Visual Control for Automated Aircraft Collision Avoidance Systems

A thesis submitted for the degree of
Doctor of Philosophy

By
Aaron Mcfadyen
B Eng (AeroAv) (1st Class Honours)

Science and Engineering Faculty
Queensland University of Technology

2015

Abstract

Unmanned Aircraft Systems (UAS) represent an important future technology, with the ability to augment or replace conventionally piloted aircraft for a range of military and commercial applications. Multiple industry sectors are rapidly embracing the technology, further diversifying the application base and broadening their operational scope. The result is an increased demand to allow unmanned aircraft regular access to unrestricted civilian airspace. Integrating unmanned aircraft in such a complex and structured environment is not trivial, and creates a set of challenging technical, regulatory and social issues that remain unresolved.

The most restrictive, and arguably most important, issue is the lack of automated See and Avoid systems for unmanned aircraft. This capability refers to the uncooperative, reactive collision avoidance conducted by pilots in response to immediate threats, both static and dynamic. It involves the pilots' visual system, recollection of regulatory procedures and pilot knowledge and skill. The encounter environment and collision avoidance approach is quite unique, so although advertised for See and Avoid, many proposed solutions do not adequately model or realistically consider the problem space. In particular range, aircraft intent information or cooperative sensor technology is typically assumed.

This thesis addresses the automated See and Avoid problem from a realistic perspective, leveraging only the information directly available from a single imaging sensor. In particular, range is not assumed or estimated, and explicit consideration is given to existing aviation standards, practice and procedures. Under such constraints, and assuming object detection, emphasis is placed on investigating how an effective vision-based decision and control strategy can be developed. The problem presents a series of key challenges and considerations. First, an avoidance decision must be made using only image features measurements, that remains consistent with aviation practice to ensure predictable behaviour. Second, a stable, effective and robust image-based control solution must be able to safely re-position the unmanned aircraft to avoid potential collision objects. Third, a resolution decision is required to cease the avoidance behaviour at an appropriate instance. Forth, the system should be practical. The associated performance limits and parameter tuning may be accomplished with appropriate simulation studies, but the system should be verified empirically. After all, if the system cannot be implemented, it is of no operational significance.

Considering the avoidance control, a set of spherical image-based controllers are derived, that require only a single point feature and a coarse range approximation for feedback control. Two general control approaches using classical partitioned image-based visual servoing and visual predictive control are developed, comparatively assessed and subsequently implemented on a small aerial platform. The effects of over and underestimating object range, added uncertainty and reference image feature placement are analysed. The visual predictive control approach is shown to better manage the See and Avoid problem constraints, whilst demonstrating sufficient robustness properties and real-time performance. The results constitute the first implementation of both schemes for the relative navigation of aerial platforms. The empirical nature of the research helps support the viability of using image-based control for real applications, and represents a practical contribution to vision-based control.

Leveraging the success of the visual predictive control approach, stability-based design approaches for general nonlinear model predictive control are then applied to the control problem. A Quasi-Infinite Horizon Visual Predictive Control (QIH-VPC) framework is used to realise two significant benefits. First, an average reduction in the control effort by 3% can be achieved compared to the standard (classical) visual predictive control approach. Second, the inclusion of the terminal region in the control structure can be used to select better reference image feature locations from a stability and feasibility perspective. To explain, the terminal region size reflects the system nonlinearity and can be used to approximate the domain of attraction for each reference image feature location. For a single point feature, the terminal region is directly mapped to an image region, providing a convenient visualisation of the system properties. The idea is general, so can be extended to other camera models, object types and image-based control applications. As determining the domain of attraction is difficult or impossible for classical image-based control, the results represent a general contribution for vision-based control.

Considering the avoidance and resolution decision strategies, a novel threshold based approach is derived for each function. The avoidance decision is based on assessing visual cues to actively re-position the object on the image surface. The specific reference position is selected by considering the properties of spiral motion, aviation right-of-way rules and the expected uncertainty on image feature measurement. The resolution decision uses an objective function to indicate an appropriate time to stop the avoidance behaviour, by coupling the platform heading and visual controller. The avoidance and resolution decisions remain decoupled, and can be tuned using two mutually exclusive decision thresholds. The complete decision strategy can be used with any imaging or bearing-only sensor, complementing existing detection and tracking algorithms, and representing a novel contribution toward See and Avoid subsystem development.

As the optimal decision threshold values cannot be determined analytically, a new probabilistic performance analysis methodology is introduced. Modified System Operating Curves (SOC) are used to simultaneously evaluate system performance and visualise the design tradeoffs. In this work, the key tradeoffs concern the decision threshold placement, but the framework allows the effects of other system parameters to be uniformly compared and visualised. Considering the attributes of a proposed system, an augmented set of collision outcomes and associated collision avoidance metrics are derived. The metrics acknowledge the fact that the system is reactive and fully automated such that action is always taken, and taking avoidance action does not imply satisfactory resolution. By remaining consistent with aviation practise, this unique probabilistic evaluation framework provides a novel tool to aid See and Avoid system analysis and development.

Combining each of the research elements and using the proposed evaluation framework, the decision thresholds can be optimised such that the See and Avoid system is over 90% effective at avoiding collisions with minimum negative outcomes (collisions). The encounter attributes that cause difficulty for the decision and control strategy are also identified. The system performance was more sensitive to the resolution decision threshold, with both liberal and conservative avoidance thresholds resulting in similar performance. Of note, the predictive control strategy remained stable regardless of the encounter type, whilst managing problem constraints.

Designing, verifying and certifying operational See and Avoid systems is a challenging task that remains in the developmental stage worldwide. The approach presented in this thesis demonstrates how vision-based decision and control for automated See and Avoid can be derived, implemented and statistically evaluated in a generalised framework. These advances constitute an important step toward certifying See and Avoid systems, and provides a good foundation in which to stem further development and research. As such, the results represent a particularly unique contribution toward the progression of automated See and Avoid systems.

Keywords: *Visual Servoing, Collision Avoidance, See and Avoid, Unmanned Aircraft, Nonlinear Model Predictive Control, System Operating Curves*

I hereby declare that this submission is my own work and to the best of my knowledge it contains no material previously published or written by another person, nor material which to a substantial extent has been accepted for the award of any other degree or diploma at QUT or any other educational institution, except where due acknowledgement is made in the thesis. Any contribution made to the research by colleagues, with whom I have worked at QUT or elsewhere, during my candidature, is fully acknowledged.

I also declare that the intellectual content of this thesis is the product of my own work, except to the extent that assistance from others in the project's design and conception or in style, presentation and linguistic expression is acknowledged.

Aaron Mcfadyen

QUT Verified Signature

28 January 2015

Acknowledgements

First, I would like to acknowledge the continual support from my supervisors and colleagues. I am grateful to Dr Luis Mejias and Prof. Peter Corke for their invaluable insights, guidance and constructive feedback throughout this research. I would also like to acknowledge the opportunities and quality research environment provided by supporting academic supervisors including Prof. Cédric Pradalier, Prof. Roland Siegwart and Prof. Duncan Campbell. Thank you also to all the students, research colleagues and administrative staff at their associated research centres including the Australian Research Centre for Aerospace Automation[1] (ARCAA) and the Autonomous Systems Lab (ASL) at Eidgenössische Technische Hochschule[2] (ETH). In particular, I would like to acknowledge the support of colleagues at ARCAA including Dr Troy Bruggemann and Duncan Greer, and researchers at ETH including Dr Stefan Richter, Dr Konstantinos Alexis, Stefan Leutenegger and Janosch Nikolic. Your involvement through inpromptu discussions and advice has helped to shape this research and improve the outcomes. I would like to extend a special thanks to Michael Burri, Ignacio Mellado-Bataller and Dr Miguel Olivares-Mendez for the time and effort spent developing the basis for the quadrotor architectures used in this research.

Second, I would like to thank the Australian Research Centre for Aerospace Automation, Queensland University of Technology[3] (QUT) and the Australian Government though the Endeavour Awards Program[4] for their financial support of this work. I am particularly grateful for the opportunity to conduct a significant portion of my research at Eidgenössische Technische Hochschule (ETH), Zurich, Switzerland.

Lastly, I would like to acknowledge my friends and family. To my friends, I am grateful for the distractions you provided, which I now see as necessary to complete such a long endeavour. To my family, John, Connie, Anna and Stephen I thank you for your constant positive encouragement and reassurance. To my girlfriend Danielle, I could not have completed this work without your support and understanding. I am very thankful for your long-term patience, kind words and acceptance of my unusually long working hours.

[1] ***Australian Research Centre for Aerospace Automation (ARCAA)*** *is a research centre of QUT located in Brisbane, Australia. Research focus is on autonomous technologies to support efficient and safe utilisation of airspace, and the development of autonomous aircraft systems for commercial applications.*

[2] ***Eidgenössische Technische Hochschule (ETH)*** *(Swiss Federal Institute of Technology) is a globally recognised tertiary education provider and research institute, located in Zürich, Switzerland.*

[3] ***Queensland University of Technology (QUT)*** *is a globally recognised tertiary education provider and research institute, located in Brisbane, Queensland, Australia.*

[4] ***Endeavour Scholarships and Fellowships*** *are internationally competitive, merit-based scholarships provided by the Australian Government that support citizens of the Asia-Pacific, the Middle East, Europe and the Americas to undertake study, research and professional development programs in Australia and for Australians to undertake these programs overseas. The Endeavour Scholarships and Fellowships program forms part of the Australia Awards initiative established in 2009, which brings together under one banner scholarships offered by the Department of Education (the Department), Department of Foreign Affairs and Trade and the Australian Centre for International Agricultural Research. The program is currently administered by the Australian Education International (AEI), Australian Government, Canberra, Australia.*

To Danielle, John and Connie.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Unmanned Aircraft Systems

An unmanned aircraft can be described as an aircraft flown without a pilot-in-command onboard. The aircraft is either remotely and fully controlled or programmed and fully autonomous. An Unmanned Aircraft System (UAS) can be described as an unmanned aircraft and all supporting components and crew required for its operation. This may include single or multiple command, control and support stations geographically dispersed from the unmanned aircraft and situated on the ground, airborne or in space [1]. Existing infrastructure such as navigational aids (global positioning satellites, terrestrial nav-aids etc.) and air traffic control surveillance equipment (radars, instrument landing systems etc.) are not considered as part of the unmanned aircraft system.

1.1.1 Benefits & Potential

The fundamental reason for developing any unmanned aircraft systems is to remove the pilot from the aircraft in order to realise a number of key operational benefits. First, the unmanned aircraft can stay airborne for longer hours (days) and travel greater distances, without inducing pilot fatigue. Second, the unmanned aircraft may be operated in close proximity or within environments that are either inhospitable or potentially harmful, without the associated human risk. Third, some platforms may be operated at considerably low cost and human resource. These benefits make unmanned aircraft an attractive alternative to conventionally piloted or manned aircraft (CPA) for typically dull, dirty and dangerous tasks. They also open up the possibility of new tasks, impossible with manned aircraft. As such, unmanned aircraft systems have since been exploited for a diverse set of existing and novel military and civilian applications.

For civilian applications, unmanned aircraft are typically viewed as a useful tool for advanced surveillance and data collection [2]. Recent developments have seen smarter platforms emerge, capable of interacting with the environment or delivering increasingly diverse payloads [3, 4]. This has seen government departments (public service) and private industry consider unmanned aircraft for a range of tasks within search and rescue, border protection, infrastructure monitoring, agriculture, law enforcement and disaster monitoring. Over the short term, precision agriculture and biosecurity applications are expected to dominate [5, 6]. Current examples including weed identification and spraying [7]-[9], crop and livestock monitoring [10, 11] and disease prevention [12]. Public services

and infrastructure inspection are expected to follow, with current examples including fire monitoring [13] and the inspection of power lines [14, 15]. In the longer term, it is anticipated that unmanned aircraft will also be used for long-haul cargo and commercial passenger transport [16].

This depth and diversity of applications are expected to provide significant benefits to the community and across multiple industries. As an example, consider the frequent collection of imagery over broad-acre commercial crops. The data can be used to better manage land use, pesticide application and general health as well as predict optimal harvest times. In turn, there is strong potential to increase productivity and improve profit margins. The economical and environmental benefits are clear and similar examples can be formulated for other natural resources and man-made assets. Alternately, a parallel argument can be extrapolated for the social benefits inherent for other applications such as law enforcement and disaster monitoring.

The perceived economical, social and environment benefits across multiple industries have seen an increased investment in research and development of commercial platforms. As a result, advanced low-cost systems are becoming available and the business case is becoming increasingly viable for potential commercial operators. This trend is evidenced by multiple certified operators currently offering UAS-based services [17]. As the adoption of this new technology continues to spread into new applications, and public acceptance of unmanned aircraft gains traction, more business opportunities and significant industry growth is expected. Current projections suggest the global UAS market value will almost double in the next 10 years[1], with an estimated worth of $89 billion (USD) [18]. This figure includes both procurement and research and development activities, and accounts for the expected relative growth in various industry sectors mentioned previously. Unmanned Aircraft Systems therefore represent an important future technology, with real potential to positively impact society.

1.1.2 Issues & Challenges

Despite the abundance of applications and associated benefits, unmanned aircraft are not without their drawbacks. A number of technological, social and regulatory barriers must be addressed to ensure unmanned aircraft not only remain useful, but can be frequently and freely operated in a legal and safe manner.

From an application specific perspective, the issues concern whether the intended task can actually be accomplished successfully. The task itself may be considerably challenging when using an unmanned aircraft. For example, this may be due to limitations of onboard sensors or operator workload. Additionally, the physical environment in which the unmanned aircraft must operate may be demanding. Dust, rain, visibility, temperature, terrain or adverse weather (tropical storms) may need to be accounted for and managed appropriately. As such, further research and development may be required to bring technology to a satisfactory level required for success. Many of these issues are

[1] Includes civilian and military unmanned aircraft systems and associated technology.

currently being addressed in the robotics community to ensure that unmanned aircraft remain useful.

From a regulatory perspective, a more general issue concerns whether the unmanned aircraft can operate both legally and safely in the intended environment [19, 20]. For indoor operations, rules and regulations do not exist as such flights are generally not considered as aviation. For outdoor operations, dedicated regulatory bodies[2] determine the rules and regulations for all aircraft operations within their respective airspace. For such operations, it is infeasible to suggest existing air traffic management systems are modified in light of this new airspace user. Unmanned aircraft must therefore comply with existing practice and procedure such that they operate seamlessly with other airspace users and existing infrastructure [21]. This means that UAS must operate at an equivalent level of safety (ELOS) to manned aircraft so as to not degrade the overall safety of the current air traffic management system. This requirement introduces a significant number of non-trivial issues that must be addressed. Multiple attempts have been made to enumerate each specific issue [22]-[24] and suggest possible solutions [25]-[27]. For many issues, this involves defining the equivalent level of safety itself and subsequently providing solutions to meet that level.

Consider again an unmanned aircraft conducting aerial work over agricultural land. Operation will be primarily in daylight and in remote non-populous locations. From a application perspective, this means the UAS may be required to operate Beyond Visual Line of Sight (BVLOS) within Class G (uncontrolled) and possibly class C or D (controlled) airspace [28]. From a regulatory perspective, this means the UAS must conform to Visual Flight Rules (VFR), respect Visual Meteorological Conditions (VMC) and maintain continuous radio communication (VHF Air-Band) [29]-[33]. The unmanned aircraft should also operate in a manner that is consistent and predictable to other airspace users and air traffic controller expectations [34, 35]. A number of complex navigation, communication, surveillance and information management tasks must therefore be replicated to at least the equivalent level of safety as manned aircraft.

Most UAS lack sufficient capability to adequately perform all of the key functions required by the regulator, and are therefore unable to operate freely within the airspace. Instead, potential UAS operators must apply for an operators certificate[3] to gain access to civilian airspace [29, 30]. This process can be rigorous, costly and time consuming with significant restrictions on where, when and how the UAS can be operated [36]-[38]. As such, the full potential of unmanned aircraft cannot be realised until appropriately certified solutions addressing strict regulatory requirements are derived.

[2]International Civil Aviation Organisation (ICAO), Federal Aviation Administration (FAA), Civil Aviation Safety Authority (CASA) and EUROCONTROL (Europe).

[3]Certificate of Authorization (United States), Unmanned Operators Certificate (Australia) or experimental certificate.

1.1.3 See and Avoid

Collision avoidance and separation assurance is a multi-layered process at the core of aviation safety. One of the key collision avoidance layers that is consistently identified as a major roadblock to civil UAS integration, is the absence of See and Avoid (SAA) capability. It is considered the last line of defence against a mid-air collision once all auxiliary layers of the collision avoidance process have failed. In short, it is a form of decentralized short term collision avoidance in which the pilot must independently identify and avoid any unplanned hazard, be it static or dynamic [39]. It involves the pilots visual system, recollection of regulatory procedures and pilot knowledge and skill. It does not rely on existing infrastructure (radar etc.), onboard surveillance equipment (TCAS, ACAS, ADS-B etc.) or air traffic services. The SAA function is thus a particular type of collision avoidance constrained by pilot ability and behaviour.

The See and Avoid task itself is typically dissected into a subset of functions including Detect, Decide and Act (DAA) or Observe, Orient, Decide and Act (OODA) [40, 41]. Taking a more generic approach, the system components can be positioned in the common collision avoidance system framework namely; Detection (Detect or Observe and Orient), Avoidance (Decide and Act) and Resolution. Detection involves the visual acquisition of a potential collision threat. Avoidance involves the decision of how to act in response to the threat and the implementation of that action. This may mean the alteration of the aircraft path or indeed no action at all. Resolution, in this work, denotes when the collision can be considered over or resolved and the aircraft is free to cease the avoidance behaviour and return to its original path. The alignment of proposed See and Avoid architectures with the traditional collision avoidance framework is depicted in Fig.1.1.

Recalling that UAS must be able to demonstrate an equivalent level of safety[4] to that of manned aircraft, the same applies to SAA capability. This means that either a pilot remains in the loop or the system performs this task autonomously. Considering that a pilot's ability to adequately See and Avoid has also been deemed questionable [42]-[44], research has focused on how an automated system may be able to aid, augment or replace the pilot completely. This presents its own set of challenging problems [45]-[48].

Defining the exact technical design and performance requirements for SAA systems remains as challenging as developing the systems themselves, and they continue to be refined simultaneously [49]-[52]. In an attempt to classify SAA systems, two criteria are commonly used based on the sensor technology employed [53, 55]. Cooperative systems rely on two-way communication between aircraft or ground equipment, where as uncooperative technology is self reliant. The later is typically divided into active and passive methods. Active systems broadcast and receive return signals where as passive systems only receive signals from the environment. Active radar based solutions are attractive, but are typically either too heavy or too expensive for many small unmanned aircraft operations [56]. Recent solutions have started to overcome such issues, but at the expense of very small

[4]Note, the equivalent level of safety (ELOS) regarding See and Avoid is in itself difficult to quantify.

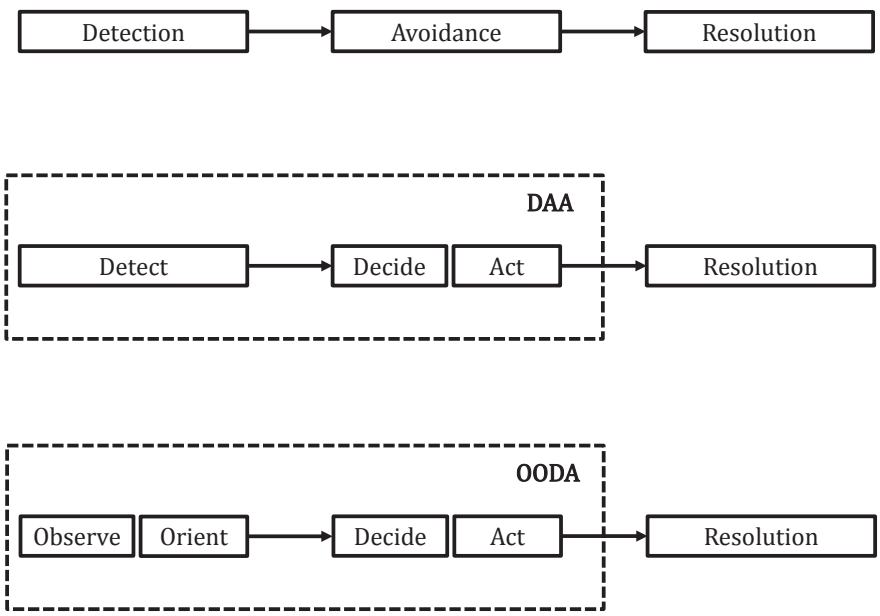

Figure 1.1: Collision avoidance system components including Detection, Avoidance and Resolution with respect to common See and Avoid architectures. These include Detect, Decide and Act (DAA) and Observe, Orient, Decide and Act frameworks (OODA).

detection ranges[5] [57]. Alternatively, and in a more realistic attempt to replicate pilot performance, passive sensors such as onboard video cameras are a more natural choice for target detection [58]. These electro-optic devices are often small, lightweight and low-cost. This means that despite typical size, weight and power (SWaP) constraints on some UAS, a collision avoidance system leveraging computer vision may be used on both small and large platforms. For larger platforms, cameras could augment additional sensors such as radar (or laser scanners) in an attempt to improve object detection (sensor fusion) [59]. For smaller platforms however, the additional sensor remains restricted by the platform constraints. In any case, the benefits of using computer vision may also extend to manned aircraft. The technology could augment existing systems by directing the pilots attention to potential collision threats earlier than if visually scanning [60]. Visual sensors may also be useful for other key safety functions including forced landing [61].

Autonomous computer vision-based See and Avoid systems can therefore be considered a key enabling technology for a broad class of unmanned aircraft. The problem is significantly complex and multiple challenges still remain in the design and implementation of the individual subsystems for detection, avoidance and resolution [62, 63]. Careful consideration must be given to any inherent limitations of visual sensing, regulatory constraints and existing aviation practise and procedure. This is key to ensuring any certified system is able to maintain strict aviation safety standards, and provide an adequate level of predictability for other airspace users and service providers.

[5]Detection ranges between 100m-300m have been demonstrated, which means a near mid-air collision may have already occurred.

1.2 Research Program

1.2.1 Outline

The proposed research program is motivated by the pending See and Avoid problem for unmanned aircraft. In search of a generic approach, efforts are directed at technology based solutions for specific components of an autonomous See and Avoid system using only onboard visual sensing. In particular, object detection and tracking is assumed and the focus of the work is on vision-based control and decision making (or logic) for the avoidance and resolution functions. This includes the **avoidance decision** (Decide) determining where to position the unmanned aircraft to avoid collision, the **avoidance control** (Act) to actual maneuver the aircraft and the **resolution decision** (Resolution) to stop the avoidance behaviour. A representation of the proposed research areas with respect to the generic collision avoidance and See and Avoid frameworks is depicted in Fig 1.2.

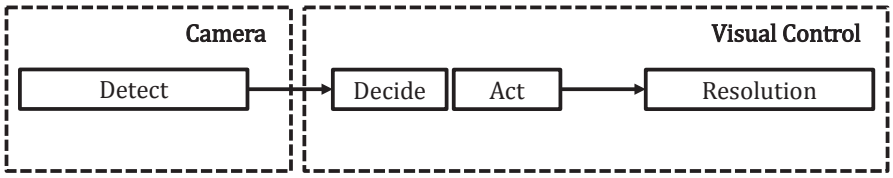

Figure 1.2: Graphical representation of the proposed research areas (*Visual Control*) with respect to the generic See and Avoid system components (*Detect, Decide, Act*). Visual object detection and tracking using computer vision is assumed (*Camera*).

Given the explicit consideration for the See and Avoid problem and the importance to address the associated issues from a realistic perspective, emphasis is placed on four key elements unique to this collision avoidance research. These include:

- The explicit consideration of visual sensing constraints inherent within the See and Avoid environment. Often these limitations are neglected or simplified such that the resulting solution is impractical in some aspects. This work uses realistic assumptions regarding the available state information obtained from computer vision in typical See and Avoid encounters.
- The explicit consideration of operational and regulatory constraints pertaining to the See and Avoid function. Often existing aviation practices and procedures are ignored in place of geometrically optimal or coordinated solutions regarding the collision avoidance and resolution decision. Although optimality is desirable, predictability is more important in such an uncooperative and shared environment relying on users to conform to regulations and advisories. This work uses existing aviation standards, practices and procedures for system design.
- The use of existing performance metrics and evaluation techniques used in the aviation community for automated collision avoidance (or hazard alerting) systems. Often performance evaluation methods (and metrics) are inconsistent with existing approaches. For example, measures of miss distance may be analysed in a subset of encounters, which may be insufficient to truly ascertain the expected performance level relative to

existing approaches. In this work, performance is evaluated using existing metrics and techniques within the aviation community, or extrapolated and extended to account for the unique structure of automated vision-based collision avoidance systems.

- The implementation of an automated closed-loop system in a proof-of-concept approach. Due to strict regularity standards, any system will require extensive testing for subsequent certification. This is difficult considering the economical and safety issues surrounding the implementation of collision avoidance systems in the real environment. This work implements a completely automated closed-loop collision avoidance system in a scaled environment to ensure safety and repeatability whilst reducing costs.

1.2.2 Research Questions

Given the proposed research areas and specific considerations, a principle research question can be formulated to fully encompass the research investigation.

Principle Question How can an automated vision-based decision and control strategy for collision avoidance and resolution, using only the available state information from monocular computer vision observations, be derived that remains consistent with existing aviation practise and procedure in See and Avoid encounters? How effective is such an approach and how can it be implemented using existing hardware systems and technology?

In order to fully address this question and satisfy the overall research objectives, the principle research question can be decomposed into three refined questions. Each question helps to frame the research by directing the investigation around key underlying concepts and issues to be addressed. The partitioning also helps to manage the research agenda.

Question 1 How can the state information obtained from computer vision alone be used directly in decision and control for collision avoidance and resolution of both static and dynamic objects?

The first component of the proposed research addresses the **avoidance decision** and **resolution decision** functions. First, the relative state information available using computer vision in the See and Avoid environment must be investigated. This includes identifying any temporal and spatial information, and associated uncertainty, that can be reliably obtained using existing object detection and tracking algorithms. Second, decision and control strategies consistent with aviation practise and standards must be investigated. This includes determining common criteria used for making autonomous avoidance and resolution decisions in existing and proposed systems. This also includes an investigation into human navigation concepts and pilot behaviour relevant to the See and Avoid environment. Third, the information must be brought together to investigate how to best design and construct a realistic and usable collision avoidance and resolution decision framework. This may require the introduction of new concepts that can appropriately align with the See and Avoid process, whilst managing any limitations and constraints on

visually acquired state information. Importantly, the investigation seeks to find a generic approach to for autonomous vision-based decision making for avoidance and resolution.

Question 2 How can visual-based control be used to safely re-position the unmanned aircraft (avoidance manuever) once an avoidance decision has been made?

The second component of the proposed research addresses the **avoidance control** function. It involves an investigation into how vision-based control (visual servoing) can be used directly to control the motion of the unmanned aircraft having acquired the object and determined the avoidance maneuver. Considerations must be given to object visibility issues, platform constraints, state information limitations and stability. The emphasis of this research component is both theoretical and practical in nature, seeking to extend visual-servoing concepts through simulation and in proof-of-concept implementations.

Question 3 What are the performance limitations of the proposed collision avoidance and resolution system, and what impact does this have on using such a concept as a general framework for designing automated vision-based systems?

The third component of the proposed research addresses the **performance** of the closed-loop collision avoidance and resolution system. First, existing metrics and performance evaluation techniques used in aviation to assess automated collision avoidance systems must be investigated. This includes identifying possible collision avoidance outcomes using automated vision-based systems. This also includes investigating how a generic **performance evaluation framework** can be derived such that other vision-based collision avoidance systems can be analysed using a common tool. Importantly, the investigation seeks to quantify the performance limitations of the proposed vision-based collision avoidance system using a common or extensible evaluation framework to enable future comparisons.

1.2.3 Scope

The scope of this research is bounded through the explicit considerations of the expected See and Avoid operational environment and the safety and economical limitations imposed on full scale implementation. This research considers:

- Pair-wise encounters involving only the unmanned aircraft and a single collision object.
- Short-term encounters with small time horizons in the order of seconds.
- Three-Dimensional (spatial) encounters in which the involved aircraft may manuever in the horizontal (lateral and longitudinal) and vertical planes.
- Daylight operations conducted in uncontrolled (Class G) airspace, in visual meteorological conditions (VMC) and under Visual Flight Rules (VFR). Collectively, this restricts operations to under 10,000', clear of cloud (weather) and full visibility ($\geq$10km). Many civilian unmanned aircraft operations, particular in agriculture, will be operated under these environmental considerations and related regulatory constraints.

Assumptions made in this research are either implicitly justified through the scope of the research, supported by relevant literature or necessary to support the research focus areas. Key assumptions include:

- The unmanned aircraft are small quadrotor (rotary wing) platforms less than 8kg. This assumption is not restrictive as the proposed research can be extended to larger platforms and/or fixed wing aircraft by altering the platform dynamics (dynamic model).
- The unmanned aircraft has adequate sensors available to determine its own state in terms of position, velocity and orientation.
- The unmanned aircraft is operating in the en-route or mission flight phases such that adequate way point navigation is available and active. Such navigation is disengaged upon object detection.
- The unmanned aircraft has only a single electro-optic sensor (video camera) to observe the conflicting aircraft or object. The camera field of view approximates a full spherical surface (spherical camera model) for simulation and a discrete section of the spherical surface for experimentation. This allows a perspective camera to be used to approximate a spherical camera section.
- Object motion is either static or dynamic. Dynamic objects are restricted to constant velocity ($\leq$250kts) and constant heading. This is a valid assumption when considering short-term encounters, associated flight rules, airspace class and existing air traffic models [64]-[67]
- Object detection and tracking is assumed in simulation and simplified in experimentation. Existing computer vision tools and algorithms [73] for detection and tracking are modified to realistically model the appearance of the object in typical See and Avoid encounters.
- Assuming initial object detection, in the event that the vision-based tracking fails, it is assumed that the See and Avoid system is suspended (disabled) and an alternate navigation system assumes control. Failure may be due to missed frames, occlusion or noise for example. If object tracking is recovered, the See and Avoid system may resume.
- Simulation studies use stochastic collision and non-collision encounter trajectories based on nominal flight paths and full nonlinear dynamic models of the unmanned aircraft. Uncertainty is included through wind (turbulence) disturbance models, actuator imperfections and sensor noise.
- Experimental studies use collision encounters of real static objects or virtual dynamic objects. Virtual objects are created by assuming a known position then approximating the object position in the image with added noise (uncertainty).
- Experimental studies are conducted indoors as a proof-of -concept approach. A motion capture system (Vicon T40) is used for ground truth and not for vision-based decision and control. This ensures a more realistic hardware/sensor arrangement as the real environment whilst circumventing regulatory and cost issues.

1.2.4 Contributions & Significance

The main contribution of this research is the further progression of automated vision-based collision avoidance and resolution systems. Consequently, a number of significant theoretical and practical contributions advancing the vision-based control (visual servoing) and aviation fields are presented. Specifically, the research makes a total of five main contributions as described below.

- **Contribution 1** An extension to Spherical Image-Based Control (S-IBVS). This includes a new partitioned control scheme for effectively tracking spirals that is verified in simulation and empirically. The work constitutes the first successful practical implementation of S-IBVS on an aerial vehicle. It is found that the controller is capable of satisfactory spiral tracking for collision avoidance.
- **Contribution 2** A novel derivation of Spherical Visual Predictive Control (S-VPC). This includes a new partitioned scheme for tracking spirals that is verified in simulation and empirically. The work constitutes not only the first successful practical implementation of S-VPC, but the only existing practical implementation of image-based Visual Predictive Control (VPC) of any form for aerial vehicles. It is found that the controller is capable of improved spiral tracking for collision avoidance compared to S-IBVS schemes, whilst remaining computationally tractable. Other benefits include better management of platform dynamics and problem constraints. Qualitatively, the approach shows improved robustness to parameter uncertainty and image processing delay, important for vision-based control.
- **Contribution 3** The first application of Quasi-Infinite Horizon Nonlinear Model Predictive Control (QIH-NMPC) to visual predictive control (QIH-VPC). This includes the subsequent identification of its benefits to spiral tracking and collision avoidance, apart from the inherent nominal stability guarantees. It is found that the control design implicitly enables better selection of reference image features for spiral tracking and collision avoidance, and decreases the control effort by 3% compared to visual predictive control.
- **Contribution 4** Novel vision-based decision strategies for collision avoidance and resolution. The work constitutes significantly unique yet very realistic approaches to decision making for See and Avoid. The leading approach combines spherical imaging and associated uncertainty, the properties of conical spiral motion and common aviation practice and procedure (rules of the air [36]). The avoidance and resolution decisions are de-coupled and rely on tuning two mutually exclusive thresholds. Through simulation, the decision strategy is shown to be over 90% effective at avoiding collisions.
- **Contribution 5** A novel performance evaluation framework suitable to vision-based collision avoidance systems, and strongly aligned to existing aviation methods. This includes the identification of additional collision avoidance outcomes for completely autonomous systems and the subsequent adaptation of System Operating Curves (SOC) in the context of vision-based systems. It is found that the performance evaluation framework may be used as a standard tool for refining existing systems and comparing between different systems.

The significance and potential impact of this research is best exemplified through the consideration of every contribution and the way in which they relate to each other.

From a visual control perspective, contributions 1-3 provide empirically evidence to support the use of image-based visual control for unmanned aircraft. This increases the diversity of potential control approaches and design choices for a range of complex navigation tasks. Although applied to the collision avoidance task, the QIH-VPC framework can be applied to other camera models, vehicles and applications. The QIH-VPC design has the potential to decrease the control effort, which is an important practical consideration given current battery limitations. The QIH-VPC control parameters can also be used to guide selection of feasible reference values. This has implications for both set point tracking (fixed and variable) and region reaching control.

From an aviation perspective, contributions 1-4 demonstrate a working autonomous vision-based See and Avoid system. Although the experimental set-up is a scaled down version of the real environment, the associated assumptions provide a reasonable approximation to the real situation. As such, the work provides ample evidence to support the use of vision-based See and Avoid systems. This is significant, as vision-based systems provide arguably the most realistic approach to See and Avoid. They can also be used to complement existing systems onboard manned aircraft or provide a level or redundancy.

From a regulatory perspective, contributions 3-5 constitute a significant step toward certification of See and Avoid systems. Certified systems do not currently exist, but are expected to be based on existing performance evaluation techniques currently used in aviation. A performance metric that is most likely to be required is a verified estimate of the probability of collision when using the See and Avoid system. This can be derived from a performance evaluation framework using System Operating Curves. This work uses such metrics derived from extensions to the System Operating Curve approach. The evaluation framework is then applied to the proposed See and Avoid system. As such, a direct attempt is made to align with the aviation community and work with the existing regulatory structure and requirements. As this is vital to ensuring See and Avoid systems can be certified, the framework may serve as a preliminary guide to evaluating vision-based See and Avoid system. It may also be used to help inform regulatory policy.

The See and Avoid problem for unmanned aircraft is considered a major roadblock to unmanned integration into civilian airspace. In the absence of this key enabling capability, the full potential of unmanned aircraft is currently unrealised. This work represents a significant effort to contribute toward such a difficult global problem on multiple levels.

Research Tools

- **Simulation Models:** A number of researchers have simulated complex encounter models based on real (radar) data to verify their collision avoidance concepts [64]-[67]. Such data is not freely available and cannot be obtained in the scope of this research. As such, Monte Carlo Simulations are used, in part, to conduct the performance analysis to validate the proposed See and Avoid approach. This is commonly used in collision avoidance evaluation prior to implementation [68, 69].
- **Software Tools:** MATLAB [70], C/C++ [71] programming languages and software tools are used for simulation development and implementation. The Robot Operating System (ROS) [72], OpenCV [73] and ACADO Optimal Control Toolkit [347] are used for practical implementation.

Publications

The following publications have been produced during this research:

- A. Mcfadyen, P. Corke and L. Mejias, "Visual predictive control of spiral motion," *IEEE Trans. Robotics*, vol. 30, no. 6, pp. 1441 - 1454, 2014
- A. Mcfadyen, L. Mejias P. Corke and C. Pradalier, "Aircraft collision avoidance using spherical visual predictive control and single point features," *Proc. IEEE/RSJ Int. Conf. Robotics and Intelligent Systems (IROS'13)*, pp. 50 - 56, 3 - 7 Nov. 2013
- A. Mcfadyen, P. Corke and L. Mejias,"Rotorcraft collision avoidance using spherical image-based visual servoing and single point features," *Proc. IEEE/RSJ Int. Conf. Robotics and Intelligent Systems (IROS'12)*, pp. 1199 - 1205 , 7 - 12 Oct. 2012
- A. Mcfadyen, A. Durand-Petiteville and L. Mejias, "Decision strategies for automated visual collision avoidance," *Proc. Int. Conf. Unmanned Aircraft Systems (ICUAS'14)*, pp. 715-725, May 2014
- A. Mcfadyen and L. Mejias,"Visual servoing approach to collision avoidance for aircraft," *Proc. Int. Congress of the Aeronautical Sciences (ICAS'12)*, Sep. 2012

1.2.5 Thesis Structure

Chapter 1 has since provided an introduction to the research area and an outline of the specific research to be conducted. The remaining sections in this chapter provide a comprehensive background to the See and Avoid problem. This includes a literature review of the associated visual sensing considerations and existing collision avoidance approaches. Focused literature reviews regarding specific vision-based control techniques are reserved for Chapters 4 and 5 respectively.

Chapter 2 provides preliminaries regarding platform dynamics and control, as well as the camera geometry and image kinematics used throughout this work.

Chapter 3 outlines the development of a collision avoidance and resolution decision strategy based on conical spiral motion, current rules of the air and the expected uncertainty on image feature measurements.

Chapter 4 presents a collision avoidance approach using Spherical Image-Based Visual Servoing (S-IBVS) and a simple avoidance and resolution decision strategy. Simulated and preliminary experimental results are presented using a small *Parrot - ARDrone*.

Chapter 5 presents a collision avoidance approach using Spherical Visual Predictive Control (S-VPC) and an advanced avoidance and resolution decision strategy. Novel Quasi-Infinite Horizon Visual Predictive Control (QIH-VPC) is also presented and used to refine the avoidance decision strategy and improve control performance. Simulated and preliminary experimental results are presented using a small *Ascending Technologies - HummingBird*.

Chapter 6 provides a probabilistic performance assessment of the proposed vision-based collision avoidance and resolution strategies using spherical visual predictive control. Extensions to collision avoidance outcomes, performance metrics and analysis techniques are also presented. This includes a novel framework to evaluate vision-based collision avoidance systems, aligned with existing aviation practise.

Chapter 7 provides conclusions regarding the research findings in the context of the proposed research questions. Recommendations and suggestions for further work are also presented.

1.3 Background

1.3.1 Visual Sensing

In the context of See and Avoid, visual sensing involves the detection and tracking of static and dynamic airborne objects that pose a potential collision threat to the unmanned aircraft. Estimation is an attempt to determine more information about how the object is moving. This section provides an important background on the detection, tracking and estimation considerations for vision-based See and Avoid systems. The review aims to provide relevant information regarding the limitations and constraints that visual sensing imposes on the avoidance and resolution functions for automated See and Avoid systems.

Detection & Tracking

The primary concern for onboard detection and tracking is understanding how an aircraft or object initially appears in the image an subsequently behaves over time. Specifically, the objects' spatial and temporal attributes when projected onto an imaging surface. In the See and Avoid environment, an aircraft initially appears as a small, low contrast point feature in the image that lacks distinguishable shape or color variation [39]. For a collision or near collision encounter, the object will remain small and relatively stationary in the image until the final moments before collision. At this point, the object expansion rate increases and the object starts to fill the camera field of view. In a non-collision encounter, the aircraft will show greater movement on the imaging surface, but will still remain relatively small. As an example, a set of grey scale images taken from an onboard camera observing another aircraft in a near collision encounter are shown in Fig 1.3.

The reason for the unique appearance of aircraft is due to a number of related factors. These include the cameras intrinsic and physical parameters, the exact environmental conditions and the object characteristics [75]-[77]. The details of how these factors influence object appearance, and subsequently effect detection and tracking systems, are described below with the aid of Fig 1.4.

Camera Parameters: First, an object can only be detected if it lies within the cameras field of view (FOV). This is defined by the maximum visible azimuth and elevation angles measured from the camera focal point or centre of projection. As the aircraft maneuvers the field of view is rotated, covering a different region of the airspace. If the field of view is small, then the object could be lost. This is significant in terms of collision avoidance, as objects outside the field of view may still pose a collision threat [79]. The expected field of view requirement for See and Avoid systems has been quoted at $\pm 110^\circ$ in azimuth and $\pm 15^\circ$ in elevation [39, 46, 51, 52, 55] and since verified in simulation [80]. As such, a single perspective camera may not be suitable. Instead, an array of cameras [81], wide angle cameras or a spherical camera may be required [82, 83]. Second, an aircraft will only become visible to the camera once it occupies at least one pixel on the imaging surface. Therefore, the cameras resolution and pixel size (spatial resolution) will significantly effect detection distance. For a fixed camera field of view and focal length, if the resolution is increased and the pixel sized reduced, then more distant objects can be detected. Each

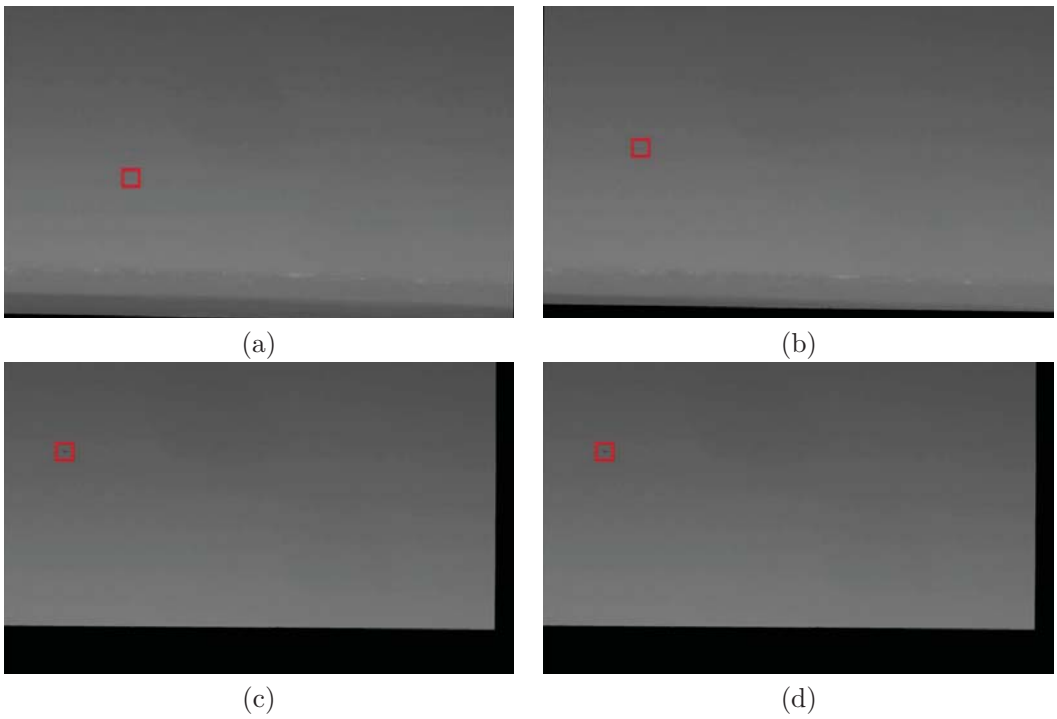

Figure 1.3: Example images taken at 15Hz from an unmanned aircraft (Silvertone Flamingo) in a near-miss scenario with a general aviation aircraft (Cessna 172R). The intruder can be seen as the dark dot within the red square in each case [95]. (a) $r = 4078$m, $t_{\text{cpa}} \approx 40s$ (b) $r = 2993$m, $t_{\text{cpa}} \approx 30s$ (c) $r = 1998$m, $t_{\text{cpa}} \approx 20s$ (d) $r = 1578$m, $t_{\text{cpa}} \approx 15s$

pixel occupies a smaller physical space and thus smaller or more distant objects can be detected. A similar effect occurs using a longer focal length and reduced field of view. Notably, the spatial resolution of current video cameras is often 0.1-1mrad compared to the human eye at approximately 0.3mrad.

Environmental Conditions: First, existing aircraft color choices are often conspicuous with respect to the background. This results in a low contrast projection of the aircraft. Second, environmental conditions such as cloud, weather (rain and haze), dust, terrain and background luminance (sunlight) can alter the appearance of the aircraft in the image. The consistently changing background (terrain) and background luminance (sunlight) will influence the entire image. This is strongly dependant of the camera sensitivity to illumination differences[6]. In some cases, the brightness of the corresponding aircraft pixel and those of background may be similar, resulting in a low signal to noise (SNR) ratio. As such, it may not be detected at all or only partially projected onto the imaging surface. Cluttered backgrounds including cloud, weather and terrain may partially occlude the object or result in contour interaction, and again result in partial or no distinct projection onto the image. This means an aircraft may not be detected as early or consistently as possible for a given camera type.

Object Characteristics: First, due to the quantisation levels used in digital image processing and the typically uniform aircraft color [78], the resulting projection onto the imaging surface lacks color variation. Second, there exists an implicit coupling between

[6]CMOS cameras are generally more sensitive to lighting conditions than CCD.

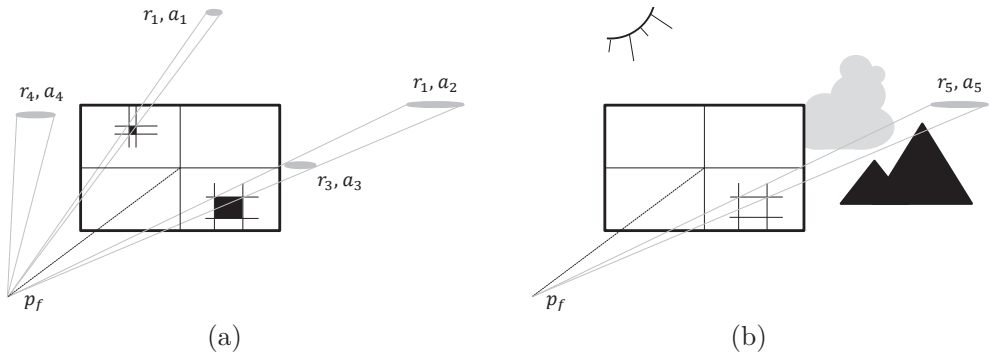

Figure 1.4: Factors affecting object detection and tracking. The objects $(a_1, a_2, a_3, a_4, a_5)$ can be seen as variable sized ellipses $(\bullet)$, positioned at different distances $(r_1, r_2, r_3, r_4, r_5)$. The objects are projected onto the imaging surface with focal point p_f. (a) Effects of field of view, spatial resolution and object size (b) Effects of sunlight (spectral reflectance and occlusion due to cloud cover and background clutter (terrain)

aircraft and pixel size that influences the detection distance. A large target may be seen further away than a small target using the same spatial resolution. In the See and Avoid environment, object scale is generally unknown. For detection and tracking, this has little impact considering all objects will initially appear pixel-sized. Third, the relative motion of the object and associated geometry affects the projected image behaviour of the object. Generally, exact knowledge of the underlying object motion is unknown. It is therefore difficult to track all object types without accurate motion models.

It should be clear that a number of complex issues need to be addressed when developing automated detection and tracking systems. As such, a large number of approaches using optical flow [84, 85], cascaded learning [86]-[88], dynamic programming [89]-[91], hidden markov models [92]-[98], particle filters [99, 100] and wavelets [101] have since been proposed. Many algorithms rely on advanced image processing techniques, incorporating both spatial and temporal filters (or filter banks) to account for the unique appearance and behaviour of the object. Often each algorithm is designed or tailored to manage a subset of particular issues, resulting in no global solution. Instead, performance is usually compared using the initial detection distance and the consistency at which correct objects can be tracked and incorrect objects rejected. Consistency is usually expressed as a probabilistic measure of correct detections, missed detections, false alarms and their corresponding rates. Using such metrics, a simple plot can be constructed to provide a rough comparison between some leading tracking algorithms[7]. This is shown in Fig 1.5.

In the context of this research, how the expected performance of existing detection and tracking algorithms may impact subsequent avoidance and resolution functions is of more importance. Improved tracking consistency provides better assurance that all collision objects can be detected. If collision objects are missed, the avoidance and resolution functions are useless. Improving the initial detection distance increases the available time for avoidance and resolution. Considering though that the object motion is generally unknown, it is difficult to determine an exact mapping from distance to time to closest

[7] A true comparison is difficult due to undisclosed performance measures, different noise characteristics, camera parameters, aircraft types and the use of real and synthetic data.

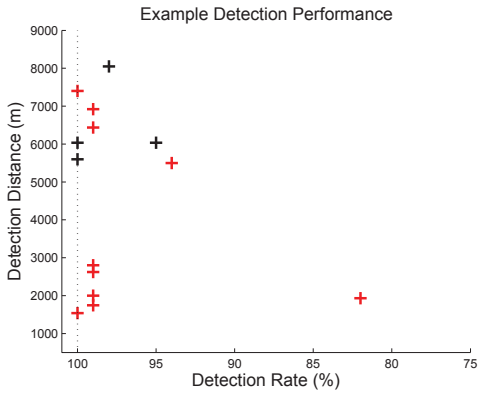

Figure 1.5: Some example detection distances and approximate detection rates for existing vision-based dim target detection algorithms using real data [84]-[106]. Results include images taken from ground (+) and airborne (+) observers

approach (or collision). An analytical approximation can be derived however using realistic assumptions on the expected See and Avoid encounter parameters [102]. This can then be used to estimate the expected time constraints imposed on subsequent avoidance and resolution functions (control).

Consider two aircraft moving with constant velocity and heading in a collision encounter. Under this assumption, common for See and Avoid, the time to collision and detection distances can be estimated. Assuming a worst case scenario, the object will approach head with a zero relative heading. Using the expected maximum and minimum allowable relative velocity[8], the relationship between detection distance and time to collision is depicted in Fig 1.6(a). A similar plot using en-route relative velocities for general aviation aircraft and unmanned aircraft is shown in Fig 1.6(b). If the object approaches from in front of the unmanned aircraft and at any non-zero relative heading, the time to collision will be to the right of the corresponding relative velocity line. The approximation suggests that current detection and tracking algorithms can provide at least 12-15 seconds for avoidance and resolution, assuming a detection distance between 2000m and 3000m. As such, existing algorithms can provide comparable or better performance than pilots in See and Avoid encounters [39, 42, 44]. Therefore, it is feasible to assume vision-based detection and tracking algorithms can be used successfully for See and Avoid systems.

Estimation

Estimation is a common problem in aerospace applications [103]. Typically aircraft motion is estimated using related observations from one or more sensors, in an appropriately designed filter. In the context of See and Avoid, estimation involves the determination of various quantities related to the relative motion of conflicting aircraft. For vision-based systems, the collection of observations used in the filter must then come from visually acquired measurements. Useful quantities for the avoidance and resolution functions include

[8]Determined by Visual Flight Rules to be approximately 100kts and less than 500kts.

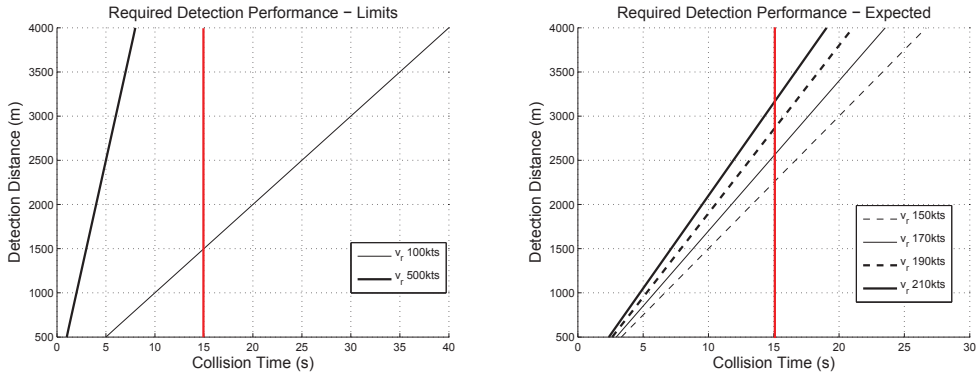

(a) Detection Range Limits

(b) Expected Detection Ranges

Figure 1.6: Example initial detection distances and available avoidance time in worst and best case (a) as well as expected (b) encounters. Worst and best case encounters correspond to a closing velocity of 500kts and 100kts respectively. The commonly accepted mean avoidance time of 15s is also shown (−)

the relative angular position (and associated rate), range (relative position and velocity) and the time to closest point of approach (or collision).

Angular Position & Velocity: Estimating relative angular position (azimuth and elevation) and angular rates can be achieved using many of the detection and tracking algorithms and filters previously discussed [84]-[98]. Inducing motion is not required to obtain accurate estimates, regardless of the encounter geometry. Stable, convergent filter behaviour can be achieved for both on line and off-line (batch) processing with acceptable computation cost [104]-[106]. Such behaviour can be attributed to the fact that the underlying process models are defined in the image space, and do not rely on specific knowledge of the relative encounter geometry. For example, the assumption the aircraft or object remains stationary in the image can be used to define relatively accurate dynamic models (or transition matrix probabilities) used in temporal filtering stages.

The accuracy of the angular estimates depends on the camera parameters, camera calibration and issues such as platform vibration. Recent algorithms [106] using real flight data suggest azimuth and elevation angles can be estimated to less than 43.6mrad (2.5°) and 3.7mrad (0.21°) RMS respectively. The azimuth and elevation angular rates can be estimated to less than 2.3mrad/s (0.13°/s) and 1.9mrad/s (0.11°/s) respectively. This is a considerable improvement on previous results suggesting RMS errors in angular position of up to 129.2mrad (7.4°) [107, 108].

Range, Position & Velocity: Estimating range, and subsequently relative position and velocity, can be achieved using stereo vision or monocular vision. However, each has significant drawbacks when realistically considering the See and Avoid environment.

Using stereo vision, two or more cameras positioned at different locations on the aircraft must observe the same object at the same time. The object position in each image must then be sufficiently different in order to triangulate position and estimate range (and velocity) [109]. Considering the small size of unmanned aircraft and the expected relative encounter geometries (scale), the approach can be considered overly optimistic for See and Avoid. The camera baseline to object range will be too small and potentially

inconsistent. The latter artefact is a result of aerodynamic (and wind) effects, and the inherent flexibility and elasticity within the aircraft structures (wings, fuselage etc.). Only recently has stereo vision been successful on small unmanned aircraft, but only at ranges up to 100m [110].

Using monocular vision, an appropriate filter must be designed that now considers the true underlying dynamics instead of those defined in the image space. As the only available observation will be angular measurements, the problem is then a bearings-only tracking or target motion analysis problem [111]. The problem has become well studied due to the unique observability [112]-[115] properties and highly nonlinear measurement process [116]. In short, range may be estimated using bearing measurements if the observer (camera) has a higher non-zero motion derivative compared to the object. This is best described as out-maneuvering the object in an attempt to acquire a unique sequence of bearing measurements [117]. Some example cases are depicted in Fig 1.7 for observable and unobservable geometries.

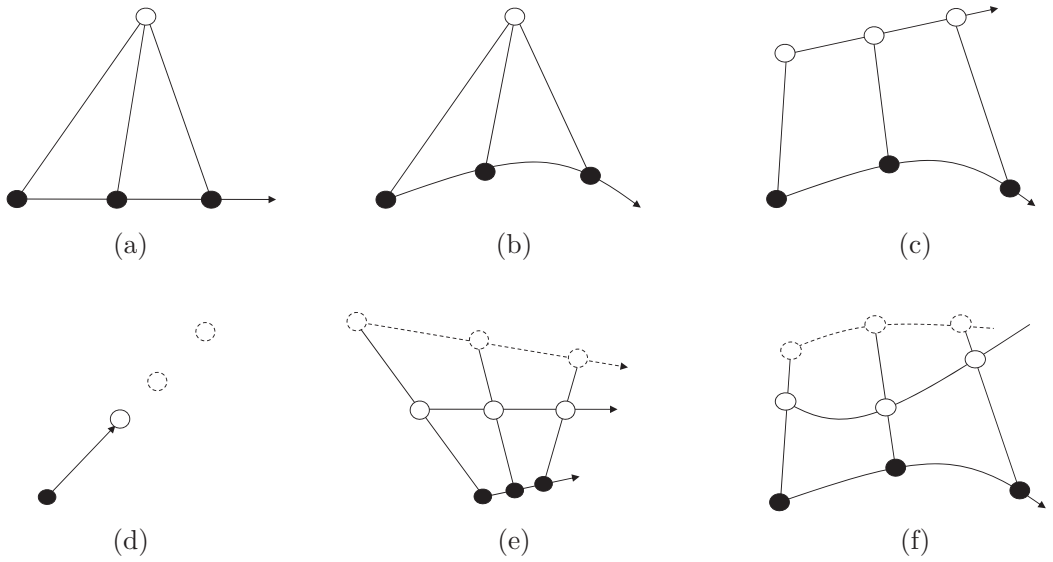

Figure 1.7: Observability examples related to bearings-only estimation. The observer (•) motion is shown with respect to the object (○) (a)-(c) Observable geometries (d)-(f) Unobservable geometries

Considering the case when both objects move with constant velocity, range is then unobservable [118, 119]. As seen in Fig 1.7(e), an infinite number of object trajectories result from the same sequence of angular observations. The observer must accelerate, effectively turning, in an attempt to acquire a unique sequence of observations as seen in Fig.1.7(c). This is commonly referred to as passive ranging, and many studies have been conducted to determine the optimal motion(s) required for full observability [120, 121]. Often the required motion cannot be realised by aircraft, so predefined maneuvers have been investigated with the use of various Extended Kalman Filters (EKF). To help address the observability issue, filters have been designed to decouple the observable and unobservable states. As a result common filters include variants of Modified Spherical

(MS-EKF) [122]-[125], Range-Parametrised (RP-EKF), [126, 127] and Log-Polar Coordinate (LPC-EKF) [127, 128]. Theoretical results are encouraging, but it remains unclear which exact maneuver to adopt given the generally unpredictable object motion. This means the encounter geometry may be unintentionally degraded and there will exist certain geometries, and noise characteristics, that render range unobservable. Despite this, some practical implementations for See and Avoid have been attempted [129, 130]. Large convergence times ($\geq$ 5s) and inaccurate range estimates (%10-30 error) were observed. Importantly, even if improvements can be made, the available avoidance time has been considerably reduced from 12-15s to 7-10s.

It is important to recall that the object motion is generally unknown, so the constant velocity dynamic models used in many of the EKF approaches above may not be representative of the actual motion [131, 132]. The result is highly uncertain (large variance) range estimates, which effects the behaviour (convergence) of the observable states and thus compromises filter stability [133]. As such, many extensions have been proposed to account for such maneuvering objects [134] without inducing predefined observer motion. Promising filter designs include methods based on multiple dynamic models [135]-[137] and particle filters [138]-[140]. However, the computational expense may not be suitable to the short encounter duration [141]. Additionally, and congruent to the constant velocity case, range observability (filter stability) is highly dependent on the encounter geometry [112, 115].

Collision Time: Estimating time to closest point of approach (or collision), without explicitly estimating relative state (range, position and velocity), can be achieved using visual looming or object expansion [142]-[144]. Time to collision is estimated, or qualitatively approximated, by processing variations in the objects expansion rate through observing the optic flow patterns created by relative motion. If a static or constant velocity object of known size is assumed, measuring the objects expansion rate can yield an accurate estimate of collision time. Alternatively, a collision can be inferred by approximating the time it will take the object to fill the camera field of view. The approach has been considered for aircraft observing large objects with multiple image feature points, or objects that are significantly close [145]-[147]. For distant or small objects (pixel sized), as in a See and Avoid encounter, the approach has not been implemented. However, the approach will likely fail, recalling that the object only demonstrates significant expansion just prior to collision.

1.3.2 Collision Avoidance Systems

In the See and Avoid context, collision avoidance systems refer to the automated technology aimed at replicating the reactive collision avoidance function of pilots. Assuming the object (aircraft, weather, terrain etc.) has been visually detected, the remaining collision avoidance functions involve the pilots decision making and control actions in response to the situation. Specifically, this includes determining if the object is a collision threat (conflict detection), deciding what (if any) avoidance maneuver to take (avoidance decision or logic), and subsequently applying that action (avoidance control). The goal being to avoid a Near Mid-Air Collision (NMAC) and safely resolve the conflict[9]. Importantly, the See and Avoid process does not include the pilots leverage of existing collision avoidance systems such as the Traffic Alert and Collision Avoidance System (TCAS I, TCAS II), Automatic Collision Avoidance System (ACAS), Air Traffic Controller (ATCO) directives or other cooperative communication devices (VHF, ADS-B etc.). This is a common point of misunderstanding even in recently proposed See and Avoid systems [148]-[151].

This section provides an important background regarding relevant conflict detection, avoidance decision and avoidance control concepts used in human navigation, manned aviation and vision-based robotics. The term concept is used to describe a generic approach as opposed to specific details. The review aims to provide relevant information regarding how such concepts may be used or extended in the design of vision-based See and Avoid systems. The discussion is restricted to aircraft and obstacle avoidance, and does not include terrain (ground) and weather avoidance systems.

Human Avoidance Systems

The human visual navigation system is a key component for effective collision avoidance [152, 153]. Given the requirement for equivalence, the human visual navigation system is an important consideration in the subsequent design of automated See and Avoid systems. Although central to the See and Avoid discussion, human collision avoidance concepts are commonly ignored or by-passed, particularly in the robotics community. This is typically under the assumption that machines (automated systems) can do better. This may be the case in some instances, but often coincides with complete negligence regarding the notion of predictability. In aviation, it is important airspace users behave in a predictable manner where possible, given the airspace is shared with manned aircraft. A system that outperforms the human see and avoid function, although useful, may degrade the overall airspace integrity (and safety). As such, human conflict detection, avoidance decision and avoidance control concepts are discussed below.

Conflict Detection: Humans can perform conflict detection using visually acquired measurements of angular position and rate. Specifically, a zero angular rate or constant angular position is used to infer a collision object [154]. Importantly, the object is assumed to be in front of the observer such that the magnitude of the relative angular position is always less than 90 degrees. The object may be stationary or moving in a

[9] A Near Mid-Air Collision (NMAC) is defined as the incursion or breach of a cylindrical protection zone of height 200' and radius 500' about each aircraft.

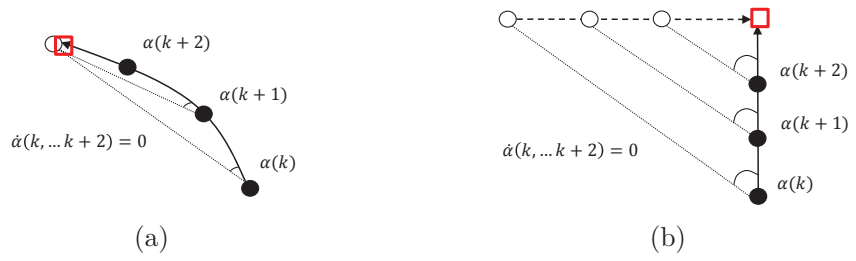

Figure 1.8: Example human conflict detection concepts between a single object (○) and observer (*bullet*) using angular (azimuth) rate $\dot{\alpha}(\cdot)$. Only collision cases (□) are shown for (a) Static Collision Object (b) Dynamic Collision Object

linear or nonlinear manner (turning) [155]. As such, the constant angle model offers a simple and useful conflict detection mechanism for objects moving in an arbitrary fashion. Additionally, the model verifies many assumptions often made in manned aviation [39]. The conflict detection concept is shown in Fig 1.8

Considering practical limitations, such as sensor noise characteristics, the observed angular rate $\dot{\alpha}$ will rarely be zero in See and Avoid encounters. It is then necessary to place a small threshold on the angular rate in an attempt to distinguish between collision and non-collision aircraft. This can be defined as

$$C(t > k) \iff \dot{\alpha}(k) \leq \epsilon, \quad \bar{C}(t > k) \iff \dot{\alpha}(k) > \epsilon \tag{1.1}$$

or, by considering some arbitrary time history n of the angular observations

$$C(t > k) \iff \dot{\alpha}(k-n, \ldots, k) \leq \epsilon, \quad \bar{C}(t > k) \iff \dot{\alpha}(k-n, \ldots, k) > \epsilon \tag{1.2}$$

where C and $\bar{C}$ denote collision and non-collision respectively, ϵ is a constant scalar threshold and k is the current time. This approach was empirically investigated in a number of See and Avoid flight trials [156] with mixed outcomes. Results showed collision objects can be correctly identified ($> 90\%$), but at the expense of multiple falsely classified non-collision objects. This suggests that additional information may be required for adequate threshold placement or enable the derivation of a continuous measure of collision threat.

Avoidance Decision: Humans demonstrate a variety of behaviours with respect to avoidance decisions [157]. This includes whether to pass in front of or behind a moving obstacle. The decision is often ambiguous and indeed difficult to derive an explicit model for the avoidance behaviour. From a See and Avoid perspective however, this ambiguity may be resolved by considering the current right-of-way rules (see Appendix A). The rules define the conditions in which to give way to nearby aircraft, and how to manage (near) head on collision encounters using lateral separation. Using vertical separation however, and except for the overtaking case, the rules-of-the air do not provide any explicit guidelines. To this end, recent pilot centric studies suggest that a decision to descend is preferable [158, 159]. Importantly, the avoidance decision is qualitative for both vertical and lateral separation. This means direction (up, down, left, right) is used

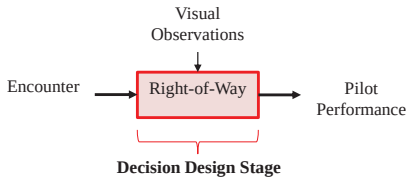

Figure 1.9: Example avoidance decision or logic for human collision avoidance systems.

and not a specific heading, velocity, altitude etc. Additionally, following the right-of-way rules suggest that precautionary avoidance maneuvers may be prevalent in manned aircraft for ambiguous collision objects.

Avoidance Control: Considering the qualitative nature of human avoidance decisions, the subsequent control uses a direct visual feedback loop. Initial models, derived from object interception (reciprocal of avoidance), suggest the object is visually steered in such a way that a non-zero angular rate is maintained [160]. This is demonstrated for both static and constant velocity objects and results in smooth control. Importantly, no specific non-zero reference position or velocity in the eye is used for control. This means explicit path planning or other waypoint based navigation is not used, and instead a reactive control strategy is employed.

Extensions to the control strategy include the consideration of small object velocity changes over the encounter [161, 162]. An intermediate anticipatory strategy was observed whereby a predictive control strategy aims to command a non-zero angular rate for a short time into the future (time horizon). Further research has showed that short prediction times (0.5-3.5s) result in more robust and accurate control [163]. This means predictive avoidance control strategies using relative angular observations may be representative of actual human avoidance behaviour.

Aircraft Avoidance Systems

Many automated collision avoidance approaches have been derived for manned aircraft [164, 165], including the well known Traffic Alert and Collision Avoidance System (TCAS) [166]. However, most of the methods are aimed at planned separation assurance and not reactive collision avoidance, as required for See and Avoid. The time at which the conflict is detected and resolved differs such that separation assurance occurs prior (minutes) to See and Avoid (seconds) [167]. Nonetheless, it is important to consider existing separation assurance concepts from a See and Avoid perspective. This is for a number of reasons. First, the basic separation assurance functions are very similar to those required for See and Avoid. After all, they are both essentially aimed at separating aircraft. This means some principles and concepts may be adapted or scaled to fit the See and Avoid constraints. Second, many existing concepts were derived with a particular focus on subsequent certification. Given See and Avoid systems will also require certification, it may be wiser to adopt similar concepts and better align with regulatory expectations and standards. As such, existing aircraft conflict detection, avoidance decision and avoidance control concepts are discussed below.

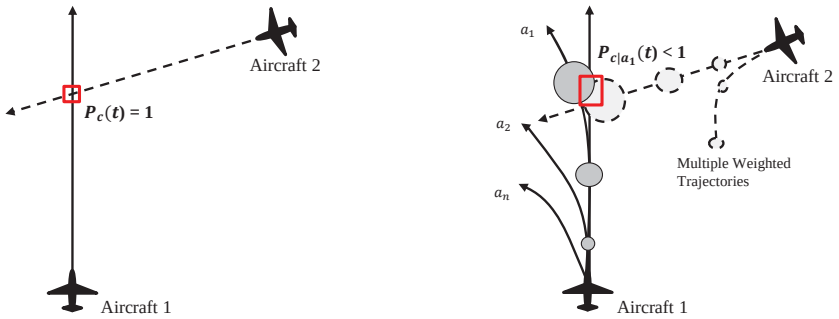

Figure 1.10: Example conflict detection concepts using state propagation to estimate conflict probability $P_c(t)$. Uncertainty in trajectories (•) and collision cases (□) are shown for (a) Nominal state propagation without avoidance maneuvers where $P_c(t) = 1$ (b) Probabilistic state propagation considering multiple avoidance maneuvers where $0 > P_{c|a_1}(t) < 1$

Conflict Detection: Conflict detection is primary accomplished by either monitoring the current relative state (position, range) [168], or by predicting the future relative state (state propagation) and acquiring an appropriate estimate of the probability of collision [169]. Although computational restrictive, state propagation is generally preferred, providing a relative measure of the probability of collision $P_c(t)$ by considering future events. State propagation can be conducted in a nominal [170] or probabilistic manner [171, 172] with respect to the predicted trajectories, and may consider one or more avoidance manuevers. As such, multiple collision probability estimates given an action a_n was taken can be obtained such that $P_c(t) \in \{P_{c|a_0}(t), \ldots, P_{c|a_n}(t)\}$. Multiple definitions for conflict probability exist including maximal [173], integrated [174, 175] and probability flow [176], along with multiple analytical [177], numerical [178]-[180] and probabilistic (Monte-Carlo) [181] approaches to approximate or calculate them. Regardless, conflict detection is then based on comparing the collision probability estimates. Considering only the current state or using nominal state propagation, then $P_c(t) = \{0, 1\}$. Considering probabilistic state propagation, then $0 \leq P_c(t), \leq 1$. Some simple examples of the differences in these conflict detection concepts are depicted in Fig 1.10.

Applying conflict probability based approaches to vision-based systems in not straight forward, often leveraging additional sensors and filters (fusion) [182]. For vision-only systems that estimate relative state, the approach can be applied directly. Recent results demonstrate that the conflict of probability estimate is highly dependent on the filter performance. Additionally, calculating the conflict probability can take up to 3s, reducing the available avoidance time [183]. For vision-based system that do not estimate relative state, state propagation cannot be performed and the direct application of conflict probability approaches is not feasible. Attempting to map the approach to the image space will be considerably difficult. Essentially, a meaningless estimate for the conflict probability may be obtained. This is due to the generally unknown motion of the aircraft, sensitivity to state uncertainty, and the fact that multiple image positions (or the entire image) could have an equally likely collision probability. Instead, by considering human

navigation models directly, the amount of time the object remains stationary in the image has been suggested as an alternate estimate of collision probability [184] such that

$$P_c(k) = \frac{1}{1 + \frac{1}{k-1}} \sum_{i=k-N}^{k-1} (\alpha(k) - \alpha(i))^2 \tag{1.3}$$

where $\alpha(\cdot)$ is the relative azimuth, k is the current time instance and N is the number of past observations considered. However, this has shown marginal performance in simulation on a select set (head-on, overtaking) of encounter geometries. It is also unclear how to select N, and although robustness to measurement uncertainty is claimed, no analysis or evidence is provided.

Avoidance Decision: Traditionally, many aircraft systems, including the certified Traffic Alert and Collision Avoidance System (TCAS), use discrete logic to make avoidance decisions [164]. The logic is derived by placing one or more thresholds on the collision probability estimates in a tree-like (nested) decision structure[10]. Multiple prescribed avoidance actions are generally considered in the collision probability estimates, but often refined to a limited set of vertical climb, descend and level-off maneuvers due to aircraft performance limitations [185, 186]. Some extensions including prescribed lateral avoidance maneuvers have also been considered [187]. The resulting decision policy is often complex with multiple stages, but can be designed to be conservative, delayed or delay-conservative [188]. The nomenclature stems from the differences in sensitivity regarding the time at which the resulting avoidance decisions are issued. In any case, the decision policy is derived using an iterative process. The logic structure and associated thresholds are specified, evaluated on simulated encounters, and then further refined using a set of statistical performance metrics [189]. Importantly, the evaluation process is based on well established signal processing theory, using derivatives of Receiver Operating Curves (ROC) to visualise performance variations [190]. Such curves are designed to capture a complete description of the decision threshold effects regardless of the system particulars, so may be extended to other avoidance decision frameworks.

Recently, Dynamical Programming (DP), Markov Decision Process (MDP) and Partially Observable Markov Decision Process (POMDP) techniques have been combined to better optimise the avoidance decision policy [191]-[193]. The optimisation problem can be solved offline [194], online [195] or in a hybrid approach [196]. The result is an optimal logic table that maps the current state to an avoidance action (or lack thereof) [197, 198]. Similar to traditional approaches, a discrete set of avoidance actions is considered. In contrast to traditional approaches, a reward structure (cost function) requires tuning instead of collision probability thresholds. Additionally, a comprehensive encounter model is required for the both logic optimisation and validation through a simulated statistical performance evaluation.

[10]Updated specific TCAS II Minimum Operational Performance Specification (MOPS) have been published by RTCA (DO-300A,DO-185B) and EUROCAE (ED-143).

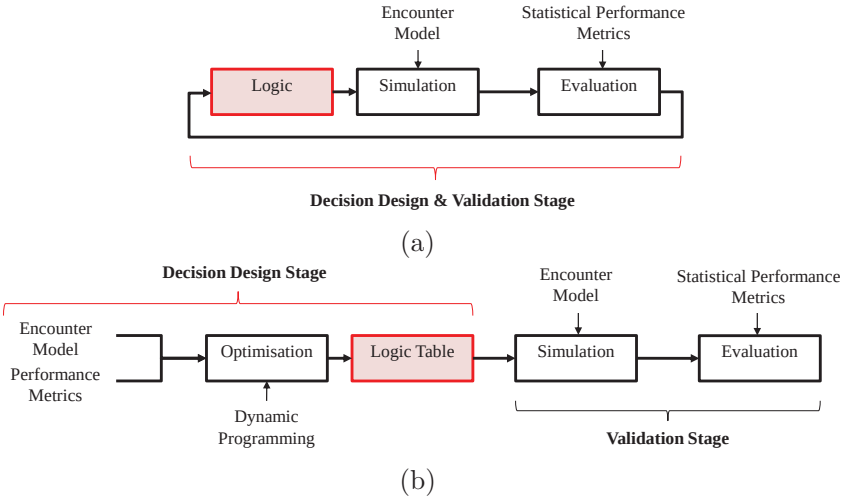

Figure 1.11: Example avoidance decision or logic design and validation approaches for aircraft collision avoidance systems. (a) Traditional conflict probabilities approach (b) Optimal logic design using Dynamic Programming and MDP/POMDP approaches.

Despite the optimality of recent approaches, the extensibility to vision-based See and Avoid systems is questionable. Computational complexity can be restrictive and comprehensive encounter models for a mixed airspace environment are not readily available. Only recently, have some approximations been derived using recorded radar data [64]-[67]. More importantly, they do not perform well when applied to angle-only sensors, inducing unwanted oscillatory avoidance behaviour [194]. Given the limited avoidance time available in See and Avoid, this is unacceptable. Traditional approaches do not suffer from these negative effects, and can still approximate the optimal logic reasonably well [188]. The inherent overhead using an iterative approach to refine the logic can also be reduced, by minimising the number of variable design parameters and thresholds. Additionally, traditional logic designs and associated performance evaluation methods were used in the certification of current systems (TCAS) [186]. They may offer an attractive framework in which to design vision-based See and Avoid avoidance logic [199].

Avoidance Control: Considering an onboard pilot is usually assumed with many aircraft collision avoidance systems, the avoidance decision can be implemented either manually or automatically. At a high level, the avoidance control is implemented in an open-loop manner whereby a prescribed avoidance maneuver is set and generally followed unless an updated maneuver is issued. The prescribed actions are quantitative in the sense they are made up of both direction (up, level-off, left etc.) and magnitude defined by an achievable rate (i.e. $\pm$1500'/s etc.) for common aircraft types [196]. At a low-level, the specific action is implemented in a closed-loop manner to establish and maintain the required vertical rate. If automated, onboard autopilots are used. If manual, the pilot regulates the corresponding rate using onboard instrumentation and available control inputs. Similar to human navigation, the avoidance control does not use path planning or other optimal control approaches. In contrast to human navigation, the localised controller regulates the direction and rate at which to move in that direction given a unique reference value.

Robot Avoidance Systems

Many automated collision avoidance approaches using a variety of sensors have been derived for a range of robotic systems including small Micro Aerial Vehicles (MAV), Autonomous Underwater vehicles (AUV), Unmanned Ground Vehicles (UGV) or humanoid robots. Although the underlying dynamics (motion models) are significantly different, it is important to consider existing robotic collision avoidance concepts from a See and Avoid perspective. After all, flying robots are effectively aircraft [200] without a history of operational rules and regulations. This allows more freedom in the development of unmanned systems, which may lead to novel automated collision avoidance approaches [201, 202]. This attribute is often overlooked in the aviation community. However, considering the limitations imposed by the See and Avoid environment and the plethora of robot collision avoidance systems, many approaches clearly violate the problem constraints so cannot be considered[11]. As such, robotic vision-based conflict detection, avoidance decision and avoidance control concepts are discussed below.

Conflict Detection: Conflict detection has been accomplished using localised maps, feature or template matching and by directly exploiting the properties of optic flow.

Localised maps of the observed scene can be artificially re-constructed (in full or in part) using dense or sparse image features, in an attempt to classify collision objects. First, the optic flow field patterns and onboard inertial sensors are used to create the map, instead of inferring collision objects directly [203, 204]. Second, the depth [205] or collision time [146] of each object is used to prioritise potential collision objects within the scene. Often, both visual and non-visual sensors (lasers, ultrasonic or stereo vision) are required, but the task can be accomplished using monocular vision [206, 207]. However, it is then assumed the observer is slow and the environment consists of large static objects with sufficient texture. Indeed, extending the approach to dynamic objects is difficult given the observability issues and image processing challenges associated with feature correspondence (data association). As such, localised map-based conflict detection will likely fail in the See and Avoid environment given the scene characteristics.

Feature or template matching based conflict detection first requires the detection of a known object or specific feature set [208]. A visual template of the object or known feature arrangements (corners, edges etc.) are used to find regions of the image that match the model [209]. The computer vision algorithms used for detection can be trained to detect a range of objects, but often require a comprehensive library of training data for adequate detection rates. Additionally, the object itself must be distinct enough (shape, color etc.) to be reliably identified amongst other objects in the scene. Once detected, range or time to collision may then be approximated given the actual object size is known (a priori) [210, 211]. As such, template matching based conflict detection will likely fail in the See and Avoid environment given the object of interest is a single point feature and the object size is unknown.

[11] Such approaches primarily include cooperative and/or intent based systems.

Direct optic-flow based conflict detection has been inspired by the navigation and collision avoidance behaviour of insects [212, 213]. It is generally accepted that insects use visual looming and optic flow field properties to qualitatively asses collision objects for navigation [214, 215]. An explicit estimation of range or time to collision is not required, and instead the direction and magnitude of the optic flow vectors associated with specific objects are used. A relatively stationary object that grows in size (positive optic flow) denotes a collision [216]. Similar to map-based approaches, multiple features from the same near or large object are usually required to obtain a reliable optic flow field pattern [217]. In contrast to map-based approaches, the whole scene is not reconstructed and the conflict detection is somewhat qualitative. When applied to See and Avoid encounters, and similar to human navigation, the same principle may be used but applied to distant objects. In this case, a collision can be detected by the lack (very small) optic flow from detected objects compared to the background.

Avoidance Decision: For conflict detection approaches using relative position and velocity, a common approach to making avoidance decisions is based on geometric optimisation. Geometric optimisation uses the current and projected geometry of the encounter to make decisions that optimise the avoidance behaviour, with respect to miss distance. Common approaches include optimal bang-bang control [218]-[221], collision cones [222, 223] and velocity (and acceleration) obstacles [224]-[230]. In all cases the objective is to define an optimal aiming point or waypoint for the aircraft to ensure collision avoidance. Often a set of avoidance maneuvers could be used, so additional criterion is used to optimise the decision and derive a single aiming point. For example, collision cones define the aiming point as the position on an arbitrary safety boundary perpendicular to the velocity vector of the object. Optimal bang-bang control determines the optimal left or right avoidance decision by considering predefined turning rates and assuming timely cooperative behaviour. Velocity obstacles define the minimum bearing change to realise a relative velocity vector outside the velocity obstacle or collision field. Importantly, the emphasis of geometrical approaches is to improve avoidance behaviour in a collision encounter by optimising the relative geometry between objects.

Alternatively, many path planning [231]-[233] and potential field approaches [234, 235] have been considered in order to include other criteria to better trade-off the competing navigation objectives and manage problem constraints. Successful path planning approaches can simultaneously ensure collision avoidance, respect aircraft performance limitations and better accommodate other desirable behaviours or objectives (such as goal reaching). Although not strictly geometrically optimal, well-tuned potential field methods can approximate an optimal avoidance path. The goal and collision object are treated as (weighted) attractive and repelling forces to negotiate a suitable aiming point. More recently, predictive potential fields have been used to better account for unmanned aircraft dynamics [236]. Importantly, the emphasis of path planning and potential field approaches is to improve overall behaviour in a collision encounter by optimising geometric and non-geometric objectives.

For conflict detection approaches explicitly using only visual cues, a common approach to making avoidance decisions is based on artificially inducing specific image feature movement through prescribed observer actions. The avoidance decision is typically based on changing the direction and magnitude of the image feature motion, and therefore the optic flow pattern. For static obstacles, it is common to move the object toward the edge of the image corresponding the side in which it is detected [237]. For example, the aircraft should move right and down if detected to the left and above the image centre. To achieve the desired motion, the direction (region) and a scaled magnitude may be used [238, 239] or a specific image feature reference position may be given. In the latter case, the position has been selected in an arbitrary fashion without justification [240] or based on achievable or desired vehicle performance [241]-[243]. More recently, motion has been derived by simultaneously attempting to avoid collision and improve observability [244]-[246] or by considering an approximate object size [247]. In most cases, the resulting unmanned aircraft motion tends to approximate a safe spiral trajectory about the object. This particular behaviour is rarely identified using such approaches, but could be exploited directly in the design of similar avoidance decision strategies.

Important to this discussion is how robotic collision avoidance system performance is evaluated. As operating standards are generally undefined, determining how safely a robot performs collision avoidance is subjective. Indeed, no rules exist for collision avoidance other than not to hit the object. As such, minimum safe distances are often chosen arbitrarily or based on the extent of the objects (size/scale), dynamics and control constraints and/or application specific requirements. The performance of the system is then analysed according to these metrics using a subset of encounters. Rarely, is the collision avoidance system analysed using statistical performance metrics as in manned aviation. This involves the comprehensive testing of the collision avoidance system through multiple high fidelity simulated encounters (Monte-Carlo, Encounter Models, Airspace Models etc.). So despite guarantees on collision avoidance, the introduction of modelling errors and uncertainties present in a realistic environment should warrant such an analysis to prove the safety case. This is particularly true of methods boasting robust range estimation strategies that underpin subsequent avoidance decisions.

Avoidance Control: Depending on the information used, the avoidance decision can be implemented at a high level using either a position-based or image-based visual closed loop control framework (visual servoing) [248]. Using relative state information allows position-based methods to be used. In this case, the unmanned aircraft can be guided to the aiming point (or way point) using direct positional feedback such as proportional navigation [249], prescribed maneuvers such as curves [250, 251], fuzzy systems [252] or optimal control approaches [253] such as minimum effort guidance [254] and model predictive control [255]. The latter has become increasingly popular due its ability to handle uncertainty, parameter variations and multiple constraints whilst remaining computationally efficient for practical implementations [256]-[259].

Using optic flow or feature position explicitly, reactive image-based methods derive feedback control directly from the image feature attributes. The approach is a natural way to

realise avoidance decisions that depend on the optic flow patterns directly [260]-[262] or to re-position objects in the image [314] without the computational over head of position-based approaches. Additionally, the approach more accurately models the avoidance control adopted in human navigation.

At a low-level, the specific action is implemented using onboard autopilots to regulate the required speed and attitude. Many low-cost high quality commercial autopilots suitable to many unmanned aircraft types (fixed/rotary) and sizes have become available that offer adequate performance. Typically, an arrangement of linear controllers such as Proportional-Integral-Derivative (PID) or Linear Quadratic Regulators (LQR) are used either coupled or de-coupled to control some or all aircraft degrees or freedom [264]-[266].

Resolution Strategies

A resolution strategy is generally required for automated systems to stop the avoidance manuever and possibly return to the original course prior to conflict detection. See and Avoid systems should therefore include an avoidance and resolution strategy to provide a complete closed loop collision avoidance system.

In manned aviation, resolution sometimes refers to the actual avoidance maneuver [164]. This is primarily due to the fact that in some cases it is not considered as part of collision avoidance, or that it is deemed less safety critical. However, this may be short-sighted for unmanned aircraft. Indeed, just because an avoidance action is issued and adopted, does not mean it will resolve the conflict. This is particularly true of fully automated systems, which may indeed continually attempt to avoid the object in some circumstances. As such, a decision strategy may be required for unmanned aircraft to be able to recognise the end of a conflict during an avoidance maneuver.

For many collision avoidance systems using relative state information, a resolution strategy is assumed or not disclosed. More precisely, the resolution decision may not be explicit and instead incorporated into the complete avoidance maneuver. To this end potential field, path planning and optimal control methods are common. Multiple waypoints may be used to guide the aircraft away from the object and then back to a predefined position, with no specific recognition required of when (or where) the collision encounter has ended.

For avoidance systems using optic flow or image feature position explicitly (including human navigation), a conflict is considered resolved when the object can be seen behind the observer. In some cases, this is inferred by the object passing outside the field of view [237]. Strictly speaking however, a conflict can only be considered resolved if the magnitude of the relative angular position is greater than 90 degrees in azimuth and non-zero in elevation [238]. Such an approach to resolution decisions provides a simple and flexible means in which to check for successful avoidance [267]. However, this places constraints on the number of cameras, their arrangement and their respective viewing angles (field of view). Additionally, it is unclear exactly how to incorporate such an approach into an automated decision strategy, and realistically asses its effectiveness. If the approach can be threshold based, then system analysis may be possible using (or extending) existing aviation performance evaluation approaches.

1.3.3 Research Implications & Summary

The previous sections have provided a comprehensive background on visual sensing and collision avoidance system principles in the context of vision-based See and Avoid systems. Particular attention was given to human, aircraft and robot navigation and control systems to provide a holistic approach to the problem whilst ensuring it can be addressed from a realistic perspective.

A number of key findings can be derived from this preliminary review, and used to help inform a useful investigation into the design and implementation of various system components. This includes the **avoidance decision**, **avoidance control** and **resolution decision** functions central to this research investigation. Important findings and their associated implications for this research program include:

- **General Decision & Control Approach**. It is very difficult to estimate the relative state (position, velocity, range etc.) for an arbitrary collision encounter in the Sense and Avoid environment. Even with prescribed motion (passive ranging), the short duration of the collision means no guarantee can be placed on filter convergence or the availability of accurate and timely estimates. Considering highly accurate relative angular positions can be obtained directly from monocular vision, it may be wiser to use these variables directly for decision and control. Considering also the alignment with human navigation models (constant angle assumption) and recent approaches to angle-only collision avoidance, effective avoidance strategies could be built by actively re-positioning the object on the image surface. For static objects, a spiral like trajectory is induced. For dynamics objects, the effects have not yet been fully analysed. The difficulty then resides in determining where or how to re-position the object in the image (avoidance decision) and which control strategy to use (avoidance control) for both static and dynamic objects.
- **Avoidance Decision** The majority of geometrically optimal and planned avoidance decision strategies do not align well to the See and Avoid problem constraints. Explicitly using visual cues to re-position the object in the image provides a simple approach to making avoidance decisions, that is well aligned to both human navigation models and the problem constraints. Specifically, the angular rate and position could be combined with the properties of spiral motion to help build a unique decision strategy. However, considering the ambiguity on some visual observations and the uncertainty in the underlying encounter geometry, avoidance decisions may then need to leverage existing rules of the air to be effective. Considering that predictable aircraft behaviour is important, consistency across all airspace users can then be better assured to help preserve airspace integrity.

  Additionally, such an approach is likely to contain one or more thresholds upon which decisions are based. The safety critical nature and short duration of typical encounters suggests a conservative decision policy, that perhaps avoids all targets in the image, should be adopted. Each object may be treated differently depending on the projected visual cues, threshold placements and decision structure. This bears strong

resemblance to traditional collision probability approaches used in manned aviation for certified systems. The difficultly in this case however, is in designing avoidance strategies that minimise the number of decision thresholds. If the latter is possible, existing statistical performance evaluation techniques used in manned aviation could then be used to simultaneously optimise the decision policy (threshold placements). and verify system performance. The findings suggests the avoidance decision strategy could be assembled using only direct visual Cues or *Optic Flow*, existing *Flight Rules*, *Spiral Motion* properties and a *Conservative Policy* or decision structure.

- **Avoidance Control** As the collision object needs to remain visible throughout the conflict, and to avoid visibility issues, a wide field of view is required. A high spatial resolution is preferred to improve detection distance to facilitate earlier target detection in a broader range of environmental conditions. Considering also the direct use of visually acquired angular observations, a natural choice for the imaging sensor is a *Spherical Camera* or spherical imaging model.

  Considering the lack of relative state information (position, velocity etc) position-based visual servoing is not suitable for avoidance control. For a reactive approach more aligned to the available information and human steering behaviour, a computational efficient *Image-Based Visual Servoing* scheme is better suited for avoidance control. The success of optimal control approach in collision avoidance control (position or image based), provides a strong argument for pursuing similar design principles. Recent improvements in computational power have allowed position based Model predictive Control (MPC) approaches to be designed that can provide optimal avoidance behaviour whilst considering a range of problem constraints. The finite prediction horizons used allow some robustness to parameters uncertainty (model mismatch) and measurement noise. Importantly, the use of short prediction horizons bears strong resemblance to recent beliefs regarding human collision avoidance control. The finding suggests that extensions to image-based visual servoing using Nonlinear Model Predictive Control (NMPC) could be used in a *Visual Predictive Control* approach that better encapsulates the control problem.
- **Resolution Decision** Most resolution strategies used to stop avoidance motion are not explicitly developed or disclosed, and instead inferred due to a planned avoidance approach. Other resolution strategies make use of the fact that an object is behind the observer if the magnitude of the azimuth angle is greater than 90 degrees. This simple criteria could be used for a vision-based resolution strategy, provided the camera viewing angle is sufficient. Alternatively, the concept may have to be adapted to meet any viewing deficiencies. In any case, the resolution strategy should be coupled with the avoidance control such that the avoidance behaviour is not stopped prematurely. Additionally, if the resolution strategy is threshold-based, it should be decoupled from the avoidance decision, to ensure the effects of each on performance can be appropriately distinguished.
- **Performance Evaluation** In order for any See and Avoid system to be certified, it must demonstrate at least an equivalent level of safety to that of manned aircraft. Despite the difficultly in completely defining this standard, a stronger safety case may

be established by adopting similar performance evaluation techniques as those used in manned aircraft systems (such as TCAS and ACAS). The finding suggests *Statistical Performance Metrics* and simulation techniques aligned to common aviation collision avoidance (and alerting) system analysis methods should be adopted. Focus could be on using or extending *System Operating Curve (SOC)* and *System Performance Curve (SP)* based methods derived from comprehensive encounter simulations (Monte-Carlo or otherwise) to fit the vision-based decision and control strategies.

- **Implementation** The regulatory, safety, repeatability and economical difficulties associated with implementation of a full scale See and Avoid system prohibit such an approach within the scope of this research. However, it is important to ensure that the decision and control strategies can be implemented in a practical system. Indeed, much work has focused on using small unmanned aircraft in preliminary flight trials. The finding suggests implementation using a *Proof-of-Concept* approach using *Scaled Environmental Conditions* that realistically represent See and Avoid encounters (where possible). Safe and easy to implement small indoor multirotor platforms could be used, as the avoidance and decision strategies can be extended to other platforms. It is then feasible to suggest a proof-of-concept implementation could be realised provided aircraft performance limitations, sensing issues and environmental conditions are realistically scaled and appropriately bounded.

The above findings simultaneously outline and justify the novelty of the particular approach to automated vision-based See and Avoid considered in this thesis. Specifically, the investigation includes the design, implementation (proof-of-concept) and evaluation (system operating curves) of new vision-based decision and control strategies, that uniquely combine existing aviation rules and spiral trajectories with novel extensions to spherical image-based visual servoing (IBVS) using nonlinear model predictive control (NMPC) schemes. The details and results of such an investigation constitute a series of significant and original contributions, and are described in the remaining chapters.

Chapter 2

Preliminaries

2.1 Notation

The following notation conventions and style are used in this thesis unless stated otherwise. For a complete list of variables and constants, including the symbolism and qualitative description, refer to Appendix E.

For any scalar $x \in \mathbb{R}^1$, $|x|$ denotes its absolute value. A vector valued function is denoted by $\mathbf{f}(\cdot)$. For any vector $\mathbf{x} \in \mathbb{R}^n$, $\|\mathbf{x}\|$ denotes the 2-norm and $\mathbf{x}_P$ denotes the weighted norm such that $\mathbf{x}_P^2 = \mathbf{x}^T\mathbf{P}\mathbf{x}$. Positive definite and positive semi-definite matrices are denoted by $\mathbf{P} \succ 0$ and $\mathbf{Q} \succeq 0$ respectively. The maximum and minimum real eigenvalues of a matrix $\mathbf{P}$ are given by $\lambda_{\max}(\mathbf{P})$ and $\lambda_{\min}(\mathbf{P})$ respectively. A skew symmetric matrix formed from the vector $\mathbf{x}$ is denoted by $\mathrm{sk}(\mathbf{x})$ and a diagonal matrix formed from the vector $\mathbf{y}$ is denoted by $\mathrm{di}(\mathbf{y})$. A rotation matrix and translation vector from frame b to frame a is denoted by ${}^aR_b$ and ${}^a\mathbf{t}_b$ respectively. Unit vectors are denoted using a tilde such that $\tilde{\mathbf{r}}$ denotes a unit range vector. Dot notation is used to represent variable derivatives where the independent variable is time such that $\dot{x} = dx/dt$.

A setpoint or reference value is denoted with an asterisk. For example, $\mathbf{u}^*(\cdot)$ defines a reference control sequence. A star denotes optimality such that the minimum cost resulting from the application of an optimal control sequence $\mathbf{u}^\star(\cdot)$ is given by $J^\star$.

2.2 Vehicle Dynamics

The focus of this thesis is a vision-based collision avoidance solution that may be amenable to multiple platform types, from multi rotors to fixed wing aircraft. As a large number of platforms exist, it is intractable to simulate, test and verify any such collision avoidance system for all cases. Instead, a representative set of vehicle dynamics is required to help determine system performance. If simulations alone were sufficient to derive system performance attributes, the choice of dynamic model is not restrictive. If the goal is to also flight test the system, only a limited set of platforms can be considered given the safety concerns and regulatory environment. To this end, a small indoor quadrotor platform is used in this thesis. This section first details the derivation of a suitable generic dynamic model for the quadrotor. Platform specific attributes are then added to the model. Although required for the experimental implementation, the vehicle dynamics are not presented as a contribution to this thesis and have been based on [268, 269].

2.2.1 Quadrotor Basics

A quadrotor is an aerial vehicle capable of quasi-stationary or hover flight similar to conventional helicopters. The platform consists of four rotors configured in a symmetrical cross pattern about the platform centre. Each rotor consists of a propeller blade directly attached to a motor located at the extremity of each arm. The platform centre contains the avionics and payload including telemetry links, autopilot and cameras.

The rotors are arranged in counter-rotating pairs such that the front and rear rotate counter-clockwise and the left and right rotate clockwise. The propellers have a fixed pitch and their axes of rotation are parallel to each other and the vertical. A vertical thrust force and horizontal drag force acting through the centre of each rotor is thus produced at each arm. As a result, a torque is also produced about each axis from the platform centre. The net force and torque on the quadrotor result in translational and rotational movement. Therefore, the motors are the only platform actuators. At a low-level, the control inputs required are variations in motor speed Ω_n for $n = \{1, 2, 3, 4\}$ rotors. The thrust force F_n from each rotor depends on the blade configuration and physical parameters, and is directly proportional to the square of the motor speed such that

$$F_n = k\Omega_n^2 \tag{2.1}$$

The torque τ_n produced at each motor depends, in part, on the propeller drag force $\bar{F}_n$ and propeller radius r_p, but can be simplified as

$$\tau_n = r_p \bar{F}_n = k_\tau \Omega_n^2 \tag{2.2}$$

At a higher level, a mixture of rotor speed variations can be used to construct four separate control inputs U_i for $i = \{1, 2, 3, 4\}$. The input controls consist of a force U_1 and three input torques U_2, U_3 and U_4. Although the platform is still under-actuated, the controls provide a more direct correspondence to changes in roll, pitch, yaw and vertical acceleration which is more useful for controlling position and attitude. A set of simplified diagrams depicting the various forces and torques acting on the platform resulting from changes in control input is given in Fig 2.1. The rotors are labelled 1-4 starting at the front rotor and moving clockwise through to the left rotor.

If all motor speeds are identical such that $\Omega_1 = \Omega_2 = \Omega_3 = \Omega_4$, a net upward force F_T is produced. The net torque $\tau_c = 0$ as the effects from each rotor cancel due to the counter-rotating arrangement. If the speeds are decreased or increased by the same amount and the pitch and roll angles are zero, the platform will descend, hover or ascend depending on the relative magnitude of F_T compared to the platform weight mg. For all other attitudes, a lateral and longitudinal force will also be induced causing movement in the horizontal plane. Setting $U_1 = F_T$ ensures the magnitude of the thrust can be directly

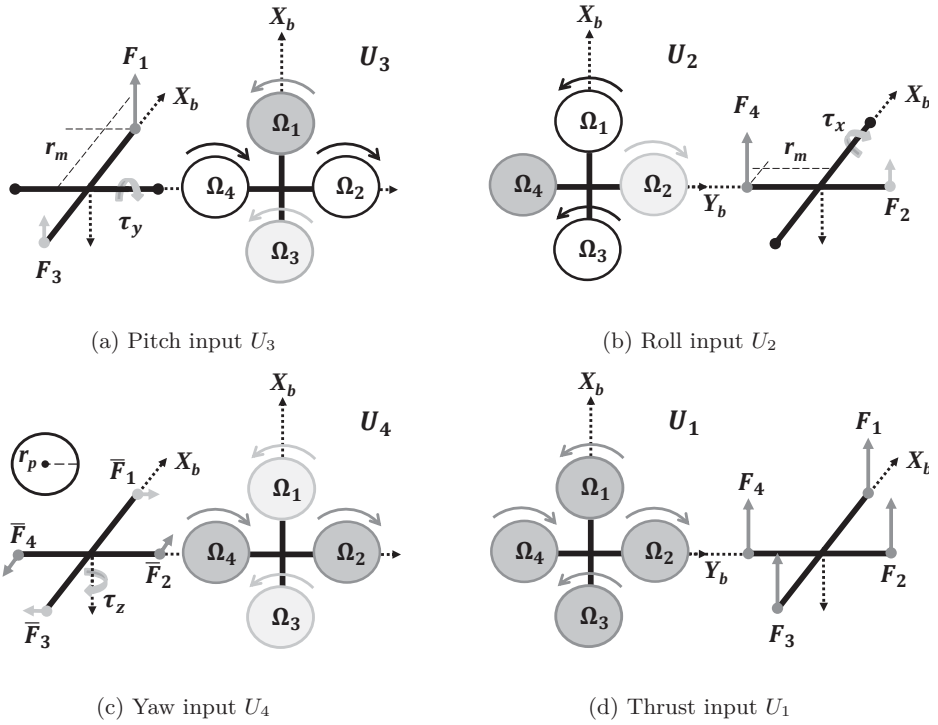

(a) Pitch input U_3

(b) Roll input U_2

(c) Yaw input U_4

(d) Thrust input U_1

Figure 2.1: Quadrotor control inputs

controlled such that

$$U_1 = F_1 + F_2 + F_3 + F_4 \tag{2.3}$$
$$U_1 = k(\Omega_1^2 + \Omega_2^2 + \Omega_3^2 + \Omega_4^2) \tag{2.4}$$

The orientation of the thrust vector is controlled by changing the platform attitude by inducing a non-zero torque. Increasing the right motor speed and decreasing the left motor speed by the same amount $\Delta\Omega$ results in a positive torque τ_x about the lateral axis whilst retaining the overall thrust such that $F_T = F_1 + (F_2 + \Delta F) + F_3 + (F_4 - \Delta F)$. An acceleration in roll angle is thus induced. Setting $U_2 = \tau_x$

$$U_2 = r_m(F_4 - F_2) \tag{2.5}$$
$$U_2 = r_m k(\Omega_4^2 - \Omega_2^2) \tag{2.6}$$

Increasing the front motor speed and decreasing the rear motor speed by the same amount $\Delta\Omega$ results in a positive torque τ_y about the longitudinal axis whilst retaining the overall thrust such that $F_T = (F_1 + \Delta F) + F_2 + (F_3 - \Delta F) + F_4$. An acceleration in pitch angle is thus induced. Setting $U_3 = \tau_y$

$$U_3 = r_m(F_3 - F_1) \tag{2.7}$$
$$U_3 = r_m k(\Omega_3^2 - \Omega_1^2) \tag{2.8}$$

The torque produced by each motor is required to maintain the propellers spin and provide the necessary thrust force. It also induces a drag force $\bar{F}_n$ perpendicular to the associated

thrust force such that each rotor contributes some torque about the body vertical axis. Increasing the right and left motor speed and decreasing the front and rear motor speed by the same amount $\Delta\Omega$ results in a positive torque τ_z about the vertical axis whilst retaining the overall thrust such that $F_T = F_1 + (F_2 + 2\Delta F) + F_3 + (F_4 - 2\Delta F)$. An acceleration in yaw angle is thus induced. Setting $U_4 = \tau_z$

$$U_4 = r_p(\bar{F}_2 + \bar{F}_4 - \bar{F}_1 - \bar{F}_3) \tag{2.9}$$

$$U_4 = k_\tau(\Omega_2^2 + \Omega_4^2 - \Omega_1^2 - \Omega_3^2) \tag{2.10}$$

The equations describing the control inputs can then represented in matrix form as

$$\begin{pmatrix} U_1 \\ U_2 \\ U_3 \\ U_4 \end{pmatrix} = \begin{pmatrix} k & k & k & k \\ 0 & -r_m k & 0 & r_m k \\ -r_m k & 0 & r_m k & 0 \\ -k_\tau & k_\tau & -k_\tau & k_\tau \end{pmatrix} \begin{pmatrix} \Omega_1^2 \\ \Omega_2^2 \\ \Omega_3^2 \\ \Omega_4^2 \end{pmatrix} \tag{2.11}$$

The model for the quadrotor control inputs derived above has been simplified. The simplifications include the following well-founded assumptions [270]-[272]:

- The effects of blade flapping resulting from deformation of the propeller moving at high velocity are neglected.
- Each rotor is identical. The associated drag force on each propeller is assumed to act at the tip of the blades to produce the rotor torque. The associated propeller velocity is such that an approximately constant thrust is induced.
- The platform is symmetrical about the zx and zy planes.
- The principles of conservation of energy apply to each rotor such that all input energy is transferred to a thrust and drag force.

2.2.2 Quadrotor Model

The physical variations between different quadrotor platforms renders the flight characteristics of each vehicle unique. Coupled with environmental factors resulting from a diverse operating region, an analysis of the forces and moments acting on a quadrotor would result in a unique, complex and cross-coupled nonlinear dynamic model. However, under some simplifying assumptions, some generic dynamic models can provide a suitable approximation to the real dynamics for a large number of similar platforms operating over a particular region of the flight envelope. To this end, a simplified generic point mass model for a quadrotor can be derived using Newtonian mechanics and the principles of linear momentum. The platform is considered to be a point object in space in which the origin of the body fixed frame $\mathcal{F}_b$ coincides with the centre of mass. The body axis is also assumed to coincide with the body principle axis of inertia. The vehicle is free to move in 6 degrees of freedom (DOF) with respect to a fixed inertial or world frame $\mathcal{F}_w$, but is assumed to remain in quasi-stationary flight (hover). In this way, the platform is assumed to be in a steady state either hovering or slowly moving such that the surrounding

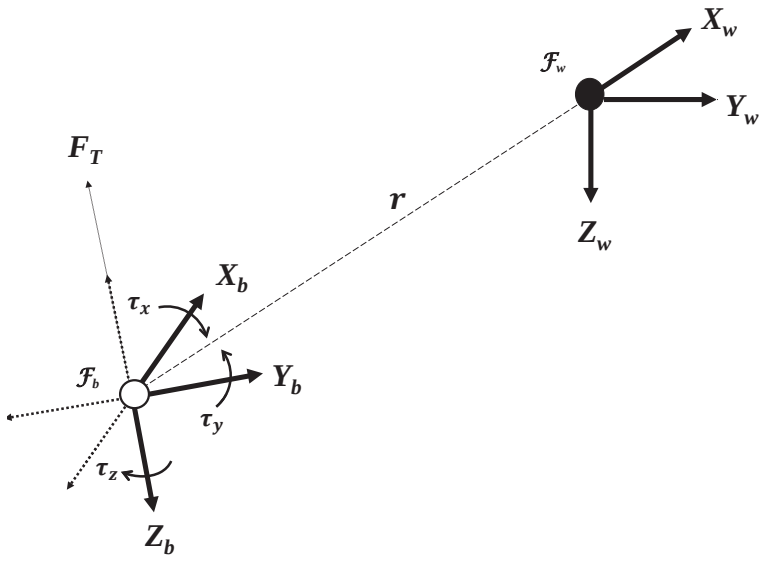

Figure 2.2: Coordinate frames for the quadrotor body $\mathcal{F}_b$ and the world frame $\mathcal{F}_w$ along with an example point in space measured with respect to the body ${}^b\mathbf{p}$ and world frame ${}^w\mathbf{p}$.

environment is stationary, aerodynamic drag and thrust forces due to the fuselage can be neglected and only small pitch and roll angles are induced. The control inputs consist of the forces U_1 and torques U_2, U_3 and U_4 derived earlier. The geometry is depicted in Fig 2.2.

The point mass model can be derived in two stages. First, the forces and moments acting on the vehicle are considered separately then merged to define the translational and angular acceleration components on the platform. Second, by applying the aforementioned simplifying assumptions, the dynamics are then reduced to a set of linear equations that can be represented in a body or world frame.

The force acting on the quadrotor platform in the world frame ${}^w\mathbf{F}$ can be defined in terms of the gravitational force $\mathbf{F}_g$, wind forces $\mathbf{F}_w$ and body forces ${}^b\mathbf{F}$ such that

$$ {}^w\mathbf{F} = {}^wR_b\,{}^b\mathbf{F} + \mathbf{F}_g + \mathbf{F}_w \tag{2.12} $$

where ${}^wR_b$ defines a rotation matrix from the body to world frames in terms of the roll ϕ, pitch θ and yaw ψ Euler angles. The force acting on the body can be defined in terms of the control input force $\mathbf{F}_c$ and angular velocity vector ω such that

$$ {}^b\mathbf{F} = \mathbf{F}_c + \omega \times m\,{}^b\mathbf{v} \tag{2.13} $$

$$ {}^b\mathbf{F} = \mathbf{F}_c - \mathrm{sk}(m\,{}^b\mathbf{v})\omega \tag{2.14} $$

where ${}^b\mathbf{v}$ is the body velocity vector and the second term defines the Coriolis forces resulting from the cross-coupling of angular velocities. As the platform remains close to the hover point, with only small pitch and roll angles, the second term can be ignored as its contribution to the resulting force is relatively small. Neglecting the effects of external

disturbances due to turbulence and wind gusts, (2.14) becomes

$$^{w}\mathbf{F} = {}^{w}R_{b}\mathbf{F}_{c} + \mathbf{F}_{g} \tag{2.15}$$

Re-writing (2.15) in terms of translational accelerations $^{w}\ddot{x}$ and expanding for the individual components

$$\begin{pmatrix} ^{w}\ddot{x} \\ ^{w}\ddot{y} \\ ^{w}\ddot{z} \end{pmatrix} = \frac{1}{m}{}^{w}R_{b}\begin{pmatrix} ^{b}F_{x} \\ ^{b}F_{y} \\ ^{b}F_{z} \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ g \end{pmatrix} \tag{2.16}$$

where m is the quadrotor mass and g is the gravitational constant. The control inputs collectively result in an applied thrust U_1 acting parallel to the vertical axis such that $U_1 = -^{b}F_{z}$. Equation (2.16) becomes

$$\begin{pmatrix} ^{w}\ddot{x} \\ ^{w}\ddot{y} \\ ^{w}\ddot{z} \end{pmatrix} = \frac{1}{m}\begin{pmatrix} C\theta C\psi & S\phi S\theta C\psi - C\phi S\psi & C\phi S\theta C\psi + S\phi S\psi \\ C\theta S\psi & S\phi S\theta S\psi + C\phi C\psi & C\phi S\theta S\psi - S\phi C\psi \\ -S\theta & S\phi C\theta & C\phi C\theta \end{pmatrix}\begin{pmatrix} 0 \\ 0 \\ -U_1 \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ g \end{pmatrix} \tag{2.17}$$

where $C = \cos(\cdot)$ and $S = \sin(\cdot)$. Expanding the above equation and simplifying

$$\begin{pmatrix} ^{w}\ddot{x} \\ ^{w}\ddot{y} \\ ^{w}\ddot{z} \end{pmatrix} = \begin{pmatrix} U_1(-\cos\psi\sin\theta\cos\phi - \sin\psi\sin\phi)/m \\ U_1(\cos\psi\sin\phi - \sin\psi\sin\theta\cos\phi)/m \\ g - U_1(\cos\theta\cos\phi)/m \end{pmatrix} \tag{2.18}$$

The torque acting on the quadrotor platform in the world frame $^{w}\tau$ can be defined in terms of the body torque $^{b}\tau$ and wind torque τ_w such that

$$^{w}\tau = {}^{w}T_{b}{}^{b}\tau_{c} + \tau_{w} \tag{2.19}$$

where $^{w}T_{b}$ defines a transformation matrix from the body to world frames in terms of the roll ϕ, pitch θ and yaw ψ Euler angles. The torque acting on the body can be defined in terms of the control torque τ_c and gyroscopic torque τ_g such that

$$^{b}\tau = \tau_{c} + \tau_{g} \tag{2.20}$$

$$^{b}\tau = \mathbf{J}\dot{\omega} + \omega \times (\mathbf{J}\omega) \tag{2.21}$$

$$^{b}\tau = \mathbf{J}\dot{\omega} - \mathrm{sk}(\mathbf{J}\omega)\omega \tag{2.22}$$

where ω is the body angular velocity and $\mathbf{J}$ is the inertia matrix. As the platform is symmetrical, $\mathbf{J} = \mathrm{di}(J_{xx}, J_{yy}, J_{zz})$ where the moments of inertia $J_{nn} = mr_{nn}^{2}$. Neglecting wind effects and assuming the algebraic sum of the motor speeds is approximately zero, such that the gyroscopic effects resulting from any imbalance have a minimum impact on

the resulting torque, the torque acting on the platform becomes

$$
\begin{aligned}
{}^{w}\tau &= {}^{w}T_b{}^{b}\tau_c && (2.23)\\
{}^{w}\tau &= {}^{w}T_b{}^{b}J\dot{\omega} && (2.24)
\end{aligned}
$$

The torque on the platform is thus a direct result of the applied control torque τ_c. Rewriting (2.24) in terms of angular accelerations and expanding for the individual components

$$
\begin{pmatrix} \ddot{\phi} \\ \ddot{\theta} \\ \ddot{\psi} \end{pmatrix} = {}^{w}T_b{}^{b}J^{-1} \begin{pmatrix} \tau_x \\ \tau_y \\ \tau_z \end{pmatrix} \tag{2.25}
$$

Substituting the torque produced by the rotors for the control inputs and expanding

$$
\begin{pmatrix} \ddot{\phi} \\ \ddot{\theta} \\ \ddot{\psi} \end{pmatrix} = \begin{pmatrix} 1 & \sin\phi\tan\theta & \cos\phi\tan\theta \\ 0 & \cos\phi & -\sin\phi \\ 0 & \sin\phi/\cos\theta & \cos\phi/\cos\theta \end{pmatrix} \begin{pmatrix} U_2/J_{xx} \\ U_3/J_{yy} \\ U_4/J_{zz} \end{pmatrix} \tag{2.26}
$$

Under the assumption the quadrotor moves about the hover position, with small pitch and roll angles, the transformation matrix can be simplified to the identity matrix $\mathbf{I}_3$. This can be seen through inspection of the transformation matrix. As such

$$
\begin{pmatrix} \ddot{\phi} \\ \ddot{\theta} \\ \ddot{\psi} \end{pmatrix} = \begin{pmatrix} U_2/J_{xx} \\ U_3/J_{yy} \\ U_4/J_{zz} \end{pmatrix} \tag{2.27}
$$

The complete set of dynamic equations can then be defined by combining (2.32) and (2.27) such that

$$
\begin{pmatrix} {}^{w}\ddot{x} \\ {}^{w}\ddot{y} \\ {}^{w}\ddot{z} \\ \ddot{\phi} \\ \ddot{\theta} \\ \ddot{\psi} \end{pmatrix} = \begin{pmatrix} U_1(-\cos\psi\sin\theta\cos\phi - \sin\psi\sin\phi)/m \\ U_1(\cos\psi\sin\phi - \sin\psi\sin\theta\cos\phi)/m \\ g - U_1(\cos\theta\cos\phi)/m \\ U_2/J_{xx} \\ U_3/J_{yy} \\ U_4/J_{zz} \end{pmatrix} \tag{2.28}
$$

Further simplifications can be made under some additional assumptions. Firstly, the yaw angle is assumed to be zero such that $\psi = 0$. Secondly, the thrust dynamics of the platform are assumed to be sufficiently fast such that the total input thrust keeps the platform from accelerating in the vertical plane. To this end, consider a virtual body frame $\mathcal{F}_b'$ depicted in Fig 2.3 that is no longer displaced by a small angle Φ, but instead remains level with the horizontal plane of the world frame. The force ${}^{b}F'$ acting along

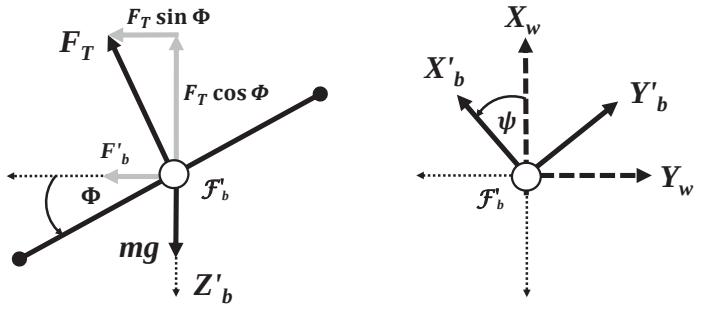

Figure 2.3: Simplification of quadrotor dynamics for a point mass model derivation

the associated axis in the horizontal plane is given by

$$^bF' = F_T \sin \Phi \tag{2.29}$$

The vertical force in the virtual body frame is the same as that in the world frame and can be expressed as

$$^wF_z = F_T \cos \Phi = mg \tag{2.30}$$

Combining (2.29) and(2.30) then

$$^bF' = mg \tan \Phi \tag{2.31}$$

Due to the symmetry of the platform, $^bF'$ can represent the force in the lateral or longitudinal direction with Φ denoting the associated roll or pitch angle. Applying the small angle theorem such that $\tan \Phi \approx \Phi$, $\sin \Phi \approx \Phi$, $\cos \Phi \approx 1$ and substituting into the translational acceleration equations, the generalised quadrotor equations of motion become

$$\begin{pmatrix} \ddot{x} \\ \ddot{y} \\ \ddot{z} \\ \ddot{\phi} \\ \ddot{\theta} \\ \ddot{\psi} \end{pmatrix} = \begin{pmatrix} U_1\theta/m \\ U_1\phi/m \\ g - U_1/m \\ U_2/J_{xx} \\ U_3/J_{yy} \\ U_4/J_{zz} \end{pmatrix} \tag{2.32}$$

Note, the reference frame for translational acceleration has been omitted in the above dynamics. As they are independent of the yaw angle, they could be considered to be the translational dynamics of the virtual body frame. They describe the forward-backward, left-right and up-down motion of the quadrotor. Without further manipulation, and considering the assumptions under which they were derived, they could be used as an approximation to the vehicle dynamics with respect to the body frame. Alternatively, they can be used as an approximation to the vehicle dynamics in the world frame provided

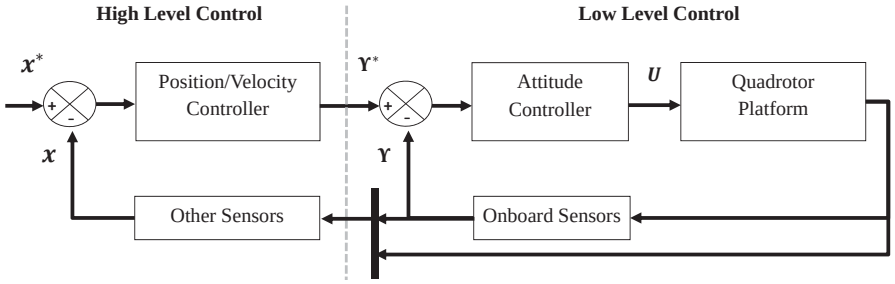

Figure 2.4: Generalised quadrotor control architecture where $\mathbf{x}^*$ denotes the reference input vector to the high level controller and Υ^* denotes the reference input vector to the low-level controller. The actual input to the quadrotor consists of motor speed commands through $U = (U_1, U_2, U_3, U_4)$.

an additional term is added to account for the yaw offset. Essentially, correcting the translational accelerations in the horizontal plane.

2.2.3 Quadrotor Control

Automatic control of most aerial vehicles is typically accomplished using a nested set of control loops. The particular arrangement of control loops may depend on the platform hardware, vehicle dynamics and actuators and the requirements of a particular application [273]-[275]. For quadrotors, a typical control architecture consists of inner low-level controller(s) and an outer high-level controller(s). The generalised quadrotor control architecture is depicted in Fig 2.5 with each controller detailed in the following sections.

Low-Level Control

The low-level controller provides high rate onboard attitude stabilisation. Specifically, the attitude controller regulates the roll, pitch, yaw rate, and thrust $\Upsilon = (\phi, \theta, \dot{\psi}, F_T)$ by controlling rotor speeds through $U = (U_1, U_2, U_3, U_4)$. The high level controller provides the roll, pitch, yaw rate and thrust references $\Upsilon^* = (\phi^*, \theta^*, \dot{\psi}^*, F_T^*)$. As many commercially available quadrotors [276, 277] come with onboard stabilisation, the attitude controller is often predetermined or fixed for a given platform. A common choice is a set of Proportional-Derivative (PD) or Proportional-Integral-Derivative (PID) controllers. Some vendors allow user access to the attitude controller parameters whilst others do not [278]. Coupled with the motor particulars and avionics arrangement, the attitude controller is generally very specific to the platform and the associated dynamics can be difficult to determine analytically. They are therefore generally unknown before the complete quadrotor system has been implemented.

The attitude dynamics can be approximated empirically using grey or black box system identification. Provided the attitude controller is sufficiently fast, each command in Υ^* can be treated independently. The controller can then be approximated as four decoupled subsystems, one for each variable in Υ. Each subsystem can be identified by measuring the response to a step command (issued during hover), then analysing the associated input-output relationship to derive an appropriate subsystem transfer function. When

required, this work leverages the experimentally identified attitude controller dynamics for a small AscTec Hummingbird platform with modified payload [268, 269]. The attitude controller dynamics for the Euler angles are defined by a set of 1$^{\text{st}}$ order transfer functions

$$\mathscr{L}\{\phi(t)\} \rightarrow \quad \Phi(s) \quad = \frac{a_\phi}{s + a_\phi}\Phi(s)^* \tag{2.33}$$

$$\mathscr{L}\{\theta(t)\} = \quad \Theta(s) \quad = \frac{a_\theta}{s + a_\theta}\Theta(s)^* \tag{2.34}$$

$$\mathscr{L}\{\dot{\psi}(t)\} = \quad \Psi(s) \quad = \frac{a_\psi}{s + a_\psi}\Psi(s)^* \tag{2.35}$$

where $\mathscr{L}$ denotes the Laplace operator and a_ϕ, a_θ, a_ψ are experimentally determined constants given in Appendix E. Taking the inverse Laplace transform of (2.33)-(2.35), the corresponding continuous time differential equations are given by

$$\dot{\phi}(t) \quad = \quad -a_\phi\phi(t) + a_\phi\phi^*(t) \tag{2.36}$$

$$\dot{\theta}(t) \quad = \quad -a_\theta\theta(t) + a_\theta\theta^*(t) \tag{2.37}$$

$$\ddot{\psi}(t) \quad = \quad -a_\psi\dot{\psi}(t) + a_\psi\dot{\psi}^*(t) \tag{2.38}$$

The above equation are of the form $\dot{x}(t) = ax(t) + bu(t)$ and can be discretised according to the method presented in [279]. Given the sampling instant k and a sufficiently small sampling period T_s, a zero-order approximation to the continuous time equations results in the corresponding difference equations

$$\phi(kT_s + T_s) \quad = \quad e^{-a_\phi T_s}\phi(kT_s) + (1 - e^{a_\phi T_s})\phi^*(kT_s) \tag{2.39}$$

$$\theta(kT_s + T_s) \quad = \quad e^{-a_\theta T_s}\theta(kT_s) + (1 - e^{a_\theta T_s})\theta^*(kT_s) \tag{2.40}$$

$$\dot{\psi}(kT_s + T_s) \quad = \quad e^{-a_\psi T_s}\dot{\psi}(kT_s) + (1 - e^{a_\psi T_s})\dot{\psi}^*(kT_s) \tag{2.41}$$

The thrust dynamics are considerably faster and can be defined as a 0$^{\text{th}}$ order system such that

$$\mathscr{L}\{F_T(t)\} = F_T(s) \quad = \quad \frac{1}{s}F_T(s)^* \tag{2.42}$$

Taking the inverse Laplace transform

$$F_T(t) \quad = \quad F_T^*(t) \tag{2.43}$$

Substituting the thrust dynamics into the equation for vertical acceleration given by (2.32), the continuous time differential equation for vertical acceleration can be expressed as

$$\ddot{z}(t) \quad = \quad g - \frac{F_T^*(t)}{m} \tag{2.44}$$

Again, taking a zero-order approximation and assuming a sufficiently small sampling time, the discrete time difference equation can be expressed as

$$\dot{z}(kT_s + T_s) = \dot{z}(kT_s) + T_s \left(g - \frac{F_T^*(kT_s)}{m} \right) \tag{2.45}$$

The complete attitude controller system dynamics can then be expressed in discrete time state space as

$$\begin{aligned}
\phi(kT_s + T_s) &= \phi(kT_s) + a_\phi T_s \phi^*(kT_s) && (2.46)\\
\theta(kT_s + T_s) &= \theta(kT_s) + a_\theta T_s \theta^*(kT_s) && (2.47)\\
\dot{\psi}(kT_s + T_s) &= \dot{\psi}(kT_s) + a_\psi T_s \dot{\psi}^*(kT_s) && (2.48)\\
\dot{z}(kT_s + T_s) &= \dot{z}(kT) + T_s(g - \frac{F^*(kT)}{m} && (2.49)
\end{aligned}$$

where the properties of the exponential series are used in the approximation such that $e^{-aT_s} \approx 1 - aT_s \approx 1$ and $1 - e^{-aT_s} \approx aT_s$. Note, the time delay (dead time) for each subsystem remains relatively small. This is particularly true for the thrust, roll and pitch dynamics with the yaw rate observed to have a slightly longer time delay. Following a similar procedure as above and including the time delays $\delta t_\phi, \delta t_\theta, \delta t_\psi, \delta t_{F_T}$ for the corresponding subsystem, the resulting discrete time difference equations become

$$\begin{aligned}
\phi(kT_s + T_s) &= \phi(kT_s - \delta t_\phi) + a_\phi T_s \phi^*(kT_s - \delta t_\phi) && (2.50)\\
\theta(kT_s + T_s) &= \theta(kT_s - \delta t_\theta) + a_\theta T_s \theta^*(kT_s - \delta t_\theta) && (2.51)\\
\dot{\psi}(kT_s + T_s) &= \dot{\psi}(kT_s - \delta t_\psi) + a_\psi T_s \dot{\psi}^*(kT_s - \delta t_\psi) && (2.52)\\
\dot{z}(kT_s + T_s) &= \dot{z}(kT_s - \delta t_{F_T}) + T_s \left(g - \frac{F_T^*(kT_s - \delta t_{F_T})}{m} \right) && (2.53)
\end{aligned}$$

Subsystem identification has been performed in this way for a number of real quadrotors resulting in sets of 0^{th} to 3^{rd} order subsystems [268, 269, 272, 284]. When considering the appropriate system order, there exists a tradeoff in complexity and performance. A higher order system may better represent the true underlying attitude dynamics at the expense of a more complex model. The first-order attitude dynamics derived above have shown to provide an adequate representation of the platform dynamics for quasi-stationary and non-aggressive (aerobatic) flight conditions [269]. Note, the dynamics pertain to the platform used in flight trial experiments presented in Chapter 5.

High-Level Control

The high-level controller provides low rate control of the quadrotor state. Specifically, the high-level controller regulates the quadrotor position $\mathbf{x}$, velocity $\dot{\mathbf{x}}$ or both via the low-level controller. Given the intended task, the reference state is provided to the controller which then outputs the control vector $\mathbf{u}$. The control vector is then fed as input to the low-level controller such that $\mathbf{u} = \Upsilon^* = (\phi^*, \theta^*, \dot{\psi}^*, F_T^*)$.

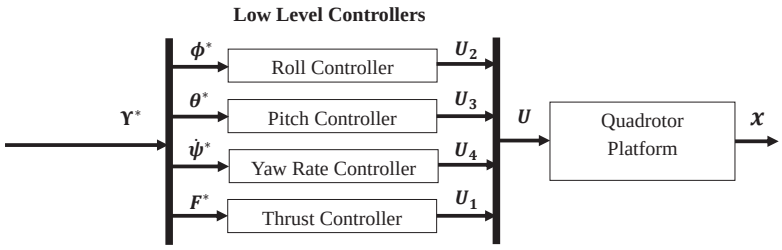

Figure 2.5: Generalised low-level quadrotor control architecture where Υ^* denotes the reference input vector to the low-level attitude controllers. The actual quadrotor input consists of motor speed commands $U = (U_1, U_2, U_3, U_4)$.

The design of the controller largely depends on the application and platform hardware. As many commercial quadrotors are intended for sports and recreation purposes, they often rely on manual input for position and velocity control. Some vendors include some high level control for position based on GPS sensors. In most cases however, a custom controller is required to provide autonomous navigation suitable to the operating environment and available sensors. When designing such a controller, a complete model of the system dynamics may be required in order to implement advanced control techniques. This can be derived by combining the point mass model and the attitude controller dynamics [280].

Consider three state vectors $\mathbf{x}_x$, $\mathbf{x}_y$ and $\mathbf{x}_z$ and associated control vectors $\mathbf{u}_x$, $\mathbf{u}_y$ and $\mathbf{u}_z$ defined in discrete time such that

$$\begin{aligned} \mathbf{x} &= (\mathbf{x}_x\ \mathbf{x}_y\ \mathbf{x}_z), \quad \mathbf{x}_x = (x\ \dot{x}\ \dot{\theta}\ \theta) \quad \mathbf{x}_y = (y\ \dot{y}\ \dot{\phi}\ \phi), \quad \mathbf{x}_z = (z\ \dot{z}\ \psi\ \dot{\psi}) \\ \mathbf{u} &= (\mathbf{u}_x\ \mathbf{u}_y\ \mathbf{u}_z), \quad \mathbf{u}_x = \theta^*, \qquad \mathbf{u}_y = \phi^*, \qquad \mathbf{u}_z = (\dot{z}^*\ \dot{\psi}^*) \end{aligned}$$

Combining the attitude controller dynamics with the point mass model for $\mathbf{x}_x$ the dynamics for the longitudinal (x) axis can be expressed as

$$\mathbf{x}_x(kT_s + T_s) = \mathbf{A}_x\mathbf{x}_x(kT_s) + \mathbf{B}_x\mathbf{u}_x(kT_s) \tag{2.54}$$

$$\begin{pmatrix} x(kT_s+T_s) \\ \dot{x}(kT+T) \\ \theta(kT+T) \\ \dot{\theta}(kT+T) \end{pmatrix} = \begin{pmatrix} 1 & T_s & gT_s^2/2 & 0 \\ 0 & 1 & gT_s & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x(kT) \\ \dot{x}(kT) \\ \theta(kT) \\ \dot{\theta}(kT) \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ a_\theta T_s \\ 0 \end{pmatrix} \Big(\theta^*(kT)\Big) \tag{2.55}$$

Combining the attitude controller dynamics with the point mass model for $\mathbf{x}_y$, the dynamics for the lateral (y) axis can be expressed as

$$\mathbf{x}_y(kT_s + T_s) = \mathbf{A}_y\mathbf{x}_y(kT_s) + \mathbf{B}_y\mathbf{u}_y(kT_s) \tag{2.56}$$

$$\begin{pmatrix} y(kT_s+T_s) \\ \dot{y}(kT+T) \\ \phi(kT+T) \\ \dot{\phi}(kT+T) \end{pmatrix} = \begin{pmatrix} 1 & T_s & gT_s^2/2 & 0 \\ 0 & 1 & gT_s & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} y(kT) \\ \dot{y}(kT) \\ \phi(kT) \\ \dot{\phi}(kT) \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ a_\phi T_s \end{pmatrix} \Big(\phi^*(kT)\Big) \tag{2.57}$$

Combining the attitude controller dynamics with the point mass model for $\mathbf{x}_z$ the dynamics for the vertical (z) axis can be expressed as

$$\mathbf{x}_z(kT_s + T_s) = \mathbf{A}_z\mathbf{x}_z(kT_s) + \mathbf{B}_z\mathbf{u}_z(kT_s) \tag{2.58}$$

$$\begin{pmatrix} z(kT_s+T_s) \\ \dot{z}(kT+T) \\ \psi(kT+T) \\ \dot{\psi}(kT+T) \end{pmatrix} = \begin{pmatrix} 1 & T_s & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & T_s \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} z(kT) \\ \dot{z}(kT) \\ \psi(kT) \\ \dot{\psi}(kT) \end{pmatrix} + \begin{pmatrix} -T_s^2/2m & 0 \\ -T_s/m & 0 \\ 0 & a_\psi T_s^2/2 \\ 0 & a_\psi T_s \end{pmatrix} \begin{pmatrix} F_T^*(kT) \\ \dot{\psi}^*(kT) \end{pmatrix} + \begin{pmatrix} gT_s^2/2 \\ gT_s \\ 0 \\ 0 \end{pmatrix} \tag{2.59}$$

Combining the system dynamics into a singe state-space representation

$$\mathbf{x}(kT_s + T_s) = \mathbf{A}(kT_s)\mathbf{x}(kT_s) + \mathbf{B}(kT_s)\mathbf{u}(kT_s) + \mathbf{G} \tag{2.60}$$

$$\begin{pmatrix} \mathbf{x}_x(kT_s+T_s) \\ \mathbf{x}_y(kT_s+T_s) \\ \mathbf{x}_z(kT_s+T_s) \end{pmatrix} = \begin{pmatrix} \mathbf{A}_x(kT_s) \\ \mathbf{A}_y(kT_s) \\ \mathbf{A}_z(kT_s) \end{pmatrix} \begin{pmatrix} \mathbf{x}_x(kT_s) \\ \mathbf{x}_y(kT_s) \\ \mathbf{x}_z(kT_s) \end{pmatrix} + \begin{pmatrix} \mathbf{B}_x(kT_s) \\ \mathbf{B}_y(kT_s) \\ \mathbf{B}_z(kT_s) \end{pmatrix} \begin{pmatrix} \mathbf{u}_x(kT_s) \\ \mathbf{u}_y(kT_s) \\ \mathbf{u}_z(kT_s) \end{pmatrix} + \begin{pmatrix} \mathbf{G}_x \\ \mathbf{G}_y \\ \mathbf{G}_z \end{pmatrix} \tag{2.61}$$

The system dynamics presented above have some important implications for control. First, control over the full vehicle state is not possible. This is natural considering that all quadrotors are under-actuated. Direct control of movement along and about the vertical axis is possible using U_1 and U_4 via F_T^* and $\dot{\psi}^*$ respectively. Indirect control of movement along the lateral and longitudinal axis is possible by regulating the pitch and roll attitude using U_2 and U_3 via ϕ^* and θ^* respectively. In any case, the specific states to regulate will largely depend on the control task and available state estimates.

Second, as the resulting dynamics are essentially a linearisation of the true nonlinear dynamics about a stable equilibrium operating point (hover), a range of linear control techniques can be used to regulate platform motion. This includes Linear Quadratic Regulators (LQR) [281, 282], Model Predictive Control (MPC) [283]-[285] and simple Proportional-Integral-Derivative (PID) controllers [286, 287]. The techniques are well studied and lead to simple control architectures that can be effectively implemented on real platforms.

Lastly, the dynamics equations have been partitioned into three subsystems in the above derivation. Due to the decoupling of the system dynamics, it is possible to partition the dynamics in a number of other ways. This means that the control structure is not limited to a single controller, but instead can be comprised of multiple controllers to provide adequate control of the platform. This is an important consideration for the work presented in this thesis. This research involves, in part, the design of a high-level controller for a particular type of vision based collision avoidance. Full control of the quadrotor is not possible with the visual information available, forcing the consideration of such partitioned control schemes. The high-level control structure can then be divided into the vision-based and non vision-based controllers.

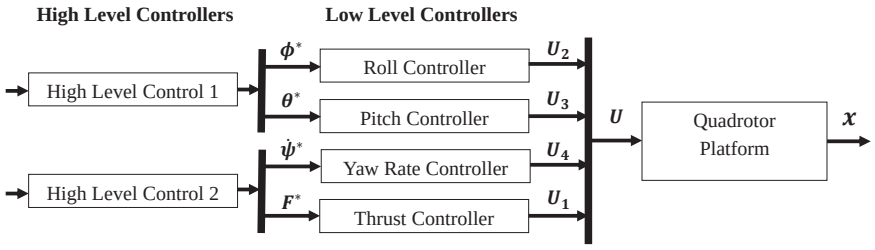

Figure 2.6: Generalised quadrotor control architecture where $\mathbf{x}^*$ denotes the reference input vector to the high level controller and Υ^* denotes the reference input vector to the low-level controller. The actual input to the quadrotor consists of motor speed commands through $U = (U_1, U_2, U_3, U_4)$.

a) Non-Visual Controllers

The non-visual controllers include simple linear Proportional-Integral-Derivative (PID) controllers and Linear Quadratic Regulator (LQR) controllers with integral action (LQRI) to control forward and lateral velocity.

PID Control: The PID controller does not require an explicit model of the quadrotor. Instead, the controller requires tuning of the three gain terms. The general form of the PID controller can be defined as

$$u^* = k_p e + k_i \int_0^{T_s} e \, dt + k_d \frac{de}{dt} \tag{2.62}$$

where k_p, k_i and k_d denote the proportional, integral and derivative gain terms respectively. The variable x is used in the general form of the controller but can be replaced by the controlled state. Consider the control of forward $\dot{x}$ and lateral $\dot{y}$ velocity. Given the decoupled nature of the quadrotor dynamics along the x and y axis, separate controllers can be derived for each state such that

$$\theta^* = k_p(\dot{x} - \dot{x}^*) + k_i \int_0^{T_s} (\dot{x} - \dot{x}^*)dt + k_d \ddot{x} \tag{2.63}$$

$$\phi^* = k_p(\dot{y} - \dot{y}^*) + k_i \int_0^{T_s} (\dot{y} - \dot{y}^*)dt + k_d \ddot{y} \tag{2.64}$$

The corresponding gain terms can be estimated using Ziegler-Nichols tuning and further refined empirically [287]. As the dynamics for forward and lateral velocity are approximate similar linear transfer functions, the PIDs for pitch and roll were tuned with the same parameters, such that

$$k_p = 0.25, \quad k_i = \frac{k_p}{t_s/3} = 0.25, \quad k_d = 0.01 \tag{2.65}$$

where t_s corresponds to a 2% settling time and the derivative gain was tuned manually. These parameters are used for the quadrotor system in Chapter 4, and are also given in Appendix E.

LQRI Control: The Linear Quadratic Regulator controllers with integral action (LQRI) require an augmented linear time-invariant (LTI) model of the dynamics about the x and

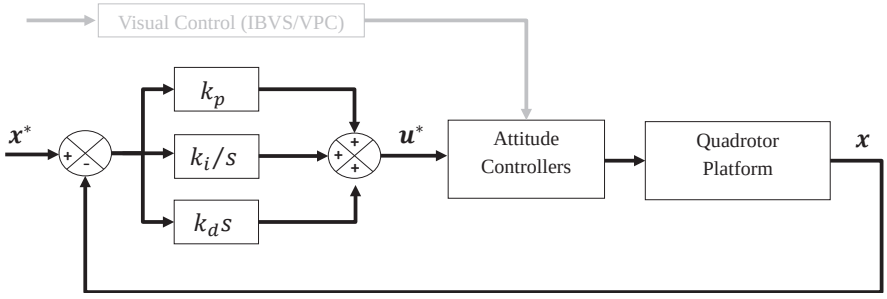

Figure 2.7: Generalised PID control architecture where $\mathbf{x}^*$ denotes the reference input vector to the PID controller and $\mathbf{u}^*$ denotes the reference input vector to the low-level attitude controllers.

y axis. Consider first the x axis dynamics. The position x is removed from the state vector $\mathbf{x}_x$ to get $\mathbf{x}_1$ such that

$$\dot{\mathbf{x}}_1(kT_s + T_s) = \mathbf{A}_1\mathbf{x}_1(kT_s) + \mathbf{B}_1\mathbf{u}_1(kT_s) \tag{2.66}$$

$$\begin{pmatrix} \dot{x}(kT_s+T_s) \\ \theta(kT_s+T_s) \\ \dot{\theta}(kT_s+T_s) \end{pmatrix} = \begin{pmatrix} 1 & gT_s & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \dot{x}(kT_s) \\ \theta(kT_s) \\ \dot{\theta}(kT_s) \end{pmatrix} + \begin{pmatrix} 0 \\ a_\phi T_s \\ 0 \end{pmatrix} \Big(\theta^*(kT_s) \Big) \tag{2.67}$$

The resulting state-space model is then augmented with an x velocity error term $\dot{x}_1^e$ such that

$$\begin{pmatrix} \dot{\mathbf{x}}_1(kT_s+T_s) \\ \dot{x}_1^e(kT_s+T_s) \end{pmatrix} = \begin{pmatrix} \mathbf{A}_1 & 0 \\ -\mathbf{C}_1 & 0 \end{pmatrix} \begin{pmatrix} \mathbf{x}_1(kT_s) \\ x_1^e(kT_s) \end{pmatrix} + \begin{pmatrix} \mathbf{B}_1 \\ 0 \end{pmatrix} \mathbf{u}_1(kT_s) + \begin{pmatrix} 0 \\ 1 \end{pmatrix} \mathbf{x}_1^* \tag{2.68}$$

which can be re-written as

$$\dot{\bar{\mathbf{x}}}_1(kT_s + T_s) = \bar{\mathbf{A}}\bar{\mathbf{x}}_1(kT_s) + \bar{\mathbf{B}}\mathbf{u}_1(kT_s) + \mathbf{x}_1^* \tag{2.69}$$

An optimal state feedback control $\mathbf{u}_1^\star(kT_s)$ can then be derived according to standard Linear Quadratic Regulator (LQR) design for LTI systems [279]. Consider a feedback controller of the form

$$\mathbf{u}_1^\star(kT_s) = \bar{\mathbf{K}}\bar{\mathbf{x}}_1(kT_s) \tag{2.70}$$

where the reference control has been omitted due to system linearisation about hover such that

$$\mathbf{u}_1^\star(kT_s) = -[\mathbf{K}\ \mathbf{K}^e][\mathbf{x}_1(kT_s)\ x_1^e(kT_s)]^T = -\mathbf{K}\mathbf{x}_1(kT_s) - \mathbf{K}^e x_1^e(kT_s) \tag{2.71}$$

The optimal constant gain matrix $\bar{\mathbf{K}}$ can be found by considering the minimisation of an infinite horizon quadratic cost function of the general form

$$J = \sum_{k=0}^{\infty} \bar{\mathbf{x}}(k)^T \bar{\mathbf{Q}}\bar{\mathbf{x}}(k) + \mathbf{u}(k)^T \mathbf{R}\mathbf{u}(k) \tag{2.72}$$

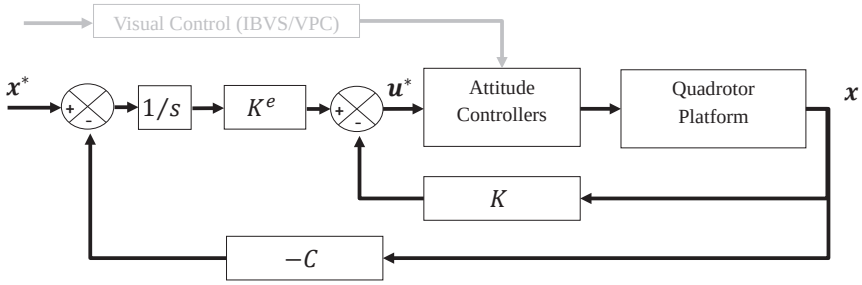

Figure 2.8: Generalised LQRI control architecture where $\mathbf{x}^*$ denotes the reference input vector to the LQRI controller and $\mathbf{u}^*$ denotes the reference input vector to the low-level attitude controllers.

where the weighting matrices $\mathbf{Q} \succeq \mathbf{0}$ and $\mathbf{R} \succeq \mathbf{0}$. The solution $\mathbf{P}$ to the associated (algebraic) Ricatti Equation

$$-\dot{\mathbf{P}} = \bar{\mathbf{A}}\mathbf{P} + \mathbf{P}\bar{\mathbf{A}} + \mathbf{C}^T\bar{\mathbf{Q}}\mathbf{C} - \mathbf{P}\bar{\mathbf{B}}\mathbf{R}^{-1}\bar{\mathbf{B}}^T\mathbf{P} = \mathbf{0} \tag{2.73}$$

can be used to find the gain such that

$$\bar{\mathbf{K}} = -\mathbf{R}^{-1}\bar{\mathbf{B}}^T\mathbf{P} \tag{2.74}$$

Including anti windup through amplitude saturation, the x velocity can be controlled via a reference pitch command such that

$$\theta^* = \mathbf{u}_1^*(t) = \max[a, \min(b, \mathbf{u}_1(t))] \tag{2.75}$$

where $\mathbf{u}_1^*$ denotes the saturated reference control input (1-vector) for the pitch attitude controller, and a and b are constants denoting the maximum and minimum angular displacements respectively. The y axis can be treated in the same manner to derive a similar state-space model and LQRI controller. In this case, the corresponding state vector is denoted by $\mathbf{x}_2$. The LQRI controllers are used for the quadrotor system in Chapter 5 and 6, and have been derived, tuned and validated using the comprehensive set of LQR design tools available in the MATLAB software package [70]. The resulting system models (parameters) and gain matrices are given in Appendix E.

b) Visual Controller

The visual control is based on both predictive and classic image-based visual servoing designs that rely primarily on visual observations to derive the control command. In both cases, the visual controllers provide the reference yaw rate and thrust to the low-level controller. The classical approach does not explicitly use the vehicle dynamics in the control scheme (Chapter 4). The predictive controller requires a dynamic model of the relevant vehicle states (Chapter 5). The details of the visual controllers are provided in Chapters 4 and 5, but the associated vehicle dynamics are presented here for reference and to avoid repetition. Removing the altitude from the state vector $\mathbf{x}_z$, the relevant

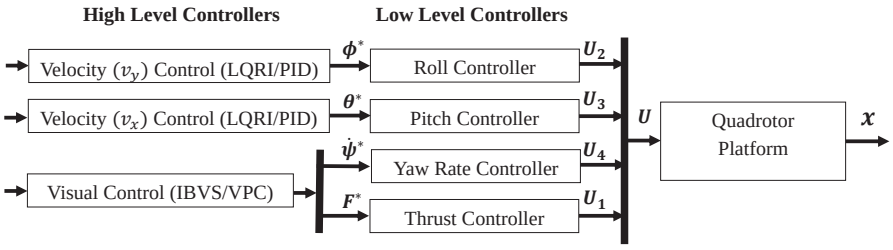

Figure 2.9: Generalised quadrotor control architecture where $\mathbf{x}^*$ denotes the reference input vector to the high level controller and Υ^* denotes the reference input vector to the low-level controller. The actual input to the quadrotor consists of motor speed commands through $U = (U_1, U_2, U_3, U_4)$.

state space model for the platform dynamics used by the visual control is given by

$$\dot{\mathbf{x}}_3(kT_s + T_s) = \mathbf{A}_3\mathbf{x}_3(kT_s) + \mathbf{B}_3\mathbf{u}_3(kT_s) \tag{2.76}$$

$$\begin{pmatrix} \dot{z}(kT+T) \\ \psi(kT+T) \\ \dot{\psi}(kT+T) \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & T_s \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \dot{z}(kT) \\ \psi(kT) \\ \dot{\psi}(kT) \end{pmatrix} + \begin{pmatrix} -T_s/m & 0 \\ 0 & a_\psi T_s^2/2 \\ 0 & a_\psi T_s \end{pmatrix} \begin{pmatrix} F_T^*(kT) \\ \dot{\psi}^*(kT) \end{pmatrix} + \begin{pmatrix} gT_s \\ 0 \\ 0 \end{pmatrix} \tag{2.77}$$

c) Controller Remarks and Considerations

▷ *Reference Frames:* The system states may be expressed in a local or world fixed frame. If the control task was based on position-based visual control, it would make sense to use dynamics expressed in the world frame. Conversely, for an image-based visual control task considered in this thesis it makes more sense to select dynamics expressed in a local frame. As such, the translational velocities are considered to be approximations to the quadrotor forward, sidewards and vertical velocity. This provides a more natural interpretation of the quadrotor motion in relation to image feature motion observed from a rigidly attached camera.

▷ *Experimental Considerations:* The high level controllers may be located onboard the platform, at a ground station or even accessed over a network. The choice largely depends on the application, platform hardware constraints and design stage. Positioning the controller at the ground station initially allows flexibility in development and testing, having only to manage the communications link to the platform. The onboard systems can then remain unchanged until the design matures and is ready for migration onto the platform avionics suite.

▷ *Simulation Considerations:* The time delay is excluded in the partitioned dynamic model presented above. Considering optimal controllers such as Model Predictive Control, it may be useful to omit the time delay in the dynamic model used for state propagation (prediction) within the controller. For simulations and validation however, including the time delay in the plant model provides a realistic environment in which to evaluate any control scheme. In this way, the model used in the controller and that representing the real system in simulation are different. This simulation approach is adopted in this thesis.

2.3 Spherical Imaging

2.3.1 Camera Model

Spherical cameras are useful for many visual surveillance and navigation applications in which a large field of view is required from a single device. The imaging surface is represented by a sphere such that a 4π steradian field of view is possible in the ideal case. Real spherical cameras are often created using a set of perspective or wide angle cameras positioned in an array such that each cameras field of view constitutes a section of the sphere. The images are then artificially stitched together to make up the spherical imaging surface. As such, most practical spherical cameras cannot represent a complete spherical image with current devices capable of up to 90% coverage [83].

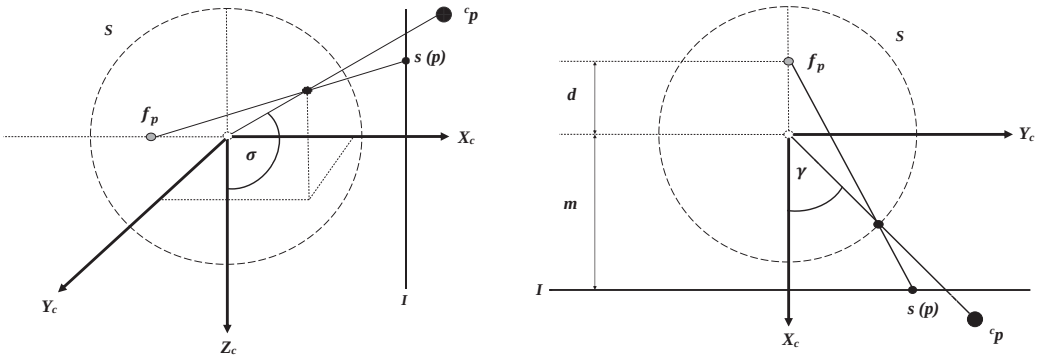

Figure 2.10: Simplified spherical projection model and ideal spherical imaging. (a) Azimuth angle re-projection (b) Colatitude angle re-projection

To obtain the image feature representation of an arbitrary projective point in space, a number of steps are required. The point is first mapped to the surface of a unit sphere S then re-projected via S onto an imaging plane I. The exact mapping depends on the cameras intrinsic parameters including focal point of re-projection and mirror type, and can be described using the unified imaging model [82]. Applying the model, an optical axis on the sphere is first chosen. A re-projection point f_p along this axis and contained within the sphere is then defined. The imaging plane is offset from the spherical centre by a distance m and positioned perpendicular to the optical axis. The spherical centre is then considered to be the origin of the camera frame. A simplified representation of the spherical camera geometry and the unified imaging model is depicted in Fig 2.10. Here, the x axis is chosen as the optical axis and the re-projection point is given by $f_p(-d, 0, 0)$. A point defined in the camera frame ${}^c\mathbf{p}(X, Y, Z)$ corresponds to the image feature representation $\mathbf{p}(x, y, z)$ where

$$\mathbf{p}(x, y, z) = (-m, \frac{Y(d+m)}{dr - X}, \frac{Z(d+m)}{dr - X}) \tag{2.78}$$

and $r^2 = X^2 + Y^2 + Z^2$ is the square of the distance between the projective point and the spherical centre. Antipodal points resulting from the use of a spherical surface are not

considered in (2.78), but would result in some minor sign changes. A more generic model can then be defined as

$$\mathbf{p}(x,y,z) = (-m, \pm\frac{Y(d+m)}{dr \mp X}, \pm\frac{Z(d+m)}{dr \mp X}) \tag{2.79}$$

The unified imaging model can also be used to map image features measured using other camera types, including wide angle and perspective, to an ideal imaging sphere [306]. As an example, consider a perspective camera with x as the optical axis and a focal length $f = 1$. Assuming the spherical centre and re-projection point are aligned such that $f_p(0,0,0)$ and the imaging plane lies tangential to the unit sphere such that $m = f = 1$, (2.78) becomes

$$\mathbf{p}(x,y,z) = (-1, \frac{Y}{X}, \frac{Z}{X}) \tag{2.80}$$

which is the model for a perspective camera. The application of the unified imaging model to perspective and wide angle imaging is important for practical implementations. A spherical camera may not be available, but the corresponding image feature representations may be required or preferred.

The image features have previously been defined in Cartesian coordinates $\mathbf{p}(x,y,z)$ but can also be represented in spherical coordinates $\mathbf{s}(\sigma,\gamma)$, providing a more natural interpretation of the image feature position when using spherical imaging. A projective point in the camera frame can be expressed in spherical coordinates as

$${}^c\mathbf{p}(X,Y,Z) = (r\sin\sigma\cos\gamma,\ r\sin\sigma\sin\gamma,\ r\cos\sigma) \tag{2.81}$$

where the azimuth $\gamma \in [-\pi, \pi)$ and colatitude $\sigma \in [0, \pi)$ angles are measured from the spherical centre. The angles can be considered similar to measuring the longitude and latitude of a point on a sphere, where a constant colatitude defines a great circle. Using the camera frame coordinates, the spherical image features can be defined as

$$\mathbf{s}(\sigma,\gamma) = (\arccos(Z/r), \arctan(Y/X)) \tag{2.82}$$

Using the Cartesian representation of the image features, $x^2 + y^2 + z^2 = 1$ and $\mathbf{p} = {}^c\mathbf{p}/r$ by virtue of the unit spherical imaging surface. The constraint results in a redundant coordinate such that the spherical image feature angles can then be expressed as

$$\mathbf{s}(\sigma,\gamma) = (\arccos(z), \arctan(y/x)) \tag{2.83}$$

The perpendicular distance from the image feature to the camera z axis is denoted by $\bar{r}$, and can also be used to determine the spherical image features such that $\sigma = \arcsin\bar{r}$. This can be inferred using Fig 2.11.

Combining the unified imaging model and the spherical image feature representation, a perspective camera could be used to represent a partial section of the spherical image. The particular section of the sphere will depend on the relative orientation of the camera optical axis with respect to the vehicle or body frame. The image features measured

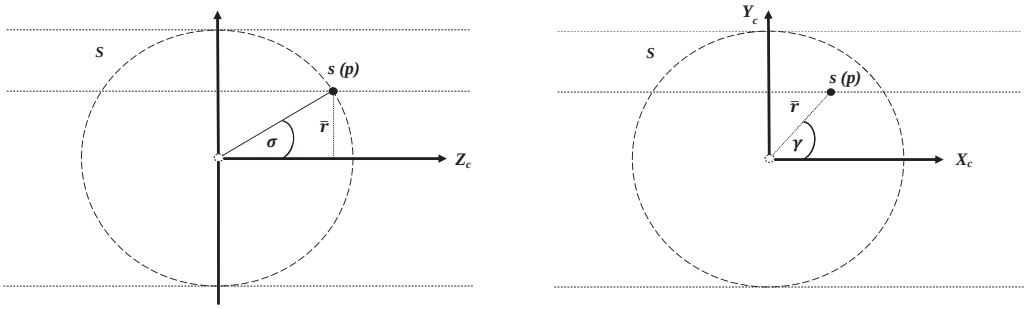

(a) (b)

Figure 2.11: Spherical image feature measurements (a) Colatitude measurement (b) Azimuth measurement

from the perspective camera $\mathbf{f}(u, v)$ could then be used to approximate the spherical images features $\mathbf{s}(\sigma, \gamma)$ in that section of the sphere. For example, assuming again that $f_p(0, 0, 0)$ coincides with the focal point of the perspective camera and assuming $m = f$, the spherical image features can be approximated as

$$\mathbf{s}(\sigma, \gamma) \approx (\arctan\left((v - v_m)/f\right) + \sigma_o, \arctan((u - u_m)/f) + \gamma_o) \tag{2.84}$$

where (u_m, v_m) denotes the centre coordinate of the perspective image and σ_o and γ_o denote fixed offset angles. The angles are included to account for camera positioning such that the correct section of the sphere is being represented. If the optical axis is switched to Z_c instead of X_c, then the spherical angles can be approximated as

$$\mathbf{s}(\sigma, \gamma) \approx (\arctan\left((v - v_m)/f\right) + \sigma_o, \arctan((v - v_m)/(u - u_m)) + \gamma_o) \tag{2.85}$$

Assuming X_c defines the optical axis and the camera and body frames are initially aligned, if the perspective camera is then rotated about the Z_c axis by ψ_c, (2.84) becomes

$$\mathbf{s}(\sigma, \gamma) \approx (\arctan\left((v - v_m)/f\right) + \pi/2, \arctan((u - u_m)/f) + \psi_c) \tag{2.86}$$

For practical implementations, this thesis makes use of perspective imaging and (2.86) to approximate a spherical camera section. A perspective camera is positioned on an aerial vehicle such that the camera is offset from the direction of forward motion to ensure coverage of key angular position important for collision avoidance. The particular field of view will be made clearer in Chapter 3 through a discussion on conical spiral motion.

2.3.2 Image Kinematics

The image kinematics define the relationship between image feature velocity $\dot{\mathbf{p}}$ or $\dot{\mathbf{s}}$ and camera velocity $\dot{\mathbf{x}} \in \mathbb{R}^6$. The camera velocity includes the translational τ and angular velocity ω components such that $\dot{\mathbf{x}} = (\mathbf{v}\ \omega)$. Using the traditional approach used for perspective imaging [288], a point in the camera frame ${}^c\mathbf{p}$ has a velocity ${}^c\dot{\mathbf{p}}$ defined by

$${}^c\dot{\mathbf{p}} = -\tau - \omega \times {}^c\mathbf{p} \tag{2.87}$$

which in scalar form becomes

$$\dot{X} = -v_x Z - \omega_y + Y\omega_z \tag{2.88}$$
$$\dot{Y} = -v_y X - \omega_z + Z\omega_x \tag{2.89}$$
$$\dot{Z} = -v_z Y - \omega_x + X\omega_y \tag{2.90}$$

Taking the derivative of the spherical image features defined in (2.82)

$$\dot{\sigma} = \frac{ZX\dot{X} + ZY\dot{Y} + Z^2\dot{Z} - r^2\dot{Z}}{r^3 \sin\sigma} \tag{2.91}$$
$$\dot{\gamma} = \frac{X\dot{Y} - Y\dot{X}}{r^2 \sin^2\sigma} \tag{2.92}$$

Substituting equations (2.88)-(2.90) into (2.91)-(2.92) and simplifying

$$\dot{\sigma} = \frac{-ZXv_x - ZYv_y - (Z^2 - r^2)v_z + r^3 Y\omega_x - r^2 X\omega_y}{r^3 \sin\sigma} \tag{2.93}$$
$$\dot{\gamma} = \frac{Yv_x - Xv_y + XZ\omega_x + YZ\omega_y - (X^2 + Y^2)\omega_z}{r^2 \sin^2\sigma} \tag{2.94}$$

Substituting equations (2.81) into (2.93)-(2.94) and rearranging, the spherical image feature velocities become

$$\dot{\sigma} = -\frac{v_x \cos\gamma\cos\sigma}{r} - \frac{v_y \sin\gamma\cos\sigma}{r} + \frac{v_z \sin\sigma}{r} + \omega_x \sin\gamma - \omega_y \cos\gamma \tag{2.95}$$
$$\dot{\gamma} = \frac{v_x \sin\gamma}{r\sin\sigma} - \frac{v_y \cos\gamma}{r\sin\sigma} + \frac{\omega_x \cos\gamma\cos\sigma}{\sin\sigma} + \frac{\omega_y \sin\gamma\cos\sigma}{\sin\sigma} - \omega_z \tag{2.96}$$

Rewriting in matrix form, the image feature velocity $\dot{\mathbf{s}}$ is given by

$$\dot{\mathbf{s}} = \mathbf{L}_s(\mathbf{s}, r)\ \dot{\mathbf{x}} \tag{2.97}$$

where $\mathbf{L}_s$ is the spherical image Jacobian[1] defined as

$$\mathbf{L}_s = \begin{pmatrix} -\cos\sigma\cos\gamma/r & -\cos\sigma\sin\gamma/r & \sin\sigma/r & \sin\gamma & -\cos\gamma & 0 \\ \sin\gamma/r\sin\sigma & -\cos\gamma/r\sin\sigma & 0 & \cos\gamma\cos\sigma/\sin\sigma & \sin\gamma\cos\sigma/\sin\sigma & -1 \end{pmatrix} \tag{2.98}$$

The optic flow equation can be represented in discrete time using small disturbance theory such that

$$\mathbf{s}(kT_s + T_s) = \mathbf{s}(kT_s) + T_s \bar{\mathbf{L}}_s \dot{\mathbf{x}}(kT_s) \tag{2.99}$$

where T_s denotes a sufficiently small sampling time and $\bar{L}_s$ denotes an approximation of the image Jacobian at $t = kT_s$. The optic flow equation given by (2.97) is important for visual control. The camera is often rigidly attached to a vehicle so the equation can be used to describe image feature motion as a result of vehicle motion. The observed feature motion can then be used in estimation schemes or directly to derive control feedback.

[1] Also referred to as the Interaction Matrix or Feature Jacobian.

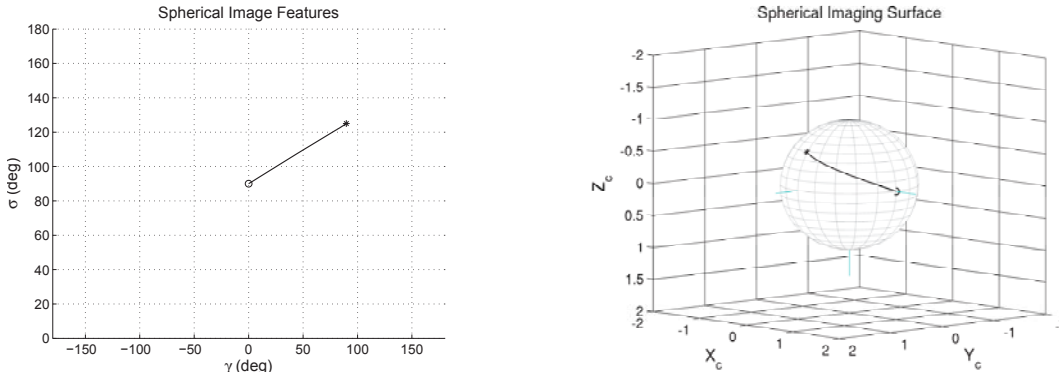

(a) (b)

Figure 2.12: Example image feature trajectory. (a) Image feature trajectory in Cartesian representation (b) Image feature trajectory on spherical imaging surface

Some example image feature motion on a planar representation of the spherical surface and the unit sphere itself is shown in Fig 2.12.

Consider now that the point object moves with a small constant velocity of magnitude v_t. If the object is far away and its motion is initially unknown and cannot be estimated, the resulting optic flow equation for the dynamic object can be expressed as

$$\dot{\mathbf{s}} = \mathbf{L}_s(\mathbf{s}, r)\, \dot{\mathbf{x}} + \frac{\partial s}{\partial t} \tag{2.100}$$

where $\frac{\partial s}{\partial t}$ represents the contribution of the object motion to the image feature velocity. The rotational motion of the object is negligible for a point feature and the only contribution to the image feature motion is a result of the object translational motion. As the projection of the translational motion is proportional to the object range such that $\frac{\partial s}{\partial t} \propto \frac{1}{r}$. The further away the object, the less influence its motion will have on the measured image feature velocity. As such, for a small time period (or horizon), a distant object will move very little in the image. This is an important consideration in this work, given the object of interest is considered to be relatively far away from the camera.

Chapter 3

Spiral Avoidance

3.1 Outline

This chapter presents an **avoidance** and **resolution decision** strategy for static and dynamic objects using only angular observations, the properties of conical spiral motion and the current right-of-way rules for manned aviation.

First, the properties of conical spirals are introduced. This includes how they can be visually tracked and the implications for avoidance decisions. Specifically, conical spirals are presented as a viable avoidance trajectory for image-based control. Second, a conservative threshold based avoidance decision strategy (logic) is presented. Aviation right-of-way rules, in conjunction with a unique representation of the observed image feature motion, are used to appropriately select a set of reference image feature positions that implicitly forces an avoidance spiral trajectory. Third, a threshold-based resolution decision concept to stop tracking the reference spiral is revised in the context of See and Avoid. Each of these elements are then brought together to define the complete vision-based avoidance and resolution decision strategy.

3.2 Conical Spirals

3.2.1 Geometry Basics

Conical spirals have been called logarithmic, equiangular and conchospirals with each name describing a particular characteristic of the resulting curve. Frequently occurring in nature and broadly studied by mathematicians, artists and scientists, they describe the set of trajectories that circumscribe the surface of a cone [289]-[291]. The idea can be reduced to two dimensions, resulting in planar spirals.

Mathematically, conical spirals require minimal parametric representation. Consider a fixed object reference frame $\mathcal{F}_o$ with its origin positioned at the apex of a stationary cone, and its z axis along the cone axis as shown in Fig 3.1. The elevation angle β defines the angle between the range vector to the apex and the positive vertical axis of the cone, such that $0 < \beta < \pi$. When $\beta \in \{0, \pi\}$ the apex is situated directly above or below the platform, and a degenerate case exists such that a conical spiral cannot be defined. The bearing angle α is defined as the angle between the projection of the range vector and platform velocity vector onto the xy-plane such that $-\pi < \alpha < \pi$. When $\alpha \in \{0, \pi, -\pi\}$ the apex is situated directly in front or behind the platform, and again a degenerate case

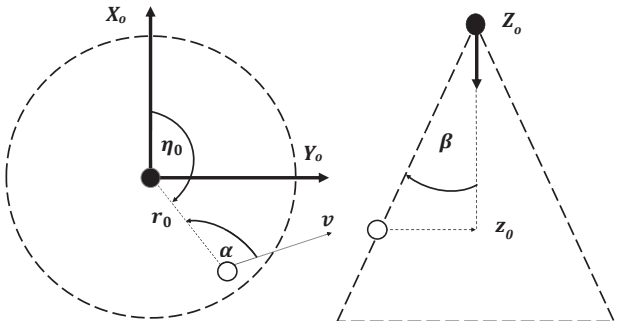

Figure 3.1: Top view and side view of a reference cone (*dashed*), its apex (*solid black dot*) and parameters defining conical spiral motion.

results such that a conical spiral is not defined. The speed is defined as the magnitude of the translational velocity vector $\dot{\mathbf{x}}_t$ such that $v = \| \dot{\mathbf{x}}_t \| = (v_x^2 + v_y^2 + v_z^2)^{\frac{1}{2}}$. The initial range and vertical displacement from the apex are defined by r_0 and z_0 respectively. The initial angular position from the apex measured from the positive x axis is defined by η_0. Importantly, fixing the conical angles $\mathbf{c}$ to a reference value such that $\mathbf{c}^* = [\beta^* \ \alpha^*]$ results in a particular spiral type, be it convergent, divergent or circular. Changing the remaining parameters simple changes the position at which that particular spiral type commences. The conical angles can then be used to represent the relative position of a maneuvering aircraft with respect to a static or dynamic apex.

3.2.2 Tracking Spirals

In order to visually track or attempt to track a conical spiral, the appropriate conical angles must be regulated directly from visual observations obtained from the unmanned aircraft. A relationship between the conical angles $\mathbf{c}$ and the spherical image features $\mathbf{s}$ is then required.

Consider a fixed object reference frame $\mathcal{F}_o$, with origin at the conical apex and orientation aligned to a fixed world frame $\mathcal{F}_w$. Now consider a spiral reference frame $\mathcal{F}_s$, attached to an aircraft moving along a conical spiral trajectory. The origin of the spiral reference frame and aircraft body frame $\mathcal{F}_b$ are initially aligned. The spiral frame inherits the body translational motion, but only its angular velocity in the z axis. Therefore, the orientation of the object frame $\mathcal{F}_o$ with respect to the spiral frame $\mathcal{F}_s$ can be defined as a rotation ${}^sR_o$ about the vertical axis of the object frame such that

$$ {}^sR_o = R(\psi) \tag{3.1} $$

where ψ denotes the vehicle yaw angle. Note, that by this definition $\alpha \neq \psi$, and is best represented in the spiral frame. Additionally, by the definition above, the relationship between the elevation angle in the spiral and object frames is given by

$$ \beta'(t) = \pi - \beta(t) \tag{3.2} $$

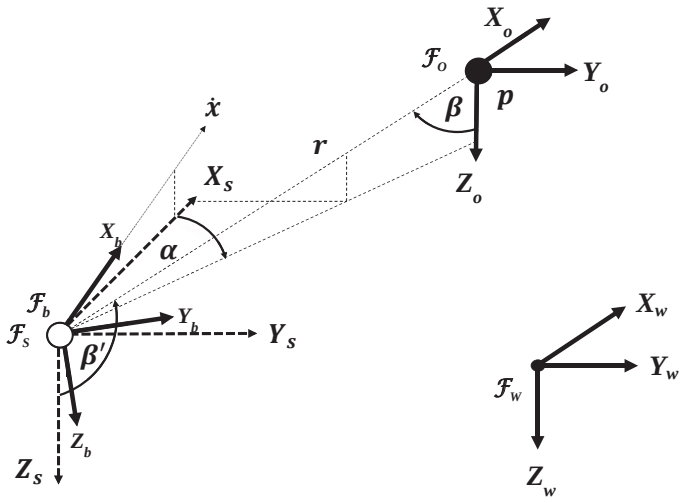

Figure 3.2: World (w), object (o), spiral (s) and body (b) coordinate frames and elevation β' and azimuth α' angles defined in the spiral frame.

where a dash denotes the elevation angle measured in the spiral frame. An aircraft can then track a conical spiral by regulating β' and α using visual observations. The relevant coordinate frames and associated angular measurement are depicted in Fig 3.3.

The introduction of the spiral reference frame allows the derivation of a convenient relationship between the image features and conical angles. Specifically, consider a spherical camera rigidly attached to an aircraft observing a point object, considered to be the conical apex. In the camera frame, the apex is denoted by ${}^c\mathbf{p}$ such that

$$ {}^c\mathbf{p}(t) = \begin{pmatrix} X(t) \\ Y(t) \\ Z(t) \end{pmatrix} = \begin{pmatrix} r(t)S\sigma(t)C\gamma(t) \\ r(t)S\sigma(t)S\gamma(t) \\ r(t)C\sigma(t) \end{pmatrix} = r(t){}^c\tilde{\mathbf{p}}(t) \tag{3.3} $$

where $C = cos(\cdot)$, $S = \sin(\cdot)$, $r(t)$ is the distance to the object and a tilde denotes a unit vector. Recalling the origin of the body and spiral frames to be equal, then in the spiral reference frame equation (3.3) becomes

$$ {}^s\mathbf{p}(t) = {}^sR_b(t){}^bR_c{}^c\mathbf{p}(t) + {}^s\mathbf{t}_c \tag{3.4} $$

$$ r(t){}^s\tilde{\mathbf{p}}(t) = r(t){}^sR_b(t){}^bR_c{}^c\tilde{\mathbf{p}}(t) + {}^s\mathbf{t}_c \tag{3.5} $$

where ${}^sR_b(t)$ and ${}^bR_c$ define rotation matrices from body to spiral and camera to body frames respectively. The camera focal point and origin of the body frame are separated by the vector ${}^b\mathbf{t}_c$, which has the same magnitude as ${}^s\mathbf{t}_c$. Dividing through by $r(t)$ and because $r(t) \gg \|{}^s\mathbf{t}_c\|$

$$ {}^s\tilde{\mathbf{p}}(t) = {}^sR_c(t){}^c\tilde{\mathbf{p}}(t) + \epsilon_s \tag{3.6} $$

where ${}^sR_c(t) = {}^sR_b(t){}^bR_c$ and $\epsilon_s = {}^s\mathbf{t}_c/r(t)$ is a small positional offset. As the object range is expected to be at least two orders of magnitude greater than ${}^s\mathbf{t}_c$, then $\|{}^o\mathbf{t}_c\| \approx \|{}^o\mathbf{t}_s\| \approx r(t)\ \forall\, t$ such that ${}^s\mathbf{t}_c \approx \mathbf{0}$ and $\epsilon_s \approx 0$. A simple simulation example can show that

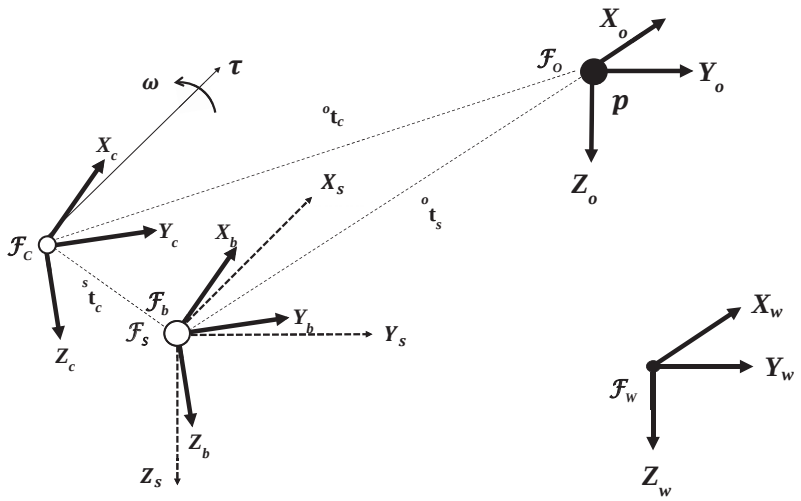

Figure 3.3: World (w), spiral (s), body (b) and camera (c) coordinate frames. Note, the camera frame displacement from the body frame has been exaggerated for clarity. It is typically an order of magnitude less than the body displacement from the object.

$\sigma \sim \mathcal{N}(0.05, 0.002)$ and $\gamma \sim \mathcal{N}(2.9, 0.37)$ measured in degrees, while $r \sim \mathcal{N}(0.007, 0.005)$ over the entire imaging surface. As such, image noise will likely dominate any uncertainty, and the assumption is justified (see Appendix B). Substituting for $\mathbf{p}$ using spherical coordinates then

$$\begin{pmatrix} S\beta'(t)C\alpha(t) \\ S\beta'(t)S\alpha(t) \\ C\beta'(t) \end{pmatrix} = {}^{s}R_c \begin{pmatrix} S\sigma(t)C\gamma(t) \\ S\sigma(t)S\gamma(t) \\ C\sigma(t) \end{pmatrix} \tag{3.7}$$

where

$$ {}^{s}R_c(t) = R(\theta)\, R(\phi) \tag{3.8}$$

$$ {}^{s}R_c(t) = \begin{pmatrix} C\theta(t) & S\theta(t)S\phi(t) & C\phi(t)S\theta(t) \\ 0 & C\phi(t) & -S\phi(t) \\ -S\theta(t) & C\theta(t)S\phi(t) & C\theta(t)C\phi(t) \end{pmatrix} \tag{3.9}$$

and $\theta(t)$ and $\phi(t)$ denote camera pitch and roll angles respectively. The conical angles can then be derived by solving (3.7) using the measured spherical image features and aircraft orientation. Depending on the aircraft dynamics, the reference spiral will typically be followed with non-zero pitch and roll. This is certainly the case for fixed wing aircraft, in which coordinated turns require non-zero yaw and roll angle, and climbing and descending require non-zero pitch. A reference pitch θ^* and roll ϕ^* angle are required according to

$$\sin\theta^* = -\cos\beta^* \cos\alpha^*, \quad \sin\phi^* = -\cos\beta^* \sin\alpha^* \tag{3.10}$$

In this case, the reference image features and reference conical angles would not be equal. However, by using de-rotated images it is reasonable to assume $\theta(t) \approx \phi(t) \approx 0\ \forall t$ so ${}^{s}R_c = I_3$ and $\mathbf{s}^* \approx \mathbf{c}^*$. In this way, the problem is generalised for any aircraft type. Reference image features that directly correspond to a particular conical spiral could first be selected. The spiral can then be tracked by directly regulating these image features as

the image kinematics are known (see §2.3.2). For the remaining sections in this chapter de-rotated images are assumed to be available, so the conical angles and image features are equal and can be used interchangeably.

3.2.3 Static Apex

For a static object, consider a spiral frame $\mathcal{F}_s$ attached to an aircraft moving with respect to the object reference frame $\mathcal{F}_o$. The object centre coincides with the origin of $\mathcal{F}_o$ and does not rotate. The object reference frame can be viewed like a fixed world frame $\mathcal{F}_w$, whose origin is shifted according to the object position. The orientation of $\mathcal{F}_s$ is controlled using a yaw angular velocity command u. The range r denotes the distance between the origins of the two frames. The heading of the vehicle ψ is measured with respect to the world frame and its speed is defined by v. An example of the geometry is shown in Fig 3.4 where $\mathcal{F}_s$ moves in the $xy-$plane of $\mathcal{F}_o$ such that $\beta = \pi/2$.

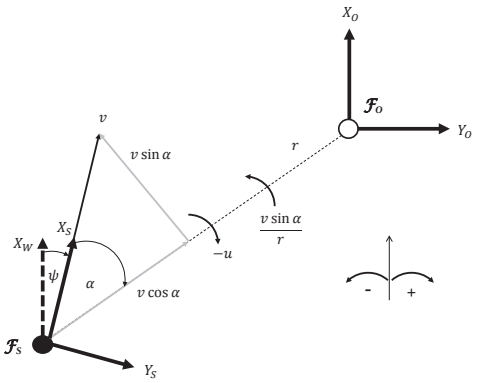

Figure 3.4: Example spiral geometry for a static obstacle (apex)

The nonlinear system of ordinary differential equations describing the vehicle motion then becomes

$$\dot{r}(t) = -v\cos\alpha(t) \tag{3.11}$$

$$\dot{\alpha}(t) = \frac{v\sin\alpha(t)}{r(t)} - u(t) \tag{3.12}$$

$$\dot{\psi}(t) = u(t) \tag{3.13}$$

Provided the sampling time T_s is sufficiently small, the system can be approximated in discrete time as

$$r_{k+1} = r_k - T_s v\cos\alpha_k \tag{3.14}$$

$$\alpha_{k+1} = \alpha_k + T_s\left(\frac{v\sin\alpha_k}{r_k} - u_k\right) \tag{3.15}$$

$$\psi_{k+1} = \psi_k + T_s u_k \tag{3.16}$$

Consider the case where the azimuth angle α is constant $\forall\, k$. Based on the equations above, it can be deduced that for $\| \alpha \| = \pi/2$ the range is constant, whereas for $\| \alpha \| > \pi/2$ the range increases. For $\| \alpha \| < \pi/2$ the range decreases, with $\alpha = 0$ denoting a degenerate

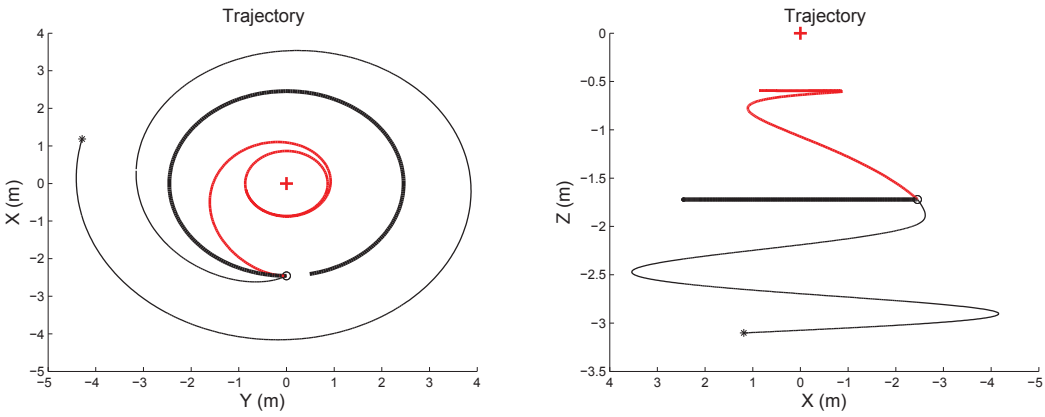

(a) xy-plane (b) xz-plane

Figure 3.5: Example trajectories of $\mathcal{F}_o$ in $\mathcal{F}_s$ for a static object with $\alpha = \pi/2$ (−), $\alpha > \pi/2$ (−) and $\alpha < \pi/2$ (−). Initial aircraft (○) and object (+) positions in the world frame $\mathcal{F}_w$ are $(0, 1)$ and $(1, 1)$ respectively. The aircraft moves with velocity $v = 1\text{m/s}$

case whereby motion is directly toward the object. Lateral separation can then be assured for all $\| \alpha \| \geq \pi/2$, if the correct avoidance direction is selected. To this end, it is trivial to see that correct avoidance direction and maximum lateral separation will occur (no crossing), by selecting the sign of the reference spiral azimuth α^* to be the same as that of the azimuth angle upon initial acquisition α_0. Specifically,

$$\alpha^* \geq \pi/2, \qquad 0 \leq \alpha_0 \leq \pi/2 \tag{3.17}$$

$$\alpha^* \leq -\pi/2, \quad -\pi/2 \leq \alpha_0 < 0 \tag{3.18}$$

Now consider the case where $\beta \neq \pi/2$. The resulting trajectory will be a 3 dimensional spiral that circumvents the surface of a cone whose opening angle is defined by β. For $\beta < \pi/2$, the object will be above the aircraft. For $\beta > \pi/2$, the object will be below the aircraft. As such, vertical separation from the object already exists or is inherently assured in each case. Similar to horizontal separation, the maximum vertical separation (no crossing) will occur by selecting the reference spiral elevation angle β^* such that

$$\beta^* > \beta_0, \qquad \pi/2 \leq \beta_0 < \pi \tag{3.19}$$

$$\beta^* < \beta_0, \qquad 0 < \beta_0 < \pi/2 \tag{3.20}$$

where β_0 is the elevation angle upon initial object acquisition. Some example trajectories are illustrated in Fig 3.5 for planar spirals about a static apex or object. Although the choice of which spiral direction to adopt can be determined using geometric intuition, it would also be possible to setup and solve an appropriate optimal control problem similar to [220] to find the appropriate direction. For reference, a simplified bang-bang optimal control version of the above problem is setup and presented in Appendix C.

3.2.4 Dynamic Apex

For dynamic objects, the object reference frame $\mathcal{F}_o$ is now moving with a constant velocity v_t, inertial heading $\Phi \in (0, 2\pi)$ and relative heading $\bar{\alpha} \in (0, 2\pi)$. Similar to the static object case, the orientation of $\mathcal{F}_s$ can be controlled using a yaw angular velocity command u, as the object does not maneuver. An example of the geometry is shown in Fig 3.6 where $\mathcal{F}_s$ moves in the $xy-$plane of $\mathcal{F}_o$ such that $\beta = \pi/2$.

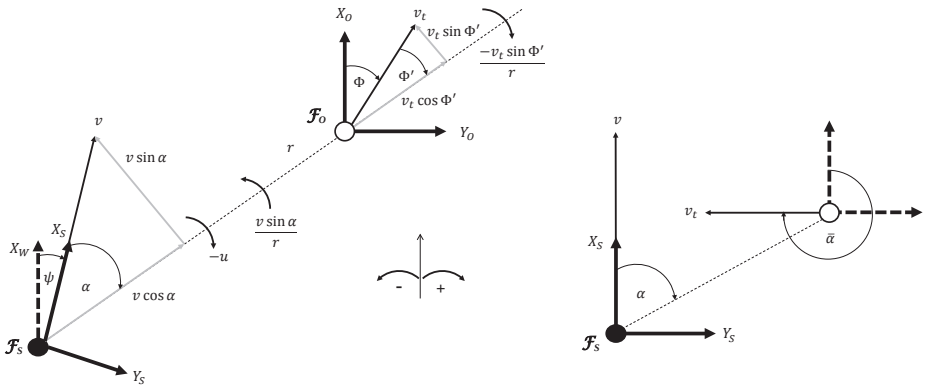

Figure 3.6: Example spiral geometry for a dynamic obstacle (apex)

The nonlinear system of ordinary differential equations describing the motion is given by

$$\dot{r}(t) = -v\cos\alpha(t) + v_t\cos(\psi(t) + \alpha(t) - \Phi) \tag{3.21}$$

$$\dot{\alpha}(t) = \frac{v\sin\alpha(t) - v_t\sin(\psi(t) + \alpha(t) - \Phi)}{r(t)} - u(t) \tag{3.22}$$

$$\dot{\psi}(t) = u(t) \tag{3.23}$$

which in discrete time can be approximated as

$$r_{k+1} = r_k + T_s(v_t\cos(\psi_k + \alpha_k - \Phi) - v\cos\alpha_k) \tag{3.24}$$

$$\alpha_{k+1} = \alpha_k + T_s\left(\frac{v\sin\alpha_k - v_t\sin(\psi_k + \alpha_k - \Phi)}{r_k} - u_k\right) \tag{3.25}$$

$$\psi_{k+1} = \psi_k + T_s u_k \tag{3.26}$$

Consider the case where the azimuth angle α is constant $\forall\, k$. Based on the equations above, the resulting spiral trajectory for a dynamic object now depends not only on α, but on the relative velocity and heading. As such, it is difficult to analyse the range rate and uniquely define the set of resulting spirals in terms of converging, diverging or circular as in the static case. In some cases, the spiral may initially diverge then converge upon attempting to circumvent the moving object. Additionally, it may be such that the intended reference spiral cannot be sustained given the relative dynamics, further complicating the analysis.

Recalling that the focus is on determining if a spiral can be used as a viable avoidance trajectory, then an alternate approach is to consider analysing the nature of the aircraft motion as it attempts to follow a reference spiral. The word nature is used to denote

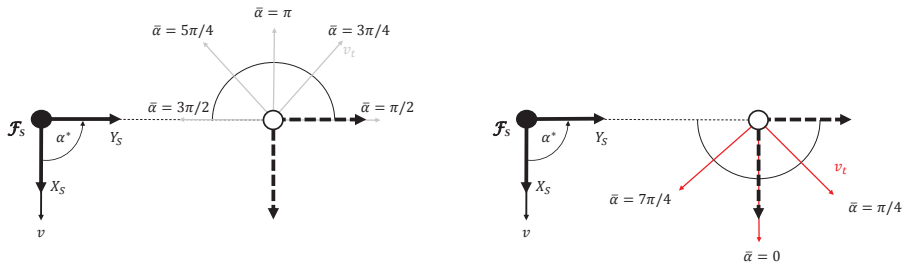

(a) Safe relative headings

(b) Unsafe relative headings

Figure 3.7: Example initial relative headings $\bar{\alpha}$ for safe and unsafe spiral trajectories for a reference spiral $\mathbf{c} = [\pi/2 \ -\pi/2]$. The object ($\circ$) and aircraft ($\bullet$) initial position are shown along with multiple safe ($-$) and unsafe ($-$) object trajectories

a qualitative description of the aircraft's ability to track a spiral and, its tendency to initially pass in front or behind the object.

Leveraging the results from the static case, consider attempting to follow a spiral such that $\alpha^* = -\pi/2$ for an arbitrary relative velocity. Consider an initial relative object heading $\bar{\alpha}_0 \in (0, 2\pi)$ such that the analysis is generalised for any initial aircraft ψ_0 and fixed object Φ heading. A simple inspection of the relative geometry suggests two types of spiral trajectory. For $\pi/2 \leq \bar{\alpha} \leq 3\pi/2$, the aircraft will attempt to pass in front of the object when $\psi(t) \geq \psi_0 + \pi$. Essentially, the aircraft has to completely turn around before passing in front of the object such that at least half a spiral has been followed. For $\bar{\alpha} < \pi/2$ or $\bar{\alpha} > 3\pi/2$, the aircraft will attempt to pass the object when $\psi(t) < \psi_0 + \pi$. This means that the aircraft will try to move in front of the object immediately, before half a spiral has been followed. Depending on the relative velocity, this may move the aircraft directly toward the object, resulting in collision. As such, although there is no guarantee that the range will not decrease as in the static case, initial lateral avoidance can be assured in some cases. Specifically,

$$\alpha^* = -\pi/2, \quad \pi/2 \leq \bar{\alpha}_0 \leq 3\pi/2 \tag{3.27}$$

For all other relative headings, collision avoidance may still be possible by attempting to follow a reference spiral where $\alpha^* = -\pi/2$. However, avoidance is not assured and the initial relative geometry is such that a collision is more likely at a sooner instance. Some example cases for safe and unsafe initial relative headings using $\alpha^* = -\pi/2$ are depicted in Fig 3.7. Of note, a similar analysis can be conducted for $\alpha^* = \pi/2$ and $\beta^* = \pi/2$. In this case, the safe and unsafe relative headings would be different and the trajectories would rotate about the object in the opposite direction.

The above analysis can be extended to any relative velocity. As evidence, consider an object moving with constant velocity v_t and constant heading Φ. An aircraft then attempts to follow a reference spiral about the object, with $\alpha^* = -\pi/2$ and $\beta^* = \pi/2$. The aircraft is moving with velocity $v = 1\text{m/s}$ and in the same plane as the object. Consider two cases in which the object moves with $v_t > v$ and $v_t < v$. For each velocity, the object

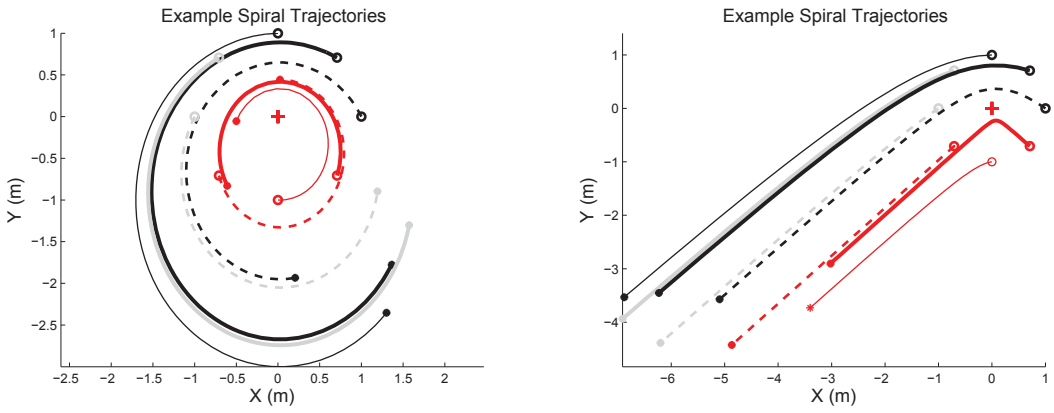

(a) $v_t < v = 0.5$m/s (b) $v_t > v = 1.5$m/s

Figure 3.8: Example spiral trajectories displayed in the object frame $\mathcal{F}_o$ for a reference spiral $\mathbf{c} = [\pi/2 \; -\pi/2]$. The object initial position in the world frame $\mathcal{F}_w$ is (1,1) and the aircraft is initially displaced by 1m (∘,∘,∘). The object position (+) is shown along with safe (−/−) and unsafe (−) aircraft trajectories.

adopts a heading $\Phi \in \{0, \pi/4, \pi/2, 3\pi/4, \pi, 5\pi/2, 3\pi/2, 7\pi/4\}$ and the aircraft is initially re-positioned to ensure $\alpha = -\pi/2$. The resulting aircraft trajectories are shown with respect to the object frame $\mathcal{F}_o$ in Fig 3.8.

For $v_t < v$ the aircraft can track the reference conical angles, resulting in distorted spiral trajectories. For $\pi/2 \leq \bar{\alpha}_0 \leq 3\pi/2$ (*black, grey*), the aircraft initially moves away from the object such that a safe avoidance maneuver is adopted. Indeed, it is not until the aircraft attempts to overtake and pass in front of the object when $\psi(t) > \psi_0 + \pi$, that a potential collision may occur. For $\bar{\alpha}_0 < \pi/2$ or $\bar{\alpha}_0 > 3\pi/2$ (*red*) the aircraft initially moves toward the object, such that an unsafe avoidance maneuver is adopted. For $v_t > v$ the aircraft cannot track the reference conical angles for all object headings, and no spiral trajectory is adopted. Essentially, the assumptions used to analyse the initial geometry are violated so no longer apply for all time. However, for $\pi/2 \leq \bar{\alpha}_0 \leq 3\pi/2$, attempting to track the reference spiral causes the aircraft to initially move away from the object, such that the correct avoidance direction is adopted. Again, for $\bar{\alpha}_0 < \pi/2$ or $\bar{\alpha}_0 > 3\pi/2$, the aircraft initially moves toward the object, such that an unsafe avoidance direction is adopted. Importantly, the analysis supports (3.27), but also suggests that the faster object influences the encounter geometry (and subsequent collision outcome) more than the slower object. The specific effects of velocity therefore require further analysis.

Extending the previous analysis, consider the effect of relative velocity in more detail. Consider again setting the reference spiral defined by $\alpha^* = -\pi/2$ and $\beta^* = \pi/2$, but varying the relative velocity such that $v = 1$m/s and $v_t \in \{0.1 : 0.2 : 1.9\}$m/s. Some example sets of safe and unsafe trajectories are illustrated in Fig 3.9 and Fig 3.10 respectively.

First, consider the subset of safe encounters consisting of relative headings such that $\bar{\alpha}_0 \in \{3\pi/2, 5\pi/4, \pi, 3\pi/4, \pi/2\}$. In all cases, upon adoption of the conical spiral trajectory, the aircraft always moves away from the object suggesting the spiral path is a viable collision avoidance maneuver. Even if the spiral cannot be tracked, the resulting azimuth angle α increases, as the aircraft cannot turn around fast enough to overtake the object. As such,

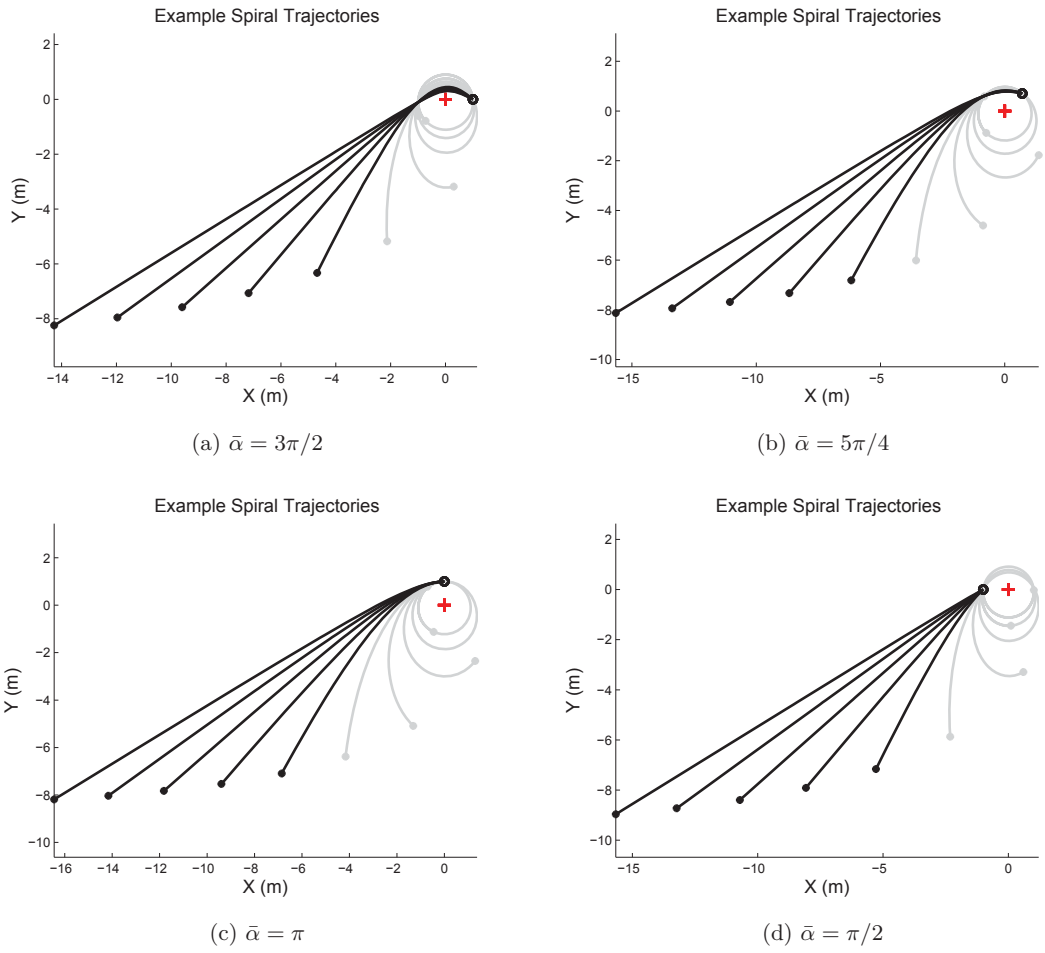

(a) $\bar{\alpha} = 3\pi/2$ (b) $\bar{\alpha} = 5\pi/4$

(c) $\bar{\alpha} = \pi$ (d) $\bar{\alpha} = \pi/2$

Figure 3.9: Example avoidance cases displayed in the object frame $\mathcal{F}_o$ for a reference spiral $\mathbf{c} = [\pi/2 \; -\pi/2]$. The object initial position in the world frame $\mathcal{F}_w$ is (1,1) and the aircraft is initial displaced by 1m (∘,∘). The object position (+) is shown along with aircraft trajectories for $v < v_t$ (—) and $v > v_t$ (—).

the vehicle moves to a safer location ensuring avoidance. For spirals than can be tracked, it is not until the aircraft attempts to completely turn around to overtake and pass in front of the object, that the encounter geometry is degraded. In all cases then, initial avoidance can then be assured, but a resolution decision is required to stop the spiral or divergent motion.

Second, consider the subset of unsafe encounters consisting of relative headings such that $\bar{\alpha}_0 \in \{0, \pi/4\}$. Now an attempt is made to track the same reference spiral for crossing objects approaching from the opposite side than the safe encounters. In all cases, there is no guarantee that the aircraft and object will not collide as the aircraft initially moves toward the object, degrading the encounter geometry before improving it. The exception is for $\bar{\alpha}_0 = \pi/4$ when $v_t > v$. The object is too fast for the aircraft to maintain the reference spiral resulting in divergent behaviour. Importantly, the relative heading is still considered unsafe for $v_t < v$, as the geometry is initially degraded. The analysis suggests that the reference spiral $\alpha = -\pi/2$ is not a viable avoidance trajectory in these cases, and a different reference spiral should have been adopted.

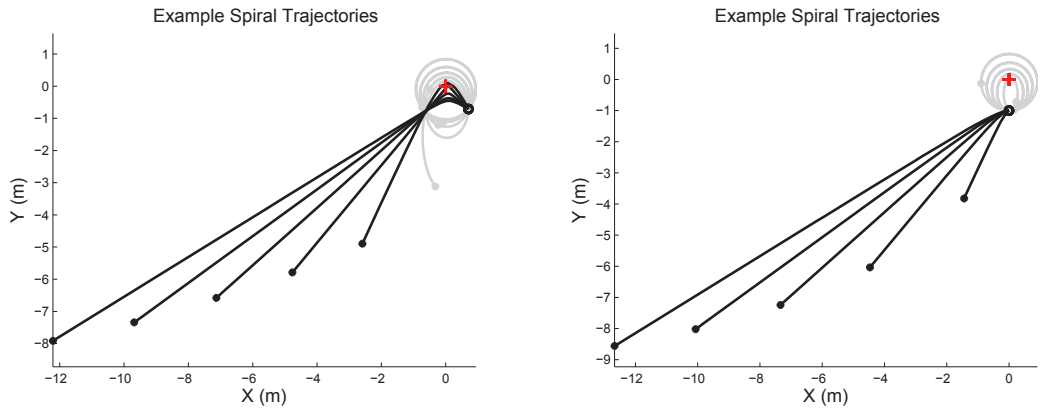

(a) $\bar{\alpha} = 0$ (b) $\bar{\alpha} = \pi/4$

Figure 3.10: Example collision cases displayed in the object frame $\mathcal{F}_o$ for a reference spiral $\mathbf{c} = [\pi/2 \ -\pi/2]$. The object initial position in the world frame is (1,1) and the aircraft is initially displaced by 1m ($\circ$). The object position ($+$) is shown along with aircraft trajectories for $v < v_t$ ($-$) and $v > v_t$ ($-$).

Now consider the inclusion of the elevation angle β. Provided $\beta \neq \pi/2$ means the object is not at the same level as the vehicle. However, unlike the static case, the object and aircraft may be ascending or descending at a constant rate. This means a similar situation to that in the lateral plane exists in the vertical plane. As such, a similar analysis to that used for lateral avoidance can then be conducted for vertical avoidance. The geometry can be redefined in terms of a relative inclination or glide angle $\bar{\beta}$, and the initial relative inclination angle $\bar{\beta}_0$ can be analysed. Importantly, to ensure polarity is correctly managed, the elevation angle needs to be shifted by $\pi/2$ in the analysis (see §3.2.2). In this case, initial vertical avoidance can be assured in some cases. Specifically,

$$\beta^* \in \{0, \pi\}, \quad \pi \leq \bar{\beta}_0 \leq 2\pi \tag{3.28}$$

This means vertical avoidance requires positioning the object directly above or below the aircraft such that $\beta \in \{0, \pi\}$. However, a degenerate conical spiral results at these elevation angles. Additionally, β is not defined in the safe avoidance range $(\pi, 2\pi)$ and instead, α is used to implicitly define an object behind the aircraft. Instead, consider attempting to provide vertical separation by selecting

$$\beta^* \neq \pi/2, \quad 0 \leq \beta_0 \leq \pi \tag{3.29}$$

Although vertical collision avoidance cannot be assured in this case, it allows feasible non-degenerate reference conical angles to be tracked. As such, a spiral trajectory about the object will still result. Importantly, assuming lateral avoidance is the primary concern and can be assured, then $\beta \neq \pi/2$ constitutes a viable avoidance spiral. In some cases, the resulting spiral may provide additional vertical separation and increase miss distance.

The analysis above indicates that collision avoidance of constant velocity objects is possible by attempting to establish and track a spiral such that $\alpha^* = \pm\pi/2$ and $\beta^* \neq \pi/2$. Even if the reference spiral cannot be achieved, attempting to establish and maintain it can

result in safe avoidance behaviour. This implies that spiral trajectories are a viable collision avoidance trajectory requiring the use of only angular position measurements. The difficulty lies in determining the orientation of the reference spiral to track ($\alpha = \pm\pi/2$), given the unknown intent of the object. Specifically, the unknown heading means the initial relative heading $\bar{\alpha}_0$ is unknown, so it is impossible to analytically determine the optimal spiral to adopt. However, an avoidance decision is still required upon initial detection to help determine which exact reference spiral to track. To this end, it is possible to introduce aviation flight rules to address the ambiguity, and help determine the specific conical spiral to adopt in an arbitrary encounter.

3.2.5 Remarks

- Spiral trajectories are presented as a viable avoidance trajectory in some cases, but indeed are not the only possible avoidance solution. Simple moving left or right can also result in an appropriate avoidance maneuver, but it does not provide the same desirable qualities offered by attempting to follow a spiral. First, as image-based control will be used to track the spiral, the conical angles can be directly identified from a single image feature and therefore tracked accordingly. Using suitable avoidance spirals allows the selection of specific reference image features to track instead of regions. In this way, a specific (unambiguious) image aiming point can be used. Second, the curved spiral trajectories attempt to move the aircraft back toward the original heading after avoidance, instead of always diverging from the initial heading. This can help alleviate unnecessary avoidance action in some cases and help resolve the conflict.
- The trajectories used in the previous analysis of both static and dynamic objects are provided as examples, and it is acknowledged that an infinite number of encounters are possible. Defining the optimal control action to take in an arbitrary encounter is not trivial, often requiring numerical methods to solve the resulting optimal control problem [218]-[221]. Applying a similar approach to determine the optimal spiral to follow is possible, but as the relative heading is generally unknown, the result cannot be used directly for non-cooperative avoidance. This is the reason a qualitative approach is taken, and used to simply highlight some benefits of using spirals for vision-based control.
- The analysis provided for each specific spiral trajectories assumed that the aircraft was initially established on the appropriate spiral, in order to determine if indeed the resulting trajectory provided avoidance. However, the initial range, control constraints and relative velocity will impact the ability to first establish the spiral before a collision. If an optimal control is used to establish the spiral, this may correspond to the aircraft's best performance in some cases. So provided the appropriate avoidance direction is chosen, any resulting collision while establishing a spiral is not an artefact of using such spirals, but a result of difficult geometry. It is likely any control scheme will have difficulty in avoiding such conflict encounters.

3.3 Avoidance Strategy

This thesis proposes that conical spirals can be used as a viable avoidance trajectory. Specifically, the aircraft has to attempt to visually establish and track a conical spiral, defined with respect to the collision object, to perform collision avoidance. Based on the preceding analysis, the difficulty for avoidance lies in determining which spiral to track, and thus which reference image features to select, given the ambiguity in object motion. For resolution, the difficulty lies in determining when to stop the avoidance behaviour. This section details how these problems can be addressed by leveraging only visual observations and aviation right-of-way rules.

3.3.1 Avoidance Decision

It has been suggested that humans detect a collision by primarily considering the relative angular rate with respect to the object. In the ideal case, a zero angular rate denotes a collision object, and for all other rates a non-collision object. Applying this theory to aircraft collision avoidance however, angular rate alone has proved insufficient to initially distinguish between collision and non collision targets. This is likely due to uncertainty on image feature measurements that may cause some non-collision objects to appear stationary and vice versa (see §1.3.2). One way to better address the ambiguity, is to asses the relative measure of convergence the object exhibits in the image, with due consideration to the expected uncertainty (variance) on the angular observations.

Consider the image feature position $\mathbf{s}(\sigma, \gamma)$ and image feature rate $\dot{\mathbf{s}}(\dot{\sigma}, \dot{\gamma})$. Consider a new metric $\check{\mathbf{s}}$ derived from the image observations that can be used to represent the relative convergence, divergence or lack thereof such that

$$\check{\mathbf{s}} = \mathrm{di}(\mathbf{s})\, \dot{\mathbf{s}}^T \tag{3.30}$$

A strongly diverging object moving away from the image centre will mean that $\check{\mathbf{s}} \gg \mathbf{0}$. A strongly converging object moving toward the image centre will mean that $\check{\mathbf{s}} \ll \mathbf{0}$. A relatively stationary object, that may exhibit some small arbitrary movement, will mean that $\check{\mathbf{s}} \approx \mathbf{0}$ where $\mathbf{0}$ is a vector of zeros. These cases can be summarised as

$$\check{\mathbf{s}} \gg \mathbf{0} \ : \ \text{Divergent}, \quad \check{\mathbf{s}} \ll \mathbf{0} \ : \ \text{Convergent}, \quad \check{\mathbf{s}} \approx \mathbf{0} \ : \ \text{Stationary} \tag{3.31}$$

The aforementioned variables are shown on an example planar representation of a spherical imaging surface in Fig. 3.11. A danger zone D, defined by any image position in the region $-\pi/2 < \gamma < \pi/2$ and $0 < \sigma < \pi$, is also introduced to denote the image region in which an object is in front of the aircraft.

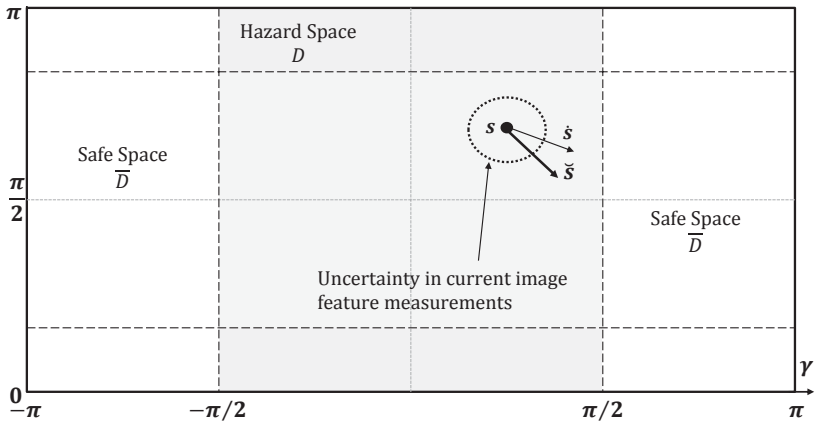

Figure 3.11: Example planar representation of the spherical imaging surface, including an example object (•), bounded uncertainty ($\cdots$) on the image feature position ξ_s and danger area D (■).

Consider also some uncertainty vectors ξ_s and $\xi_{\dot{s}}$ for the image feature observations, whose elements consists of the variance ξ^1 on the image feature angular positions and rates such that $\xi_s = (\xi_\sigma \; \xi_\gamma)$ and $\xi_{\dot{s}} = (\xi_{\dot{\sigma}} \; \xi_{\dot{\gamma}})$. The variables can be used to define uncertainty ellipses about the image features with respect to position and velocity such that

$$\mathbf{s}\,\mathrm{di}(\xi_s)\,\mathbf{s}^T = 0, \qquad \dot{\mathbf{s}}\,\mathrm{di}(\xi_{\dot{s}})\,\dot{\mathbf{s}}^T = 0 \tag{3.32}$$

In a similar fashion as above, consider an uncertainty threshold η such that

$$\eta = \mathrm{di}(\xi_s)\,\xi_{\dot{s}}^T \tag{3.33}$$

Now consider comparing the feature behaviour $\check{\mathbf{s}}$ against a positive avoidance threshold η at the initial confirmed detection instant t_d (or k_d) according to

$$\check{\mathbf{s}}(t_d) \leqq \eta, \qquad \eta \geqq \mathbf{0} \tag{3.34}$$

The comparison evaluates the objects image behaviour whilst considering the expected measurement uncertainty. It helps qualitatively distinguish between the actual object behaviour and that induced by noise. To explain, consider the nominal case with perfect sensing and setting $\eta = \mathbf{0}$. Evaluating if $\check{\mathbf{s}}(t_d) = \eta$ determines if a stationary object is in the centre of the image, or a dynamic object is indeed stationary in the image. These are the conditions known to lead to collision. For imperfect sensing, setting $\eta > \mathbf{0}$ and assessing if $\check{\mathbf{s}}(t_d) \leqq \eta$ suggests the object is either relatively stationary in the image with arbitrary motion bounded by η, or very close to the image centre. In this case, the object may be considered as a more significant collision threat than if $\check{\mathbf{s}}(t_d) > \eta$. If the threshold is large such that $\eta \gg \mathbf{0}$, then almost all stationary, converging or diverging objects would be considered a major collision threat using the same assessment.

[1]The symbol ξ is used instead of the standard σ^2 to denote variance to avoid confusion with the colatitude angle σ.

If the avoidance threshold is then used to denote the confidence (of variance) in the visual observations based on expected uncertainty, it represents a single parameter that can be tuned based on the degree of conservativeness that is desired. If large, the implication is that the camera is perhaps of lower quality of the ambient conditions are causing difficulties in object detection and tracking. If small, the opposite might be implied. Alternatively, the threshold value may be set based on the cameras measured performance during calibration, or updated during flight.

Once the object motion has been qualitatively assessed using the avoidance threshold, a means to map the observed object behaviour to a specific avoidance decision is required. In particular, a reference avoidance spiral must be selected. The exact spiral can be defined implicitly through a set of reference image feature positions where $\alpha^* = \pm\pi/2$ and $\beta^* \neq 0$, as outlined previously. Recalling that defining the direction of the spiral (or polarity of the reference image features), is impossible without knowledge of the object heading. Another method is thus required.

Aviation flight rules, and in particular the right-of-way rules, provide a convenient baseline solution to build such avoidance decisions or logic. By combining these rules with the avoidance decision threshold η, discrete logic can be built for vertical and lateral avoidance that can resolve the ambiguity in determining the direction of the reference conical spiral. Specifically, the logic can be used to select an appropriate set of reference image features $\mathbf{s}^*(\sigma^*, \gamma^*)$ to track for avoidance. The resulting logic for lateral avoidance is given in algorithm 1, and a set of example cases depicted in Fig 3.12. The resulting logic for vertical avoidance is given in algorithm 2, and a set of example cases depicted in Fig 3.13. To help describe how the logic works, the lateral and vertical avoidance decisions for the example cases given in each set of figures are described below. In all cases, the term right-of-way, and its context in aviation, is taken from Australian Civil Aviation Regulations (CAR), Section 162 and summarised in Appendix A [29].

Algorithm 1 Avoidance Decision Strategy - Azimuth

```
if γ ∈ D then                                   ▷ Danger Area
    if γ̆ < η_γ then                             ▷ Convergent, Centreline or Static Features
        γ* = −π/2                               ▷ **Allow Crossing
                                                ***Right-of-Way Rules, CAR 162(1-4)***
    else                                        ▷ Divergent Features
        if γ < 0 then                           ▷ Left Centreline
            γ* = −π/2                           ▷ **No Crossing
        else                                    ▷ Right Centreline
            γ* = π/2                            ▷ **No Crossing
        end if
    end if                                      ▷ Static or Centerline Features
else                                            ▷ Outside Danger Area
    γ* = γ                                      ▷ **No Movement
end if
```

** Aircraft Action

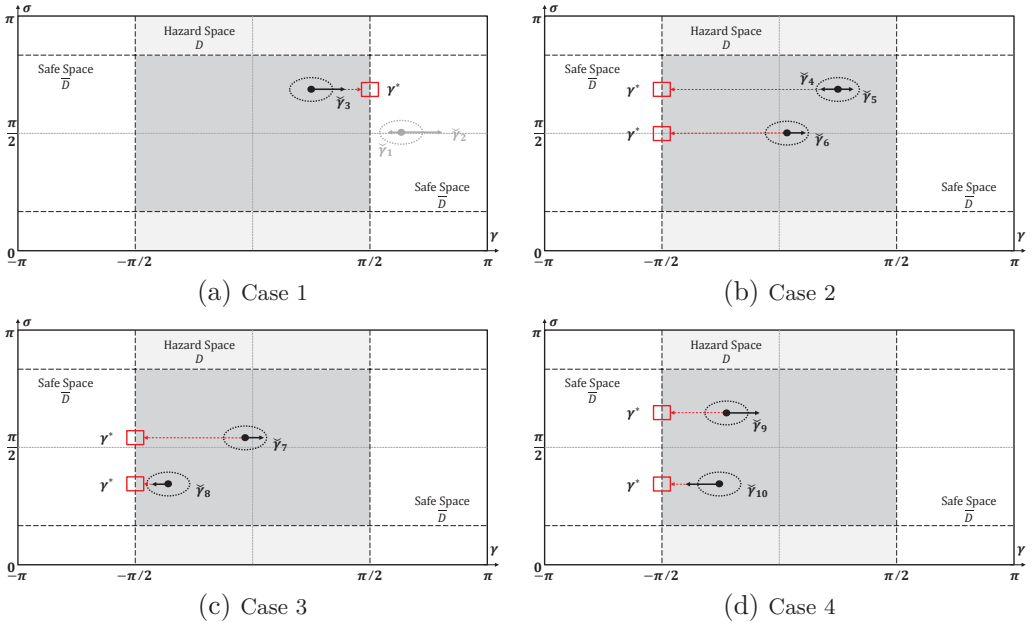

Figure 3.12: Example collision avoidance cases. The image feature position (•) is shown along with $\breve{\gamma}$ (⟶) and the resulting avoidance decision or logic outcomes that set the reference azimuth angle α^* (□).

The following descriptions apply to the each of the **lateral avoidance** cases in Fig 3.12.

▷ *Case 1:* For $\breve{\gamma}_1$ and $\breve{\gamma}_2$, the object appears behind the aircraft and is therefore not a collision object regardless of its behaviour. This is because it is no longer the primary responsibility of the aircraft. For $\breve{\gamma}_3 > \eta_\gamma$, the object is diverging in the right half plane. It is likely a non-crossing, non-collision object (either static or dynamic). The object is then allowed to pass to the right of the aircraft by setting $\gamma^* = \pi/2$.

▷ *Case 2:* For $\breve{\gamma}_4 \leq \eta_\gamma$ and $\breve{\gamma}_5 \leq \eta_\gamma$, the object is relatively stationary in the right half plane. It is likely a crossing collision object (dynamic), and right-of-way must be given. The object is then allowed to pass in front by setting $\gamma^* = -\pi/2$. For $\breve{\gamma}_6 \leq \eta_\gamma$, the object may be static and directly in front of the aircraft, or dynamic and just prior to collision. This constitutes a near-head on encounter, so the aircraft must turn right. As such, the object is allowed to pass in front by setting $\gamma^* = -\pi/2$.

▷ *Case 3:* For $\breve{\gamma}_7 \leq \eta_\gamma$ and $\breve{\gamma}_8 \leq \eta_\gamma$, the object is relatively stationary in the left half plane. It is likely a crossing collision object (dynamic), and right-of-way must be given. In this case however, the aircraft has right-of-way which means the object is required to pass behind. As such, the object is forced to pass behind by setting $\gamma^* = -\pi/2$.

▷ *Case 4:* For $\breve{\gamma}_9 > \eta_\gamma$ the object is converging in the left half plane. It is likely a crossing non-collision object (dynamic). For and $\breve{\gamma}_{10} > \eta_\gamma$, the object is diverging in the left half plane. It is likely a non-crossing, non-collision object (either static or dynamic). In both cases however, applying the right-of-way rules means the object is required to pass behind the aircraft. As such, the object is forced (or allowed) to pass behind by setting $\gamma^* = -\pi/2$[2].

[2]Although induced collision are possible, the logic remains consistent with aviation practise.

Algorithm 2 Avoidance Decision Strategy - Colatitude

Set $\sigma = \sigma - \pi/2, \quad \sigma_o^* = 35\pi/180$
if Object Above Horizontal **then**
  if $\breve{\sigma} > \eta_\sigma$ **then** ▷ Divergent Features
    $\sigma^* = \pi/2 + \sigma_o^*$ ▷ **No Crossing
  else if $\breve{\sigma} < -\eta_\sigma$ **then** ▷ Convergent Features
    $\sigma^* = \pi/2 - \sigma_o^*$ ▷ **Allow Crossing
  else ▷ Static or Centerline Features
    if Overtaking **then** ▷ Overtaking
      $\sigma^* = \sigma$ ▷ **No Movement
      CAR 162(4)
    else Not Overtaking
      $\sigma^* = \pi/2 - \sigma_o^*$ ▷ **Allow Crossing
    end if
  end if
else Below Above Horizontal
  Reciprocal Logic
end if

** Aircraft Action

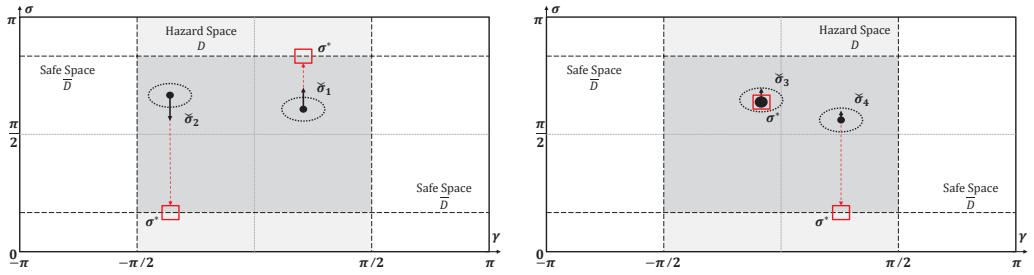

(a) Case 5 (b) Case 6

Figure 3.13: Example collision avoidance cases. The image feature position (•) is shown along with $\breve{\sigma}$ (⟶) and resulting avoidance decisions or logic outcomes with respect to colatitude reference σ^* (□)

The following descriptions apply to the each of the **vertical avoidance** cases in Fig 3.13.

▷ *Case 5:* For $\breve{\sigma}_1 > \eta_\sigma$, the object is diverging in the top half plane. It is likely a non-crossing, non-collision object (either static or dynamic). The object is then allowed to pass above the aircraft by setting $\sigma^* = \pi/2 + \sigma_0$. For $\breve{\sigma}_2 > \eta_\sigma$, a similar situation occurs as for $\breve{\sigma}_1$, but the object is converging. As such, the object is then allowed to cross in front of the aircraft by setting $\sigma^* = \pi/2 - \sigma_0$

▷ *Case 6:* For $\breve{\sigma}_3 \leq \eta_\sigma$ and $\breve{\sigma}_4 \leq \eta_\sigma$, the objects are relatively stationary in the top half plane. They are likely collision objects (dynamic). No right-of-way rules exist for the vertical dimension, other than not to change altitude for an overtaking encounter. If an overtaking encounter can be identified[3], then $\sigma^* = \sigma(k_d)$. In non-overtaking cases, recent studies suggest that a mix of pilot reactions occur including both climb and descend [158, 159]. As such, in this thesis, objects are forced to pass in front by setting $\sigma^* = \pi/2 \pm \sigma_0$ accordingly.

[3]Overtaking may be indicated by slow qualitative object growth as shown by $\breve{\sigma}_3$

There are some important points to notice when applying the avoidance logic, including:

- Only a single threshold needs to be tuned for lateral and vertical avoidance, whilst the underlying logic applying aviation right-of-way rules does not need to change. The avoidance structure is then able to capture the expected avoidance behaviour in manned aviation, with minimal design parameters.
- Lateral and vertical avoidance logic is separate. Lateral avoidance logic is unique for the left and right half planes. The vertical avoidance logic is reciprocal for the bottom and top planes. This means that a mixture of the above cases exist. An object may be considered a significant collision object in the lateral plane, and so adopt the relevant avoidance maneuver. In the vertical plane however, it may exhibit enough apparent image motion to suggest that only precautionary avoidance is required.
- An avoidance action is always taken. Even if the object is qualitatively classified as a non-collision object in both lateral and vertical planes such that ($\check{\mathbf{s}} > \eta$), precautionary action is taken. For true non-collision objects, this can improve miss distance. For incorrectly assessed collision objects, it may provide an element of robustness by providing an avoidance action that alters the initial geometry. However, this does not imply that collision avoidance can be assured.
- The avoidance rules may not result in geometrically optimal behaviour, or even ensure collision avoidance. However, abiding by them ensures the behaviour of unmanned aircraft in a collision encounter is more predictable from a pilots' perspective. Using such aviation flight rules further adds to the approaches alignment with human behaviour in a See and Avoid conflict.
- Although not intentional, the lateral avoidance logic can account for the object also taking action. Indeed, if the object employs the same avoidance logic, then the right of way rules render the object motion complementary to that of the aircraft. For example, if the object observes the aircraft on the right, the object gives way and the aircraft has right of way. As such, $\gamma^* = -\pi/2$ for both vehicles.
- The colatitude angle used in the avoidance logic needs to be shifted by $\pi/2$ such that the vertical centreline corresponds to $\sigma = 0$. This is to ensure an appropriate assessment of $\check{\sigma}$ can be made, similar to the lateral avoidance logic. The reference colatitude then needs to account for this such that $\sigma^* = \sigma_0 \pm \pi/2$ accordingly.
- Many reference colatitude angles could be chosen. The logic selects a colatitude reference offset from the horizontal image centreline by a fixed amount denoted by σ_o. The idea is to vertically separate the aircraft to improve miss distance, but rely on horizontal separation to assure collision avoidance. As such, it is unclear how to best select a specific value for $\sigma_o \neq 0$.

  One way to approach the problem is to consider selecting a value based on improving control performance. For example, selecting a value by considering the nonlinearity of the corresponding reference image position, and the expectations on stability and feasibility for any visual controller. In this way, there is better assurance the reference

feature will be tracked from a control perspective. In turn, consistent collision avoidance control can be better assured, thus improving safety. This is the approach taken in this thesis, and leads to some novel visual control schemes detailed in Chapter 5.

3.3.2 Resolution Decision

For a completely automated system, the aircraft needs to be able to determine when it is safe to leave the avoidance path and resume normal flight. Given the nature of the avoidance spiral trajectories, a resolution decision or stopping criteria is required to prevent the aircraft from continually spiralling around the object ($0 \leq v_t < v$) or attempting to track an unachievable spiral ($v \leq v_t$). The reason a resolution decision is required for each case is straight forward. If the aircraft continually spirals the object, a new collision opportunity is presented at each rotation resulting in unsafe geometry. Additionally, the aircraft moves further away from its initially intended trajectory. If the aircraft cannot track the spiral, the aircraft will continually diverge, again moving away from the original path. A simple and easy approach to the problem can be based on monitoring the aircraft heading from the initial detection instance $\psi(t_d) = \psi_0$. Once the aircraft returns to its initial heading (for the first time) upon adoption of the avoidance trajectory, any further motion will force the aircraft to turn back toward the object. As such, a feasible resolution strategy can be expressed as

$$|\bar{\psi}(t)| \leq |\bar{\psi}_0|, \quad t > t_d \ : \ \text{Stop Avoidance} \tag{3.35}$$

where $\bar{\psi}$ denotes the aircraft heading bounded in the region $(-\pi, \pi)$. This ensures the resolution strategy is valid for reference conical angles resulting in left or right spirals.

Assuming the spiral has been established and subsequently maintained such that $\alpha = \pm\pi/2$, the resolution strategy suggest the aircraft remains directly beside the object until returning to its initially heading. The avoidance manuever is then ceased, and the aircraft will be displaced from its original path and tracking the same heading. Assuming the spiral cannot be maintained such that $\alpha = \pm\pi/2$, but the correct avoidance direction has been selected, the resolution strategy behaves differently. Specifically, the aircraft may not be directly beside the object when the aircraft returns to its initial heading as in the previous case. However, the aircraft will again be displaced from its original path but tracking the same heading. Some example cases that illustrate the concept are depicted in Fig 3.14

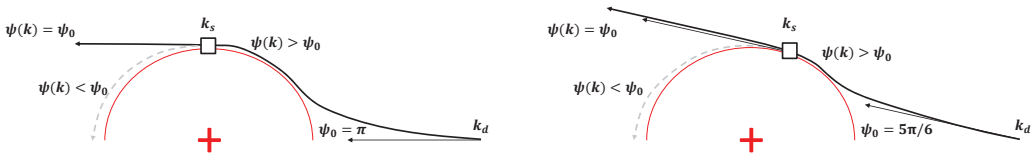

(a) $\psi_0 = \pm\pi$ (b) $\psi_0 = 5\pi/6$

Figure 3.14: Example cases for resolution decision or logic. The original (−) and resolved (−) trajectories are shown along with object position (+) and reference conical spiral (−). The avoidance stops at t_s such that $\psi(t_s) = \psi_0$ (□)

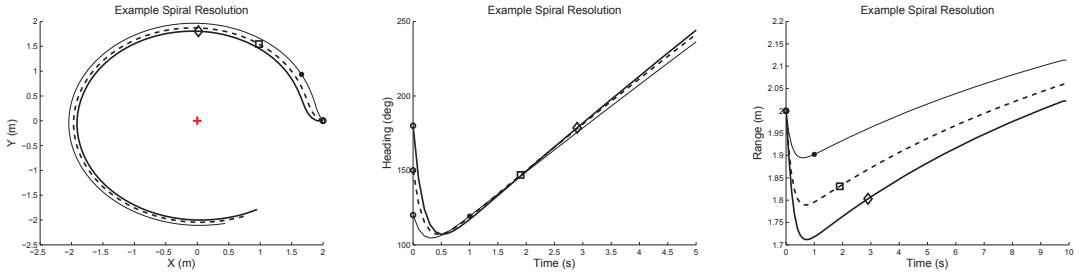

(a) Aircraft Position (b) Aircraft Heading (c) Range

Figure 3.15: Example resolution cases for a static object, displayed in the object frame $\mathcal{F}_o$ for a reference spiral $\mathbf{c} = [\pi/2 \ -\pi/2]$. The object position (+) in the world frame is (1,1) and the aircraft is initially displaced by 1m (∘). The aircraft state (−/−−/−) and corresponding resolution instance (◇/□/●/) are given for $\psi_0 \in \{3\pi/2, 5\pi/3, 11\pi/6\}$ respectively.

Considering the analysis in the previous section, and maintaining consistency, it is important to consider the proposed resolution strategy in light of both static and dynamic objects. Some example cases for static and dynamic objects are depicted in Fig 3.15 and Fig 3.16 respectively. In each case, the plots assume that the correct reference conical spiral $\alpha^* = \pm\pi/2$ has been determined but still needs to be established.

First, consider the case for static objects. The avoidance manuever is always stopped after the point of minimum separation (range). This point also corresponds to the point at which the spiral is established in the static case. The exact amount of time after this point depends on the amount of manoeuvring the aircraft has to do to initially establish the spiral. This is because the same amount of time is then required to re-establish the aircraft's initial heading once on the spiral. Therefore, the bigger the difference between the reference azimuth and relative heading, the further away the resolution time is from the point of minimum separation. Put simply, the more maneuvering the aircraft has to perform to establish the spiral, the more time the aircraft spends tracking it before resolution. In the best case, if the initial relative aircraft heading is equal to the reference azimuth, the spiral is established and the stopping time and point of minimum separation coincide. In the worst case, if the initial relative aircraft heading creates a head-on encounter a maximum resolution time is observed. Importantly, in all cases, the avoidance is stopped before the aircraft attempts to spiral back towards the object, resulting in successful resolution.

Second, consider the case for dynamic objects. The avoidance manuever is always stopped before the point of minimum separation (range). The exact amount of time before this point also depends on the amount of manoeuvring the aircraft has to do to initially establish the spiral, but in a reciprocal manner to that of the static case. Specifically, the less maneuvering that the aircraft has to do to establish the spiral, the closer the resolution is to the minimum separation distance. This is because the amount of manuevering the aircraft has to do is now related to the closing velocity. The larger the closing velocity, the less the aircraft maneuvers to establish the spiral and then return to its initial heading. This is because the object motion now reinforces the aircraft's attempt to establish the spiral. Once established, the aircraft then has to turn quickly to maintain the spiral. This

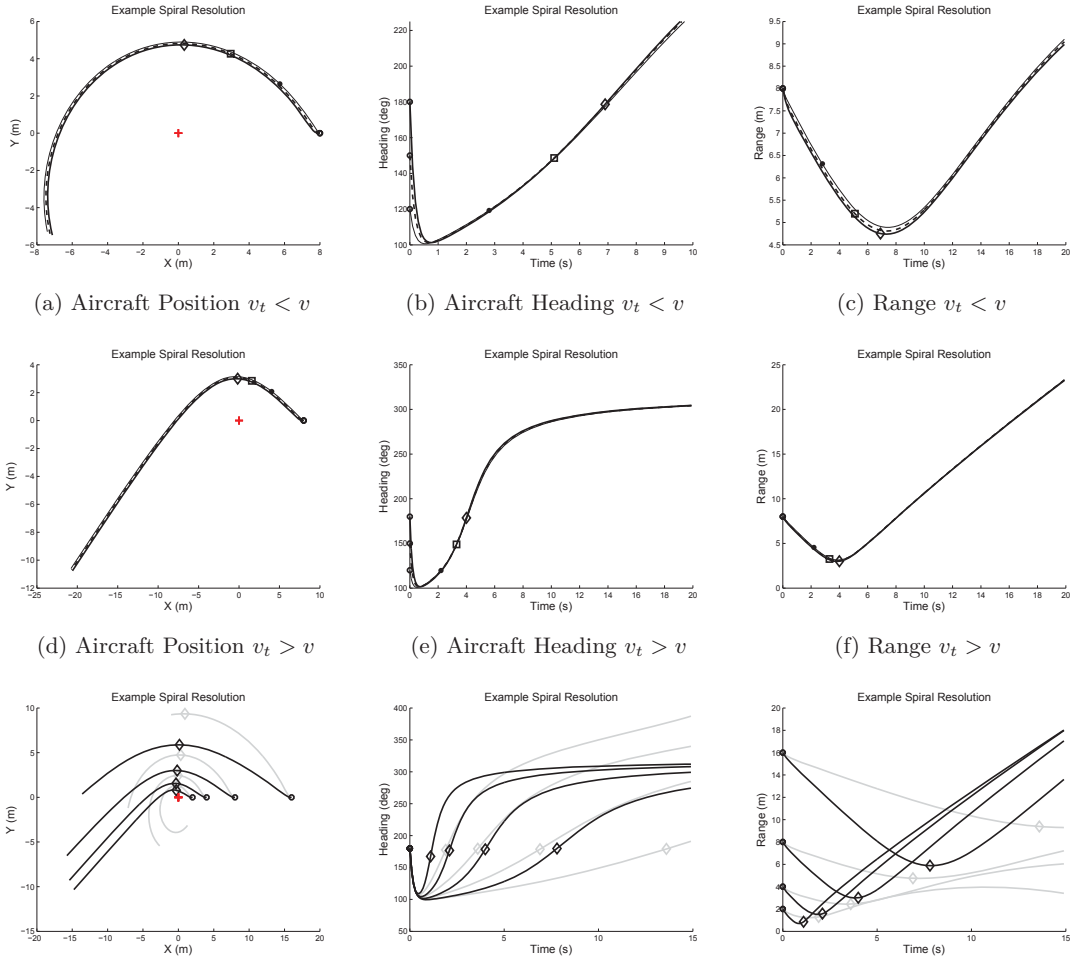

(a) Aircraft Position $v_t < v$ (b) Aircraft Heading $v_t < v$ (c) Range $v_t < v$

(d) Aircraft Position $v_t > v$ (e) Aircraft Heading $v_t > v$ (f) Range $v_t > v$

(g) Aircraft Position $\psi_0 = 3\pi/2$ (h) Aircraft Heading $\psi_0 = 3\pi/2$ (i) Range $\psi_0 = 3\pi/2$

Figure 3.16: Example resolution cases for a dynamic objects displayed in the object frame $\mathcal{F}_o$ for a reference spiral $\mathbf{c} = [\pi/2 \ -\pi/2]$. The object initial position in the world frame is (1,1) and the aircraft is initial displaced by 1m (∘,∘). The aircraft state (−/−−/−) and corresponding resolution instance (⋄/□/•/) are given for $\psi_0 \in \{3\pi/2, 5\pi/3, 11\pi/6\}$ respectively. (a)-(f) $r_0 = 2$ (g)-(i) $r_0 \in \{2, 4, 8, 16\}$

decreases the time spent tracking the spiral before the initial heading is re-established. Therefore, for a given relative velocity, the bigger the difference between the reference azimuth and relative heading, the closer the resolution time is to the point of minimum separation. The effect is seen in Fig 3.16(a)-(c) for $v_t < v$, where the closing velocity increases as the relative heading approaches a head on ($\psi_0 = 3\pi/2$) encounter. The result is further amplified in Fig 3.16(d)-(f) where $v_t > v$, and the resolution instance in all cases tends toward the time of minimum separation. To further highlight the result, Fig 3.16(g)-(i) shows the resolution outcomes for multiple velocities and initial range values r_o with $\psi_0 = 3\pi/2$. Again, in all cases, the avoidance is stopped before the aircraft attempts to spiral back towards the object, resulting in successful resolution.

There are some important points to notice when applying the resolution criteria, including:

- The resolution criteria is only defined with respect to lateral avoidance. As an attempt at vertical separation may also be included in the collision avoidance strategy, vertical avoidance is also ceased upon satisfying the resolution criteria. Note however, it is also

possible to define a vertical resolution strategy in some cases. For example, vertical resolution could be inferred by monitoring the orientation of the vertical control. When the polarity switches, the maximum vertical separation is achieved, as any subsequent control will force the aircraft back toward the object. This is because $0 < \beta < \pi$ which forces the aircraft to move back toward the object the closer it becomes.

- The resolution concept presented here is not explicitly coupled to the control strategy used to regulate spiral motion. It is a simple conformance based decision. This means there may exist cases in which the conical angles have not been established, but the resolution criteria is met. As such, the avoidance behaviour may cease prior to establishing a safe spiral or at an inappropriate instant. This is particularly true at the point of initial detection. Wind effects may briefly displace the aircraft beyond the initially heading ψ_0, triggering the resolution. In this case, no action is taken and the collision avoidance may fail.

  In this thesis the basic resolution concept presented here is maintained but treated in a novel way depending on the vision-based controller. Specifically, for classical image-based control presented in chapter 4, the exact resolution strategy presented here is employed and analysed practically. For predictive image-based control presented in chapter 5, a novel method to couple the visual control with the resolution decision is derived to help prevent resolving the collision prematurely.

3.3.3 General Avoidance Strategy

Combining the avoidance and resolution strategies, the automated collision avoidance approach can now be defined in a series of sequential steps as summarised below:

1. Apply a set of appropriate image processing techniques to detect a potential collision object. The projection of the object onto the imaging surface is then approximated by a single point feature $\mathbf{s}$ consisting of a colatitude σ and azimuth angle γ such that $\mathbf{s}(\sigma, \gamma)$. De-rotate the images using the aircraft attitude (and camera parameters) such that the image features approximate the true conical angles $\mathbf{c}(\beta, \alpha)$ (see §3.2.2).
2. At the confirmed detection instance t_d, asses the image feature motion parameter $\check{\mathbf{s}}(t_d)$ against the avoidance threshold parameter η, to determine the appropriate placement of the reference image features $\mathbf{s}^*(\sigma^*, \gamma^*)$ for visual control. The reference image features are determined using the avoidance logic based on conical spiral motion and aviation rule of the air (see §3.3.1). Set the heading resolution parameters according to the current heading such that $\bar{\psi}^* = \bar{\psi}(t_d) = \bar{\psi}_0$ where $\bar{\psi}(\cdot)$ denotes the aircraft heading shifted in the region $(-\pi, \pi)$ (see §3.3.2).
3. Apply image-based control (visual servoing) to reposition the aircraft such that the actual image features converge toward the reference image features. Specifically, define the image feature error $\mathbf{e}(t)$ and input control vector $\mathbf{u}(t)$ for $t \geq t_d$ such that

$$\lim_{t \to \infty} \mathbf{e}(t) = \mathbf{0} \tag{3.36}$$

   subject to (s.t)

$$\dot{\mathbf{s}}(t) = \mathbf{f}(\mathbf{s}(t), \mathbf{u}(t)) \tag{3.37}$$

$$\mathbf{u}(t) \in \mathbb{U} \subset \mathbb{R}^2, \qquad \forall t \geq t_d \tag{3.38}$$

$$\mathbf{s}(t) \in \mathcal{S}^2, \qquad \forall t \geq t_d \tag{3.39}$$

   where $\mathbb{U}$ defines the aircraft control constraint domain, $\mathbf{f}(\cdot)$ is a model of the combined image kinematics and aircraft dynamics and $\mathcal{S}^2$ defines the set of points on a 2-sphere. The latter can be thought of as 2 dimensional manifold in Euclidean space representing the spherical imaging surface. Two approaches to solving the image-based visual control problem are presented in this work, including novel classical and predictive control schemes (see §4 and §5 respectively)
4. At the confirmed resolution instance t_s, stop the image-based control and resume normal flight control. This may mean the aircraft maintains the heading and altitude at t_s or attempts to resume the original flight path defined at $t < t_d$. The resolution instance is determined by monitoring the aircraft heading $\bar{\psi}(t)$ with respect to $\bar{\psi}^*$ and determining an appropriate time to stop the spiral avoidance trajectory (see §3.3.2, §4.4.1 and §5.5.2).

Chapter 4

Classic Visual Control

4.1 Outline

This chapter presents a new practical spherical image-based control strategy for avoidance control that extends existing classical image-based visual servoing concepts.

First, classical image-based control approaches are reviewed in the context of See and Avoid. Specific attention is paid to practical schemes that account for limited visual information and controllable degrees of freedom whilst considering aircraft applications. Second, novel partitioned spherical image-based visual control (S-IBVS) schemes are derived and simulated results are presented. Third, a practical image-based control scheme is implemented using a small *Parrot - ARDrone* in a scaled indoor collision avoidance environment in a proof-of-concept approach. A simple resolution strategy is also presented and validated empirically. The control schemes are then analysed, highlighting their advantages and drawbacks.

4.2 Fundamentals & Prior Work

The use of visual acquired information to control a robot to perform a specific task is referred to as visual servoing. Broadly speaking, there are two classes of visual servoing labelled *Position-based Visual Servoing (PBVS)* and *Image-based Visual servoing (IBVS)*.

Position-based approaches require an estimate of the object pose (position and orientation) in order to derive feedback control in the task space. The approach can be computational demanding, sensitive to noise and highly dependant on camera calibration. Additionally, multiple image feature points of the same known object are often required to obtain quality pose estimates [292]. Collectively, these features violate the See and Avoid problem constraints (see §1.3). As such, position-based visual servoing does not provide a suitable control approach for See and Avoid systems.

Image-based visual servoing (IBVS) approaches do not require an estimate of the object pose, as feedback control is derived directly from the image features themselves. The approach provides a reactive computationally efficient control solution with inherent robustness to range (depth) uncertainty, noise and camera calibration errors. Additionally, the object of interest does not need to be known exactly and a single feature point can be used for at least partial control [293]. Each of these features are inherent in the See and Avoid environment, with the control approach bearing strong resemblance to how a

pilot may use visual cues to navigate and react to collisions. As such, image-based visual servoing provides a more natural control approach for See and Avoid systems.

In the classical approach, image-based visual servoing involves regulating a set of image features $\mathbf{s}(t)$ toward a set of reference values $\mathbf{s}^*(t)$ using the image feature error $\mathbf{e}(t)$ directly. Often, a fixed set of reference image features are chosen, and an exponential decrease in the image feature error $\mathbf{e}(t) = \mathbf{s}(t) - \mathbf{s}^*$ is desired such that

$$\mathbf{e}(t) = e^{-\lambda t}, \quad \dot{\mathbf{e}}(t) = -\lambda \mathbf{e}(t) = \dot{\mathbf{s}}(t) \tag{4.1}$$

where $\lambda > 0$. For a static object, the optic flow equation can then be re-written using the image feature error such that

$$\dot{\mathbf{e}}(t) = \mathbf{L}_s \, \dot{\mathbf{x}}(t) \tag{4.2}$$

where $\dot{\mathbf{x}}(t)$ denotes the spatial velocity of the camera (and thus robot) consisting of translation and rotation components, and $\mathbf{L}_s$ is the image Jacobian relating the camera velocity to image feature velocity (see §2.3.2). By inverting the above equation, a general velocity controller can be obtained such that

$$\mathbf{u}(t) = -\lambda \hat{\mathbf{L}}_s^+ \mathbf{e}(t) \tag{4.3}$$

where $\mathbf{u}(t)$ is the control vector and λ can now be viewed as a constant positive gain term. The Moore-Penrose pseudo inverse of the image Jacobian $\hat{\mathbf{L}}_s^+$ is an approximation to the true inverse of the image Jacobian. The word approximate is used because the true image Jacobian (and thus its inverse) remains unknown due to camera calibration errors, noisy image measurements and in some cases an unknown object range (depth) value. If the object is moving, then the control law above often requires additional terms and stability issues may arise depending on the object motion [293]. In most cases, locally asymptotic stability is often ensured using Lyapunov stability criterion.

By inspection, it is clear that the control largely depends on the structure (elements) of the image Jacobian and the approximation of its inverse $\hat{\mathbf{L}}_s^+$, as λ serves to simply amplify the resulting control. The image Jacobian (and its inverse) itself depends on the range (depth) value used in the approximation, the number of image features as well as type of image features and coordinate frame (and camera model) in which they are expressed.

The type of image features used for image-based control largely depends on the characteristics of the observed object. In some cases, multiple features types may be feasible, with different feature types providing varying degrees of coupling between the resulting controls. Lines [297, 298], moments [299] and the projection of spherical targets [300]-[302] have been used as image features to try to decouple motion, but such schemes will not work in See and Avoid encounters. Indeed, given the object characteristics, only point features are suitable. Point features can however be represented in a number of different coordinate frames including Cartesian [296], polar [303, 304] and spherical [305, 306]. Each of these representations can also offer different advantages with respect to decoupling the resulting control in different degrees of freedom (DOF). Polar coordinates have

been exploited to improve the servoing performance with respect to the camera retreat problem inherent using Cartesian coordinates. The improvement in the servoing task comes at the expense of imperfect de-coupling between the x and y axis translation and rotation. By using a combinations of polar and Cartesian image feature representations, a hybrid scheme has shown to provide performance similar to the polar form with worsened cross coupling in the x and y axis motions [304].

Using a spherical image features representation offers similar benefits to that of the polar representation, with decoupled motion in the z translational and angular velocity. The use of spherical coordinates is thus an appropriate image feature representation for a classical visual servoing scheme for spiral tracking. The forward x velocity is constant and the y velocity is assumed to be zero. The x and y rotational velocities are also assumed to be zero, given the control is based on de-rotated images (and features). This alleviates any issues regarding coupled motion in the x and y axis. The decoupled z axis velocity then allows each image feature component to directly regulate the vertical and lateral motion of the aircraft, and thus track a spiral. Representing image features in spherical coordinates is of course a natural choice for spherical cameras, but they can also be used with perspective and wide angle cameras to approximate a spherical section (see §2.3).

Combined with the type of image features, the number of available image features determines the controllable degrees of freedom. Many generic image-based control approaches assume that all camera degrees of freedom can be independently controlled, such that rotation and translation can be regulated about each camera axis. This is a valid assumption for many systems such as robotic arms, and requires at least four image feature points to obtain an overdetermined system for which a unique control solution exists. For other systems, such as aerial vehicles, each degree of freedom cannot be controlled independently and the system is often under actuated. For example, horizontal translation may be controlled by regulating pitch and roll angular rates (attitude). Therefore, partitioned control schemes are common whereby only some degrees of freedom are visually controlled. The specific degrees of freedom depend on the number (and type) of image features used in the visual control scheme.

Given the diversity of possible image features, a large number of unique image-based control schemes have since been derived for aircraft. Some notable examples include translational and hover control of multirotors [307]-[310] and visually guiding aircraft to land [311]-[313]. Although novel, the object of interest is often assumed to be a planar object or 3D structure with multiple distinct features. As such, these approaches cannot be adapted directly for the See and Avoid problem. This is because the fundamental difference in the See and Avoid problem is the existence of a single point feature, allowing image-based control of only two degrees of freedom. Considering the use of spiral trajectories for See and Avoid, then both vertical and lateral motion are required which confines which degrees of freedom should be controlled. This limitation can also help guide the structure of any partitioned control strategy.

To this end, variations of the classic image-based control structure have been used to derive 2D [314, 315] and 3D [237, 240] controllers for spiral tracking and collision avoidance control. Preliminary flight test results are reported in some cases, the majority of which focus on a scaled version of the See and Avoid problem. These works were based on similar collision avoidance methods that do not explicitly consider the image Jacobian [241]-[243], but leverage only bearing observations. More recently, the work has moved away from range independence, and focused on optimal control approaches for the simultaneous avoidance and estimation of range (or time to collision) [244]-[246]. Although novel, the approach may only be successful if the object intent is known and range recovered. Again, this moves away from the realistic See and Avoid environment.

The reference range value used in the servoing scheme impacts the resulting control, as the image Jacobian is directly dependant on the object range r. As this is often unknown, a common assumption is to consider a fixed reference value r^* within the Jacobian (and its inverse) [293]. The range term then acts in a similar way to the gain parameter λ in the control law, scaling the control output accordingly. The motivation is that as the actual camera pose approaches the desired pose, the control becomes increasingly accurate. However, stability issues can arise, resulting in recent studies on the general robustness of image-based control to r^*. The results suggest that if the geometry is unknown, care must be taken in the selection of r^* to ensure stable control [316]. As the analysis was conducted under a specific type and number of image features (planar), it is unclear what the exact implications are for a control strategy using single point features and a subset of controls.

To this end, it is important to consider how the choice of r^* affects the resulting control from a collision avoidance perspective. In particular, how the resulting control affects the tracking of an avoidance spiral and the implications for design. Given the aforementioned robustness properties, and assuming a fixed gain value is chosen, it may be such that the selection r^* is not critical for safe avoidance. Instead, consistently over or under estimating the parameter may reinforce or aid the avoidance behaviour.

The analysis above suggests that classical image-based control, and specifically partitioned control frameworks, offer a viable control structure for spiral tracking and collision avoidance. Some benefits include reactive control, inherent robustness (to added noise and parameter uncertainty) and a simple flexible control structure. These features render the partitioned image-based control approach amenable to the See and Avoid problem environment, however it has not yet been investigated.

The remaining sections of this chapter detail the derivation, simulation, implementation and analysis of a set of image-based controllers using classical image-based approaches, that are suitable for See and Avoid. Specifically, two novel spherical image-based controllers are presented. The first is designed under simplifying assumptions. The second is designed under the same assumptions, but augmented to better compensate for their undesirable effects. The controllers are then used to design a preliminary practical closed loop vision-based collision avoidance system for static objects.

4.3 Spherical Image-based Visual Servoing

To derive a partitioned image-based visual controller using only a single point, consider first the relationship between the camera velocity and image feature error

$$\dot{\mathbf{e}}\left(t\right) = \mathbf{L}_s^* \dot{\mathbf{x}}\left(t\right) + \frac{\partial \mathbf{e}(t)}{\partial t} \tag{4.4}$$

where $\dot{\mathbf{x}}\left(t\right)$ is the camera translational and angular velocity, $\dot{\mathbf{s}}\left(t\right)$ is the image feature velocity and $\mathbf{L}_s^*$ is the spherical image Jacobian using a reference range value r^* (see §2.3.2). As the object intent is generally unknown, the last term is included to account for optic flow induced by the object motion. As only a single feature point consisting of two angular measurements is observed, only two degrees of freedom can be controlled. Typically the Jacobian is partitioned into translational and rotational parts, but arbitrary partitioning is possible [248]. As such (4.4) can be expressed as

$$\dot{\mathbf{e}} = \mathbf{L}_{xy}^* \dot{\mathbf{x}}_{xy}(t) + \mathbf{L}_z^* \dot{\mathbf{x}}_z(t) + \frac{\partial \mathbf{e}(t)}{\partial t} \tag{4.5}$$

where $\mathbf{L}_{xy}^*$ is the Jacobian made up from the rotational and translational components about the x and y axis, $\mathbf{L}_z^*$ is the Jacobian made up from the rotational and translational components about the z axis and the vectors $\dot{\mathbf{x}}_{xy}$ and $\dot{\mathbf{x}}_z$ are given by

$$\dot{\mathbf{x}}_{xy} = \left(v_x \; v_y \; \omega_x \; \omega_y\right)^T \qquad \dot{\mathbf{x}}_z = \left(v_z \; \omega_z\right)^T \tag{4.6}$$

where the independent variable t has been omitted. Assuming a constant desired image feature set $\mathbf{s}^*$, the image feature error can be defined as

$$\mathbf{e}(t) = \mathbf{s}(t) \ominus \mathbf{s}^* \tag{4.7}$$

where $\ominus$ denotes a modulo-2π subtraction required to bound the error such that $\mathbf{e}(t) \in \{-\pi, \pi\}$. Assuming an exponential decrease in the feature error, then substituting (4.7) into (4.5) and rearranging

$$\dot{\mathbf{x}}_z(t) = \hat{\mathbf{L}}_z^{-1} \left(-\lambda \mathbf{e}(t) - \mathbf{L}_{xy}^* \dot{\mathbf{x}}_{xy}(t) + \frac{\partial \mathbf{e}(t)}{\partial t} \right) \tag{4.8}$$

where λ is a positive gain value and $\hat{\mathbf{L}}_z^{-1}$ is the approximate inverse image Jacobian corresponding to the z axis velocity components using r^*. Assuming the velocity about the z axis can be controlled directly, then

$$\mathbf{u}_3(t) = \hat{\mathbf{L}}_z^{-1} \left(-\lambda \mathbf{e}(t) - \mathbf{L}_{xy}^* \dot{\mathbf{x}}_{xy}(t) + \frac{\partial \mathbf{e}(t)}{\partial t} \right) \tag{4.9}$$

where $\mathbf{u}_3(t)$ denotes the reference vertical velocity and yaw rate controls for a low-level attitude controller (see §2.2.3). If the object velocity was known, then the induced optic flow could be directly accounted for in the control law using the partial derivative term. As it is generally unknown however, some simplifying assumptions are required.

Consider that the induced optic flow from object motion is only dependent on the object translational velocity. This is a valid assumption considering only a point feature is observed. As such

$$\frac{\partial \mathbf{e}(t)}{\partial t} \propto \frac{1}{r(t)} \tag{4.10}$$

Assuming that the relative displacement is expected to be much greater than the magnitude of the relative velocity, then the term will provide only a small contribution to the optic flow. As such, consider first making the assumption that the term can be ignored. This means the control should provide zero offset for static objects and an unpredictable non-zero offset for moving objects. In this case, the control law can then be simplified such that

$$\mathbf{u}_3(t) = \hat{\mathbf{L}}_z^{-1}\left(-\lambda \mathbf{e}(t) - \mathbf{L}_{xy}^{*}\dot{\mathbf{x}}_{xy}(t)\right) \tag{4.11}$$

To track a spiral, the aircraft is assumed to move with a fixed forward velocity v_x and zero lateral velocity v_y. Additionally, the image features can be obtained from de-rotated images such that the roll ϕ and pitch θ angles will not change significantly. This means their corresponding contribution to the control law can also be ignored. The resulting control law can then be expressed as

$$\mathbf{u}_3(t) = \begin{pmatrix} -\frac{\cos\sigma(t)\cos\gamma(t)}{\sin\sigma(t)} \\ \frac{\sin\gamma(t)}{r^*\sin\sigma(t)} \end{pmatrix} v_x + \begin{pmatrix} \frac{-\lambda r^*}{\sin\sigma(t)} & 0 \\ 0 & \lambda \end{pmatrix} \mathbf{e}(t) \tag{4.12}$$

which can be expanded and expressed as

$$v_z(t) = -\lambda \frac{r^* e_\sigma(t)}{\sin\sigma(t)} - \frac{\cos\sigma(t)\cos\gamma(t)}{\sin\sigma(t)} v_x(t) \tag{4.13}$$

$$\omega_z(t) = \lambda e_{\gamma(t)} + \frac{\sin\gamma(t)}{r^*\sin\sigma(t)} v_x(t) \tag{4.14}$$

where r^* denotes a fixed range value and e_σ and e_γ are the colatitude and azimuth feature errors respectively. The control is decoupled as the vertical velocity depends on the colatitude error, whilst the yaw velocity depends on azimuth error. The aircraft can therefore independently change its lateral and vertical displacement to track a spiral trajectory, using only a single point feature. Additionally, v_x can be treated like a parameter and fixed in the control law given that the platform velocity is constant.

▷*Remark* For rotary wing aircraft, the input controls can be mapped directly to a yaw rate command and transformed into a collective (or thrust) command. For fixed wing aircraft, the input controls can transformed into a roll and pitch command. As such, the controller is applicable to both aircraft types. In any case, the direct application of the controller assumes image feature errors defined with respect to de-rotated images to ensure the roll and pitch attitude are zero. Of note, the above controller can be reformulated using different degrees of freedom, namely yaw and pitch velocity [315]. This control law however is not as extensible to rotary wing platforms and does not retain the same control decoupling.

4.3.1 Static Objects

Consider an aircraft attempting to visually track a conical spiral about a static object. The object is approximated as a point feature in the image, and assumed to represent a fixed conical apex. Consider applying the spherical image-based control law

$$\mathbf{u}_3(t) = \hat{\mathbf{L}}_z^{-1}\left(-\lambda\mathbf{e}(t) - \mathbf{L}_{xy}^{*}\dot{\mathbf{x}}_{xy}(t)\right) \tag{4.15}$$

$$\mathbf{u}_3^{*}(t) = \max[\mathbf{a}, \min(\mathbf{b}, \mathbf{u}_3(t))] \tag{4.16}$$

where $\mathbf{u}_3^{*}$ denotes the saturated reference control vector for the attitude controllers and $\mathbf{a}$ and $\mathbf{b}$ are constant vectors denoting the upper and lower amplitude saturation limits respectively. The control limits can be chosen based on aircraft performance and control constants. Given the simplifying assumptions used to derive the controller, it is important to verify an appropriate spiral track for collision avoidance can be maintained. To this end, a set of simulation studies are used to visualise the controller performance. The analysis is conducted with respect to reference feature placement $\mathbf{s}^*$, added noise and the reference range parameter r^* used in the control. Simulation parameters are given in Appendix E.2 and the results are detailed in the following sections.

Reference Image Features

Consider varying the reference image features $\mathbf{s}^*$ such that the reference spirals include convergent, divergent and circular motion. The reference range used in the control is set to the actual range value such that $r^* = r(t)$. This ensures only the effects of varying the reference image features are observed. The objective is then to verify that the controller can provide suitable spiral tracking performance under ideal conditions and parameters.

Using a fixed reference colatitude, the reference azimuth γ^* is varied, and the resulting platform and image feature trajectories are depicted in Fig 4.1. Using a fixed reference azimuth, the reference colatitude σ^* is varied, and the resulting platform and image features trajectories are depicted in Fig 4.2. The feature error, velocity control and feature trajectory for an example case is shown in Fig 4.3

In all cases, the platform is able to track the image features and maintain the expected spiral. For divergent spirals and circular motion, the image features can be tracked for all time. For convergent spirals, the reference image features can be tracked until such time that the control limits are reached. This is expected as convergent spirals eventually require an infinite number of rotations about the object at successively smaller radii. This is not a concern for two reasons. First, convergent spirals would not be selected as an avoidance maneuver if possible. Second, if a convergent spiral is required due to camera limitations, the spiral can be initially tracked, which forces the aircraft away from the object. Eventually, the control constraints would prevent the aircraft from tracking a spiral of decreasing radius. This means the controller is able to initially guide the aircraft

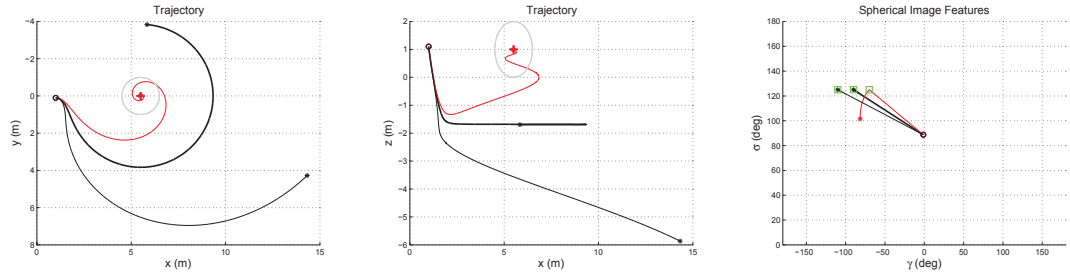

(a) xy-plane (b) xz-plane (c) Feature Trajectory

Figure 4.1: Aircraft and image feature trajectories in the world frame $\mathcal{F}_w$ for a static object with $\sigma^* = -125\pi/180$ and variable reference azimuth (□) such that $\gamma^* = -70\pi/180$ (−), $\gamma^* = -\pi/2$ (−) and $\gamma^* = -110\pi/180$ (−). Initial aircraft (○) and object (+) positions are $(1, 0.1, -1)$ and $(5.5, 0, -1)$ respectively.

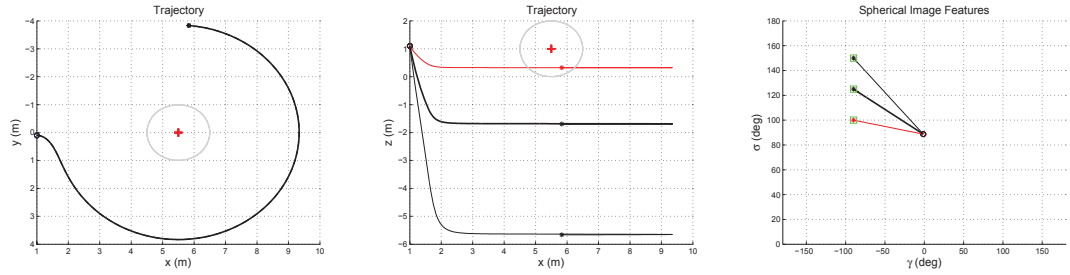

(a) xy-plane (b) xz-plane (c) Feature Trajectory

Figure 4.2: Aircraft and image feature trajectories in the world frame $\mathcal{F}_w$ for a static object with $\gamma^* = -\pi/2$ and variable reference colatitude (□) such that $\sigma^* = 100\pi/180$ (−), $\sigma^* = 125\pi/180$ (−) and $\sigma^* = 150\pi/180$ (−). Initial aircraft (○) and object (+) positions are $(1, 0.1, -1)$ and $(5.5, 0, -1)$ respectively.

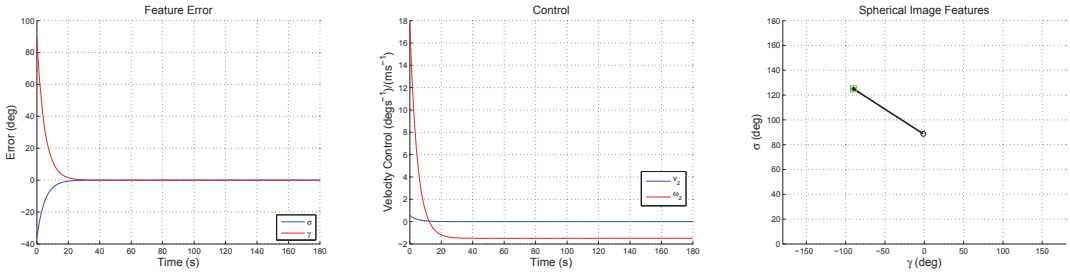

(a) Feature Error (b) Control (c) Feature Trajectory

Figure 4.3: Image feature trajectories and control for a static object with $\gamma^* = -\pi/2$ and $\sigma^* = 125\pi/180$ (□). Initial aircraft and object positions are $(1, 0.1, -1)$ and $(5.5, 0, -1)$ respectively.

along an appropriate avoidance trajectory[1]. This in an important observation from a practical perspective.

Considering an example case using the intended reference azimuth angle for collision avoidance $\gamma = \pm\pi/2$, the image feature motion is straight forward, the control is smooth and the feature error exponentially decreases. The results are as expected, and verify that the controller is suitable for spiral tracking and collision avoidance of static objects in the nominal or ideal case.

[1]The chosen reference image features represent the types of spiral that can be tracked for a static object. For negative azimuth angles, a similar set of trajectories result but rotating in the opposite direction. For colatitude angles below the horizontal, the reference conical spiral is simply inverted (see §3.2.3).

Robustness

Consider adding uncertainty to the system, such that the robustness properties of the controller can be analysed. The reference image features $\mathbf{s}^*$ are fixed, and set to the intended reference angles for collision avoidance. The reference azimuth angle $\gamma = \pm\pi/2$ is selected for circular motion, and an arbitrary colatitude angle is used such that $\sigma^* \neq \pi/2$. Uncertainty is then added to the system in the form of additive noise, imperfect actuation and variable reference range parameter r^*. The objective is then to verify that the controller can provide suitable spiral tracking performance under realistic conditions and uncertain or unknown system parameters.

First, consider adding image feature uncertainty and imperfect actuation to the system in the form of additive white Gaussian noise $q(t) \sim \mathcal{N}(0, 0.19^2)$. The variance on the image feature position and velocity controls are equivalent, and approximately 2 degrees. Importantly, the variance on the image feature estimates is realistic and comparable to that expected using recent aerial object tracking systems (see §1.3.1). The reference range used in the control is set to the actual range value such that $r^* = r(t)$. This ensures only the affects of added noise are observed.

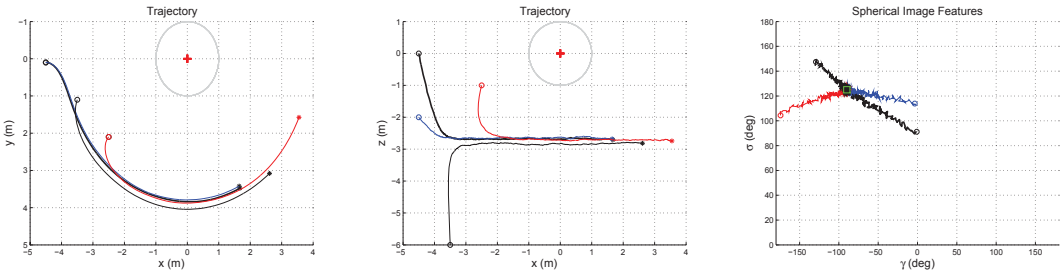

(a) xy-plane (b) xz-plane (c) Feature Trajectory

Figure 4.4: Aircraft and image feature trajectories in the object frame $\mathcal{F}_o$ for a static object with $\gamma^* = -\pi/2$ and $\sigma^* = 125\pi/180$ (□). Four separate encounters are shown (−/−/−/−). Initial aircraft position (∘) is varied such that $2 < r_0 < 8$. The object (+) position in the world frame is $(5.5, 0, -1)$.

A set of example encounters were simulated using varied initial conditions and depicted in Fig 4.4. The initial displacement r_0 was varied such that $2 < r_0 < 8$. The heading ψ_0 was varied such that the initial image features $\mathbf{s}_0$ appeared in the regions defined by $\sigma_0 \in (\pi/2, \pi)$ and $\gamma_0 \in (-\pi, 0)$. This is because alternate reference image features would be used for the other image regions. In all cases, the platform is able to establish and maintain the expected spiral track, shown by the convergence of each trajectory to a planar spiral. The image feature image feature motion is straight forward, residing in a small neighbourhood about the reference image feature position. The results verify that the controller is suitable for spiral tracking and collision avoidance of static objects in the presence of image feature noise and imperfect actuation.

Second, consider varying the reference range parameter used in the controller such that a constant $r^* \neq r(t)$ is used, and all remaining encounter parameters are fixed. This aspect is important from a See and Avoid perspective, as range is assumed to be unknown throughout the encounter. A reference value is still required for control but, only a coarse estimate would be available based on the expected initial object detection distances (see §1.3.1). To this end, consider two cases in which the reference value is considerably overestimated and underestimated respectively.

For the underestimated case the reference range value is such that $r^* \ll r(t)$ throughout the encounter. Consider the structure of the control given in (4.13) and (4.14). In this case, the vertical control will under compensate for the colatitude error by directly decreasing the amplification of the error term. The control is then unchanged as the contribution from the forward velocity does not depend on r^*. The yaw rate control is also reduced as the contribution from the forward velocity is increased. This is because any image feature motion induced by forward velocity that naturally tends toward the reference image feature position is overestimated (larger). The contribution from the azimuth error term remains unchanged, and therefore has less influence on the control. The result is then a weaker yaw rate control. Collectively, less aggressive control would result and the reference spiral will be approached slowly. Considering the nature of spiral motion, a small unpredictable offset in the reference image features may also remain, and therefore the reference spiral may not be tracked. The attenuated control will likely undershoot the reference image features and force a slightly convergent spiral. The analysis suggests that underestimating r^* may be problematic for collision avoidance, forcing the aircraft toward the object.

For the overestimated case the reference range value is such that $r^* \gg r(t)$ throughout the encounter. Again consider the structure of the control given in (4.13) and (4.14). In this case, the vertical velocity control will over compensate for the colatitude error by directly increasing the amplification of the error term. The control is then unchanged as the contribution from the forward velocity does not depend on r^*. The yaw rate control is also amplified as the contribution from the forward velocity is reduced. This is because any image feature motion induced by forward velocity that naturally tends toward the reference image feature position is underestimated (smaller). The contribution from the azimuth error term remains unchanged, and therefore has greater influence on the control. The result is then a stronger yaw rate control. Collectively, more aggressive control would result and the reference spiral will be approached faster. Again, considering the nature of spiral motion, a small unpredictable offset in the reference image features may still remain, and therefore the reference spiral may not be tracked (similar to $r^* \ll r(t)$). However, the amplified control will likely overshoot the reference image features and force a slightly divergent spiral. The analysis suggests that overestimating r^* can still provide adequate motion for collision avoidance. There is then evidence to suggest that an accurate range value is not required for collision avoidance. In fact, given the expected maximum detection distances, a reference range value twice as large still results in stable and effective collision avoidance control.

Despite the collision avoidance benefits when overestimating r^* in the control, the exact reference spiral may not be obtainable in sufficient time or indeed at all. From a safety perspective, better assurance the actual reference spiral can be tracked is required. To address this, the offset in image feature error anticipated when $r^* \neq r(t)$ can be considered as a small variable steady state error. This is a common problem in many real control systems, and can be addressed using classical linear control techniques. As such, consider introducing a small integral control term on the image feature error such that (4.13) and (4.14) become

$$v_z^*(t) = -\frac{\cos\sigma(t)\cos\gamma(t)}{\sin\sigma(t)} v_x(t) - \lambda \frac{r^* e_\sigma(t)}{\sin\sigma(t)} + \lambda_i \int^w e_\sigma(t)\, dt \tag{4.17}$$

$$\omega_z^*(t) = \frac{\sin\gamma(t)}{r^*\sin\sigma(t)} v_x(t) + \lambda e_{\gamma(t)} + \lambda_i \int^w e_\gamma(t)\, dt \tag{4.18}$$

where λ_i denotes a fixed integral gain term such that $0 < \lambda_i \ll \lambda$. The notation $\int^w$ is used to express that the integration can be restricted to a finite time history or performed for all time provided anti-windup action is included. The augmented spherical image-based control law can then be expressed

$$\mathbf{u}_3(t) \;=\; \hat{\mathbf{L}}_z^{-1}\left(-\lambda \mathbf{e}(t) - \mathbf{L}_{xy}^* \dot{\mathbf{x}}_{xy}(t)\right) + \lambda_i \int^w \mathbf{e}(t)\, dt \tag{4.19}$$

To verify the augmented controller and the previous range parameter analysis, a set of example encounters were simulated using the original and augmented controller and depicted in Fig 4.5. An over and under estimated range parameter was used for the original controller such that $r(t)/2 \leq r^* \leq 2r(t)$. Only the over estimated range parameters were used in augmented controller with $\lambda_i = 0.01$. The initial conditions in all cases remain unchanged, and represent a head on encounter in which the object appears 4.5m ahead of the aircraft. This ensures only the effects of highly uncertain system parameters and the augmented controller are observed.

The results confirm the analysis presented in the previous section, and reinforce the use of an overestimated range parameter in the visual controller for collision avoidance. In this case the aircraft trajectory rapidly diverges from the object and more direct image feature motion is observed. The augmented controller is also shown to actively compensate for the range uncertainty, ensuring the reference spiral can be tracked. The control is smooth and the image feature error shows better convergence.

As further evidence to support the use of the augmented controller, image feature noise and imperfect actuation is added to the simulation and the results depicted in Fig 4.6. Again, expected aircraft behaviour is obtained and the addition of the integral term has helped maintain the image features within a small neighbourhood of the reference values.

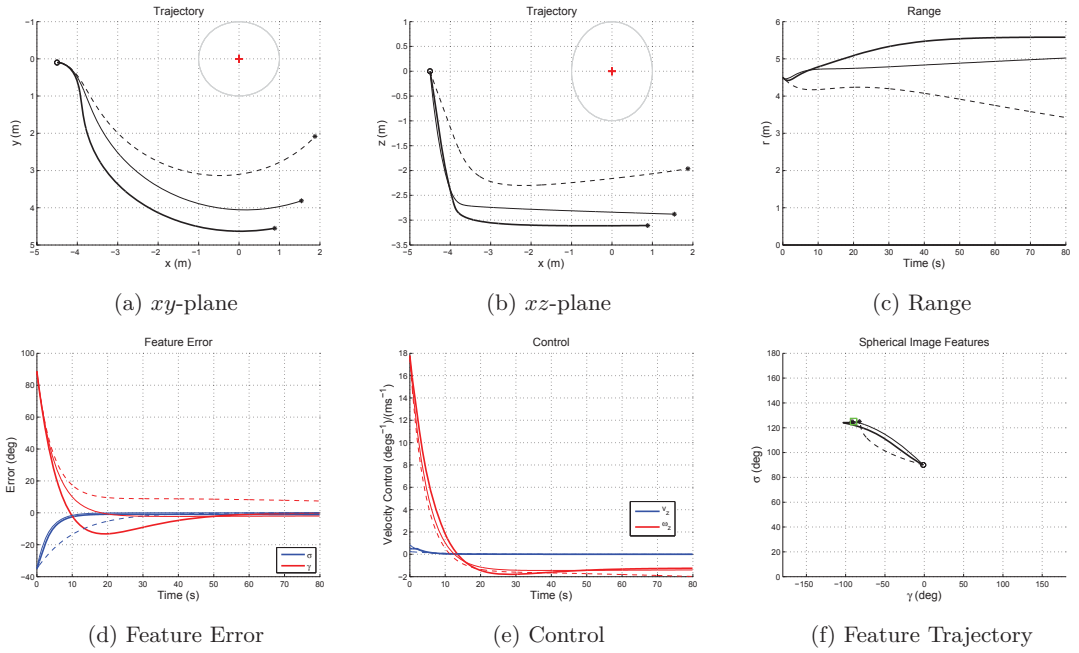

(a) xy-plane (b) xz-plane (c) Range

(d) Feature Error (e) Control (f) Feature Trajectory

Figure 4.5: Example aircraft behaviour in the object frame $\mathcal{F}_o$ for a static object with $\gamma^* = -\pi/2$ and $\sigma^* = 125\pi/180$ (□). Initial aircraft and object positions in $\mathcal{F}_w$ are $(1, 0.1, -1)$ and $(5, 0, -1)$ respectively. Using the original controller $r^* = 2$ $(--)$ and $r^* = 7$ $(-)$ is used. Using the augmented controller only $r^* = 7$ $(-)$ is used.

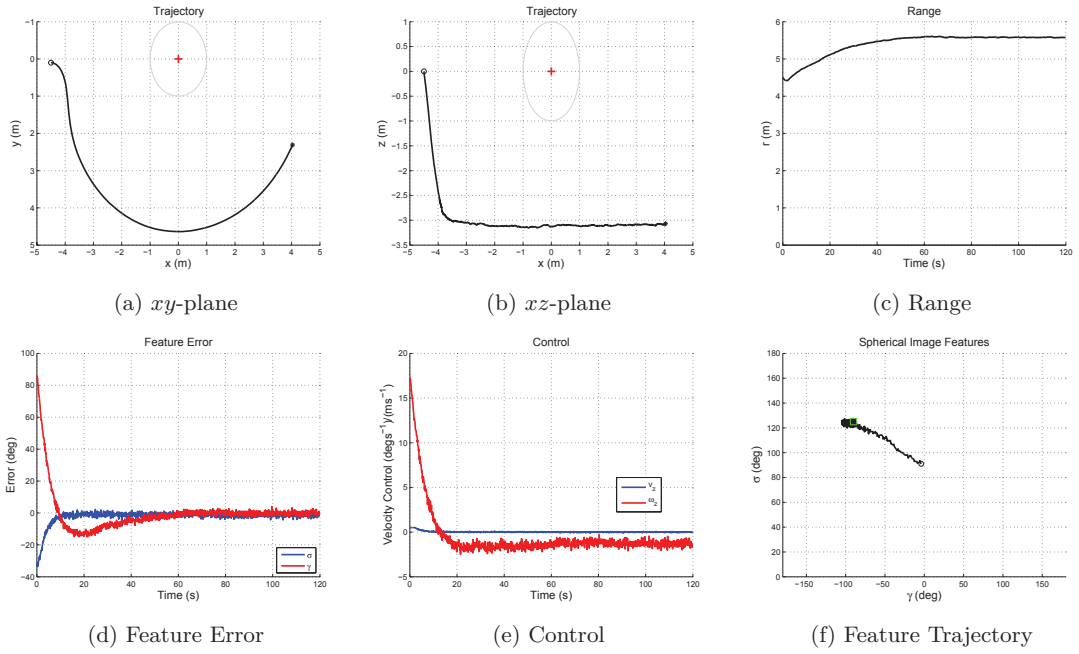

(a) xy-plane (b) xz-plane (c) Range

(d) Feature Error (e) Control (f) Feature Trajectory

Figure 4.6: Example aircraft behaviour in the object frame $\mathcal{F}_o$ for a static object with $\gamma^* = -\pi/2$ and $\sigma^* = 125\pi/180$ (□). Initial aircraft and object positions in $\mathcal{F}_w$ are $(1, 0.1, -1)$ and $(5, 0, -1)$ respectively. Additive noise $q(t)$ and the augmented controller with $r^* = 7$ $(-)$ is used.

4.3.2 Dynamic Objects

Consider an aircraft attempting to visually track a conical spiral about a dynamic constant velocity object. The object is approximated as a point feature in the image, and assumed to represent a moving conical apex. Consider applying the augmented spherical image-based control law

$$\mathbf{u}_3(t) = \hat{\mathbf{L}}_z^{-1}\left(-\lambda\mathbf{e}(t) - \mathbf{L}_{xy}^{*}\dot{\mathbf{x}}_{xy}(t)\right) + \lambda_i \int^{w} \mathbf{e}(t)\,dt \tag{4.20}$$

$$\mathbf{u}_3^{*}(t) = \max[\mathbf{a}, \min(\mathbf{b}, \mathbf{u}_3(t))] \tag{4.21}$$

where $\mathbf{u}_3^{*}$ denotes the saturated reference control vector for the attitude controllers and $\mathbf{a}$ and $\mathbf{b}$ are constant vectors denoting the upper and lower amplitude saturation limits respectively. As before, the control limits can be chosen based on aircraft performance and control constants.

Although a static object was assumed in the derivation of the controller, it is extensible to the dynamic object case. To explain, reconsider the initial assumptions ignoring the projected image feature velocity from the object motion. The object motion is unknown so it is difficult, if not impossible, to provide direct compensation. However, as the induced feature motion is expected to be relatively small, and considering the properties of spiral motion, a small variable image feature offset will likely result. The same integral control used to help compensate for the inaccurate reference range value may then suffice to help account for the unknown image feature motion.

To verify this hypothesis, a set of example encounters were simulated using the original and augmented controller and depicted in Fig 4.7. A worst case head-on encounter is considered, where the object moves with equal forward velocity to that of the aircraft but different heading. Based on the previous analysis, $r^{*} = 7$ and the integral gain is such that $\lambda_i = 0.01$. The remaining simulation parameters are the same as those used for the static case. This ensures only the effects of considering a dynamic object are observed.

The results confirm the extensibility of the augmented controller to the dynamic object case. As with the static case, the more aggressive augmented control enables the aircraft to approach the reference spiral quicker than the original controller. The augmented controller then actively compensates for the uncertain object motion, forcing the rapid convergence of the image feature error, and ensuring the reference spiral is tracked. Additionally, the resulting trajectory provides added benefits from a collision avoidance perspective. Specifically, a greater minimum lateral and vertical displacement (range) can be obtained.

As further evidence to support the use of the augmented controller, a crossing encounter is now included in addition to added image feature noise and imperfect actuation. The results are depicted in Fig 4.8. Again, expected aircraft behaviour is obtained and the addition of the integral term has helped compensate for the image feature motion induced by the non-zero object velocity. The control remains stable and the image feature error remains close to the reference value.

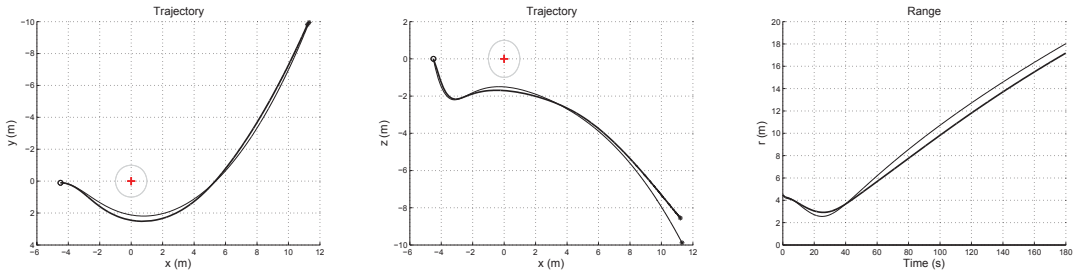

(a) xy-plane (b) xz-plane (c) Range

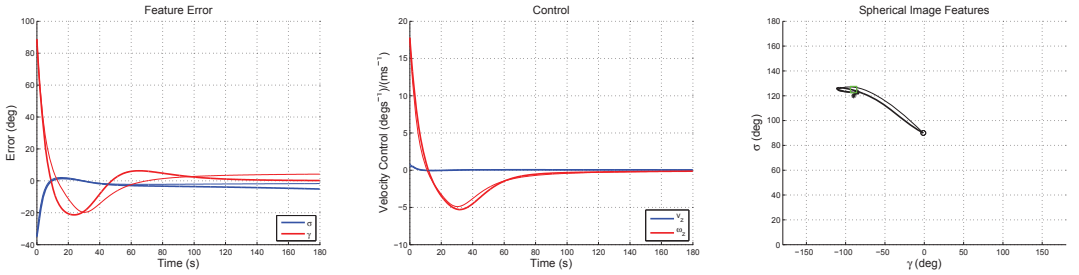

(a) Feature Error (b) Control (c) Feature Trajectory

Figure 4.7: Example aircraft behaviour in the object frame $\mathcal{F}_o$ for a head-on dynamic object using the original ($-$) and augmented ($-$) controller with $\gamma^* = -\pi/2$ and $\sigma^* = 125\pi/180$ (□). Initial aircraft and object positions in $\mathcal{F}_w$ are $(1, 0.1, -1)$ and $(5, 0, -1)$ respectively. The object moves with forward velocity $v = 0.1$m/s and $r^* = 7$.

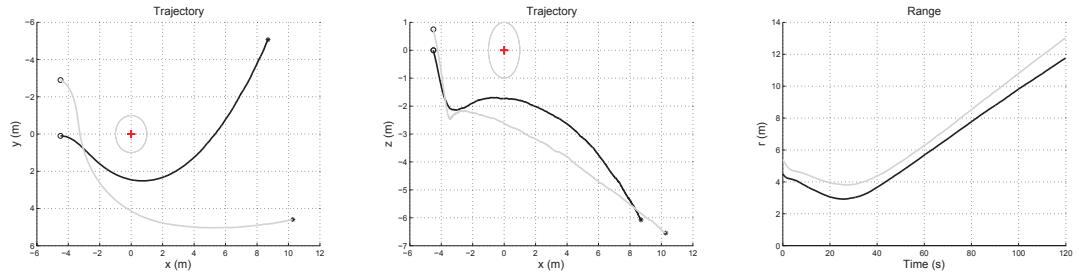

(a) xy-plane (b) xz-plane (c) Range

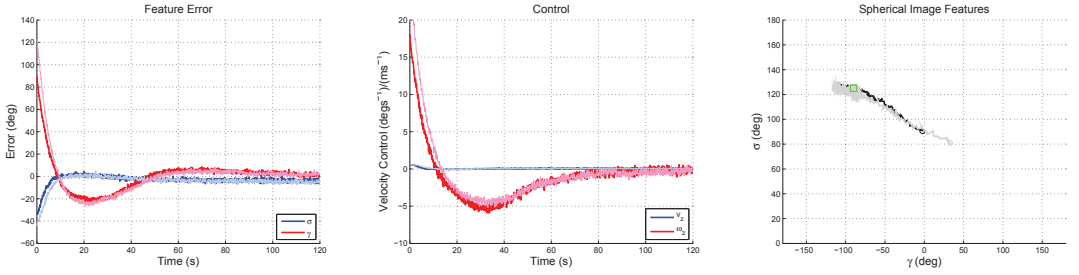

(a) Feature Error (b) Control (c) Feature Trajectory

Figure 4.8: Example aircraft behaviour in the object frame $\mathcal{F}_o$ for a head-on ($-$) and crossing ($-$) object with $\gamma^* = -\pi/2$ and $\sigma^* = 125\pi/180$ (□). Initial aircraft and object positions in $\mathcal{F}_w$ are $(1, 0.1, -1)$ and $(5, 0, -1)$ respectively. The object moves with forward velocity $v = 0.1$m/s and $r^* = 7$.

4.4 Collision Avoidance

The spherical image-based visual control laws derived in the preceding sections represent a suitable set of controllers for tracking a reference spiral. As such, they can be used directly in a closed-loop collision avoidance system when coupled with an appropriate avoidance and resolution decision strategy. This section details the implementation of such a collision avoidance system on a small aerial platform in a proof of concept approach.

First, modified avoidance and resolution decision strategies based on previously established concepts are presented (see §3). This is required to ensure the collision avoidance system is amenable to practical platform constraints and limitations. Second, results from real flight tests in a scaled collision avoidance environment are presented and analysed for a static object. They are the first recorded flight tests using such a spherical image-based visual servoing control scheme.

4.4.1 Avoidance & Resolution Decision

Consider a quadrotor system in which a perspective camera is used to approximate a spherical section. The camera is rigidly attached to the platform with its optical axis aligned to the platform x axis. Given the restricted field of view, the ideal reference image features $\mathbf{s}^*(125\pi/180, \pm\pi/2)$ cannot be used. Instead, reference values need to chosen that such that $|\gamma^*| < \pi/2$, resulting in a convergent spiral path. This is not ideal from a collision avoidance perspective. However, provided the right spiral direction is chosen, the aircraft will initially move in favour of collision avoidance. The reference image features can then be selected such that

$$|\gamma^*| > |\gamma(t_d)| \quad \wedge \quad \operatorname{sgn}\gamma^* = \operatorname{sgn}\gamma(t_d) \tag{4.22}$$

$$|\sigma^* - \pi/2| < |\sigma(t_d) - \pi/2| \quad \wedge \quad \operatorname{sgn}(\sigma^* - \pi/2) = \operatorname{sgn}(\sigma(t_d) - \pi/2) \tag{4.23}$$

where $\gamma(t_d)$ and $\sigma(t_d)$ are the initial azimuth and colatitude angles upon initial object detection instance t_d. Importantly, if the reference image features are chosen in this way, a resolution decision is mandatory to stop the controller from forcing the aircraft back toward the object in both vertical and lateral planes. To this end, the initial heading upon detection can be used directly to stop lateral avoidance (see §3). Additionally, the sign reversal of the z velocity command can be used to stop vertical avoidance. The refined resolution decision strategy for the lateral and vertical control can then be defined as

$$\omega_z^*(t) = \begin{cases} 0, & |\bar{\psi}(t)| < |\bar{\psi}(t_d)| \\ \omega_z^*(t), & \text{else} \end{cases}, \quad v_z^*(t) = \begin{cases} 0, & \operatorname{sgn} v_z^*(t) \neq \operatorname{sgn} v_z^*(t_d) \\ v_z^*(t), & \text{else} \end{cases} \tag{4.24}$$

where $\bar{\psi}$ denotes the aircraft heading bounded in the region $(-\pi, \pi)$. To help verify the refined resolution strategy, some example simulated encounters for a static and head-on dynamic object case when $\psi(t_d) = 0$ are shown in Fig 4.9 and Fig 4.10 respectively. In each case, successful resolution is demonstrated.

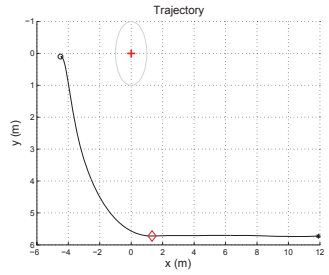

(a) xy-plane

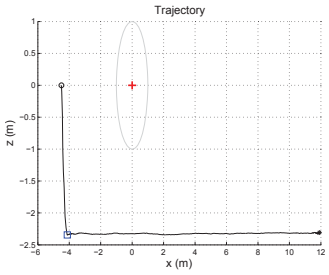

(b) xz-plane

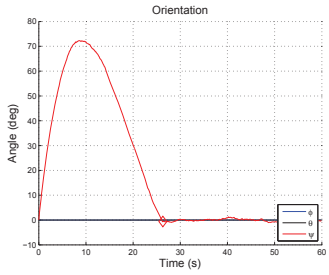

(c) Orientation

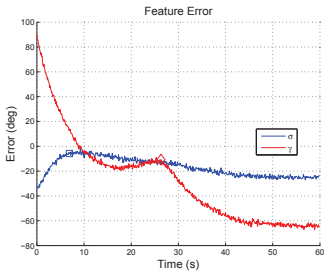

(a) Feature Error

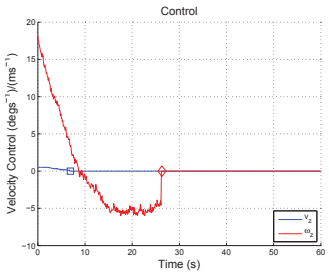

(b) Control

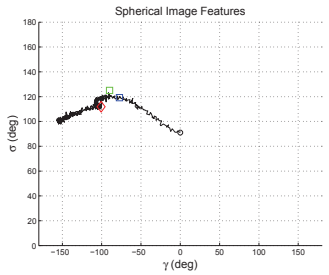

(c) Feature Trajectory

Figure 4.9: Example aircraft vertical (□) and lateral (◊) resolution instances for a static object with $\gamma^* = -\pi/2$ and $\sigma^* = 125\pi/180$ (□). Initial aircraft and object positions in $\mathcal{F}_w$ are $(1, 0.1, -1)$ and $(5, 0, -1)$ respectively, and shown in the object frame $\mathcal{F}_o$. The aircraft moves with velocity $v = 0.1$m/s and $r^* = 7$.

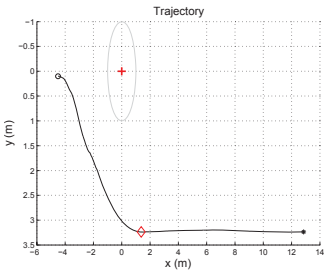

(a) xy-plane

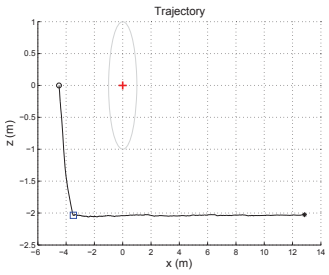

(b) xz-plane

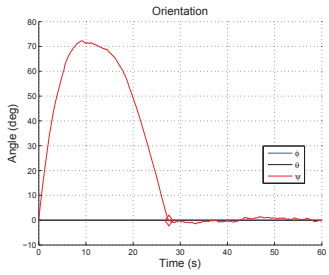

(c) Orientation

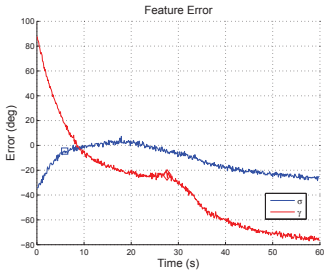

(a) Feature Error

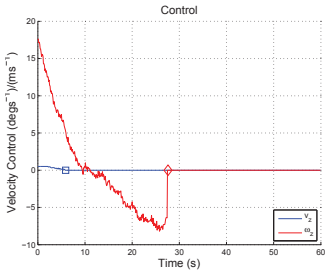

(b) Control

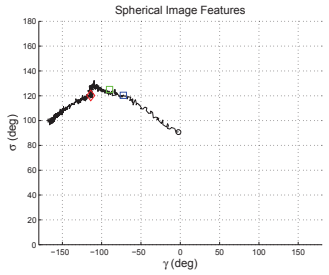

(c) Feature Trajectory

Figure 4.10: Example aircraft vertical (□) and lateral (◊) resolution instances for a dynamic object with $\gamma^* = -\pi/2$ and $\sigma^* = 125\pi/180$ (□). Initial aircraft and object positions in $\mathcal{F}_w$ are $(1, 0.1, -1)$ and $(5, 0, -1)$ respectively, and shown in the object frame $\mathcal{F}_o$. The aircraft and object move with forward velocity $v = 0.1$m/s and $r^* = 7$.

4.4.2 Implementation - Experimental System 1

Control Architecture

A small commercially available *Parrot - AR Drone* quadrotor was used to implement the spherical image-based visual controller and refined avoidance and resolution strategy. The vertical velocity and yaw rate were controlled directly using the S-IBVS scheme or a simple set of PID controllers via $\mathbf{u}_3^*$. The PID controllers are used to maintain constant heading and altitude when the collision has been resolved. Another set of PID controllers were used to control forward and lateral velocity via $\mathbf{u}_1^*$ and $\mathbf{u}_2^*$ respectively. To this end, the velocities were measured using a T40 Vicon, whilst the roll, pitch and yaw angles and associated rates were measured from both the Vicon and onboard IMU's. Reference roll and pitch angles were then derived via the velocity PID's and used by the on-board attitude controller to regulate forward and lateral velocity (see §2.2.3). The visual and non-visual controllers were implemented in C/C++ using the MAVwork framework (C++ API) in a custom ground station command and control module [287], and transmitted wirelessly to the quadrotor. The control architecture is depicted in figure 4.11.

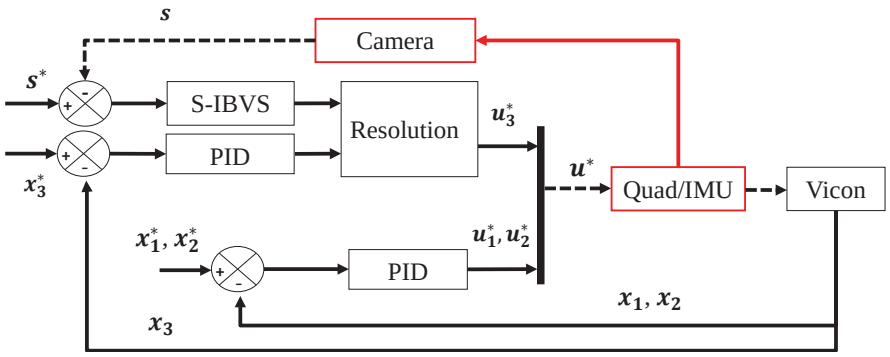

Figure 4.11: ARDrone quadrotor control architecture. The onboard (□) and ground station components (□) are shown along with wired (−) and wireless connections (− −).

Vision System

An onboard forward facing camera was used to collect images of a cone shaped fluorescent coloured collision object of height 40cm and radius 20cm. The camera was operating at 10Hz with a resolution 320×240pix and focal length of 206.6pix. Images were sent to the ground station and processed to obtain a point feature representation of the cone. To this end, a simple image processing algorithm using the OpenCV library [73] was used for robust detection and tracking. Specifically, a combination of RGB normalization, color segmentation and blob detection was used to consistently track the collision object. An outline of the image processing algorithm is given by algorithm 3.

Of note, the detection and tracking algorithm is susceptible to missed detections and false alarms (tracking the wrong object) depending on the background environment and object characteristics. The detection algorithm, including blob area and color thresholds, were tuned empirically to improve tracking performance in an arbitrary background. This was done to ensure the visual controller was being assessed and not the image processing details.

Algorithm 3 Image Processing Algorithm - Color Blob Detection

▷ Split each image I into red R, green G and blue B channels

▷ Normalise each channel and convert to greyscale images R', G' and B'

$$R'_{x,y} = \frac{r_{x,y} \cdot 255}{r_{x,y} + g_{x,v} + b_{x,v}}, \quad G'_{x,y} = \frac{g_{x,y} \cdot 255}{r_{x,y} + g_{x,y} + b_{x,y}}, \quad B'_{x,y} = \frac{b_{x,y} \cdot 255}{r_{x,y} + g_{x,y} + b_{x,y}},$$

▷ Threshold each greyscale image using user defined thresholds τ_R, τ_G and τ_B to obtain a set of binary images $\mathcal{R}$, $\mathcal{G}$ and $\mathcal{B}$.

$$\mathcal{R}_{x,y} = \begin{cases} \mathcal{R}_m, & R'_{x,y} > \tau_R \\ 0, & \text{else} \end{cases}, \quad \mathcal{G}_{x,y} = \begin{cases} \mathcal{G}_m, & G'_{x,y} > \tau_G \\ 0, & \text{else} \end{cases}, \quad \mathcal{B}_{x,y} = \begin{cases} \mathcal{B}_m, & B'_{x,y} > \tau_B \\ 0, & \text{else} \end{cases}$$

▷ Merge each channel and find all non-white blobs $\mathcal{C}$ within a user specified area $\mathcal{A}$

$$\mathcal{A}_{\text{min}} < \mathcal{A} < \mathcal{A}_{\text{max}}$$

▷ Assuming a single blob is detected, find its centre $\mathcal{C}_{\mathcal{X}_c,\mathcal{Y}_c}$ using the minimum $\mathcal{X}_m, \mathcal{Y}_m$ and maximum $\mathcal{X}_M, \mathcal{Y}_M$ dimensions of the blob

$$\mathcal{X}_c = (\mathcal{X}_m + \mathcal{X}_M/)2, \quad \mathcal{Y}_c = (\mathcal{Y}_m + \mathcal{Y}_M/)2$$

▷ Find the spherical image feature representation $\mathbf{s}(\sigma, \gamma)$ of the blob centre using the image centre I_{x_c,y_c} and focal length f

$$\sigma = \pi/2 - \arctan(\mathcal{Y}_c - y_c)/f, \quad \gamma = \arctan(\mathcal{X}_c - x_c)/f$$

Results & Analysis

Consider an instance similar to the previous simulation results. The quadrotor was allowed to fly toward a static object with forward velocity $v_x = 0.25$m/s and relative heading $\bar{\psi} \approx 0\,\text{deg}$, corresponding to a head-on encounter. When $r \approx 4.5$m, the visual control was activated, denoting the initial detection instance t_d and setting $\bar{\psi}^*$. An avoidance decision was made, moving the image features to a desired position $\mathbf{s}(\pm 10\pi/180 + \pi/2, \pm 25\pi/180)$ depending on the initial object image position. The choice of reference image features is an artefact of the practical limitations including the camera field of view constraints. As such, a large number of flight trials were undertaken using the spherical image based controller, with and without the integral term, using variable reference range value r^* and ignoring the optic flow from the forward velocity. The objective was not to tune the controller for ideal behaviour for the said encounter, but to verify the simulation results presented previously (were possible), and demonstrate spherical image-based control for collision avoidance.

In order to help demonstrate the robustness of the approach, results for a challenging case are presented here. The object is first detected to the right and below the de-rotated image centre. The controller without integral action is used with $r^* = 2$, $\lambda = 0.5$ and the optic flow from forward velocity is ignored. This is a difficult case, as the controller can no longer compensate for the poor range estimate or induced optic flow produced by

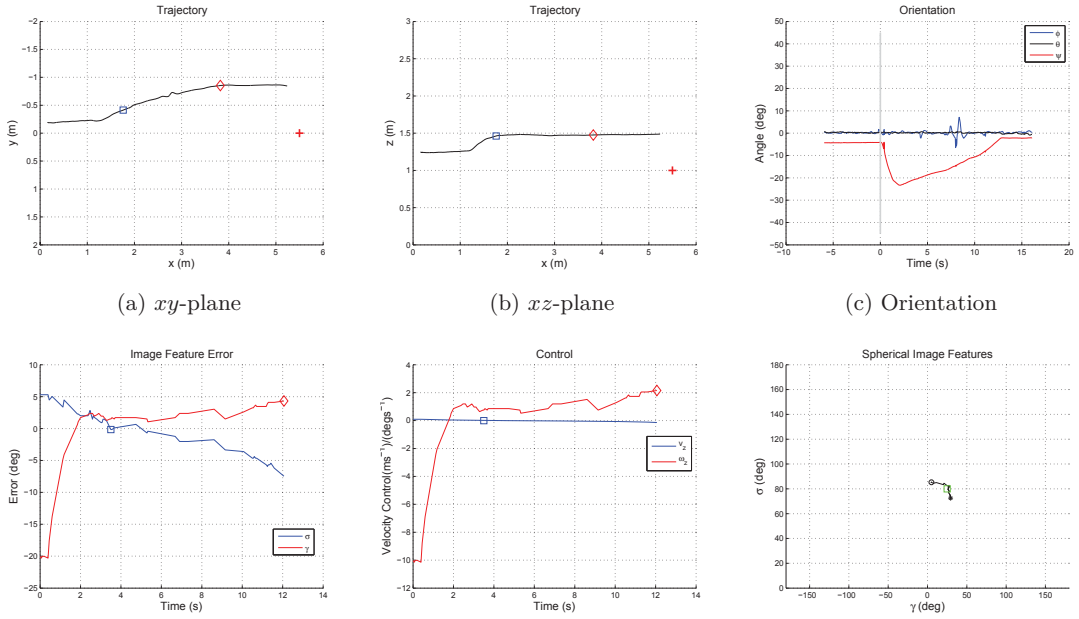

(a) xy-plane (b) xz-plane (c) Orientation

(a) Feature Error (b) Control (c) Feature Trajectory

Figure 4.12: Example real aircraft trajectories and image features for a static object with $\gamma^* = 25\pi/180$ and $\sigma^* = 100\pi/180$ (□), including lateral (◊) and vertical (□) resolution instances. Initial aircraft and object positions in $\mathcal{F}_w$ are approximately$(0, 0.25, -1)$ and $(5.5, 0, -1)$ respectively. The aircraft moves with velocity $v = 0.25$m/s with $r^* = 2$.

the forward velocity. However, as the reference image features cannot be selected such that $\gamma^* = \pm\pi/2$, the controller is more suitable than that including the integral term. To explain, an example set of results depicting the quadrotor position, orientation, input control and image feature motion is presented in Fig 4.12 and analysed below.

Despite the smooth vertical velocity control and general decrease in image feature error, the colatitude image feature error shows small discontinuities through small inconsistent increases. This can be attributed to small fluctuation in vertical height and imperfect image de-rotation (errors). Additionally, the variable image processing rate can adversely affect the control in a similar manner. In some cases, the image processing frequency drops to as little as 2Hz, forcing each control to be sustained for longer periods. Having not accounted for optic flow from forward velocity, the weaker z velocity commands have reduced the vertical separation.

The azimuth feature error decreases smoothly as the image features move toward the reference value. Recalling that because the azimuth angle is not dependant on the image de-rotation, the same discontinuities as for the colatitude feature error are not observed. The azimuth angle then moves past the desired position before attempting to settle with a positive offset error. This occurs as the optic flow from forward velocity has been neglected, forcing the controller to issue stronger ω_z commands. The effect is indeed analogous to overestimating r^*, justifying the controller selection. From a collision avoidance standpoint this is not a problem as the positive offset results in an attempt to follow a spiral of increased radius, and thus allows for an increased lateral separation at a faster rate. Using the augmented image-based controller with integral term in this case is not

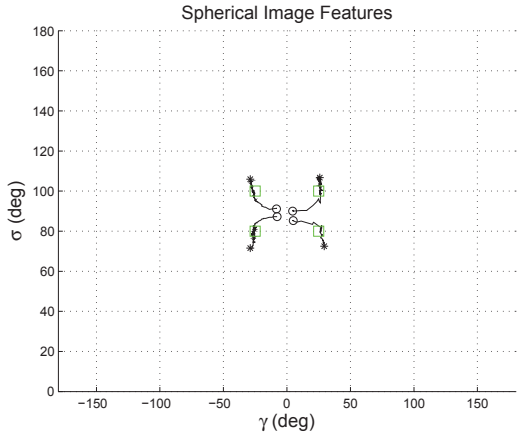

Figure 4.13: Image feature trajectories resulting from each of the possible avoidance decision outcomes. Initial (∘), final (∗) and reference (□) image feature locations are shown.

advantageous. This is because the desired azimuth $|\gamma^*| < \pi/2$, so a tighter spiral would be tracked at reduced lateral separation. Important to note, that if the reference azimuth could have been selected such that $\gamma^* = \pm\pi/2$, then the augmented controller would have been chosen and the optic flow from forward velocity included.

A successful resolution decision in the lateral and vertical planes is also shown, as the quadrotor ceases the avoidance behaviours at appropriate times. The image feature trajectory is allowed to move away from the reference value, as the quadrotor maintains its heading and altitude after the corresponding stopping time t_s. The quadrotor is allowed to return to its initial heading upon detection $\bar{\psi}(t_d)$, helping to reduce the deviation from the intended flight plan before object detection.

A successful avoidance decision for the lateral and vertical planes is also shown, as the reference image features are selected correctly. As further evidence of the avoidance functionality, the image feature trajectories from four separate near head-on encounters are shown in Fig 4.13. The avoidance criterion is successfully evaluated in each case, and the reference image feature chosen appropriately based on the initial image feature position. In the examples shown, each of the four possible outcomes are depicted, consisting of reference image features in each corner of the image.

Remark 1 Initially the assumption of zero roll and pitch velocity was removed and they were included in the control. It was found that the control law was quite sensitive to inaccurate velocity measurements, causing the quad to issue more aggressive and rapidly fluctuating z velocity commands. Removing the optic flow measured from pitch and roll, and adopting the control law derived in this chapter alleviated this problem. Alternatively, an appropriate filter could have been used to smooth the velocity variables, but at the risk of introducing additional system delay.

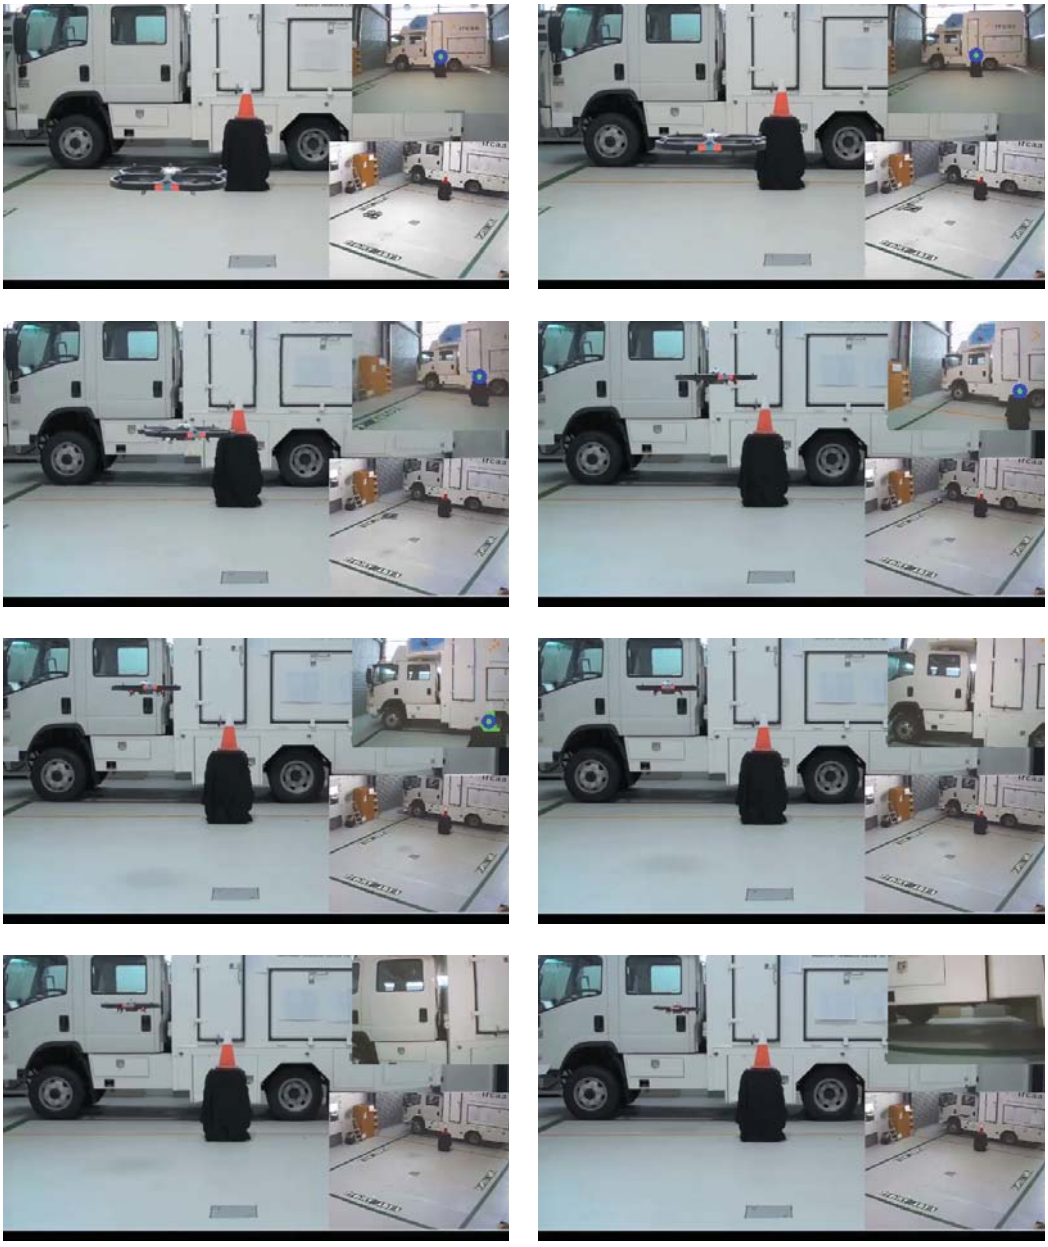

Figure 4.14: Example flight trial using spherical image-based visual control for collision avoidance of a small orange fluorescent cone (○) initially positioned in front of the platform. The controller regulates the image feature motion to the reference position , avoiding the object by spiralling to the left. The encounter evolves from $t = 0$s (*upper left*) to $t < 20$s (*lower right*).

4.5 Summary

This chapter introduced a set of novel spherical image based visual controllers for spiral tracking and collision avoidance, using principles from classic image-based visual servoing. Control strategies were derived that attempt to account for unknown object motion and thus range. The results were used to improve the control design through appropriate parameter selection, resulting in a conservative spiral tracking approach. A simple avoidance and resolution decision strategy was then included for collision avoidance. A practical implementation for static objects was demonstrated using a small quadrotor. Minor modifications to the avoidance and resolution decision strategies were included to account for practical constraints. The key attributes, benefits and drawbacks of using the spherical image-based visual control framework for collision avoidance include:

- The avoidance controller provides a flexible reactive control solution to spiral tracking and collision avoidance. The approach is computationally tractable, robust to range parameter uncertainty, easy to implement and amenable to multiple platform types (fixed and rotary wing). The effects of underestimating and overestimating range were also validated, providing insight into some design considerations. In particular, it was shown that an accurate range value is not necessarily required for effective collision avoidance, given the expected maximum detection distance. This helps to circumvent the difficult range estimation problem.
- Although a rough selection of the controller gain terms leads to satisfactory performance, some tuning is still required. As such, it may be difficult to obtain a general controller for an arbitrary aircraft and encounter scale. As such, stability can only be guaranteed in the single case. Additionally, the platform constraints cannot be explicitly included in the controller design, and instead rely on either tuning the gain appropriately or saturating the input. Given the success of recent optimal control approaches, merging similar concepts with visual control could alleviate such issues, with the possibility of improved controller performance and spiral tracking (control effort and stability).
- A simple image processing algorithm was used to detect an inconspicuous collision object. The algorithm was capable of robust object detection and tracking in a range of environments, but not completely free of false alarms (incorrect object tracking) and missed detections. The low frequency of such events did not cause significant issues for the controller. Often false alarms and missed detections were only temporary (< 3 frames), such that the avoidance control could easily be recovered. Although improvements can be made, it is noted that object detection and tracking is not the focus of this work.
- The reference colatitude angle σ^* has not yet been justified. Indeed, any reference such that $\sigma^* \neq \pi/2$ can be used. A means to better select the reference colatitude angle is required. This may be derived from a collision avoidance perspective, or a control perspective in which better assurance the image feature can be maintained is provided (see §5.4.5).

- The avoidance decision approach has not been fully analysed. Only a few cases are presented in both the simulated and practical environment using only a single threshold value (see §3). To fully analyse the avoidance decision, a large number of trials would be required under realistic conditions to provide a probabilistic performance assessment. This would be best accomplished using Monte-Carlo methods in a high fidelity simulation environment.
- The resolution decision approach has not been fully analysed. Only a few cases are presented in both the simulated and practical environment. Indeed, only a simple strategy that is not explicitly coupled to the visual control has since been presented. Considering the lateral resolution, a number of flight trials demonstrated the sensitivity of the resolution decision to environmental conditions, causing the avoidance to stop prematurely. Small variations in heading near the time of initial detection can cause the avoidance to stop before establishing a spiral, and indeed prevent sufficient separation. If the resolution decision could be coupled to the visual control such that resolution could only occur once established (or nearly established) on the spiral, a more robust resolution strategy would result. Additionally, a threshold based approach, similar to the avoidance decision, would conform to existing aviation collision avoidance system designs (see §5.5.2). The derivation of an appropriate threshold and subsequent performance analysis can then be accomplished using existing probabilistic performance analysis techniques (see §6).

Chapter 5

Visual Predictive Control

5.1 Outline

This chapter presents new spherical visual predictive control strategies for avoidance control. The control approaches extend existing visual predictive control concepts both theoretically and experimentally, resulting in novel contributions for general visual control applications.

First, optimal visual control approaches are reviewed in the context of See and Avoid. Specific attention is paid to image-based schemes with supporting practical implementations. Second, two novel Spherical Visual Predictive Control (S-VPC) schemes are derived. The first is based on existing Visual Predictive Control (VPC) schemes, while the second introduces new stability-based design concepts to derive a Quasi-Infinite Horizon Visual Predictive Control (QIH-VPC) approach. Third, a practical predictive control scheme is implemented using a small *AscTec - Hummingbird* in a scaled indoor collision avoidance environment in a proof-of-concept type approach. An improved resolution strategy is also presented and validated empirically. The control schemes are then analysed, highlighting their advantages and drawbacks.

5.2 Fundamentals & Prior Work

Despite the flexibility and simplicity of many classical image-based visual control schemes, they are not without their common drawbacks. Typically, classical approaches require tuning a gain term to simultaneously satisfy problem constraints and ensure sufficiently fast convergence. The gain term is often empirically derived, and can lead to sub-optimal behaviour that only guarantees stability in this single case. Additionally, the classical control structure cannot explicitly consider vehicle dynamics and actuator limitations in the control design. Such controller attributes were present in the partitioned spherical image-based controller derived in the previous chapter.

Optimal control based approaches have emerged to help address some of these issues. Control is derived by minimizing an objective function J over an infinite time horizon such that

$$\mathbf{u}^{\star}(\cdot) = \underset{\mathbb{U}}{\operatorname{argmin}}\ J(\dots, \dots) \tag{5.1}$$

where $\mathbf{u}^\star(\cdot)$ denotes the applied optimal control sequence and $\mathbb{U}$ denotes the set of possible controls (sequences), defined over the corresponding horizon. The objective function however is quite general and could contain not only image feature error, but a combination of image features, control effort, time and other application specific criteria. This allows flexibility in the control specifications whilst explicitly considering problem constraints. Included in the set of optimal control based approaches are tactical control designs such as path planning [321, 322], and various linear matrix inequalities [323, 324] formulations using linear differential inclusions (LDI) or otherwise. Although global in nature they require a unique solution to exist, which is not always guaranteed. They are often designed with robustness in mind and by considering worst case scenarios, result in overly conservative control schemes with larger computational overhead. This is particularly relevant when the online computation of a minimax problem is required [324]. Many existing approaches also assume a position-based visual control framework or require additional information such as object position. Such approaches have not been extended to spiral tracking, but are likely to be unsuitable considering the requirement for multiple image features from the same object.

Nonlinear model predictive or receding horizon control (NMPC) describes a particular class of optimal control based approaches, solved over a finite time or prediction horizon T_p instead [325]. As opposed to a global solution, NMPC allows the control law to be re-calculated periodically, incorporating a feedback mechanism into the system to help compensate for external disturbances. Generally, accurate process models are required for improved control. As good approximations for both image and robot kinematic models can often be derived a priori, the NMPC framework is well suited to visual control. In this case, the term visual predictive control (VPC) is commonly used and the optimal control problem to be solved becomes

$$\mathbf{u}^\star(\cdot) = \underset{\mathbb{U}}{\operatorname{argmin}}\ J(\mathbf{e}(t),\ \dots)\ |_t^{t+T_p} \tag{5.2}$$

Solving a finite horizon as opposed to an infinite horizon optimal control problem reduces the computational overhead. Significantly complex or large systems can then be solved provided the horizon time is reasonable. In VPC, the state vector may consist only of image features or a combination of vehicle states and image features. The nonlinear optimization problem is then solved over the resulting state space such that control, state and visibility constraints are managed accordingly. The system dynamics are often linearised, and a common objective function consisting of a weighted sum of quadratic terms in state and control is used. As with classical visual control, both position-based and image-based visual predictive control schemes can be derived.

For image-based schemes, object pose is not reconstructed and differences in various schemes often revolve around image feature representation, constraint inclusion and the objective function structure [326]-[328]. Recent theoretical studies on unconstrained [329] and constrained [330] predictive schemes have suggested the prediction horizon can further reduce the undesired coupling effects. Additionally, the increased prediction horizon can

result in straight forward camera motion, and more direct image feature trajectories in the presence of noise, range uncertainty and large camera displacements. In the special case using a quadratic objective function and prediction horizon of one sample, the control resulting from (5.2) is equal to that derived for the classical control scheme (4.3), as $\hat{L}_s^+$ approximates a least squares solution to the inverse image kinematics [337]. In any case, the additional benefit is evident in the explicit management of control limits, state constraints and visibility issues in a single optimisation. Of note, the visual predictive control framework has not been practically applied to aircraft control for any application.

As with classical image-based control, stability of visual predictive control is a major concern. It is then important to relate the stability attributes of general nonlinear model predictive control strategies to visual control. To this end, ensuring closed-loop stability for nonlinear model predictive control is difficult, but can be addressed in a number of ways. In most cases stability can be achieved by assuming the general form of the objective function and selecting suitable design parameters. These include the prediction horizon T_p, control horizon T_c and objective function weighting matrices [331]. The approach is typically used for practical applications and although it has proved to be effective for VPC [332]-[335], defining the parameters often requires experience. They are chosen to provide a trade-off between computational complexity and stability. Typically the prediction and control horizon are equal and tuning is predominantly on the prediction horizon length. If too short, stability issues may arise. If too long the computational expense increases significantly and the control horizon can be used to reduce complexity. To this end, move blocking approaches are commonly used. By considering a single constant control over multiple steps within the prediction horizon, the complexity of the optimization problem is considerably reduced. This has been studied in [336] and applied to VPC in [337]. Unfortunately, stability can no longer be guaranteed, constraint satisfaction may be difficult and the resulting constant control may be sub-optimal or infeasible. Tuning guidelines exist [338, 339], but due to the variability in system dynamics, no formal rules have been established for general NMPC or VPC. For visual control, the nonlinearity of the image dynamics results in considerable differences in system behaviour for each operating point, further complicating parameter selection. A loss of performance may then result as a compromise for ensured stability.

Alternatively, design methods exist that guarantee nominal closed-loop stability, without explicitly tuning system parameters. By altering the control structure, they rely only on the existence of a feasible solution at the initial time. Feasibility at each subsequent time is then guaranteed resulting in stable behaviour. In a first approach, a terminal equality constraint can be added [340, 341] such that (5.2) becomes

$$\mathbf{u}^\star(\cdot) = \underset{\mathbb{U}}{\operatorname{argmin}}\ J(\mathbf{e}(t),\ \dots)\ |_t^{t+T_p}\ \text{ s.t }\quad \mathbf{e}(t+T_p) = 0 \tag{5.3}$$

The reference state must be reached within a finite prediction horizon, which often leads to a restricted region of operation. So feasibility, and thus stability, becomes an issue for small prediction horizons. To ensure a satisfactory region of operation, a large prediction

horizon is required which brings increased computational burden. Additionally, exact numerical satisfaction of the terminal equality requires an infinite number of calculations. Approximations must be made for tractability, which can compromise stability results. One example is the addition of a terminal penalty term to the cost function to approximate the zero terminal constraint. The idea is to penalise deviation from the reference state at the end of the prediction horizon, to help force hypothetical zero terminal region constraint. Approaches based on this concept have been successfully applied to recent VPC schemes for robot manipulators [342, 343]. Again, this visual predictive control framework has not been practically applied to aircraft control for any application.

A second approach is to apply a terminal penalty term and terminal region to the control structure, forcing the final state to lie in a bounded region about the reference value at the end of the horizon [344, 345]. Within this region, convergence is then guaranteed. Transforming (5.2) into this framework then

$$\mathbf{u}^{\star}(\cdot) = \underset{\mathbb{U}}{\operatorname{argmin}}\ J(\mathbf{e}(t),\ \ldots, \mathbf{e}(t+T_p))\ |_{t}^{t+T_p}\ \text{ s.t}\quad \mathbf{e}(t+T_p) \in \Omega_{\varsigma} \tag{5.4}$$

where $\mathbf{e}(T_p)$ denotes a terminal penalty term and Ω_{ς} defines the terminal region. By including the terminal region constraint, the approach is less restrictive than the first. The requirement to satisfy a zero terminal constraint in finite time is removed, decreasing the potential for feasibility issues when using small prediction horizons.

The terminal region has some other convenient properties. The existence, shape and magnitude of the terminal region can also be used to discriminate between particularly nonlinear regions of the state space and help estimate the domain of attraction (region of operation). A large terminal region indicates the system dynamics are relatively linear about the associated reference state, resulting in a greater region of convergence. The reverse is true for small terminal regions. If applied to image-based control then, the terminal region could represent a particular image region, and its attributes would then be with respect to particular reference image features and image dynamics. The terminal region could then be used to help select better reference image features from a control perspective. Additionally, if servoing to a specific region of the image was required, better reference features within the region could then be selected, offering an alternate region-reaching control approach.

The terminal penalty term is used to penalise deviation from the reference states at the end of the horizon, similar to previously mentioned predictive schemes. However, the penalty term is now explicitly dependent (coupled) to a specific terminal region, and serves to virtually extend the horizon to infinity by approximating an upper bound on the objective function. As a result, the approach is commonly referred to as quasi-infinite horizon nonlinear model predictive control (QIH-NMPC).

Importantly, the terminal penalty term and terminal region for QIH-NMPC depend on multiple system specific factors, and cannot be chosen freely. They are in general difficult to determine, but can be approximated offline. This improves the design process by allowing analysis of some system attributes before controller implementation. Additionally,

by moving the majority of the computational burden offline, the online optimization can be solved sufficiently fast with standard computing power and modern solvers (optimisation algorithms) [347]. Of note, stability based design techniques using the quasi-infinite horizon approach have not been applied to visual predictive control of any form in any application.

Other stability-based designs exist for nonlinear model predictive control, but they are typically overly complex (Feedback Linearisation NMPC), have higher computational burden through increases degrees of freedom (Dual-Mode NMPC) or require additional tuning parameters that are difficult to determine (contractive NMPC) [325, 331]. As such, their extension to visual predictive control is difficult, and likely to be unsuccessful for real applications such as spiral tracking and collision avoidance.

The analysis above suggests that visual predictive control, including the quasi-infinite horizon framework, can offer some significant benefits over classical image-based control approaches. These benefits include improved robustness (added noise and parameter uncertainty), explicit constraint handling (including state, control and visibility constraints), improved reference image feature selection and nominal stability guarantees. Each of these features are particularly important for spiral tracking and aircraft collision avoidance. The visual predictive control approach therefore lends itself naturally to the See and Avoid problem environment, however it has not yet been investigated.

The remaining sections of this chapter detail the derivation, simulation, implementation and analysis of improved image-based controllers using predictive image-based approaches, that are suitable to See and Avoid. Specifically, two novel spherical visual predictive controllers are presented. The first is designed using quasi-infinite horizon approaches and validated in simulation. The second is designed from a practical perspective and validated empirically. The controllers are then used to design an improved practical closed loop vision-based collision avoidance system for static and dynamic objects.

5.3 Spherical Visual Predictive Control

Control Structure

Visual predictive control is based on well established nonlinear model predictive control [325]. It requires a reference state vector, a process model, an objective function and a solver for the resulting optimization problem. At each sampling instance, an optimal control sequence of camera velocities $\mathbf{u}^\star(\cdot)$ is found that minimizes the objective function J under a set of nonlinear visibility, state and control constraints. The objective function is calculated over a finite prediction horizon T_p using a process model $\mathbf{f}(\cdot)$ that may include a combination of image kinematics and vehicle dynamics. The first of the control sequence is applied before the process is repeated each time a new measurement arrives. Essentially, a constrained, finite horizon, nonlinear optimal control problem is solved at each sampling instance.

Typically the objective function takes the general form, consisting of a weighted sum of quadratic terms in state and control. The specific weights are typically fixed and defined by the matrices $\mathbf{Q}$ and $\mathbf{R}$ for state and control respectively, such that

$$J(\mathbf{e}(t), \mathbf{u}(t)) = \int_{t}^{t+T_p} \mathbf{e}(\tau)_Q^2 + \mathbf{u}(\tau)_R^2 \, d\tau \tag{5.5}$$

However, it is possible to manipulate the objective function structure online. Terms may be added, removed or altered, reflecting their relative importance to the intended application. Any objective function weighting matrices may be constant or time varying over the prediction horizon or successive optimizations. Recently, a progressively incremented positive definite state penalty matrix $\mathbf{Q}(\tau)$ is used to penalise states more heavily toward the end of the horizon [337]. No penalty on control was used, and is essentially handled by the control constraints. This visual predictive control problem can be defined in continuous time as a single online optimisation such that

$$\mathbf{u}^\star(\cdot) = \underset{\mathbb{U}}{\operatorname{argmin}} \; J(\mathbf{e}(t)) \tag{5.6}$$

where

$$J(\mathbf{e}(t)) = \int_{t}^{t+T_p} \mathbf{e}(\tau)_{Q(\tau)}^2 \, d\tau \tag{5.7}$$

s.t

$$\dot{\bar{\mathbf{s}}} = \mathbf{f}(\bar{\mathbf{s}}, \bar{\mathbf{u}}), \qquad \bar{\mathbf{s}}(t) = \mathbf{s}_0 \tag{5.8a}$$

$$\bar{\mathbf{e}} = \bar{\mathbf{s}}(\tau) - \mathbf{s}^* \tag{5.8b}$$

$$\mathbf{Q}(\tau) = 2\mathbf{Q}(\tau - T) \tag{5.8c}$$

$$\bar{\mathbf{u}}(\tau) = \rho, \qquad \tau \in [t, t + T_p] \tag{5.8d}$$

$$\rho \in \mathbb{U} \tag{5.8e}$$

where T defines the sampling time and internal variables used in prediction are distinguished using a bar. This is because the internal variables will not in general be equal to their actual future values. To this end, the trajectory of the predicted image features $\bar{\mathbf{s}}(\cdot)$ results from adopting the control sequence $\bar{\mathbf{u}}(\cdot)$, according to the process model $\mathbf{f}(\bar{\mathbf{s}}, \bar{\mathbf{u}})$ and starting at the true initial image features $\mathbf{s}_0$. The predicted image feature error $\bar{\mathbf{e}}$ is determined by the smallest difference between the reference and actual image features.

Applying the control structure to spherical image features, the image feature error then corresponds to a modulo 2π subtraction denoted by $\ominus$ such that the $\bar{\mathbf{e}} \in (-\pi, \pi)$. Applying the control structure to spiral tracking using a single point (and a spherical image features), the process model $\mathbf{f}(\bar{\mathbf{s}}, \bar{\mathbf{u}})$ represents the image kinematics defined by the optic flow equation (5.9). Applying the framework to aircraft spiral tracking, the process model must then include a combination of image kinematics and vehicle dynamics. To keep the approach general, consider first a process model consisting only of the image kinematics for a static object. Recalling the image kinematics define the relationship between image feature velocity $\dot{\mathbf{s}}(t)$ and camera translational and angular velocity components $\dot{\mathbf{x}}(t)$

$$\dot{\mathbf{s}}(t) = \mathbf{f}(\mathbf{s}(t), \mathbf{u}(t)) = \mathbf{L}_s(\mathbf{s}(t), r(t))\, \dot{\mathbf{x}}(t) \tag{5.9}$$

where $\mathbf{L}_s$ is the full spherical image Jacobian and $\dot{\mathbf{x}}(t) \in \mathbb{R}^6$. We can then express $\dot{\mathbf{x}}(t)$ in terms of the controlled velocities $\mathbf{u}(t)$ and the remaining velocities $\mathbf{v}(t)$ such that $\dot{\mathbf{x}}(t) = [\mathbf{u}(t)\ \mathbf{v}(t)]$ where $\mathbf{u} \in \mathbb{R}^m$, $\mathbf{v} \in \mathbb{R}^n$ and $m+n = 6$. Recalling, an aircraft can track a spiral trajectory with a fixed forward velocity v_x and variable vertical velocity v_z and yaw rate[1] ω_z. All three velocities cannot be directly controlled using visual feedback, as only a single point is being observed. As such, setting $\mathbf{u} = \mathbf{u}_3 = [v_z\ \omega_z]$ and $\mathbf{v} = [v_x\ v_y\ \omega_x\ \omega_y]$, the image kinematics can be re-written as

$$\dot{\mathbf{s}}(t) = \mathbf{L}_z(\mathbf{s}(t), r(t))\mathbf{u}_3(t) + \mathbf{L}_{xy}(\mathbf{s}(t), r(t))\mathbf{v}(t) \tag{5.10}$$

where

$$\mathbf{L}_z(\mathbf{s}(t), r(t)) = \begin{pmatrix} \frac{\sin\sigma(t)}{r(t)} & 0 \\ 0 & -1 \end{pmatrix} \tag{5.11}$$

$$\mathbf{L}_{xy}(\mathbf{s}(t), r(t)) = \begin{pmatrix} \frac{-\cos\sigma(t)\cos\gamma(t)}{r(t)} & \frac{-\cos\sigma(t)\sin\gamma(t)}{r(t)} & \sin\gamma(t) & -\cos\gamma(t) \\ \frac{\sin\gamma(t)}{r(t)\sin\sigma(t)} & \frac{-\cos\gamma(t)}{r(t)\sin\sigma(t)} & \frac{\cos\gamma(t)\cos\sigma(t)}{\sin\sigma(t)} & \frac{\sin\gamma(t)\cos\sigma(t)}{\sin\sigma(t)} \end{pmatrix} \tag{5.12}$$

and $\mathbf{L}_z$ and $\mathbf{L}_{xy}$ represent partitions of the spherical image Jacobian such that $\mathbf{L}_s = (\mathbf{L}_z\ \mathbf{L}_{xy})$. This is the same partitioning used in the classical image-based control approach, and ensures decoupled control in the vertical xz and lateral xy planes in the camera frame (§4.3). As in the case for classical image-based control, the forward velocity can then be

[1] These may be used directly or transformed into angular displacement commands, such as commonly done for fixed wing aircraft.

regulated by an external controller, and a fixed reference range value can be assumed such that (5.10) reduces to

$$\dot{\mathbf{s}}(t) = \mathbf{L}_z^*(\mathbf{s}(t))\mathbf{u}_3(t) + \mathbf{L}_{xy}^*(\mathbf{s}(t))\mathbf{v}(t) \tag{5.13}$$

where $\mathbf{L}_z^*$ and $\mathbf{L}_{xy}^*$ denote approximations of the partitioned image Jacobian taken for $r(t) = r^*$. Again, assuming the forward velocity v_x remains relatively constant and the image features are taken from de-rotated images, then $\mathbf{v}(t)$ can be considered constant as $v_y(t) \approx \omega_x(t) \approx \omega_y(t) \approx 0$. The velocity $\mathbf{v}(t)$ may then be passed as a predetermined vector parameter $\mathbf{v}$. Re-writing (5.13) in standard ODE form, the local process model can then be defined as

$$\dot{\mathbf{s}}(t) = \mathbf{f}(\mathbf{s}(t), \mathbf{u}_3(t)) \tag{5.14}$$

$$\mathbf{f}(\mathbf{s}(t), \mathbf{u}_3(t)) = \mathbf{L}_z^*(\mathbf{s}(t))\mathbf{u}_3(t) + \mathbf{L}_{xy}^*(\mathbf{s}(t))\mathbf{v} \tag{5.15}$$

There are four important attributes of the aforementioned visual predictive control structure. First, the process model does not take into account that the object may be moving. However, recalling that the induced optic flow from the object motion $\frac{\partial \mathbf{s}}{\partial t} \propto \frac{1}{r}$, the model mismatch will be small. One of the benefits of using the model predictive control structure is the inherent robustness to such model mismatch or parameter uncertainty. Therefore, this property can be directly leveraged to simplify the control design for both static and dynamic objects[2]. Second, move blocking is used as a single constant control ρ makes up the control sequence for the entire prediction horizon. Although computationally efficient, move blocking results in potentially sub-optimal solutions, having not considered the full richness of control actions over the prediction horizon [336]. Third, although local stability is ensured, it is difficult to infer the domain of attraction for a given reference or consider any global stability properties. Large prediction horizons increase the computational complexity, whilst small horizons will inherit similar stability properties as classical image-based approaches. Fourth, no visibility constraints have been imposed such that the image features are not restricted to a particular image region. They can however be added in a similar manner to the control constraints defined by (5.8e). The same can be said for the state constraints, and can be important if specific vehicle dynamics were to be included in the process model defined by (5.8a).

Remark The predictive control scheme presented in this section shall be referred to as classical visual predictive control (C-VPC), and will be used for comparative purposes. When required, the objective function will be augmented with a weighted quadratic term in control $\mathbf{u}(\tau)_R^2$ where $\mathbf{R} \succ 0$ to ensure a fair comparison can be made such that

$$J(\mathbf{e}(t), \mathbf{u}(t)) = \int_t^{t+T_p} \mathbf{e}(\tau)_{Q(\tau)}^2 + \mathbf{u}_3(\tau)_R^2 \, d\tau \tag{5.16}$$

[2]If the range is small, there is the possibility of significant steady state error or offset, to which an additional integral controller could be added as in the classical image-based approach.

Control Solution

Solving the online optimization problem (5.6) requires the solution of a constrained nonlinear finite horizon optimal control problem. The computational complexity lies in the solution of the minimization problem and depends, in part, on the structure of the objective function and process dynamics. For a quadratic cost function structure, the solution to the minimisation problem is often found using numerical methods including gradient decent, interior point (barrier) methods and variants of Sequential Quadratic Programming (SQP) [348]-[350] including Active-Set strategies [351]-[354]. The latter approach, as the name suggests, attempts to successively solve a series of Quadratic Programmes (QP) in an ordered manner. The problem is first broken down into a number of manageable quadratic programs, and then solved in a hierarchical fashion. The field is well studied and many solution methods, or solvers, have been proposed. For this work, an important requirement for the solver is to enable real-time operation and manage the visual predictive control structure given the inclusion of both image kinematics and aircraft dynamics.

Fortunately, solvers using the Active-Set strategy have recently become commercially available in software packages such as MATLAB's Control Systems Toolbox [70]. Such tools are useful for development and simulation, but extensions to real platforms are limited. As a result, freely available solvers using Active-Set strategy have recently become available, such as the ACADO Toolkit [347]. The package includes a MATLAB interface for development and automatic code generation capability to realise optimised re-configurable code for embedded applications. The solver can also meet the requirements for real-time operation of complex systems, given appropriate selection of controller and solver parameters. A variant of the quadratic program solver qpOASES is used in the ACADO package, and is capable of managing infeasible initial states, searching for alternate solutions when required [355, 356]. Additionally, the solver can also be tuned to behave like the active set routine in MATLAB, with the added benefit of faster computational time. In this way, the appropriate solver can then be selected to best match the problem structure. To this end, this work leverages both MATLAB and ACADO to solve various visual predictive control problems (in simulation and implementation), exploiting their benefits for a given control structure whilst ensuring fair performance comparisons can be made. Of note, developing a new nonlinear optimal control problem solver or routine is out of scope of this work.

5.4 Quasi-Infinite Horizon Visual Predictive Control

Transforming the visual predictive control problem into the quasi-infinite horizon nonlinear control framework requires the offline calculation of some additional parameters. First, the control problem is setup whereby the image kinematics are shifted and linearised about the reference point. Second, a stabilising locally linear controller is used to derive a terminal penalty matrix $\mathbf{P}$ to augment the objective function. Third, a terminal region Ω_ς about the operating point is calculated and added as an extra nonlinear constraint. The terminal region is calculated in a similar manner to that outlined in [345], but presented in a new and simple framework that can be directly and easily implemented using common software tools. The online optimization problem defined in (5.6)-(5.8) can then be reformulated with the inclusion of these additional controller parameters.

5.4.1 Problem Setup

Considering the spiral tracking application, the reference image features $\mathbf{s}^*$, range r^* and forward velocity v_x^* are first selected. The control constraint domain $\mathbb{U} \subset \mathbb{R}^2$ must then be defined, and is typically a function of the vehicle limitations such as turn or climb rate. Assuming a constant forward velocity, the reference controls $\mathbf{u}_3^*$ can then be calculated using circular motion equations or otherwise. If $\mathbf{u}_3^* \notin \mathbb{U}$ then the reference spiral is infeasible and new reference image features, range or forward velocity must be selected until the control constraints are satisfied. This is not a drawback of the controller, but results from the vehicle's physical constraints and chosen application.

The image kinematics representing the process model $\mathbf{f}(\mathbf{s}(t), \mathbf{u}_3(t))$ must be shifted about the reference operating point such that $\mathbf{f}(\mathbf{0}, \mathbf{0}) = 0$. Re-writing equation (5.15) for the shifted model

$$\begin{aligned}\mathbf{f}(\mathbf{s}(t), \mathbf{u}_3(t)) = \mathbf{L}_z(\mathbf{s}(t) + \mathbf{s}^*)(\mathbf{u}_3(t) + \mathbf{u}_3^*) \\ + \mathbf{L}_{xy}(\mathbf{s}(t) + \mathbf{s}^*)\mathbf{v}\end{aligned} \tag{5.17}$$

where the associated control constraint domain must also be shifted such that $\mathbb{U} = \{\mathbf{u}_3^{\min} \leq (\mathbf{u}_3(t) + \mathbf{u}_3^*) \leq \mathbf{u}_3^{\max}\}$. Recalling that $\mathbf{v}$ can be approximated as a constant maintained by external controllers, this system can then be linearized about $\mathbf{s}(t) = \mathbf{0}$ and $\mathbf{u}_3(t) = \mathbf{0}$. Using small disturbance theory, the resulting state-space representation for the process model is

$$\dot{\mathbf{s}}(t) = \mathbf{A}_s(t) + \mathbf{B}_s(t)[\mathbf{u}_3(t)\ \mathbf{v}]^T \tag{5.18}$$

where $\mathbf{A}_s(t) = \partial\mathbf{f}(\mathbf{s}, \mathbf{u}_3)/\partial\mathbf{s}$ and $\mathbf{B}_s(t) = \partial\mathbf{f}(\mathbf{s}, \mathbf{u}_3)/\partial\mathbf{u}_3$. The reference control and image features for this system are now zero vectors as a result of the shifted process model. Equations (5.17) and (5.18) have been derived with consideration to the necessary conditions required for general QIH-NMPC control design as stated in Appendix E.4.

5.4.2 Terminal Penalty Term

The terminal penalty matrix $\mathbf{P}$ is used to apply a quadratic terminal cost $\bar{\mathbf{s}}(t+T_p)^2_P$ within the objective function. The term is used to penalise deviation from the reference state at the end of the prediction horizon, and provide an upper bound on the objective function $\forall t > T_p$ where a fictitious locally linear state feedback controller is assumed. This linear controller is derived for the shifted process model defined by (5.18) and used directly to calculate $\mathbf{P}$.

First, a static gain matrix $\mathbf{K}_v$ must be defined about the reference value for a linear state feedback controller such that $\mathbf{u}_3(t) = \mathbf{K}_v\mathbf{s}(t)$. The choice of controller is somewhat arbitrary, however by solving a linear quadratic regulator (LQR) type problem, an optimal linear controller may be found. The associated Riccati equation

$$\mathbf{A}^T\mathbf{S} + \mathbf{S}\mathbf{A} - (\mathbf{S}\mathbf{B})\mathbf{R}^{-1}(\mathbf{B}^T\mathbf{S}) + \mathbf{Q} = 0 \tag{5.19}$$

is first solved for $\mathbf{S}$, where the state and control weighting matrices are defined by $\mathbf{Q}$ and $\mathbf{R}$ respectively. Then $\mathbf{K}_v = -\mathbf{R}^{-1}(\mathbf{B}^T\mathbf{S})$ is found such that asymptotic stability of the closed loop system ($\mathbf{A}_k = \mathbf{A} + \mathbf{B}\mathbf{K}_v$) is guaranteed in a region about the reference values. Note the sign reversals that are required due to the use of an LQR and the definition for $\mathbf{u}(t)$. As $\sigma^* \notin \{0,\ \pi\}$, the linearised system is both controllable and observable such that the pair $\mathbf{A}$, $\mathbf{B}$ is stabilisable. The existence of a suitable LQR controller is thus ensured provided $\mathbf{Q} \succeq 0$, $\mathbf{R} \succ 0$.

Second, the maximum real eigenvalue of the closed loop system $\lambda_{\max}(A_k)$ is found by solving the Lyapunov equation for $\mathbf{P}$ given by

$$(\mathbf{A}_k + \kappa\mathbf{I})^T\mathbf{P} + \mathbf{P}(\mathbf{A}_k + \kappa\mathbf{I}) = -(\mathbf{Q} + \mathbf{K}_v^T\mathbf{R}\mathbf{K}_v) \tag{5.20}$$

The constant κ is chosen such that $0 \leq \kappa < -\lambda_{\max}(\mathbf{A}_k)$ to ensure (5.20) admits a unique positive definite solution. To ensure a larger terminal region whilst retaining good control performance, the poles of the closed loop system are shifted such that $\kappa = -0.9\lambda_{\max}(\mathbf{A}_k)$. It was shown in [331] and verified for visual control that κ should be chosen near, but not equal to $-\lambda_{\max}(\mathbf{A}_k)$. If $\kappa \ll -\lambda_{\max}(\mathbf{A}_k)$ a smaller terminal region would result. If $\kappa = -\lambda_{\max}(\mathbf{A}_k)$, then the terminal penalty matrix will be large and may degrade control performance. Note that other values for κ are suitable, but the limits on κ for acceptable performance for visual predictive control are yet to be determined.

5.4.3 Terminal Region

The terminal region Ω_ς is an ellipse on the spherical imaging surface defined such that $\mathbf{s}^T\mathbf{P}\mathbf{s} \leq \varsigma$ where $\varsigma \geq 0$. The method for calculating the terminal terminal region is based on the methodology of [331, 345] and extended to provide an explicit formulation amenable to direct implementation in MATLAB. To find the terminal region Ω_ς a series of optimization problems are solved off-line, defined by algorithm 4 and using equations (5.21)-(5.25). Each optimization problem can be solved using the global optimization function *MultiSearch*. It requires nomination of a local solver and associated solver method. To this end, the *fmincon* function with sequential quadratic programming (SQP) solver can be used.

The terminal region is invariant under the locally linear controller, which itself is fictitious in the sense it is never implemented and only used to derive these controller parameters. Feasibility of the open loop optimal control problem at the initial time implies the image features will lie within the terminal region at the end of the prediction horizon, ensuring closed-loop asymptotic stability. An example is shown in Fig 5.1 where an arbitrary terminal region Ω_ς and feasible region M are drawn over the spherical surface $\mathcal{S}^2$. Strictly speaking, $\Omega_\varsigma \neq M$ but it is reasonable to assume that the larger the terminal region, the larger the domain of attraction such that $\Omega_\varsigma \subseteq M$. Importantly, a larger terminal region will generally result in a larger terminal penalty matrix. The penalty matrix must not be too large such that control performance is degraded. This is an important trade-off when selecting the prediction horizon. A smaller prediction horizon may be used for a larger terminal region to decrease computational overhead without degrading performance.

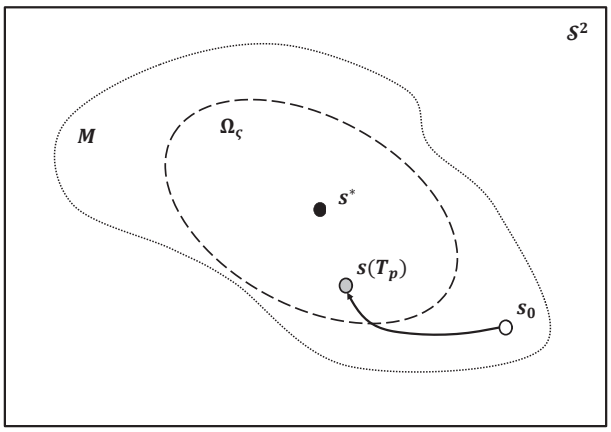

Figure 5.1: Example feasible initial image feature $\mathbf{s}_0$ moving from the feasible region M to the terminal region Ω_ς over the prediction horizon T_p.

Algorithm 4 Terminal Region Ω_ς Calculation

Set $j = 1$

PROBLEM A
Solve (5.21) subject to (5.24)
set $\varsigma_j = \mathbf{s}^\pi \mathbf{P} \mathbf{s}^\pi$

PROBLEM B
Set $\mathbf{s}^\varsigma = \mathbf{s}^\pi$
while $\mathbf{s}^\varsigma > 0$ **do**
  Solve (5.21) s.t (5.24)-(5.25)
  Solve (5.22) s.t (5.23)-(5.25)
  Set $j = j + 1$
  set $\varsigma_j = \frac{2}{3}\varsigma_{j-1}$

end while
Set $\varsigma = \mathbf{s}^{\varsigma_{j-1}} \mathbf{P} \mathbf{s}^{\varsigma_{j-1}}$

$$\mathbf{s}^\pi = \underset{\mathbf{s}}{\operatorname{argmax}} \, \|\mathbf{s}\| \tag{5.21}$$

$$\mathbf{s}^\varsigma = \underset{\mathbf{s}}{\operatorname{argmin}} \left\{ \kappa \cdot \mathbf{s}^T \mathbf{P} \mathbf{s} - \mathbf{s}^T \mathbf{P} \varrho(\mathbf{s}) \right\} \tag{5.22}$$

$$\varrho(\mathbf{s}) = \mathbf{f}(\mathbf{s}, \mathbf{u}_3) - \mathbf{A}_s \mathbf{s} \tag{5.23}$$

$$\mathbf{0} \leq \mathbf{K}_v \mathbf{s} - \mathbf{u}_3^{\min} \tag{5.24a}$$

$$\mathbf{0} \leq \mathbf{u}_3^{\max} - \mathbf{K}_v \mathbf{s} \tag{5.24b}$$

$$0 \leq \varsigma_j - \mathbf{s}^T \mathbf{P} \mathbf{s} \tag{5.25}$$

Remark Finding the true maximal terminal region is difficult, but the approach outlined in this thesis can be considered a suitable approximation having verified the results of [345] with the above algorithm. Note also that solving the first optimization in problem B is generally not required, but its solution is used as an initial value to speed up the second optimization in problem B. Additionally, when using *fmincon* in MATLAB, (5.21) must be transformed to a minimization problem.

5.4.4 Online Implementation

The online quasi-infinite visual predictive control (QIH-VPC) optimisation problem can now be defined as

$$\mathbf{u}_3^\star(\cdot) = \underset{\mathbb{U}}{\operatorname{argmin}}\ J(\mathbf{s}(t), \bar{\mathbf{u}}_3(\cdot)) \tag{5.26}$$

where

$$J(\mathbf{s}(t), \bar{\mathbf{u}}_3(\cdot)) \;=\; \bar{\mathbf{s}}(t+T_p)_P^2 \;+\; \int_t^{t+T_p} \bar{\mathbf{s}}(\tau)_Q^2 + \bar{\mathbf{u}}_3(\tau)_R^2\, d\tau \tag{5.27}$$

s.t

$$\dot{\bar{\mathbf{s}}} \;=\; \mathbf{f}(\bar{\mathbf{s}}, \bar{\mathbf{u}}_3), \qquad \bar{\mathbf{s}}(t) = \mathbf{s}_0 \tag{5.28a}$$

$$\bar{\mathbf{u}}_3(\tau) \;\in\; \mathbb{U}, \qquad \tau \in [t,\, t+T_p] \tag{5.28b}$$

$$\bar{\mathbf{s}}(t+T_p) \;\in\; \Omega_\varsigma \tag{5.28c}$$

where $\mathbf{Q} \succ 0$, $\mathbf{R} \succ 0$ and $T_p = T_c$ to avoid move blocking. The terminal penalty matrix $\mathbf{P} \succ 0$ is calculated as outlined above (see §5.4.2). As the framework considers a shifted process model, $\mathbf{s}$ can be used instead of $\mathbf{e}$ in J. Notice also, that because a spherical camera model is used, visibility constraints do not need to be explicitly considered in the control design. Importantly, this allows a simpler derivation of the quasi-infinite horizon visual predictive control approach.

5.4.5 Circular Motion

Given the expectation that the avoidance decision regarding reference azimuth angle forces circular motion, the QIH-VPC control structure should be considered for this special case of spiral motion (see §3.3.1). For circular motion, the platform remains equidistant from the object which means that the range parameter is fixed, and minimal model-mismatch results. As such, using simple circular motion equations, $\mathbf{u}_3^* = [0\ \omega^*]$ where v_x^* defines the nominal forward velocity such that $v_x^* = \omega^* r_2^*$ and $r_2 = r^* \cos(\sigma^* - \pi/2)$ defines the radius of the circle on the lateral xy-plane. By definition, the reference colatitude cannot be placed at the polar caps of the spherical surface, so $\sigma^* \notin \{0, \pi\}$. Therefore, the required conditions to formulate the QIH-NMPC structure are satisfied (see §E.4).

Under the above considerations, an analysis is conducted with respect to reference image feature selection, robustness (noise and parameter uncertainty) and control effort. Additionally, a performance comparison to the classical visual predictive approach is provided. Simulations were conducted using the parameters defined in table E.4 and constraints defined in table E.6 unless stated otherwise.

Reference Image Features

The terminal region behaviour with variable reference colatitude angle σ^* is first compared. Assuming circular motion, the reference azimuth angle is such that $\gamma^* = \pi/2$ and the reference range value is fixed such that $r^* = 2$. The results are depicted in Fig 5.2, where results for $\sigma^* < \pi/2$ are omitted due to the observed symmetry about $\sigma^* = \pi/2$.

Fig 5.2(a) shows the terminal region centred about $\sigma^* = 5\pi/6$. The black ellipse bounds the terminal region according to the control constraints. This essentially corresponds to solving only problem A in algorithm 4. The smaller area enclosed by the red ellipse shows the actual terminal region, considering the system non-linearity at the reference colatitude angle. This essentially corresponds to solving problem B in algorithm 4. The red ellipse is also included in Fig.5.2(b) for comparative purposes.

Fig 5.2(b) shows the terminal regions centred about $\sigma^* \in \{\pi/2, 2\pi/3, 5\pi/6\}$ depicted in red. For $\pi/2 \leq \sigma^* < 130\pi/180$, the terminal region is dictated by the control constraints. If we increase the control constraint domain for a given σ^*, a larger terminal region will result. In the polar regions however, the terminal region is no longer defined primarily due to control constraints. Because of the increased nonlinearity of the system, a number of iterations are required to find the terminal region for which the fictitious linear controller guarantees convergence.

Fig 5.2(c) further depicts the terminal region variation with σ^* using the terminal region constant ς, which can be considered a proxy for the terminal region area. The terminal region constant increases until $\sigma^* \approx 128\pi/180$ due to the increased domain for the locally linear controller, before rapidly decreasing as we approach the polar regions. This is because the terminal region provides a quantitatively representation of the degree of system nonlinearity for a given reference. The result is important as it suggests selecting σ^* near $125\pi/180$ may ensure a larger domain of attraction and thus larger feasible set for $\mathbf{s}_0$. As verification, a set of simulations for $\sigma^* = 125\pi/180$ and varied $\mathbf{s}_0$ is shown in Fig. 5.2(d). In each simulation the initial range is fixed, and the resulting camera trajectory rapidly converges to the circular path along a conical spiral trajectory. As evidence, the associated camera trajectory, track error and control inputs are shown in Fig 5.3. In the context of collision avoidance, the smooth control inputs force the aircraft to safely maneuver below and around the object.

Given the above analysis, the terminal region could then be used to select suitable reference image feature values for spiral tracking, before implementation. Consider an application such as collision avoidance. Using the rules of the air, and with respect to the avoidance strategy presented earlier (see §3.3.1), it is only required that $\sigma^* \neq \pi/2$. Selecting $\sigma^* = 125\pi/180$ could ensure a larger set of feasible initial states and therefore provide better assurance that the reference features, and hence avoidance, will be achieved. The terminal region can then guide the selection of reference image features based on expected control performance.

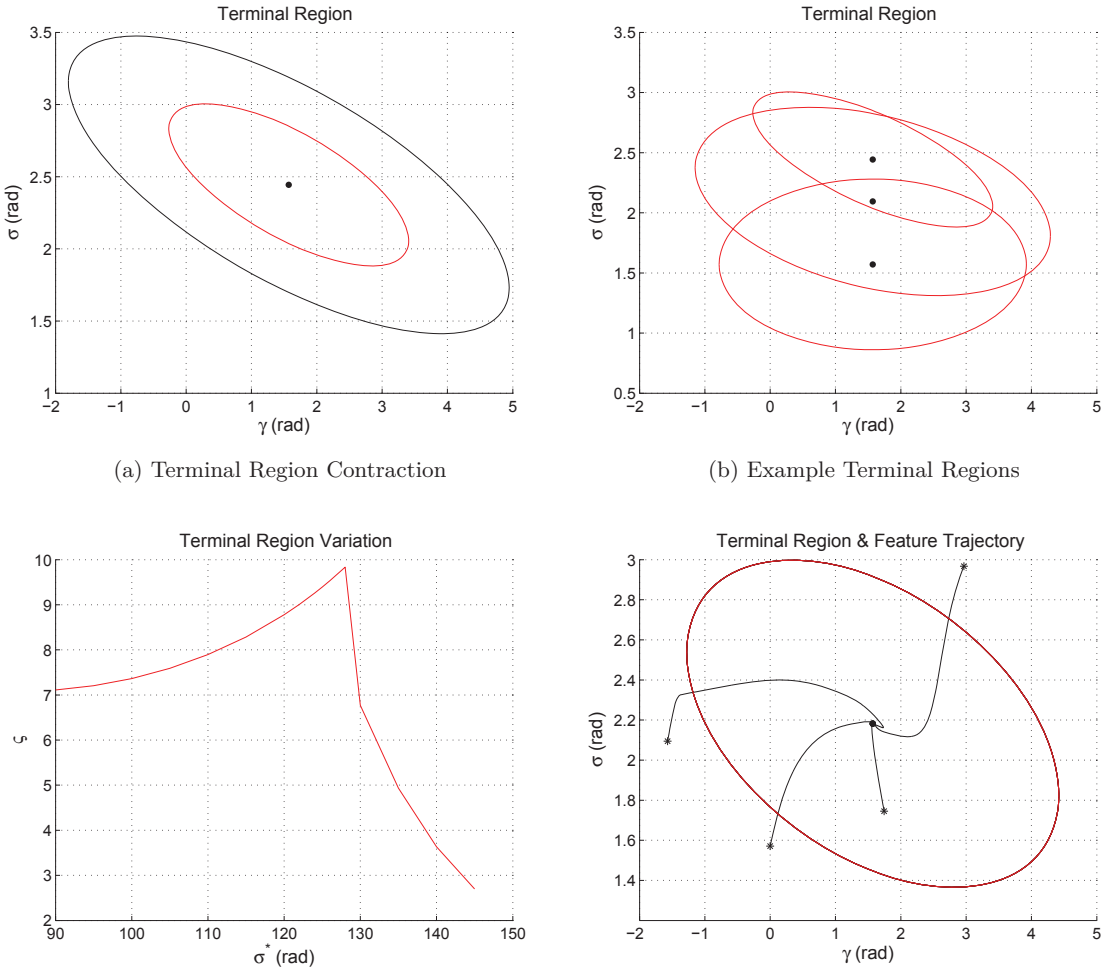

(a) Terminal Region Contraction

(b) Example Terminal Regions

(c) Terminal Region Constant

(d) Example Feature Trajectories

Figure 5.2: Terminal region variation (–) for variable σ^* (•) (a) Terminal region contraction for $\sigma^* = 140\pi/180$ (b) Terminal region differences for $\sigma^* \in \{\pi/2, 2\pi/3, 5\pi/6\}$ (c) Terminal region constant ς for $\sigma^* \in (\pi/2, 5\pi/6)$ (d) Image feature trajectories for $\sigma^* = 125\pi/180$, $r_0 = 2$ and varied $\mathbf{s}_0$.

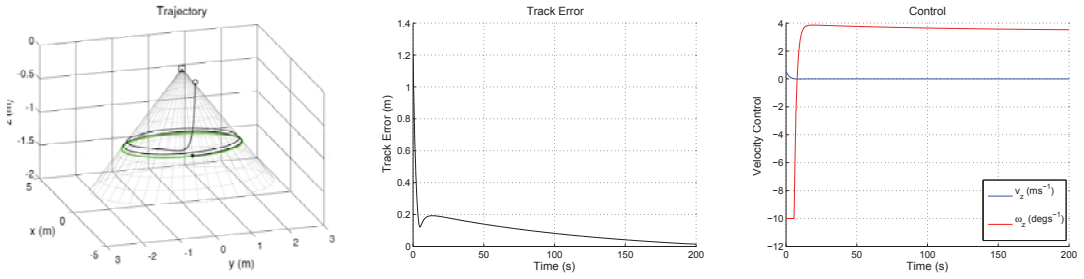

(a) Trajectory

(b) Track Error

(c) Control

Figure 5.3: Example QIH-VPC simulation for $\sigma^* = 125\pi/180$ including platform trajectory, track error and control. The track error is defined as the minimum 2-norm of the difference between any point on the reference circle (–) and the current position.

Remark The exact camera and image feature trajectory will of course depend on the choice of state and control weighting matrices, but it can be safely assumed that these are constant predefined matrices chosen based on some desired vehicle performance.

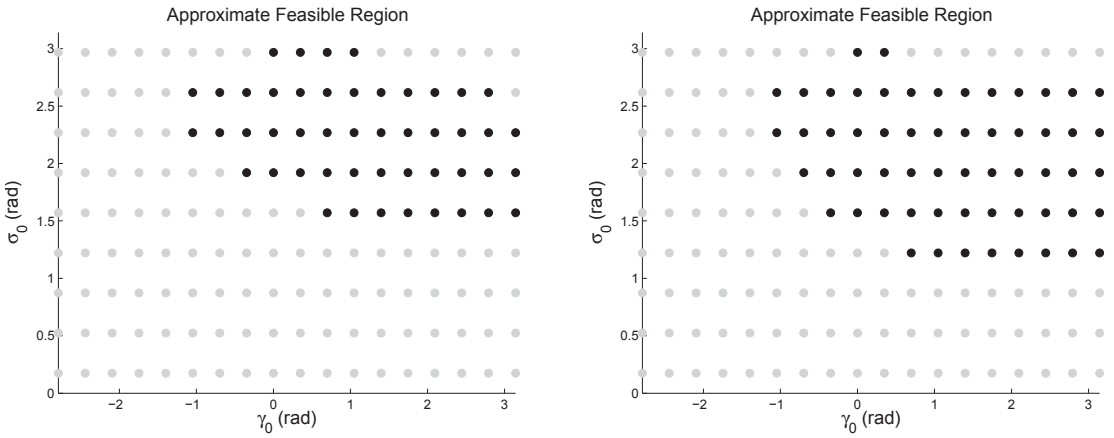

(a) $T_p = 2$ (b) $T_p = 15$

Figure 5.4: Feasibility analysis for $\sigma^* = 125\pi/180$ and simulation time of 30s. Feasible (•) and infeasible (•) $\mathbf{s}_0$ are shown for variable prediction horizons T_p.

Although analytically determining the feasible region may be possible in the nominal case, the difference between the actual system and model suggest it would be best to sample $\mathbf{s}_0$ over $\mathcal{S}^2$ and simulate. The resulting state at the end of the first iteration, and whether it lies in the terminal region, can be used to determine feasibility. Such results are depicted in Fig 5.4 for short and long horizons by sampling $\mathbf{s}_0$ at $20\pi/180$ intervals over $\mathcal{S}^2$ and simulating. The sampled set is defined as $\hat{\mathcal{S}}^2$. Note the slightly larger feasible initial state space, and thus stable region of convergence, for the longer horizon. It was noted in [293] that it is difficult to determine the region of convergence for classical image-based visual servoing. These results provide some valuable insight into this important issue.

Robustness

The quantitative robustness properties in the context of unknown time varying range parameter and bounded uncertainty on image feature measurements and actuators are investigated. This is important to consider since the approach is designed for the ideal case. Even for circular motion, $r(t) \neq r^*$ unless already established on the circular path. Additionally, sensor uncertainty and imperfect actuation will be present in a real system. Based on previous results, the reference colatitude angle is selected such that $\sigma^* = 125\pi/180$. The results are depicted in Fig 5.5 when r^* is varied, and in Fig 5.6 when r_0 is varied and additive noise is included. Results for $\sigma^* \neq 125\pi/180$ are omitted as similar terminal region behaviour and control performance were observed.

Fig 5.5(a) depicts the terminal region variation for $r^* \in \{2, 4, 8, 16\}$. The variation in the terminal region constant, and thus terminal region area, is shown in Fig 5.5(b) for $r^* = (0.5, 125)$. The existence and shape of the terminal regions are important, suggesting a controller designed with a smaller r^* may be suitable when $r(t) > r^*$. This is shown by the successively smaller terminal regions as r^* decreases. However, the elements of the terminal penalty matrix increase as the reference range parameter decreases. This may

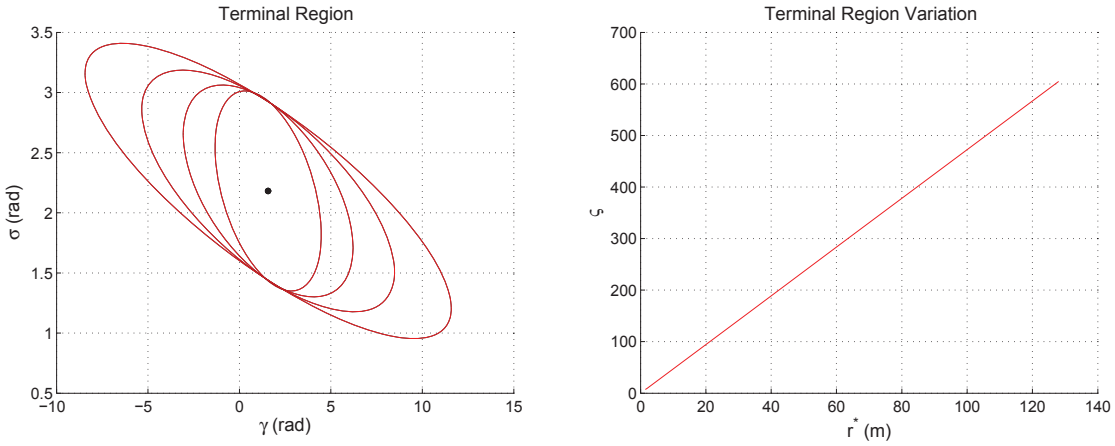

(a) Terminal Region Variation (b) Terminal Region Constant

Figure 5.5: Terminal region variation (−) with parameter uncertainty for $\sigma^* = 125\pi/180$ (•) (a) Terminal region variation with r^*. The terminal regions from the smallest to the largest correspond to $r^* \in \{2, 4, 8, 16\}$ respectively (b) Terminal region constant variation for $r^* \in [0.5,\ 125]$

reduce performance through aggressive control action, but is not a major issue given the control limits will still be maintained.

Fig 5.6(a) depicts some example image feature trajectories for variable r_0, using a fixed reference range value $r^* = 2$. In some cases, the initial object range is up to twice that used in the process model. The QIH-VPC controller is shown to handle this large parameter uncertainty, and converge to the circular trajectory. The result suggests that the control structure can alleviate the explicit dependence on over (or under) estimating the reference range parameters in the controller, as in the classical control case.

Fig 5.6(b) depicts some example image feature trajectories when the initial range is varied, and uncertainty on image feature measurements and imperfect actuation is included. Uncertainty in the form of white noise $q(t) \sim \mathcal{N}(0, 0.02^2)$ is added to the image feature measurements and actuator commands. The noise model has the same characteristics as that used for the analysis of the classical image-based controllers derived previously (see §4). In all cases, the image features converge to the desired location, and remain within a small neighbourhood of the corresponding reference value. The results suggest that the desired spiral can be tracked reasonably well. Of note, initial results suggest that added noise can be better managed than in the classical image-based approach, similar to general visual predictive control [337].

Fig 5.7 extends the above results regarding variable initial range to include a feasibility analysis. Some example feasibility statistics for $r_0 = [1.5,\ 4]$ sampled at 0.5 intervals for all $\mathbf{s}_0 \in \hat{\mathcal{S}}^2$ are shown. The approximate percentage of feasible initial features for each initial range is shown. Interestingly, the results are comparable for each initial range, suggesting that the initial range has minimal impact on the feasible, and thus stable, initial image feature set. For spiral motion, this suggests similar tracking performance can be obtained under vastly different initial states.

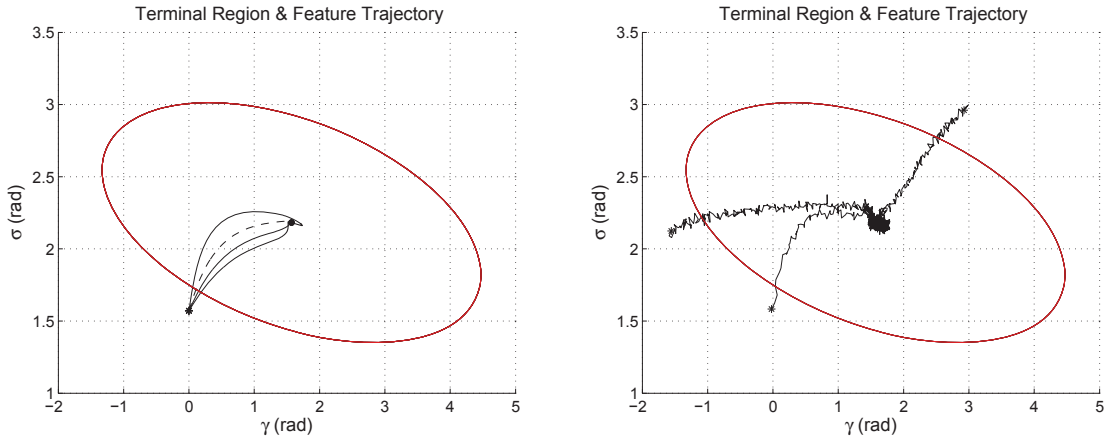

(a) Variable Initial Range

(b) Additive Noise

Figure 5.6: Robustness characteristics for $\sigma^* = 125\pi/180$ (•) (a) Example image feature trajectories for $r^* = 2$ and $r_0 \in \{1, 2, 3, 4\}$. The dashed line shows $r_0 = 2$ (b) Example image feature trajectories for variable $\mathbf{s}_0$ and additive measurement and process noise $q(t)$

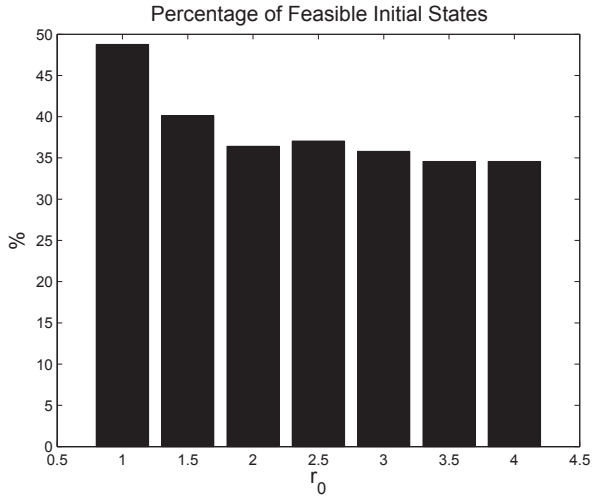

Figure 5.7: Example feasible region statistics for $\sigma^* = 125\pi/180$ and simulation time 30s

Collectively, the analysis suggests that although the approach is not designed for robustness, uncertainty and model mismatch are managed well, and performance is not significantly degraded. The control structure demonstrates similar desirable attributes to that of the general model predictive control structure. Specifically, due to the control structure optimising over a finite time, model-mismatch and additive noise (including feature measurements and actuator output) are not propagated through the process model for all time. The image feature trajectory does not diverge significantly, as the controller has an opportunity to re-optimise and better adjust the control output at each iteration. Although, similar robustness characteristics are exhibited by classical image-based visual servoing, the reference range and gain term must be tuned to ensure both stable and desirable spiral tracking behaviour (see §4). The predictive control structure thus provides some significant benefits over the classical approach to ensure adequate spiral tracking, and thus collision avoidance, using a real system.

Remark One way to guarantee robust stability, without explicitly considering uncertainty in the control design, is to shrink the terminal region along with an appropriately calculated prediction horizon [344]. Other control design such as H_∞ control and Linear Matrix Inequalities (LMI) explicitly consider a bounded uncertainty on the unknown model parameters but require the solution of a computationally expensive minimax problem online. This makes such approaches difficult to implement in practice. In visual servoing, only recently has this been considered [324], still with unsatisfactory computation expense for implementation.

Performance Comparison

A quantitative performance comparison is made between the QIH-VPC and the C-VPC control structure for tracking circular motion. Importantly, the classical VPC scheme must now include a control penalty term $\bar{\mathbf{u}}(\tau)^2_R$ in the corresponding objective function for a fair comparison. As such, the objective function defined in (5.16) is used, and the remaining elements of the control structure are retained. Based on previous results, the reference colatitude angle and range are selected such that $\sigma^* = 125\pi/180$ and $r^* = 2$ respectively. Comparative results regarding the tracking behaviour and control performance and depicted in Fig 5.8 and Fig 5.9 respectively.

Fig 5.8 depicts an example set of results for an initial range $r_0 = 2$ and image feature location $\mathbf{s}_0(\pi/2, 0)$. A realistic case is assumed such that additive noise $q(t)$ on image feature measurements and control output is also included. First, the QIH-VPC provides comparable performance to the C-VPC scheme. Similar image feature trajectories are observed with a marginal reduction in track error for the QIH-VPC scheme. The track error is defined as the minimum 2-norm of the difference between any point on the reference circle and the current position. Second, similar robustness qualities are displayed for both predictive schemes, managing model mismatch and added uncertainty. Importantly, similar results were obtained for different initial conditions and comparable noise attributes.

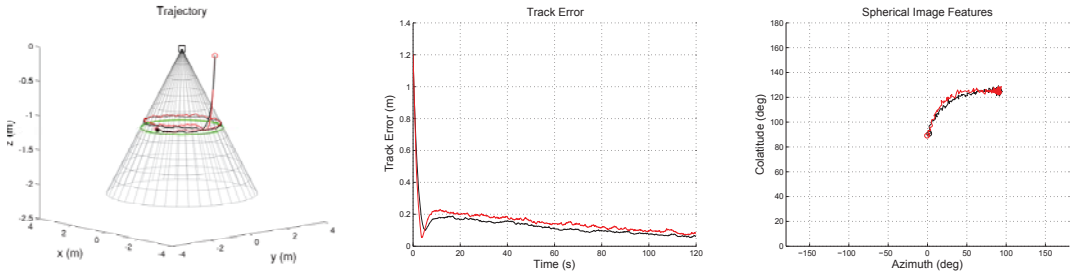

(a) Trajectory (b) Track Error (c) Feature Trajectory

Figure 5.8: Comparison of circular motion tracking using C-VPC (−) and QIH-VPC (−). Platform trajectory, track error and image feature trajectory with $\mathbf{s}^*(125\pi/180, \pi/2)$, $r^* = 2$ and $r_0 = 2$ are shown.

Fig 5.9 depicts the relative total required control effort for both schemes, analysed with respect to variable initial state. A long prediction horizon $T_p = 15$ (samples) is used for a total simulation time of 30s. The magnitude of the total control effort for each $\mathbf{s}_0 \in \hat{\mathcal{S}}^2$ for the QIH-VPC and C-VPC schemes is first calculated, then the ratio of the two is depicted using a two dimensional (2D) surface plot. A reduction in control effort for the QIH-VPC scheme results in a ratio less than unity, shown using the blue end of the color spectrum. Interestingly, the average control effort ratio is 0.9716 ($\approx 3\%$ reduction), showing a general reduction in control effort for the QIH-VPC scheme. This translates into the ability to efficiently navigate through all regions of the state space, including particularly non-linear spherical sections such as the polar regions.

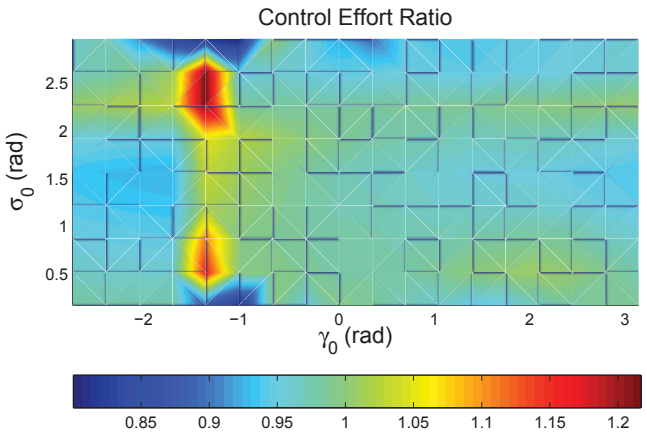

Figure 5.9: Comparison of total control effort for $\mathbf{s}^*(125\pi/180, \pi/2)$, $T_p = 15$ and simulation time 30s. The ratio of control effort for QIH-VPC over C-VPC is shown. A value less than unity depicts reduced control effort using QIH-VPC at the corresponding $\mathbf{s}_0$. The average ratio is 0.9716.

Of note, the computational effort is of less impact, as it remains small and suitable for near real-time implementation for both schemes. Using MATLAB 2011b and the ACADO Toolkit running on an Intel Core 2 Duo CPU T8100 at 2.10GHz with 2GBytes RAM, the average computational time for the online optimisation (5.26) is under 42ms. Using MATLAB's Active-Set constrained optimisation solver, the computation time is significantly greater, in some cases up to 0.5s.

5.4.6 Spiral Motion

Given that the avoidance decision regarding reference azimuth angle may not be fixed, the QIH-VPC control structure should be considered for non-circular spiral motion for completeness (see §3.3.1). For spiral motion, the deviation from a circular path results in increasing model-mismatch due to the highly variable range parameter. Thus for divergent or convergent spiral motion, the reference control is no longer constant as $\omega_z^*(t) = f(r(t))$ such that $\mathbf{u}^* = f(t)$, which violates the conditions required to formulate the QIH-NMPC structure (see §E.4). One way to manage this is of course to include range as a state, but its derivative (and thus variation) cannot be predicted until the spiral is established. As such, a stable locally linear feedback controller cannot be found for non-circular motion. This results in the terminal region approximating a point; a terminal equality constraint. This is known to lead to stability, but is restrictive in the sense that it forces $\mathbf{s}(T_p) = \mathbf{s}^*$. In this case, a terminal penalty term does not make sense and feasibility issues can arise. For this control approach, modifications to the control structure include

$$J(\mathbf{s}(t), \bar{\mathbf{u}}(\cdot)) \quad = \quad \int_{t}^{t+T_p} \bar{\mathbf{s}}(\tau)_Q^2 + \bar{\mathbf{u}}(\tau)_R^2 \, d\tau, \qquad\qquad \bar{\mathbf{s}}(t+T_p) = 0 \tag{5.29}$$

Alternatively, a non-zero terminal penalty term could be used without a terminal region, provided an appropriate prediction horizon is selected to aid stability [325]. This is similar to the classical VPC approach in the stability sense, however the end states are not progressively weighted and move blocking is not employed. For this control approach, modifications to the control structure include

$$J(\mathbf{s}(t), \bar{\mathbf{u}}(\cdot)) \quad = \quad \bar{\mathbf{s}}(t+T_p)_P^2 \; + \; \int_{t}^{t+T_p} \bar{\mathbf{s}}(\tau)_Q^2 + \bar{\mathbf{u}}(\tau)_R^2 \, d\tau, \qquad \bar{\mathbf{s}}(t+T_p) \in \mathcal{S}^2 \tag{5.30}$$

The terminal weighting matrix used in (5.30) can be calculated according to Section 5.4.2, but this assumes the reference control is constant for all time. One option is to then use the reference control input at some nominal reference range value. The result is a terminal weighting matrix derived from a linear controller with a very restrictive domain. This can lead to significantly large terminal weighting matrix element values, which can result in overly aggressive control. To avoid such control issues, the terminal weighting matrix can be selected to be twice that of the state weighting matrix. Using this approach, a performance comparison of the two alternate visual predictive control structures using (5.29) and (5.30) and the classical visual predictive approach (5.6)-(5.8) is provided. Simulations were conducted using the parameters defined in table E.4 and constraints defined in table E.6 unless stated otherwise.

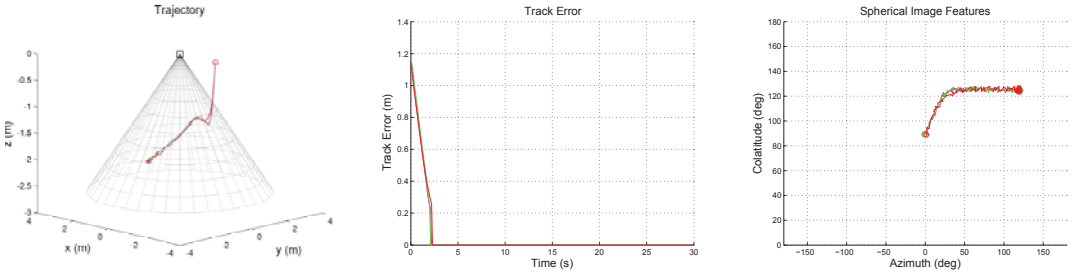

(a) Trajectory (b) Track Error (c) Feature Trajectory

Figure 5.10: Comparison of spiral motion tracking using C-VPC (–), terminal penalty term VPC (–) and terminal equality constraint VPC (–). Platform trajectory, track error and image feature trajectory with $\mathbf{s}^*(125\pi/180, 120\pi/180)$, $r^* = 2$ and $r_0 = 2$ are shown.

Performance Comparison

Similar to the preceding section, a quantitative performance comparison is made between predictive controllers using a terminal equality constraint, terminal penalty term without terminal constraint and the classical VPC scheme. To remain consistent with the preceding analysis, the reference image features and nominal reference range value are selected such that $\mathbf{s}^*(125\pi/180, 120\pi/180)$ and $r^* = 2$ respectively.

Fig 5.10 depicts an example set of results for an initial range $r_0 = 2$ and image feature location $\mathbf{s}_0(\pi/2, 0)$. A realistic case is assumed such that additive noise $q(t)$ on image feature measurements and control output is also included. The difference between the VPC controller using a terminal equality constraint and terminal penalty term without constraint are minimal, with the later offering potentially more flexibility in the feasible set for $\mathbf{s}_0$. Additionally, the performance of both VPC controllers is comparable to that of the C-VPC scheme. To this end, the VPC scheme with terminal penalty term shows slightly faster convergence, requiring less time to reach the surface of the reference cone and reduced track error. The track error in this case is now defined as the minimum 2-norm of the difference between any point on the reference cone and the current position.

Of note, it is reasonable to suggest that the C-VPC scheme could be tuned to behave the same as the predictive control scheme with penalty term, by altering the time varying state weighting matrix. However, using the predictive control scheme with penalty term avoids move blocking to ensure the full richness of controls can be adopted. Additionally, the predictive controller can of course manage circular motion presented in the previous section. It might be useful to adopt a predictive control scheme with a penalty term (and no terminal constraints) for spiral tracking and collision avoidance, given the requirement to manage unknown dynamic objects and potentially variable reference azimuth angles.

Remark Divergent spirals may be followed for all time provided the object remains visible. For convergent spirals, a limit cycle will be reached depending on the control constraints. This means the object will be continually circled at a fixed radius, albeit small if the control constraints are liberal.

5.5 Collision Avoidance

The spherical visual predictive control laws derived in the preceding sections represent an improved set of controllers for tracking a reference spiral. They can be used directly in a closed-loop collision avoidance system when coupled with an appropriate avoidance and resolution decision strategy. This section details the implementation of such a collision avoidance system on a small aerial platform in a proof of concept approach.

First, a practical and flexible avoidance controller including the aerial platform dynamics is presented. Second, improved avoidance and resolution decision strategies based on previously established concepts are presented (see §3). This is required to explicitly couple the control and resolution strategies, and ensure the collision avoidance system is amenable to practical platform constraints and limitations. Third, results from real flight tests in a scaled collision avoidance environment are presented and analysed. They are the first recorded flight tests using spherical visual predictive control of any form.

5.5.1 Avoidance Control

System Dynamics

To experimentally apply the visual predictive control framework, the aircraft dynamics must be augmented with the image dynamics to derive the process model. This ensures an accurate model of the integrated visual system can be used for prediction in the control strategy. Using the point mass model and black box identification for a small *AscTec Hummingbird* quadrotor presented earlier (see §2.2), a simplified set of decoupled linear equations were found to represent the quadrotor dynamics. A subset of the complete dynamic equations relevant to the visual control, and defined in the body frame, are given by

$$\dot{v}_x = -F_T\,\theta^*/m \tag{5.31}$$

$$\dot{v}_y = -F_T\,\phi^*/m \tag{5.32}$$

$$\dot{v}_z = -g + F_T^*/m \tag{5.33}$$

$$\dot{\psi} = \omega_z \tag{5.34}$$

$$\dot{\omega}_z = -a_{\dot{\psi}}\omega_z + a_{\dot{\psi}}\omega_z^* \tag{5.35}$$

where the controls F_T^*, ω_z^*, θ^* and ϕ^* define the reference thrust, yaw rate, pitch and roll commands for the low-level attitude controllers. The empirically derived yaw damping constant $a_{\dot{\psi}}$, mass m and acceleration due to gravity g are given in Appendix E. Now consider the simplified optic flow equations for colatitude and azimuth angles given by

$$\dot{\sigma} = \frac{-\cos\sigma\cos\gamma}{r^*}v_x + \frac{\sin\sigma}{r^*}v_z \tag{5.36}$$

$$\dot{\gamma} = \frac{\sin\gamma}{r^*\sin\sigma}v_x - \omega_z \tag{5.37}$$

The image kinematics can then be combined with the partitioned quadrotor dynamics. Augmenting (5.33) - (5.35) with (5.36 - 5.37), the resulting nonlinear process model can then be defined as

$$\dot{\mathbf{z}} = \mathbf{f}(\mathbf{z}, \mathbf{u}_3) \tag{5.38}$$

$$\dot{\mathbf{z}}(t) = \mathbf{A}_{\mathrm{P}}(t) + \mathbf{B}_{\mathrm{P}}(t)[\mathbf{u}_3(t)\,\mathbf{v}]^T \tag{5.39}$$

$$\begin{pmatrix} \dot{v}_z \\ \dot{\psi} \\ \dot{\omega}_z \\ \dot{\sigma} \\ \dot{\gamma} \end{pmatrix} = \begin{pmatrix} -g \\ \omega_z \\ -a_{\dot{\psi}}\omega_z \\ v_z \sin\sigma / r^* \\ -\omega_z \end{pmatrix} + \begin{pmatrix} 0 & 1/m & 0 \\ 0 & 0 & 0 \\ 0 & a_{\dot{\psi}} & 0 \\ 0 & 0 & -\cos\sigma\cos\gamma/r^* \\ 0 & 0 & \sin\gamma/r^*\sin\sigma \end{pmatrix} \begin{pmatrix} F_T^* \\ \omega_z^* \\ v_x \end{pmatrix} \tag{5.40}$$

where $\mathbf{z}$ defines a combined state vector of image features and vehicle states, and the subscript P is used to avoid confusion from the general model of the vertical quadrotor dynamics derived earlier (see §2.2.3). The control vector $\mathbf{u}_3$ now includes the reference thrust command F_T^* instead of the reference vertical velocity v_z^*, as in the classical image based control approach (see §4). The model could of course be adapted to use a reference vertical velocity v_z^* instead, if required by the low-level controllers.

There are three important observations regarding the process model. First, the camera and body velocities are considered to be equal, such that the origin of both frames coincide and the x axis are aligned. This is not the case for the real system, and the image features will inherit a small offset not accounted for in the process model. However, the approximation can be considered suitable assuming a small fixed camera displacement (see §3.2.2 and §B), and given that the VPC structure can manage small model-mismatch well. Second, as with the classical image-based control approach, v_x can be passed as a parameter and controlled independently of the visual controller. Third, it would appear a redundant state ψ, representing the platform heading, is included in the model. However, the state is retained in order to derive a novel resolution decision strategy presented in later sections.

Control Structure & Solution

The specific control structure used for the practical system must be aligned to the practical constraints, not only with respect to the platform state, control and camera visibility constraints, but the available software tools. The control structure should also be flexible enough to allow the inclusion of additional collision avoidance functionality, such as resolution decision strategies.

Consider the visual predictive controllers derived in the preceding section, including QIH-VPC and VPC strategies using only a terminal penalty term or a zero terminal constraint. It was shown that the QIH-VPC approach can be used to ensure stability and select better reference image features for circular motion. However, the approach is not directly extensible to general spiral (non-circular) motion without further modifications (see §5.6). For general spiral motion, there does not exist, even for a shifted system, a zero control

that ensures $\mathbf{f}(\mathbf{z}, \mathbf{u}) = 0$. To further complicate the matter, the yaw angle cannot be fixed for all time when tracking a spiral, so a monotonically decreasing objective function cannot be assured. Additionally, the ACADO Toolkit version[3] used to solve the minimisation problem does not currently support nonlinear constraints for embedded applications [356]. The VPC strategy with terminal penalty term showed slightly improved performance over C-VPC, and allows greater flexibility in the choice of reference image features over the QIH-VPC approach. Importantly, it also allows the inclusion of additional states (such as yaw ψ) and associated constraints, which will become important to derive a novel resolution decision strategy (see §5.5.2). The approach can also be implemented using the ACADO Toolkit for real applications. Of note, none of the aforementioned strategies have been successfully implemented on an aerial vehicle of any type.

Given the above considerations, the VPC strategy with terminal penalty term is chosen for implementation. Having defined the process model (5.40), the control structure including the objective function J_s, state constraint domain $\mathbb{Z}$ and control constraint domain $\mathbb{U}$ can now be defined. The control problem to be solved at each sampling time is given by

$$\mathbf{u}_3^\star(\cdot) = \underset{\mathbb{U}}{\operatorname{argmin}}\ J_s(\mathbf{z}(t), \bar{\mathbf{u}}_3(\cdot)) \tag{5.41}$$

where

$$J_s(\mathbf{z}(t), \bar{\mathbf{u}}_3(\cdot)) \quad = \quad \bar{\mathbf{z}}(t+T_p)_P^2 \ + \int_t^{t+T_p} \bar{\mathbf{z}}(\tau)_Q^2 + \bar{\mathbf{u}}_3(\tau)_R^2 \, d\tau \tag{5.42}$$

s.t

$$\dot{\bar{\mathbf{z}}} \ = \ \mathbf{f}(\bar{\mathbf{z}}, \bar{\mathbf{u}}_3), \qquad \bar{\mathbf{z}}(t) = \mathbf{z}_0 \tag{5.43a}$$

$$\bar{\mathbf{u}}_3(\tau) \ \in \ \mathbb{U}, \qquad \tau \in [t,\, t+T_p] \tag{5.43b}$$

$$\bar{\mathbf{z}}(\tau) \ \in \ \mathbb{Z}, \qquad \tau \in [t,\, t+T_p] \tag{5.43c}$$

The associated constraint domains are defined by simple box constraints such that

$$\mathbb{U} \in \mathbb{R}^2 \,|\quad \mathbf{u}_{\min} \leq \mathbf{u}(t) \leq \mathbf{u}_{\max} \tag{5.44}$$

$$\mathbb{Z} \in \mathbb{R}^5 \,|\quad \mathbf{z}_{\min} \leq \mathbf{z}(t) \leq \mathbf{z}_{\max} \tag{5.45}$$

The visibility constraints defined within $\mathbb{Z}$ are somewhat handled with the application of a spherical camera, so could be used to avoid the polar caps and ensure the spherical image Jacobian is always well defined. As such, the domain limits of the spherical image features can be used as the upper $\mathbf{s}_{\max}$ and lower $\mathbf{s}_{\min}$ constraints. Combining these limits with maximum and minimum yaw and vertical rates for the platform, the remaining components of $\mathbf{z}_{\max}$ and $\mathbf{z}_{\min}$ can be obtained. The state constraints then ensure controls are issued such that the quadrotor state lies inside a desirable region. The

[3] An updated ACADO Toolkit version is now available with increased functionality (Version 1.2.1beta, January 17, 2014).

control constraints bound the quadrotor thrust and yaw rate commands to an admissible (or desirable) region based on the control authority of the platform.

The weighting matrices are defined such that $\mathbf{P} \succ 0$, $\mathbf{Q} \succ 0$ and $\mathbf{R} \succ 0$. In some case the matrices can be optimally tuned online with external software such as that presented in [357], but the approach has not been applied to visual predictive control for systems with fast dynamics. Instead, the parameters are chosen with consideration to the intended application. For collision avoidance, and given the reactive nature of the See and Avoid task, the primary concern is establishing and subsequently following the reference spiral as soon as possible. The control weighting matrix $\mathbf{R}$ can be selected to be relatively small, having included corresponding control constraints in the optimisation. Specifically, the elements of the control weighting matrix $\mathbf{R}$ can be selected to be at least an order of magnitude smaller that those within the state weighting matrices $\mathbf{Q}$ and $\mathbf{P}$. The terminal penalty matrix $\mathbf{P}$ can be selected such that $\mathbf{P} \succeq \mathbf{Q}$, in an attempt to align the cost function to that expected using the QIH-VPC design. The choice makes sense, as placing an equal or larger penalty on the state at the end of the prediction horizon will help to force rapid convergence to the reference spiral.

Similar to the C-VPC strategy, local asymptomatic stability can be achieved through appropriate selection of the prediction horizon [325, 339, 341]. The prediction horizon T_p needs to be large enough to avoid potential stability issues, and small enough to help reduce computational expense.

Remark As the state $\mathbf{z}$ includes the image feature vector $\mathbf{s}$, care has to be taken to appropriately bound the image feature error when calculating the objective function. When using the ACADO Toolkit, it is easier to use the bounded initial feature error $\mathbf{e}_0 = \mathbf{s}_0 \ominus \mathbf{s}^*$.

5.5.2 Avoidance & Resolution Decision

Consider again that a perspective camera is used to approximate a spherical section. The camera is rigidly attached to the platform with its optical axis offset from the platform x axis by a fixed angle $\psi_c = -135\pi/180$. Despite the restricted field of view, the ideal reference image features $\mathbf{s}^*(125\pi/180, \pm\pi/2)$ and those corresponding to a subset of divergent spirals can be used, as they reside within the visible region. The reference image features can then be selected according to the avoidance decision strategy based on aviation rules of the air (see §3.3.1). Therefore, provided the right spiral direction is chosen, the aircraft will initially move in favour of collision avoidance.

Unlike the classical image-based control approach, if the reference image features are chosen in this way, a resolution decision is not necessarily mandatory to ensure safe avoidance. The controller will force the aircraft around or away from the object in both vertical and lateral planes (see §3). However, in order to design a useful automated system, a resolution decision will be needed to stop the spiral behaviour. The same strategy as used for the image-based controller could be used, but would obviously inherit the same problems. In particular, the visual control is not coupled to the resolution decision, so

there is the possibility that the avoidance behaviour ceases prior to establishing a safe spiral, close to the initial detection instance (see §4.5). If the same resolution concept using the aircraft heading is to be used however, a modified approach is required to address this issue.

Interestingly, the visual predictive control structure provides a natural framework in which to provide the required coupling between the visual control and resolution decision. By including the yaw angle in the process model, the image features and heading can be indirectly coupled such that a trade-off exists in the objective function. Depending on the relative magnitude of the elements of the state weighting matrices, the controller will implicitly prefer a spiral trajectory (defined by the image features), or the initial aircraft heading upon detection. The objective function value could then be considered as a proxy to indicate how well both collision avoidance and resolution criteria are satisfied. As such, the minimum objective function value can then be used directly as a basis for a new resolution decision strategy.

To explain, consider first a visual predictive controller using the objective function J_s defined previously, in which only the elements of the state weighting matrices that correspond to the image features (colatitude and azimuth) are non-zero. The elements are also equally weighted and denoted by Q_s. The control and terminal weighting matrices can then be scaled with respect to Q_s using the constant scaling factors $\lambda_R \geq 0$ and $\lambda_P \geq 0$ respectively such that

$$\mathbf{Q} = \begin{pmatrix} 0 & \cdots & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \cdots & Q_s & 0 \\ 0 & \cdots & 0 & Q_s \end{pmatrix}, \; \mathbf{P} = \begin{pmatrix} 0 & \cdots & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \cdots & \lambda_P Q_s & 0 \\ 0 & \cdots & 0 & \lambda_P Q_s \end{pmatrix}, \; \mathbf{R} = \begin{pmatrix} \lambda_R Q_s & 0 \\ 0 & \lambda_R Q_s \end{pmatrix} \tag{5.46}$$

As expected, the reference spiral is tracked in all cases as the control structure closely resembles that used when only the image kinematics were considered in the process model. The minimum of the objective function occurs when the spiral is established, then remains at the minimum if the spiral is maintained. Of course, the objective function value cannot provide any information regarding the relative location of the platform on the spiral, having not induced the yaw state. In this case, it would then be difficult to use the objective function to indicate an appropriate time to stop the spiral.

Now consider including a non-zero weighting on the yaw state such that any deviation from the yaw angle at the initial time of object detection incurs a penalty. This can be expressed in the state weighting matrices directly by including the relevant terms, or using an augmented objective function J_ψ such that

$$J_\psi = J_s + \lambda_\psi \int_t^{t+T_p} (\psi(\tau) \ominus \psi^*)_{Q_s} \, d\tau \tag{5.47}$$

where $\lambda_\psi \geq 0$ is a positive yaw scaling factor and $\psi^* = \psi(t_d)$ is the platform heading upon

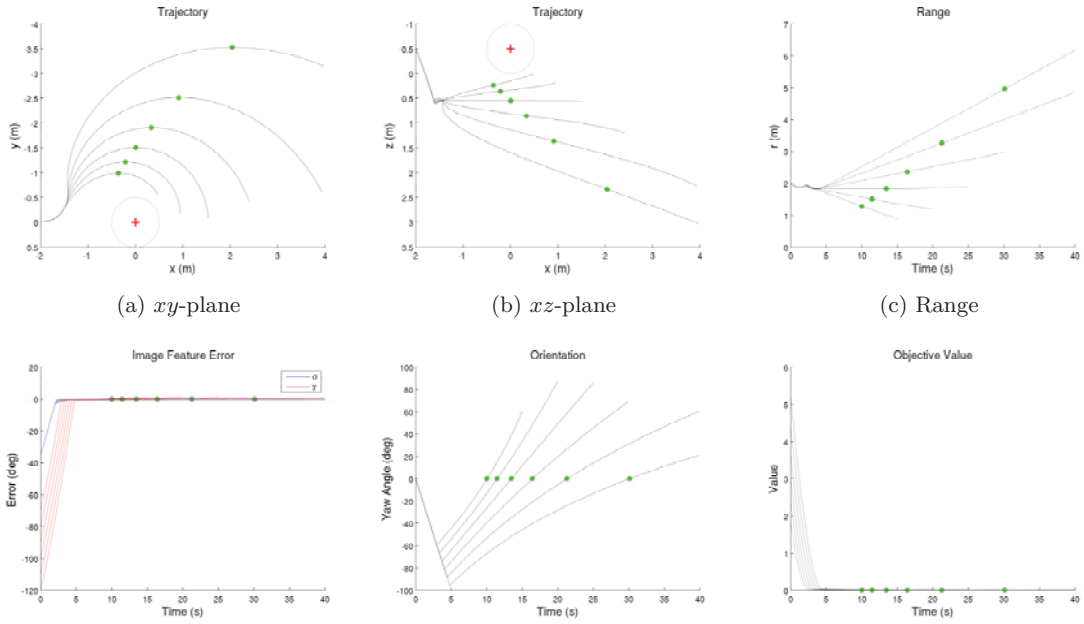

(a) xy-plane (b) xz-plane (c) Range

(d) Feature Error (e) Orientation (f) Objective Function

Figure 5.11: Example trajectories, objective function, feature error and orientation using J_ψ in the controller for $\sigma^* = 125/\pi/180$ and $\gamma^* \in (70\pi/180, 120\pi/180)$ at 10 degree intervals. The auxiliary weighting factors are such that $\lambda_\psi = 0.01$, $\lambda_P = 2$ and $\lambda_R = 1 \times 10^{-9}$. The time corresponding to the objective function minimums (•) are shown for simulation times $t \leq 40$s.

the time of initial object detection t_d. The objective function is now directly dependant on the platform heading and image features, and defined with respect to image features weighting Q_s. The minimum of the objective function would then be obtained when established on the spiral, and the platform has retuned to its initial heading. Therefore, the objective function could be used directly to indicate an appropriate time to stop the avoidance behaviour.

Fig 5.11 depicts a set of example simulations to demonstrate this concept. The augmented objective function is used in the control scheme with $Q_s = 1$, $\lambda_\psi = 0.01$, $\lambda_R = 1 \times 10^{-9}$ and $\lambda_P = 2$, while only the reference azimuth angle is varied. First, the objective function minimum occurs before the platform spirals back toward the static object, denoting an appropriate time to cease avoidance behaviour. Second, having selected a small yaw scaling factor, the general shape of the reference spiral is maintained. However, the exact reference spiral will never be achieved even in the ideal case. This is because the inclusion of the platform heading in the objective function now induces conflicting goals in the minimisation problem. Specifically, maintaining the initial heading or tracking the reference spiral. Fortunately, by formulating the objective function with respect to Q_s, the tradeoff in competing objectives can be managed using the single scaling factor λ_ψ. For $\lambda_\psi = 0$, motion is on the reference spiral such that $J_s = J_\psi$. For $\lambda_\psi \neq 0$, motion is no longer on the true reference spiral as the trajectory becomes inclined and forces a tilted spiral. As λ_ψ increases further, the resulting trajectories become more inclined, separation is reduced and the minimum objective function value occurs at an earlier

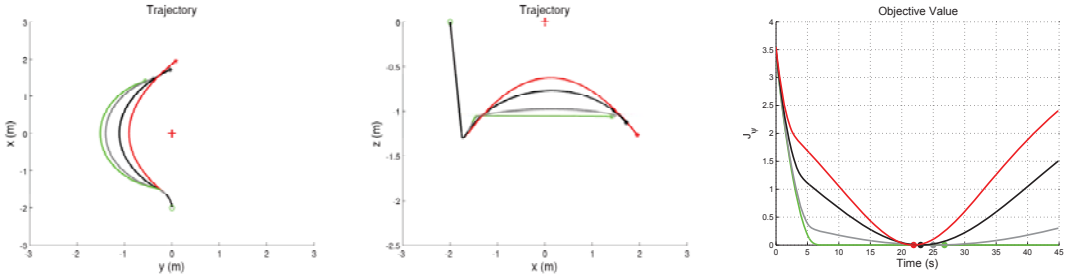

(a) xy-plane (b) xz-plane (c) Objective Function

Figure 5.12: Example simulations using J_ψ in the visual predictive controller and varying λ_ψ from 1 (–), 0.5 (–), 0.1 (–) to 0.0 (–). The corresponding minimum objective function values (●/●/●/●) are also shown for $\mathbf{s}(135\pi/180, \pi/2)$ and $\mathbf{s}_0(\pi/2, 0)$.

instant. Fig 5.12 depicts a set of example simulations to demonstrate the effects. The augmented objective function is used in the control scheme with $Q_s = 1$, $\lambda_R = 1 \times 10^{-9}$, $\lambda_P = 2$ and $\lambda_\psi \in \{0, 0.1, 0.5\}$. The reference image features and initial conditions are the same for all cases.

This situation is not ideal for a collision avoidance strategy that relies on tracking specific spirals to ensure avoidance. To address this issue, and realise the benefits of using the heading in the process model and objective function, a simple solution is proposed. The image feature based objective function $J_s^\star$ can be used in the controller, whilst $J_\psi^\star$ can be calculated outside the controller and used to indicate when to stop the avoidance behaviour. The resolution decision then maintains the explicit coupling to the image features, and the controller does not force the platform to deviate from the reference spiral path. For the nominal case, the minimum value for the augmented objective function $J_\psi^\star$ will be zero when both criteria are met. This will not be possible in practise, due to added uncertainty and the possibility that the spiral may not be completely established for dynamic objects. A small threshold ϵ needs to be applied to $J_\psi^\star$ to indicate when to cease the avoidance behaviour.

Determining an appropriate decision threshold ϵ is not trivial. Consider using the augmented objective function in which the scale (magnitude) of the weighting matrices is varied, but the relative magnitude of the state weighting matrix elements remains the same. Naturally, as the scale of the weighting matrices increases, the absolute value of the objective function will increase. However, there should be no effect on the time at which the objective function is a minimum, only the initial rate at which it converges to that point. This suggests that, in the absence of noise, the threshold value can be kept near zero regardless of the the scale of the weighting matrices. Consider now the inclusion of added measurement noise and imperfect actuation. Their affects will be amplified in the process model, causing larger variations in the objective function value. In this case, even if the resolution decision threshold is small, there is an increased risk that the threshold may be met at an inappropriate instance. As evidence, Fig 5.13 demonstrates the effect of varied weighting matrix scales and additive noise. The augmented objective function is used in the control scheme with $Q_s = 1$, $\lambda_R = 1 \times 10^{-9}$, $\lambda_P = 2$ and $\lambda_\psi = 0.1$.

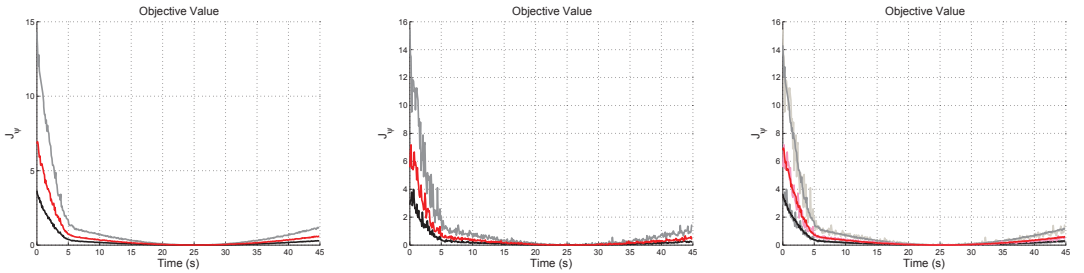

(a) Low Variance Noise (b) High Variance Noise (c) Low & High Variance Noise

Figure 5.13: Example objective function for simulations using varied weighting matrix scale and feature noise characteristics $\xi_s = 2\pi/180$ (*solid*) and $\xi_s = 5\pi/180$ (*faded*). The weighting matrix scale is increased by a factor 1 (−), 2 (−) and 4 (−) for $\mathbf{s}(135\pi/180, \pi/2)$ and $\mathbf{s}_0(\pi/2, 0)$

The reference image features and initial conditions are the same for all cases. The results suggest that an initial resolution decision threshold $\epsilon < 2$ may provide suitable performance, whilst allowing for flexibility in the specific scale of the weighting matrices. It is worth mentioning that for a given system, the weighting matrices will in general be fixed, which means the decision thresholds could be optimised for the expected encounter environment.

The complete resolution decision strategy can now be defined. Recalling, t_s denotes the time after initial detection t_d that the resolution decision has been applied and avoidance behaviour ceases. For $t_d \leq t \leq t_s$, the visual predictive controller provides the control input $\mathbf{u}_3^\star$ to the low-level controllers. For $t < t_d$ and $t > t_s$, control input is derived from a set of PID controllers $\mathbf{u}_3^{\text{PID}}$ that can be used to maintain the heading and altitude at the corresponding time. The resolution decision strategy can then be defined such that

$$\mathbf{u}_3^*(t) = \begin{cases} \mathbf{u}_3^{\text{PID}}(t), & \forall\ J_\psi^*, \qquad 0 \leq t < t_d \\ \mathbf{u}_3^\star(t), & J_\psi^* > \epsilon, \quad t_d \leq t < t_s \\ \mathbf{u}_3^{\text{PID}}(t), & \forall\ J_\psi^*, \qquad t_s \leq t < \infty \end{cases} \tag{5.48}$$

Designing the resolution decision strategy in this way has some additional benefits. The approach is simplistic and relies on tuning a single threshold ϵ that is mutually exclusive from that used in the avoidance decision η. Similarly, both thresholds are also independent of any thresholds used for object detection. As such, each collision avoidance system component can be tuned independently. To this end, probabilistic analysis techniques used in aviation collision avoidance systems can be used to simultaneously optimise each threshold and proving a comprehensive performance evaluation.

5.5.3 Implementation - Experimental System 2

Control Architecture

A small custom built *Ascending Technologies - Hummingbird* quadrotor was used to implement the spherical visual predictive controller and improved avoidance and resolution strategy. The thrust and yaw rate were controlled directly using the VPC scheme or a simple set of PID controllers via $\mathbf{u}_3^*$. The PID controllers are used to maintain constant heading and altitude when the collision has been resolved. An arrangement of LQRI controllers was used to control forward and lateral velocity via $\mathbf{u}_1^*$ and $\mathbf{u}_2^*$ respectively. To this end, the quadrotor position was measured using a Vicon motion capture system, whilst the roll, pitch and yaw angles and associated rates were measured from both the Vicon and onboard IMU's. The positional state information was then used to estimate the vertical, forward and lateral platform velocity using a set of Kalman filters, replacing a typical GPS/INS system found on outdoor platforms. Reference roll and pitch angles were then derived via the velocity LQRI controllers, and used by the on-board attitude controller to regulate forward and lateral velocity (see §2.2.3). The linearised state-space models for the x, y and z axis dynamics were derived about hover, so a set of feed forward terms using the current roll and pitch angles is required to adjust the control commands and minimize any coupling effects. A comprehensive list of the controller, filter and state space model parameters are given in Appendix E.3.

The visual and non-visual controllers were implemented using a ROS framework in a custom ground station command and control module [269], and transmitted to the quadrotor. Importantly, the visual predictive controller was developed and refined used C/C++ such that it could be used directly in MATLAB/SIMULINK, defined within a ROS node or re-built using MATLAB'S code generation toolbox. The control architecture is depicted in figure 5.14.

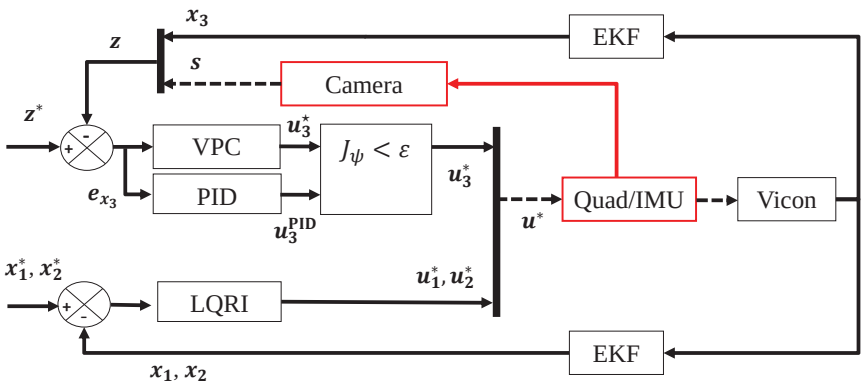

Figure 5.14: Hummingbird quadrotor control architecture. The onboard (□) and ground station components (□) are shown along with wired (—) and wireless connections (——).

Vision System

An onboard camera was used to collect images of a cylindrical shaped object collision object of height 50cm and radius 20cm. The camera was offset from the quadrotor x axis such that $\psi_c = -135\pi/180$ and operating at approximately 10Hz. The camera resolution was 752×480pix, with focal lengths of 623.2pix and 625.2pix in x and y respectively. Images were sent to the ground station and processed to obtain a point feature representation of the cylinder. To this end, a simple image processing algorithm using the OpenCV library [73] was implemented in ROS (using python script), and used for robust detection and tracking. Specifically, a machine-learning based approach using Haar feature-based cascade classifiers was used to consistently track a fixed image of a human face attached to the collision object [358]. An outline of the face detection algorithm is given by algorithm 5.

Algorithm 5 Image Processing Algorithm - Face Detection

▷ Load the ROS image I_R, convert to OpenCV image I and then to a binary image I'

$$I'_{x,y} = \begin{cases} 255, & I_{x,y} \geq 255 \\ 0, & \text{else} \end{cases}$$

▷ Create a smaller image I^* by scaling I' by a factor of 2 using bilinear interpolation

▷ Normalizes the brightness and increases the contrast of I^*. Find the histogram of the image $\mathcal{H}$, normalise the histogram and apply its integral $\mathcal{H}'$ as a transform (look-up table) to I^*.

$$I^\star_{x,y} = \mathcal{H}'\, I^*_{x,y}\,, \qquad \mathcal{H}'_i = \sum_{0 \leq j < i} \mathcal{H}(j)$$

▷ Apply Haar object detection using the classifier cascade for face detection, returning each face as a region $\mathcal{C}^\star(\cdot)$

$$\mathcal{C}^\star(\mathcal{X}^\star, \mathcal{Y}^\star, w^\star, h^\star)$$

$$(\mathcal{X}^\star, \mathcal{Y}^\star) = \text{Upper Left Corner}, \quad w^\star = \text{Width}, \quad h^\star = \text{Height}$$

▷ Assuming a single face is detected, find the region centre $\mathcal{C}_{\mathcal{X}_c, \mathcal{Y}_c}$ in the original image

$$\mathcal{X}_c = 2\left(\mathcal{X}^\star + (w^\star/2)\right), \quad \mathcal{Y}_c = 2\left(\mathcal{Y}^\star + (h^\star/2)\right)$$

▷ Find the spherical image feature representation $\mathbf{s}(\sigma, \gamma)$ of the region centre using the image centre I_{x_c, y_c} and focal lengths f_x and f_y

$$\sigma = \pi/2 - \arctan(\mathcal{Y}_c - y_c)/f_y, \quad \gamma = \arctan(\mathcal{X}_c - x_c)/f_x$$

Of note, the detection and tracking algorithm is susceptible to missed detections and false alarms (tracking the wrong object), depending on the background environment and relative viewing angle. The detection algorithm was pre-trained using the built in OpenCV classifier for frontal faces [73], to ensure satisfactory tracking performance in an arbitrary background. Additionally, the algorithm remains computational inexpense through the successive application of a smaller number of Haar features (lines, edges and four-rectangles) to detect a facial region. The algorithm was chosen to ensure the visual controller was being assessed and not the image processing details.

Results & Analysis

To validate the visual predictive control approach for spiral tracking and collision avoidance, two sets of results are presented. The first set of experimental results are used to analyse the visual predictive control approach for general spiral tracking of a static object. The second set of experimental results are used to analyse the visual predictive control with the added resolution decision strategy for collision avoidance of both static and dynamics objects. The objective of the experiments was not to tune the controller(s) for ideal behaviour for the said encounter(s), but to verify the simulation results presented previously (where possible), and demonstrate the benefits of using spherical visual predictive control for collision avoidance.

a) Spiral Tracking

Consider an instance similar to the previous simulation framework for spiral tracking. In this case, the visual predictive control scheme uses the combined process model, but does not include the augmented cost function or resolution decision. The optimisation problem defined by (5.41) is solved online, and only a single objective function is used (or calculated). The quadrotor was allowed to fly about the static object with forward velocity $v_x = 0.2$m/s and $r^* = 2$. The object initially appears level with the platform and displaced to its right such that $\mathbf{s}_0(\pi/2, -140\pi/180)$ and $r_0 \approx 1.5$m.

The visual control was activated immediately, moving the image features to a desired position $\mathbf{s}^*(80\pi/180, 120\pi/180)$, corresponding to a divergent reference spiral. The choice of reference colatitude is an artefact of the practical camera field of view constraints, and based on the preceding analysis, would have chosen such that $\sigma^* = 125\pi/180$ in a real See and Avoid system. The choice of reference azimuth is to ensure a large initial angular displacement from the reference values, whilst allowing a divergent spiral to be followed. A large number of flight trials were undertaken using a real object and a virtual object. For the real object, the camera (and image precessing) remains in the closed loop control. For a virtual object, the camera is excluded, and the image features are calculated by assuming a known object position and included additive measurement noise $q(t) \sim \mathcal{N}(0, 0.02^2)$. This helps to distinguish between controller performance issues and those not addressed in the thesis, such as model predictive control of delayed systems.

Fig. 5.15 depicts the quadrotor position, orientation, input control and image feature motion for an example set of results using a real object. In this case, induced delay from the image processing (face detection) is as much as 100ms. The reference spiral is tracked well in the vertical plane, as the colatitude error rapidly converges to zero. The added noise, induced delay and range parameter uncertainty (model mismatch) have not significantly effected the controllers ability to track the reference colatitude angle. The reference spiral is not tracked as well in the lateral plane, due to oscillations in yaw angle which result in consistent deviations about the reference azimuth angle. These oscillations are caused by large yaw velocity controls that not only switch polarity, but often reside at the upper control limit. This can be primarily attributed to the induced image processing delay. Each control command issued tends to be fixed for longer periods

of time than required, resulting in persistent overshoot of the reference image features. The control then attempts to overcompensate on subsequent iterations. One way to indirectly compensate for the delay and reduce the overshoot, is to reduce the yaw velocity limits through a more restrictive control constraint domain. Alternatively, the prediction horizon could be increased, but at the expense of increased computational complexity. Importantly, the image delay does not cause instability, keeping the quadrotor in a stable flight configuration.

Fig. 5.16 depicts the quadrotor position, orientation, input control and image feature motion for an example set of results using a virtual object. In this case, having removed the camera from the closed loop control, there is no longer any induced delay from the image processing. The reference image features are tracked well in both the vertical plane and lateral planes, indicated by the rapid convergence of the image feature error. The platform motion then approximates the desired divergent conical spiral trajectory, as the control remains well conditioned. The yaw velocity control no longer switches polarity, forcing the smooth establishment of the reference spiral. The effect is also seen by the gradual change in yaw angle. The observation also helps justify the aforementioned statements regarding the effects of image delay. Again, having compensated for unmodelled disturbances and range parameter uncertainty (model mismatch), the predictive controller ensures a stable flight configuration is maintained.

Collectively, the results suggest that even for a relatively slow frame rate (10Hz), the visual predictive control scheme does not fail, just results in slightly degraded spiral tracking performance. If the delay could be compensated for, using a smith predictor or other recent estimation techniques for visual control [359], the results would then approach those observed for the virtual object. From a collision avoidance perspective, and in a real See and Avoid encounter, it is likely that such a filter and a faster frame rate will be required, given the additional image processing required for aircraft detection and tracking. The development of such extensions is beyond the scope of this work.

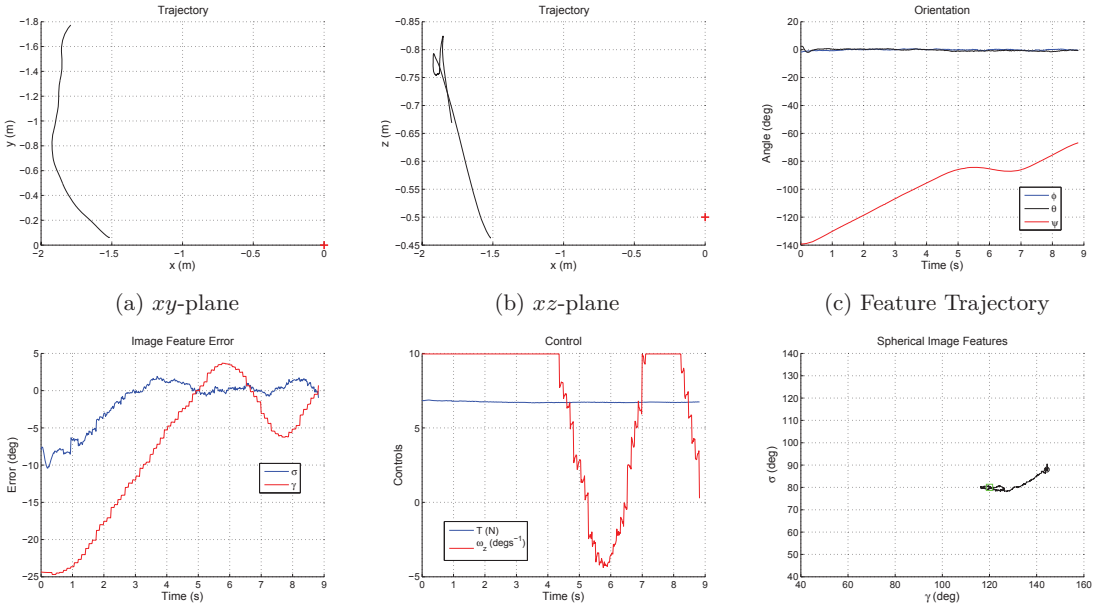

(a) xy-plane (b) xz-plane (c) Feature Trajectory

(d) Feature Error (e) Control (f) Feature Trajectory

Figure 5.15: Example quadrotor behaviour in the world frame $\mathcal{F}_w$ for a static object with $\gamma^* = -120\pi/180$ and $\sigma^* = 80\pi/180$ (□). Initial aircraft and object (+) positions in $\mathcal{F}_w$ are $(-1.5, -0.1, -0.5)$ and $(0, 0, -0.5)$ respectively. A real object is used (including image processing) in the closed loop control.

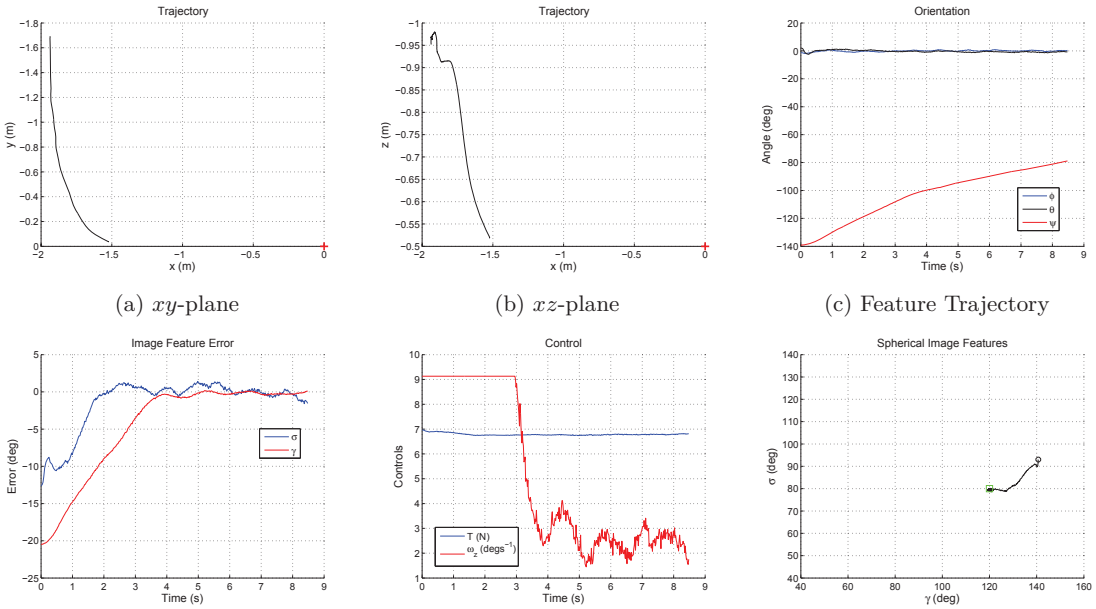

(a) xy-plane (b) xz-plane (c) Feature Trajectory

(d) Feature Error (e) Control (f) Feature Trajectory

Figure 5.16: Example quadrotor behaviour in the world frame $\mathcal{F}_w$ for a static object with $\gamma^* = -120\pi/180$ and $\sigma^* = 80\pi/180$ (□). Initial aircraft and object (+) positions in $\mathcal{F}_w$ are $(-1.5, -0.1, -0.5)$ and $(0, 0, -0.5)$ respectively. A virtual object is used in the closed loop control, including added image feature noise $q(t) \sim \mathcal{N}(0, 0.02^2)$.

b) Collision Avoidance

Consider an instance similar to the spiral tracking example above, but extending the framework for collision avoidance. Again, the visual predictive control scheme uses the combined process model, but now includes the augmented cost function J_ψ in the control framework. The optimisation problem defined by (5.41) is solved online, and the alternate objective function (5.47) is calculated offline and used in the resolution decision strategy. The quadrotor was allowed to fly toward the object with forward velocity $v_x = 0.2$m/s, $r^* = 2$ and time to collision between 10 and 20 seconds. The object initially appears level with the platform and directly ahead such that $\mathbf{s}(\pi/2, 0)$ and $r_0 \approx 2.0$m. Clearly, by scaling these dimensions by a factor of 100, a realistic See and Avoid encounter results.

The visual control was activated immediately, moving the image features to a desired position, corresponding to a circular $\mathbf{s}^*(70\pi/180, \pi/2)$ or divergent $\mathbf{s}^*(80\pi/180, 110\pi/180)$ reference spiral. The choice of reference colatitude is now due to practical constraints regarding fight area, and would have been chosen such that $\sigma^* = 125\pi/180$ in a real See and Avoid system. The choice of reference azimuth is now consistent with the general avoidance strategy outlined previously (see §3.3.1). A large number of flight trials were undertaken using a virtual object for various resolution decision thresholds, and both static and dynamic objects. A virtual object is required in this case in order to better replicate collision encounters, given the fixed camera orientation and field of view limitations. Specifically, head-on and crossing encounters originating at any initial image feature position can be modelled.

Fig. 5.17 depicts the avoidance behaviour and augmented objective function for two sets of six separate static encounters. In each case, a divergent reference spiral is required such that $\mathbf{s}^*(80\pi/180, 110\pi/180)$ and the weighting matrix values are such that $Q_s = 1$, $\lambda_R = 0.01$ and $\lambda_P = 1$ to force a moderate response. Therefore, the objective value should be greater than that in simulation due to the non-zero control penalty. Different resolution decision thresholds ϵ_1 and ϵ_2 are used for each set of encounters.

In (a)-(c), a conservative resolution decision threshold is chosen such that $\epsilon_1 = 0.8$. In each case, the platform quickly converges to the reference spiral, before ceasing the avoidance behaviour and leaving the spiral path prior to returning to the initial heading upon detection. This can be seen directly in the platform trajectories where $\psi(t_s) < \psi(t_d)$, or inferred by the increasing objective function value. In (d)-(f), a liberal resolution decision threshold is chosen such that $\epsilon_2 = 0.6$. Similarly, the platform quickly converges to the reference spiral, but now ceasing the avoidance behaviour and leaves the spiral path very near to the initial heading upon detection. Again, this can be seen directly in the platform trajectories where $\psi(t_s) \approx \psi(t_d)$, or inferred by the delayed increase in objective function.

Fig 5.18 depicts the avoidance behaviour and augmented objective function for a set of static and dynamic object encounters. In each case, a circular reference spiral is required such that $\mathbf{s}^*(70\pi/180, \pi/2)$ to align with the general avoidance decision strategy. The terminal weighting matrix values are now twice as large as those used with ϵ_1 and ϵ_2 such that $\lambda_P = 2$, to force a faster response. Based on the previous results, and considering

the increased terminal penalty term, a more conservative resolution decision threshold $\epsilon_3 = 1.9$ is chosen.

In (a)-(c), a set of seven static object encounters are shown. The platform quickly converges to the reference spiral, and considering that the conservative nature of the threshold $\epsilon_3 > 2\epsilon_2$, the avoidance behaviour ceases at an earlier instant for the majority of cases. This can be seen directly in the platform trajectories where $x(t_s) > 0.5$m. For the remaining cases however, the avoidance behaviour is no longer stopped in some cases, and the spiral path is continued. This can be seen directly in the platform trajectories and the objective function value such that $J^\star_\psi > \epsilon_3$. For clarity, the marker denoting t_s is placed at the end of the encounter. In such cases, and provided $\gamma^* \geq 90°$, continuing the flight would force the platform to spiral the object continuously until $J^\star_\psi \leq \epsilon_3$. Otherwise, for $\|\gamma^*\| < \pi/2$, a collision may result if the platform dynamics and constraint domain allow large velocities. In (d)-(f), a set of ten dynamic (constant velocity) object encounters are shown. Five different object trajectories are simulated using the virtual object. They include planar head-on an crossing encounters, with the remaining also including fixed climb or descent rates. Each encounter is simulated at two separate object velocities $\mathbf{v}_t$ such that $\|\mathbf{v}\| > \|\mathbf{v}_t\|$, resulting in ten collision encounters. Although not included in the process model, this is essentially corresponds to adding a small disturbance to the translational velocities used in the image kinematics. In each case, the platform attempts to track the reference image features and establish the appropriate spiral. In all but two cases, the avoidance behaviour ceases at an appropriate time due to the use of the augmented objective function. After which, the objective value rapidly increases as the relative motion forces faster divergence from the reference image features. Importantly, having not accounted for the object motion in the dynamic model, the reference image features can be tracked. This observation further highlights the approach's robustness to model mismatch, without requiring accurate range parameter estimates in the process model or additional integral control.

Combining the results, three important observations can be made. First, the results provide empirical evidence to support such a threshold based visual predictive control approach to collision avoidance and resolution. Most encounters are appropriately resolved and no collisions occur, even for liberal resolution decision thresholds and varied weighting matrices. Second, the resolution decision threshold represents a degree of freedom that can be designed according to safety considerations. For example, a system that requires strict safety margins to be maintained may warrant a conservative threshold. A system that is less concerned with false alarms may use a liberal threshold, relying on the fact that there may be another opportunity (pending relative velocity and heading) to satisfy the resolution criteria. Third, given the diversity of potential object behaviours, it is difficult to ensure that the resolution decision threshold will always be met. However, assuming fixed weighting matrices for given airframe[4], the threshold can be further optimised for the expected encounter environment.

[4]The weighting matrices may depend on some desired aircraft performance specifications or limitations.

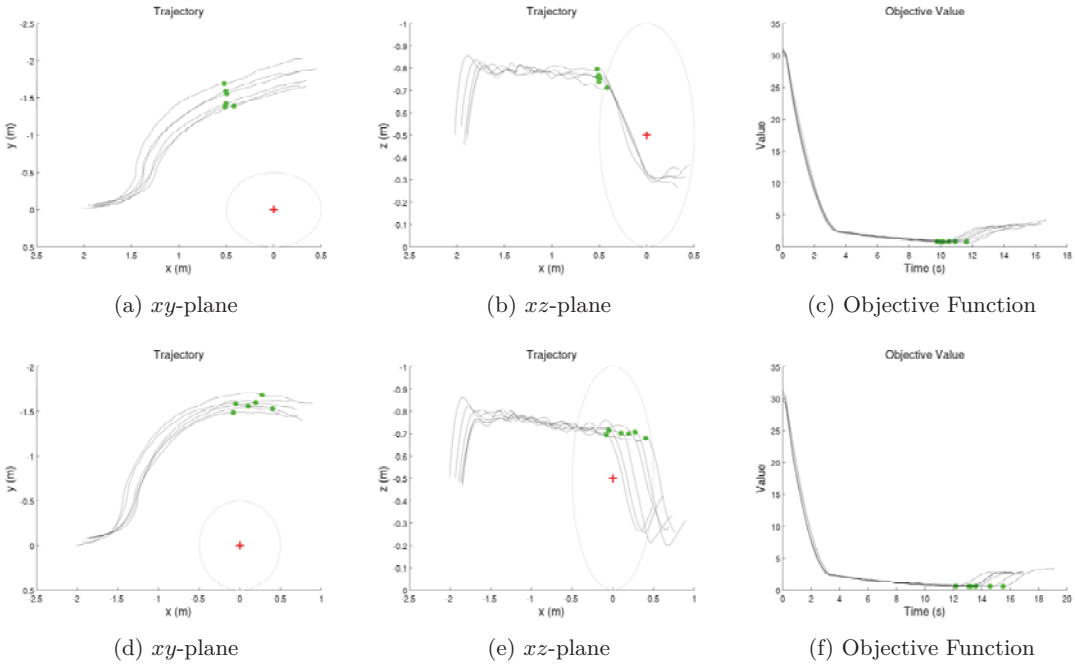

(a) xy-plane (b) xz-plane (c) Objective Function

(d) xy-plane (e) xz-plane (f) Objective Function

Figure 5.17: Example real aircraft trajectories and objective function for collision avoidance of static objects using a conservative (a)-(c) and liberal (d)-(f) resolution decision threshold *epsilon*. Initial aircraft and object (+) positions in $\mathcal{F}_w$ are $(-1.5, -0.1, -0.5)$ and $(0, 0, -0.5)$ respectively, with $r^* = 2$ and $\mathbf{s}^*(80\pi/180, 110\pi/180)$ (□). A virtual object with added image feature noise $q(t) \sim \mathcal{N}(0, 0.02^2)$ is used.

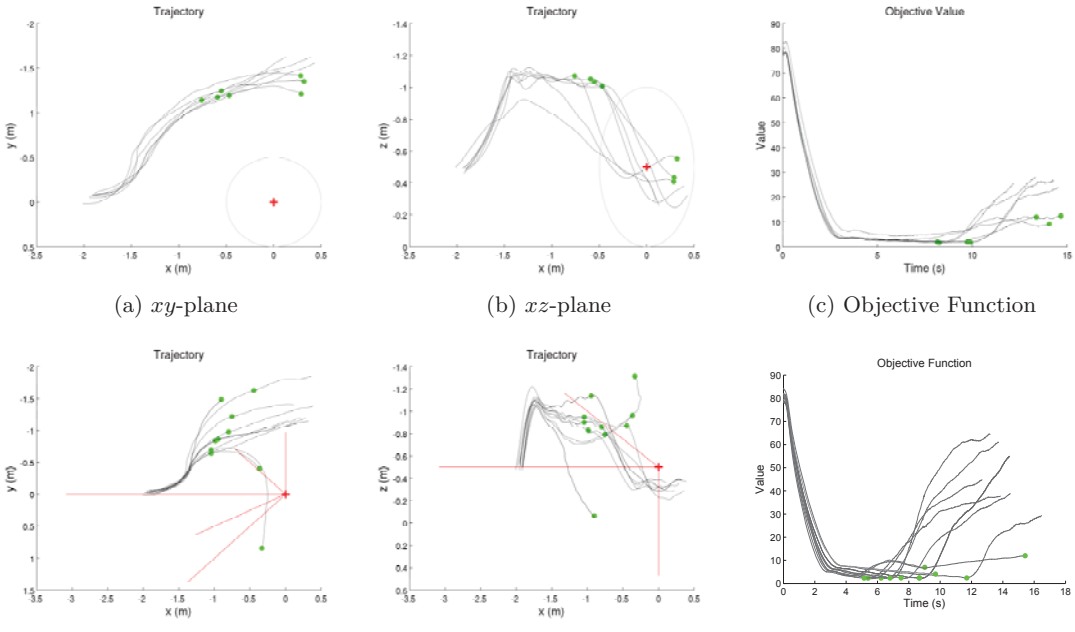

(a) xy-plane (b) xz-plane (c) Objective Function

(d) xy-plane (e) xz-plane (f) Objective Function

Figure 5.18: Example real aircraft trajectories and objective function for collision avoidance of static (a)-(c) and dynamics (d)-(f) objects using a more conservative decision threshold $epsilon_3 = 1.9$. Inital aircraft and object (+) positions in $\mathcal{F}_w$ are $(-1.5, -0.1, -0.5)$ and $(0, 0, -0.5)$ respectively, with $r^* = 2$ and $\mathbf{s}^*(70\pi/180, \pi/2)$ (□). There are five different object trajectories (−), including head-on and crossing with the remaining including climb or descent. Each are simulated at two separate $\mathbf{v}_t$ values, resulting in 10 collision scenarios. A virtual object with added image feature noise $q(t) \sim \mathcal{N}(0, 0.02^2)$ is used.

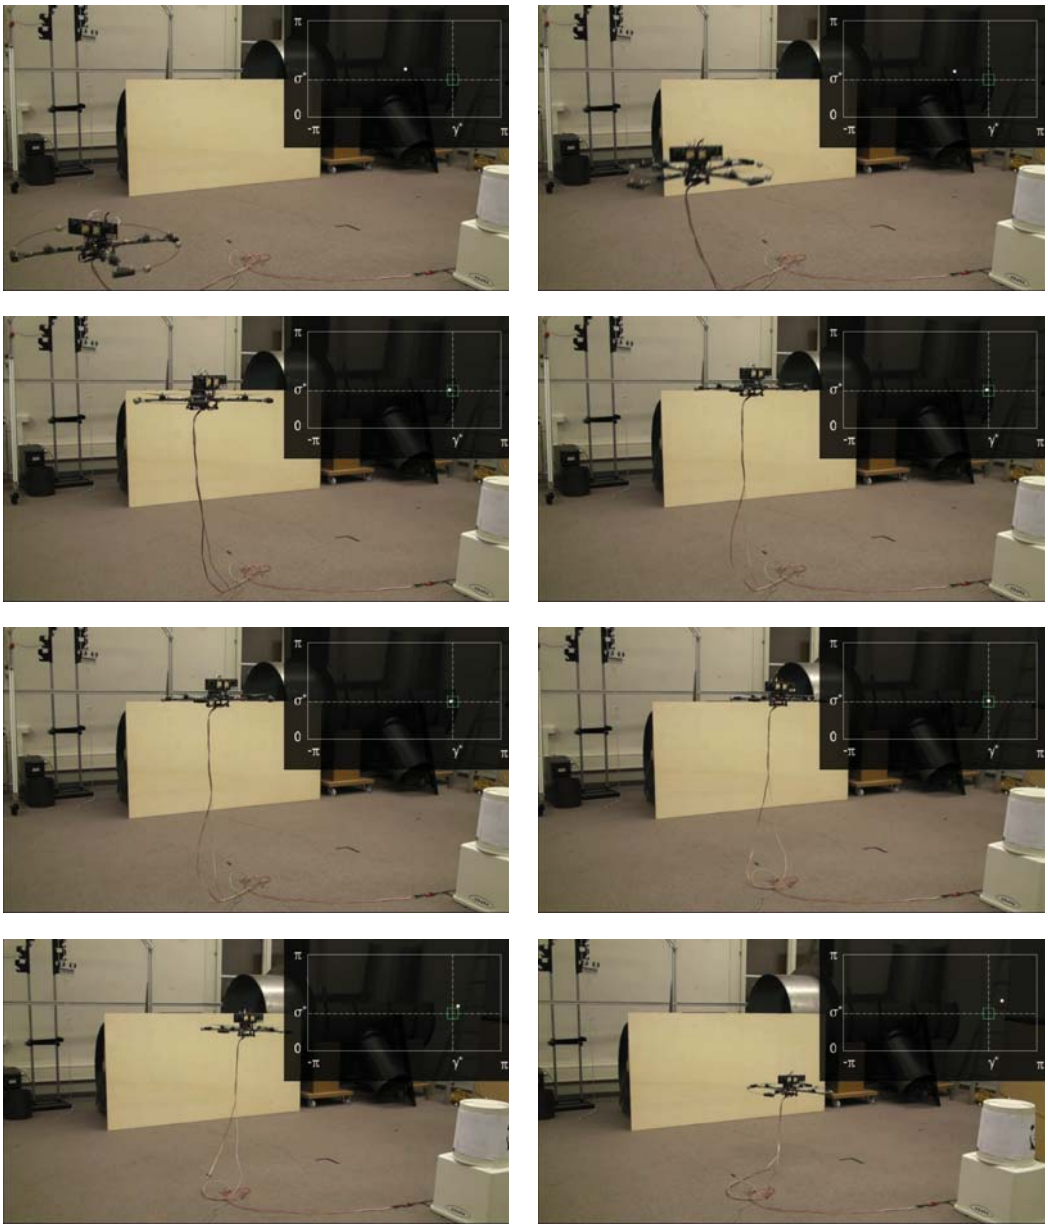

Figure 5.19: Example flight trial using spherical visual predictive control for collision avoidance of a virtual object positioned at the white cylinder. The controller regulates the image feature motion (*insert*) to the reference position $\mathbf{s}^*(80\pi/180, \pi/2)$(□), avoiding the object by spiralling to the left. The image features then diverges (○) as the quadrotor returns to its initial altitude and heading. The tether provides power to the platform and restricts motion in the confined environment. The encounter evolves from $t = 0$s (*upper left*) to $t < 20$s (*lower right*)

5.6 Summary

This chapter introduced a set of novel spherical visual predictive controllers for spiral tracking and collision avoidance, using principles from nonlinear model predictive control. Visual predictive control strategies using terminal equality constraints, terminal penalty terms and a quasi-infinite horizon framework were derived and compared to existing predictive approaches. The results were used to design a practical visual predictive control framework for spiral tracking. The structure of the controller was then exploited to include a novel threshold based resolution strategy for collision avoidance. A practical implementation for static and dynamic objects was demonstrated using a small quadrotor. Only minor modifications to the avoidance decision strategy were required to account for practical constraints. The key attributes, benefits and drawbacks of using the visual predictive control framework for collision avoidance include:

General Visual Predictive Control

- The general nonlinear model predictive control structure lends itself naturally to the vision-based collision avoidance problem, providing an excellent framework in which to uniquely manage the difficult problem attributes. Specifically, the framework is capable of explicitly managing platform and camera constraints, modest image processing delay and added uncertainty through imperfect actuation and image feature noise. This in turn allows consideration of both static and dynamic objects, without explicitly considering the projected object motion in the process model. This relaxes the section of the reference range parameter used in the controller compared to the classical image-based approach. Of note, this is only possible for objects that are considered to be relatively far away as in the See and Avoid environment.
- The predictive controllers can be implemented in a computationally efficient (and tractable) manner using modern solvers, without the need for move blocking. Even for the quasi-infinite horizon framework, any computational complexity is managed offline, and the online minimisation problem can be solved sufficiently fast. This is an important consideration for general image-based control, given its sensitivity to system latency through image processing delays, and the desire for optimal performance [288].

Quasi-Infinite Horizon Visual Predictive Control

- This work represents the first application of the quasi-infinite nonlinear model predictive control design to visual control of any form. The direct application of the stability-based control framework to image-based visual control is not straight forward, or indeed applicable to any system. Coupling a spherical camera and single point features allows such a stability based-design technique to be applied. The calculation of the terminal penalty matrix and terminal image region can be challenging, but a simplified methodology that can be directly implemented using modern tools has been provided. The approach provides a detailed outline of how to implement and solve each of the optimisation problems in a manner consistent to MATLAB requirements.

- The resulting control frameworks for circular (with terminal region constraint) and spiral (terminal equality constraint) motion provide comparable spiral tracking performance with respect to recently proposed predictive control schemes. For the circular motion case, reduced track error and a 3% reduction in control effort averaged over the entire state space is observed. Of note, comparisons between various predictive schemes are often made with respect to the classical image-based approach only and not competing predictive schemes.
- The inclusion of the terminal region constraint provides some unique advantages regarding the design of reference image features, stability and feasibility. The size of the terminal region gives a quantitative representation of the degree of system nonlinearity at specific reference image features, so could be used to help determine the domain of attraction for both predictive and classical image-based visual servoing schemes. Considering feasibility implies stability for quasi-infinite horizon approaches, stable reference image features may then be selected a priori. When applied to the general case of circular motion and collision avoidance, the reference colatitude angle σ^* associated with a larger terminal region can then be selected a priori, such that there is better assurance that the true conical spiral can be tracked. Of note, and using a similar argument, the terminal region could be used to help design suitable reference image features for region-reaching image-based visual servoing[5]
- Although nominal closed-loop stability can be assured for circular motion, the same stability guarantees cannot be assumed for a real system. For aircraft spiral tracking and collision avoidance, model mismatch and added uncertainty result in deviation from the nominal or ideal behaviour. However, by designing for the nominal case and using an appropriate prediction horizon, local asymptomatic stability can still be achieved as in standard model predictive control approaches.

Practical Visual Predictive Control

- This work represents the first practical implementation of visual predictive control of any form for control of aerial vehicles. Using a terminal penalty term and no terminal constraints, the feasibility of the resulting control structure is demonstrated for systems with fast dynamics. Stable behaviour can be assured using moderate prediction horizons, whilst enforcing platform constraints. Improvements are still required to better manage image processing delay, either though improved processing algorithms, or the inclusion of additional filters. Of note, a number of simulated and some practical visual predictive control implementations have been proposed, but are typically for robot manipulators and using simplifying assumptions such as multiple feature points.

[5]For example, a divergent spiral or one that exists entirely above or below the apex may be required. Therefore, the reference image features are not unique, but instead exist in a particular region of the image. Recent region-reaching controllers [318] will fail when using a single point feature, due to singularities at the region boundary, so selection of specific image features may still be required. Selecting reference image feature in the desired image area associated with a large terminal region reduces the likelihood of feasibility issues.

- Although the control framework was applied to qaudrotor control, the approach is general and can be extended to multiple platform types (fixed and rotary wing) and camera models (perspective, wide angle etc.). The appropriate system dynamics and image kinematics can be combined in the process model in a similar manner (see §5.5.1), and the same resolution strategy can be employed.
- The novel threshold based resolution decision strategy using two separate objective functions provides an improved solution, coupling the visual control and aircraft heading. Multiple flight trials demonstrated the effectiveness of the strategy for both static and dynamic objects. Unlike the initial approach to resolution (see §4.4.1), avoidance was never stopped prematurely and failure to meet the particular resolution decision threshold ϵ does not imply collision. Alternatively, the derivative of the objective function could have be used, but tended to be noisy in the real system. A low pass filter could be applied, but this means additional tuning parameters and possible added delay. In either case, achieving the resolution decision threshold cannot be assured given the diversity of object motion, however tuning a single parameter reduces the development work required.
- As the decision points of the collision avoidance system have been collapsed into two mutually exclusive decision thresholds (for resolution ϵ and avoidance η), the complete system naturally lends itself to probabilistic performance analysis techniques used in aviation (see §1.3.2). Not only do these techniques offer a more comprehensive evaluation of the system performance than assessing particular case studies, they also allow the simultaneous tuning of system thresholds. As such, they can be used to determine optimal thresholds $\epsilon^\star$ and $\eta^\star$ using Monte-Carlo type simulations or otherwise (see §6). Of note, this represents the first time a automated vision-based See and Avoid system has been formulated with realistic and explicit regard to existing aviation practise, procedures and certification standards.

Chapter 6

Performance Evaluation

6.1 Outline

This chapter presents a preliminary probabilistic performance evaluation strategy to assess the complete vision-based See and Avoid system. The evaluation technique extends existing approaches used in aviation, resulting in a general framework in which to asses and tune the avoidance and resolution decision strategies.

First, the possible collision avoidance outcomes are highlighted for the proposed automated See and Avoid system. Traditional collision avoidance outcomes are augmented, by considering the inclusion of automated avoidance, control and resolution decisions. Second, extensions to existing probabilistic collision avoidance performance evaluation strategies used in manned aviation are derived. Specifically, it is shown how system operating curves can be used to capture, visualise and optimise the effects of threshold (and parameter) selection on system performance for vision-based avoidance. Third, a series of Monte-Carlo simulations are presented in order to demonstrate the effectiveness of the complete system, using such system operating curves. Of note, this is the first time a vision-based collision avoidance system has been cast into a framework amenable to such probabilistic analysis techniques, and subsequently analysed. This is an important consideration when certifying operational systems.

6.2 Collision Avoidance Outcomes

To asses the performance of an arbitrary collision avoidance system, it is important to uniquely define the possible system outcomes. The outcomes themselves are related to the end state, or relative state, of the involved aircraft. The specific outcome of a random encounter is not only a function of the relative geometry, but the complex interactions between system parameters. Such parameters include, but are not limited to, sensor uncertainty and the decision points or logic used in the collision avoidance system. It is then intractable to enumerate each separate outcome, and instead it makes sense to categorise the resulting collision state according to a discrete set of outcome types.

Considering a coarse definition for the collision avoidance outcome categories, a simple binary classification could be used in which the result is either collision or no collision. However, considering that a collision avoidance system is being used, this does not provide

sufficient detail that relates to the functionality of the collision avoidance system. Therefore, additional categories are required that acknowledge the inclusion of the collision avoidance system.

Consider first that the collision avoidance system can be modelled as an alerting system. This means that an avoidance decision or alert is automatically issued, but the control and subsequent resolution decision is assumed by either a pilot or other operator. In this case, the possible outcomes can now be linked directly to an avoidance system. Generally speaking, the possible outcomes include

UA (FA)	:	Unnecessary Alert (False Alarm)
CA (CD)	:	Correct Avoidance (Correct Detection)
MA (MD)	:	Missed Avoidance (Missed Detection)
IC	:	Induced Collision

A **UA** results when an avoidance action is taken and the collision is avoided, but no action was required given the nominal aircraft trajectories. A **IC** results when an avoidance action is taken and a collision is created (induced), that would not have occurred given the nominal aircraft trajectories. A **CA** results when an avoidance action is required and the collision was successfully avoided. A **MA** results when an avoidance action was required but the collision was unsuccessfully avoided.

Importantly, each outcome category assumes that avoidance implies resolution so there is no automatic resolution decision to be made. Arguably, if a pilot was in the loop this may be the case and the assumption is valid. As such, the outcomes are commonly used for TCAS like systems that provide avoidance advisories. For a fully automated system though, it should be recognised that just because an avoidance action was taken, this does not mean that the avoidance behaviour was stopped or resolved. This is certainly the case for the collision avoidance approach presented in this work, and is likely the case for other similar vision-based systems. Therefore, this assumption cannot be overlooked for a completely automated See and Avoid, and an additional set of outcomes is then required to classify the possible end states.

Consider a See and Avoid system that now includes an automated avoidance and resolution decision. When an avoidance action is taken, then there is the possibility that the action is stopped or continued (attempted) for all time. Operator intervention may then be required to manually terminate the avoidance behaviour. Consider also that all objects are considered potential collisions based on the short duration of the expected encounter and initial detection distance. In this case an action is always taken, but will vary depending on the avoidance decision. The notion of taking decisive action upon detection, is supported by recent lines of thought in military applications regarding adaptive action [362, 363]. Taking initial action may help obtain a greater understanding of the situation, such that the avoidance response can be better adapted. In a simple approach, consider two types of action. The first could be considered action that is intended to avoid a perceived collision. The second is intended to increase separation between aircraft in a precautionary manner.

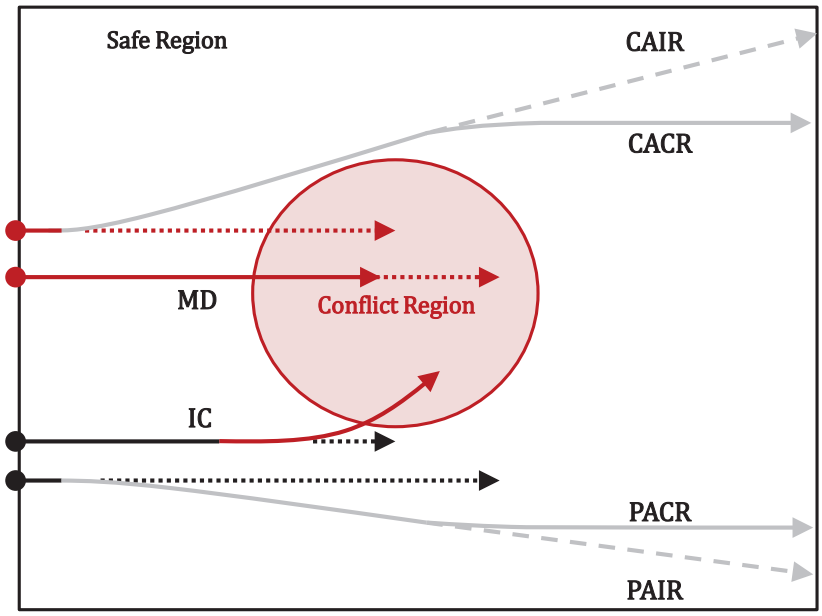

Figure 6.1: Automated collision avoidance outcomes

The precautionary action may also facilitate the acquisition of additional information to help improve future control decisions. For these systems, the possible outcomes include

PACR	:	Precautionary Avoidance & Correct Resolution
PAIR	:	Precautionary Avoidance & Incorrect Resolution
CACR	:	Correct Avoidance & Correct Resolution
CAIR	:	Correct Avoidance & Incorrect Resolution
MA	:	Missed Avoidance
IC	:	Induced Collision

A **PACR** results when precautionary action is taken and then successfully resolved. A **PAIR** results when precautionary action is taken but is incorrectly (or unsuccessfully) resolved. For both cases, avoidance was not required to maintain separation. A **CACR** results when collision avoidance action is taken and then successfully resolved. A **CAIR** results when collision action is taken but is incorrectly resolved. A **MA** and **IC** take the same meaning as before, and as they denote collision cases, augmenting them with a resolution outcome does not make sense. Of note, unnecessary actions would no longer constitute a possible outcome set under the assumption action is always taken. Instead, unnecessary action could be replaced with incorrect action. Incorrect action means that precautionary avoidance was taken in a true collision encounter and vice versa. An example of the proposed collision avoidance outcome categories for an automated See and Avoid system are depicted in Fig 6.1.

6.3 System Operating Curves

Given the set of possible outcomes, it is then possible to define a set of metrics that provide a quantitative representation of the system performance. Taking a statistical approach, the metrics can be based on the observed counts of each outcome type for a larger number of arbitrary encounters. The metrics are typically defined in terms of a probability of positive and negative outcomes. A set of outcomes may be combined (summed) to define each metric. For example, the probability of a collision may be the sum of the induced collisions and missed avoidances. Depending on how these metrics are combined and subsequently visualised, they may also be used to represent the variation in performance with system parameters. Recalling, the outcome of an arbitrary encounter depends on the system parameters, including any thresholds, and so influences the performance metrics.

In signal detection theory, the idea is to identify a target signal amongst noise. The system metrics consist of a probability of correct detection $P(\mathbf{CD})$ and false alarm $P(\mathbf{FA})$. For a given set of system parameters, detection thresholds and noise characteristics, the metrics are evaluated and plotted as a single point on a Receiver Operating Curve (ROC). The system parameters are then altered and the metrics are re-evaluated to obtain a set of points, resulting in a curve [360]. The curves have been leveraged directly to evaluate automated vision-based aircraft detection performance and trade-offs, given the problem is almost identical [95, 96].

The decision components in collision avoidance and alerting systems bears strong similarity to the object or signal detection problem. As such, receiver operating curves have since been recognised as an effective way to evaluate collision avoidance and alerting system performance [168, 189, 361]. In this domain, the term System Operating Curve (SOC) is used, and the metrics are derived by considering a combination of various outcome types. The sum of desirable outcomes is used to estimate the probability of successful avoidance $P(\mathbf{SA})$ and unsuccessful avoidance $P(\mathbf{UA})$. Again, each set of system parameters map to a particular point along the curve. Altering the avoidance (or alerting) decision threshold τ moves the operating point along the curve. Therefore, the curve can be used to directly evaluate system performance and visualise the tradeoff between the desirable and undesirable system outcomes. Additionally, they can be used to evaluate the system sensitivity to different encounter types or uncertainty characteristics. The use of System Operating Curves and Receiver Operating Curves in evaluating the performance of collision avoidance system components are depicted in Fig 6.2, with example curves shown in Fig 6.3.

For vision-based See and Avoid systems, it then makes sense to leverage System Operating Curves to optimise system parameters and demonstrate performance in a manner consist with aviation practise. The added benefit is that the detection and avoidance components can be analyses using the same principles, and easily combined to derive the overall system performance. Interestingly, most autonomous collision avoidance systems are not analysed and evaluated in this way. Instead, performance is typically analysed with respect to stability and feasibility having assumed the correct avoidance decision has been made.

Figure 6.2: System Operating Curves for See and Avoid Systems

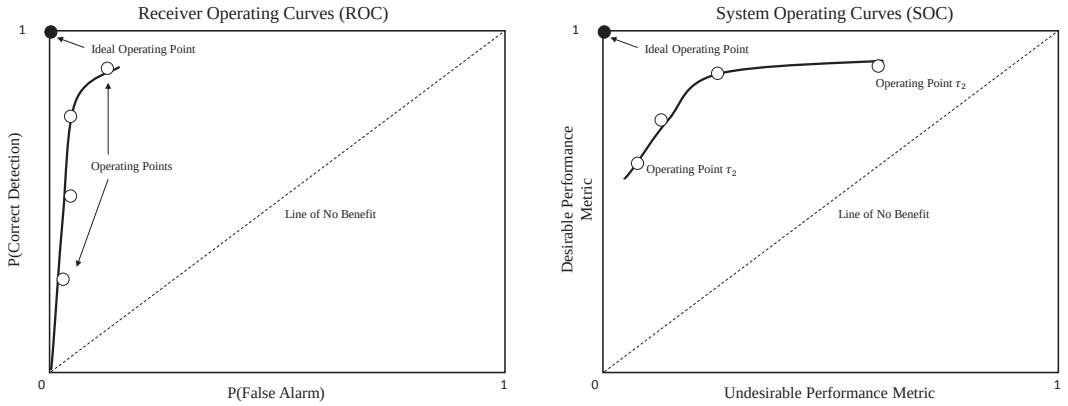

(a) ROC (b) SOC

Figure 6.3: Example Receiver Operating Curves (ROC) and System Operating Curves (SOC) used for performance evaluation in Signal Detection and Alerting Systems.

Often simplifying assumptions are made to ensure the decision is suitable. For example, the nominal case is considered, object intent is known, or a cooperative encounter in which both aircraft take action is assumed. Such assumptions are unrealistic for many See and Avoid encounters and certainly when considering a vision-based system.

Given the structure of the collision avoidance system presented in this work, System Operating curves and the traditional performance metrics cannot be applied directly. Starting from first principles, the appropriate performance metrics need to be defined. First, determine the probability that an encounter E exists (collision or non-collision) given the current state $\mathbf{x}$. Then, the probability that the object is detected D, given an encounter exists, can be defined such that

$$P(\mathbf{CD}) = P(D|E)P_{NA}(E|\mathbf{x}) \tag{6.1}$$

$$P(\mathbf{CD})' = P(D'|E)P_{NA}(E|\mathbf{x}) \tag{6.2}$$

where P_{NA} denotes a probability along the nominal trajectory, where no action is taken, and $P(\mathbf{CD})' = P(\mathbf{MD})$ denotes a missed detection. Assuming perfect detection, then $P(D|E) = 1$ and $P(\mathbf{CD})' = 0$. The probability that an encounter actually occurs is difficult to determine considering that unmanned aircraft are a new airspace user, and no traffic models exist. By extrapolating radar data, and considering a number of assumptions, an integrated air traffic model can be approximated [64, 65]. Such models could

be used to approximate $P_{NA}(E|\mathbf{x})$ but pertain to a particular air space system and are not freely available. Alternatively, quoted figures from historical data may be used in the interim to estimate this value [19].

Second, determine the probability of specific outcome types given the collision avoidance system is included, and based on how it is operating. This enables the performance metrics to be directly linked to the key system parameters, the avoidance η and resolution ϵ decision thresholds. Using the collision avoidance outcomes defined for the proposed See and Avoid system, the performance metrics can be stated as the probability of each outcome types such that

$$P(\mathbf{CACR}) = P'_{mc}P(\mathbf{CD})P_{NA}(C|E)P_{CA}(C'R|\eta^*, \epsilon^*) \tag{6.3}$$
$$P(\mathbf{CAIR}) = P'_{mc}P(\mathbf{CD})P_{NA}(C|E)P_{CA}(C'R'|\eta^*, \epsilon^*) \tag{6.4}$$
$$P(\mathbf{PACR}) = P'_{mc}P(\mathbf{CD})P_{NA}(C'|E)P_{PA}(C'R|\eta^*, \epsilon^*) \tag{6.5}$$
$$P(\mathbf{PAIR}) = P'_{mc}P(\mathbf{CD})P_{NA}(C'|E)P_{PA}(C'R'|\eta^*, \epsilon^*) \tag{6.6}$$
$$P(\mathbf{MA}) = P_{mc}P_{NA}(C|E) + P'_{mc}P(\mathbf{CD})P_{NA}(C|E)P_{CA}(C|\eta^*, \epsilon^*) \tag{6.7}$$
$$P(\mathbf{IC}) = P'_{mc}P(\mathbf{CD})P_{NA}(C'|E)P_{PA}(C|\eta^*, \epsilon^*) \tag{6.8}$$

where P_{CA} denotes a probability along the avoidance trajectory and P_{PA} denotes a probability using precautionary action. The notation C and R denote a collision and resolution event respectively, such that R' implies unsuccessful resolution and C' denotes successful avoidance. The probability of a system malfunction P_{mc} is also included such that a malfunction denotes no action is taken. Clearly, the system metrics demonstrate the complex interactions between the system parameters and the encounter geometry. For this work, the concern is the performance and optimisation of the avoidance and resolution decision thresholds. As such, assuming the system does not malfunction and the detection system is capable of detecting all objects reliable, then P($\mathbf{CD}$) ≈ 1 and $P_{mc} = 0$. To further simplify the metrics, $P_{NA}(C|E) < 1$ can be considered a constant scaling factor that will reduce each probability considering that the chance of an encounter of any sort (collision or non-collision) is rare. Indeed, it is this term that contributes to very low collision probabilities for some systems[1]. Simplifying the performance metrics in this way, then

$$P(\mathbf{CACR}) = P_{CA}(C'R|\eta^*, \epsilon^*), \qquad P(\mathbf{CAIR}) = P_{CA}(C'R'|\eta^*, \epsilon^*) \tag{6.9}$$
$$P(\mathbf{PACR}) = P_{PA}(C'R|\eta^*, \epsilon^*), \qquad P(\mathbf{PAIR}) = P_{PA}(C'R'|\eta^*, \epsilon^*) \tag{6.10}$$
$$P(\mathbf{MA}) = P_{CA}(C|\eta^*, \epsilon^*), \qquad P(\mathbf{IC}) = P_{PA}(C|\eta^*, \epsilon^*) \tag{6.11}$$

[1]The probability of an encounter is assumed to be unity, and the probability of a collision or non-collision can be considered to be equally likely. Realistically, these probabilities would depend on the location, airspace and time the aircraft is operating.

Combining the performance metrics such that System Operating Curves can be used is a matter of summing the positive and negative outcomes. The probability of desirable $P(+)$ and undesirable system behaviour $P(-)$ can then be expressed as

$$P(+) = P(\mathbf{CACR}) + P(\mathbf{CAIR}) + P(\mathbf{PACR}) \tag{6.12}$$

$$P(-) = P(\mathbf{MA}) + P(\mathbf{IC}) \tag{6.13}$$

Note the inclusion of P(**CAIR**) in the desirable performance outcome, as the primary concern is to avoid, given the See and Avoid environment. Note also that P(**PAIR**) has been excluded for the undesirable metric. This is indeed a negative outcome, but to keep focus on the the system ability to avoid a collision, the induced collision is of more importance. Of course, the metric can be included if required. Provided each of the metrics is available, or can be derived, a modified System Operating Curve can then be constructed. The term modified is used to distinguish the different axis values compared to traditional System Operating Curve.

Using the above framework, a modified Operating Curve could then be constructed for each of the decision thresholds and optimised accordingly. Holding the avoidance threshold fixed and varying the resolution decision threshold enables the system performance to be analysed with respect to the resolution strategy, and vice versa. Importantly, this is possible because the two points, and thresholds, are mutually exclusive and decoupled (see §3.3.1 and §5.5.2). The impact of altering each decision threshold can then be visualised such that the design of the thresholds can be optimised. In each case, the ideal operating point is the upper left hand corner of the plot. At this point, the decision threshold is such that no collisions occur and the encounter is always resolved, for both precautionary and avoidance actions. Due to system uncertainty, this point cannot be achieved. Instead, the best possible achievable performance can be expected in a region about the ideal operating point. Importantly, as sensors improve the entire curve will be affected, shifting closer to the ideal operating point. As such, the curve can still be used to evaluate the relative effect of the decision thresholds on overall performance. The use of the modified System Operating Curves in evaluating the performance of See and Avoid system components is depicted in Fig 6.4 with an example curve shown in Fig 6.5.

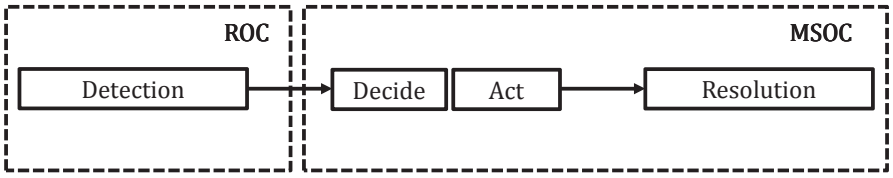

Figure 6.4: System Operating Curves for See and Avoid Systems

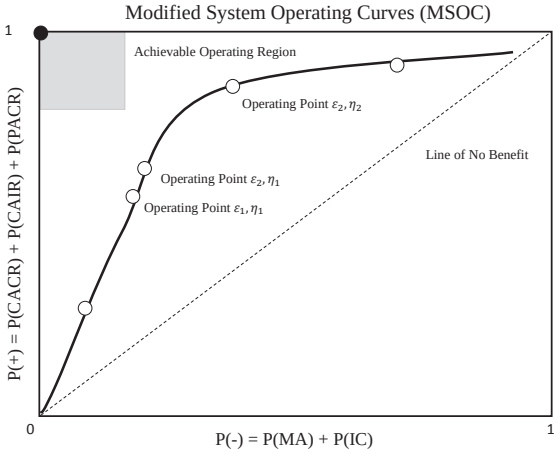

Figure 6.5: Modified System Operating Curve

The performance metrics are not always straight forward to derive. Typically, they must be approximated from counts of the observed collision outcomes when subject to a wide range of diverse encounter scenarios. Due to cost and safety aspects, this is currently only possible through simulation (typically Monte-Carlo) as the number of encounter scenarios needed would make flight trials infeasible. As such, this is the approach taken in this work to assess the system performance and optimise the decision thresholds. Of note, this is the first time a See and Avoid system has been derived and evaluated in this manner, and represents an important step toward certification. The simulation details and results are outlined in the following sections.

Remark As an alternate to System Operating curves, Detection Error Tradeoff (DET) can also be used to visualise the same system performance.

6.4 Monte-Carlo Trials

6.4.1 Simulation Environment

A simulation framework was developed to model a scaled down See and Avoid environment. To remain consistent with the practical implementation in the previous chapter, and ensure a realistic system model, the same small *AscTec Hummingbird* quadrotor platform is used in simulation. A dynamic model of the platform, including time delays is used to model the platform motion (see §2.2). A set of Kalman filters were used to estimate the platform velocity, and a set of Linear Quadratic Regulators (LQRI) controllers were used to control forward and lateral velocity. A set of PID controllers were used to regulate the vertical rate and yaw rate to zero if the resolution decision is activated. Otherwise the platform continues the avoidance behaviour. The visual predictive controller used to guide spiral motion and collision avoidance is based on the simulated and practical results obtaining in the previous chapter (see §5.5). Importantly, the process model used in the simulation environment includes experimentally derived delay terms [269], unlike the process model used in the visual predictive controller. The weighting matrix elements Q_s, λ_R and λ_P were kept small and remained fixed for all simulations. The prediction horizon $T_p = 10$ allows a solution to the minimisation problem to be found sufficiently fast, without compromising stability and control performance. Lastly, the sampling period for object detection and image processing T_i was assumed to be longer than that for the predictive controller T_v.

A target trajectory generator was created to generate either a collision encounter or non-collision encounter. The collision encounters were generated by selecting a random collision time and velocity, to then calculate the initial object position. A non-collision encounter was derived by first calculating a collision encounter as before, then adding a random non-zero translational velocity to the object. In this way, no collision is possible under the nominal trajectories. All encounters were bounded by a maximum and minimum time to closest approach t_{cpa}, initial range r_0 and object velocity v_t.

The simulation model includes environmental disturbances from turbulence $\mathbf{w}_g(t)$ and ambient wind $\mathbf{w}_a(t)$, added image feature noise $q(t)$ and imperfect actuation $\mathbf{q}_c(t)$ as sources of uncertainty. Each source is modelled using a Gaussian distribution of one or more variables. The variability in image feature (azimuth and colatitude) measurements is assumed to be equal, so $q(t)$ is added to each measurement. Imperfect actuation is modelled by additive noise on the control inputs, which results in small thrust offsets and angular displacements in roll, pitch and yaw rate. Turbulence is modelled by additive noise on the platform orientation. Ambient wind is modelled by additive noise on the translational velocities of each platform. Of note, the mean ambient wind is considered zero. This is a valid assumption as the ambient wind is usually known such that the platform heading is offset to maintain desired ground track. What cannot be accounted for is variability in the wind speed and direction, therefore the variance of the ambient wind is non-zero. A complete list of the simulation parameters is given in Appendix E.

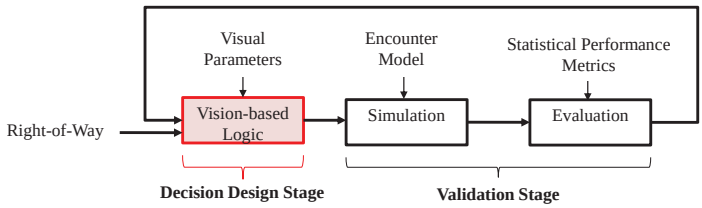

Figure 6.6: Example collision avoidance performance analysis methodology for vision-based collision avoidance systems.

The simulation environment was implemented using MATLAB Simulink and the ACADO Toolkit. A series of Monte-Carlo type simulations for a large number of collision and non-collision encounters were conducted, in order to asses the utility of the vision-based collision avoidance systems. Specifically, the focus is on the effectiveness of the avoidance and resolution decision strategies. There are two aims of the simulation trials. First, to determine the viability of using such an approach to automated See and Avoid. Considering the controller has been validated both in simulation and empirically, the idea is to quantify the effectiveness of the avoidance and resolution strategies. Second, to determine the value of using the proposed performance evaluation framework as a generalised means to asses performance, and visualise the design tradeoffs. In this case, the tradeoffs consist of analysing and interpreting the relative effects of the avoidance and resolution thresholds on collision avoidance performance.

To achieve these goals, an iterative simulation process is used. The avoidance or resolution decision threshold is fixed whilst the other is varied for a series of 1000 collision and 1000 non-collision encounters. The process is then repeated, switching the decision threshold that is varied whilst the other is fixed. The process is aligned to existing methods in which to analyse performance of threshold based systems (see §5.5), and is depicted in Fig 1.3.2. The results of the analyses are detailed in the following sections of this chapter.

6.4.2 Avoidance System

The complete collision avoidance system is first analysed from qualitative perspective. A subset of the collision and non-collision encounters are examined to ensure that each of the system components are functioning as expected, irrespective of the specific thresholds used for avoidance and resolution. First, the resolution decision is analysed to ensure the avoidance behaviour can be stopped. Second, control performance is analysed to ensure feasible aircraft trajectories are applied. Third, the avoidance decision is analysed in light of the rules of the air in which they are based upon.

An example set of 15 collision and non-collision encounters are shown in Fig 6.7. The the object and quadrotor trajectories are shown, along with the instance in which the resolution decision was activated. This is shown by the green square superimposed over the quadrotor trajectory. If the resolution decision was not activated, the green square appears at the end of the quadrotor trajectory. To this end, only a few scenarios remain unresolved for both collision and non-collision encounters.

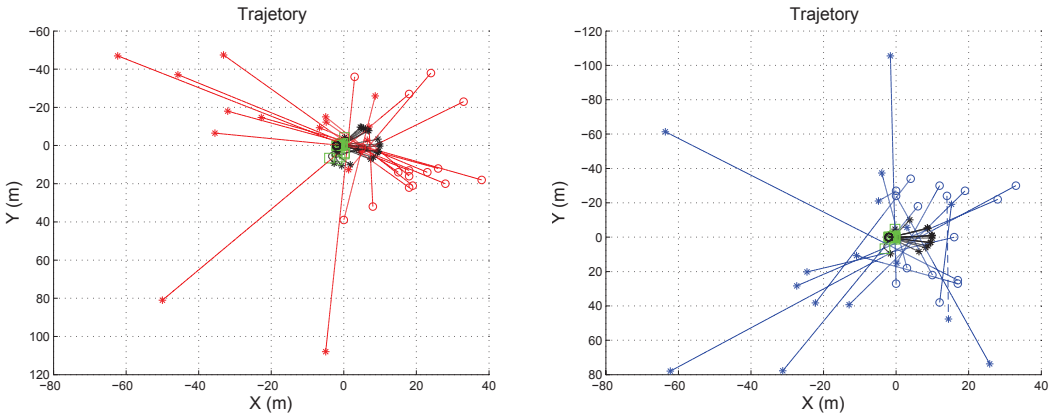

(a) Collision Encounters (b) Non-Collision Encounters

Figure 6.7: Subset of 15 quadrotor trajectories (−) subject to collision objects (−) and non-collision objects (−). The initial quadrotor position (○) in $\mathcal{F}_w$ is $(-2, 0, -10)$ whilst the initial object position (○/○) is varied. The resolution instance is also shown (□).

The quadrotor trajectories appear to be stable and avoidance decision are made, but the plots are too unclear to provide a more comprehensive analysis. To better asses the avoidance decision and control behaviour, consider removing the object trajectories and focusing on the quadrotor behaviour. Additionally, consider further examining three examples taken from the same 15 encounters. These addition plots are depicted in in Fig 6.8(a)-(c) and Fig 6.8(d)-(f) for the collision and non-collision sets respectively.

First, consider the example collision encounters. Fig 6.8(a) shows the smooth and stable quadrotor trajectories resulting from the avoidance decision. The specific avoidance decision depends on the projected feature motion and avoidance decision thresholds η. Fig 6.8(b)-(c) shows the trajectories of each aircraft for different avoidance decisions. For the blue encounter, the object appears on the right and is allowed to cross, essentially giving way to the object. For the black encounter, the object has approached from the left and the vehicle has allowed the object to pass behind it. For the red encounter, an

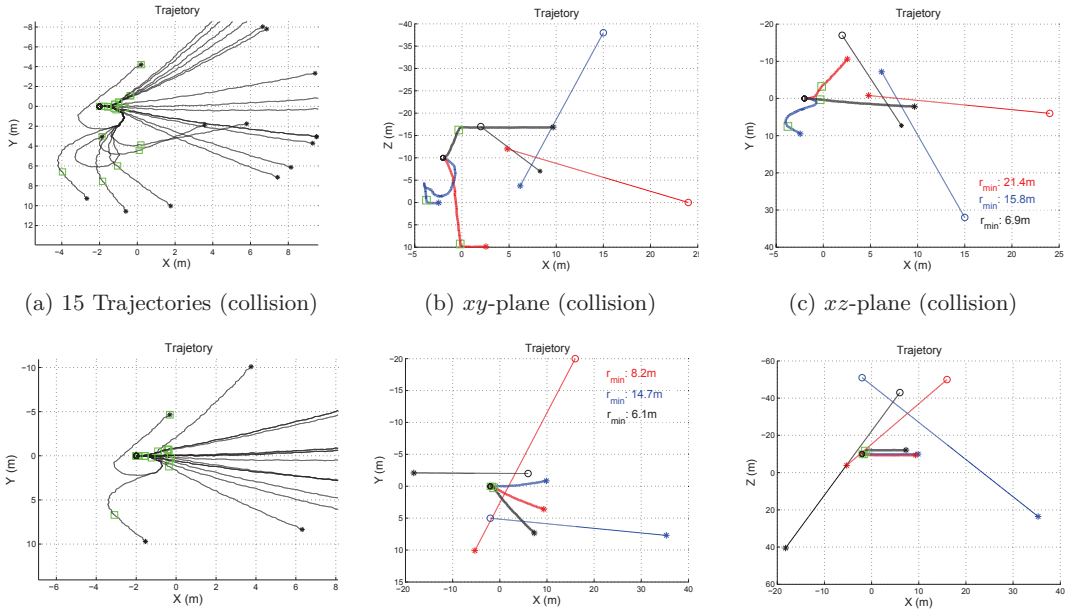

(a) 15 Trajectories (collision) (b) xy-plane (collision) (c) xz-plane (collision)

(a) 15 Trajectories (non-collision) (b) xy-plane (non-collision) (c) xz-plane (non-collision)

Figure 6.8: Example avoidance behaviour and quadrotor trajectories for a subset of collision and non-collision encounters. The initial quadrotor position in $\mathcal{F}_w$ is $(-2, 0, -10)$. The color coded quadrotor (*thick*) and corresponding object (*thin*) trajectories are for 3 example encounters. The resolution instance (□) minimum separation are also shown.

almost head-on scenario occurs and the vehicle avoids the object on the right. In all cases the appropriate avoidance decision was adopted based on the guidelines outlined by aviation rules of the air . A successful resolution decision was made in all but one case, but every collision was still avoided. As such, all three encounters can be classified as **CACR** outcomes. The single unresolved encounter can be classified as a **CAIR** outcome, and there are no **MA** outcomes.

Second, consider the example non-collision encounters. Fig 6.8(d) shows the smooth and stable quadrotor trajectories, this time resulting from the precautionary avoidance. Fig 6.8(e)-(f) shows the trajectories of each aircraft for different precautionary avoidance decisions. For the blue encounter, the object appears on the right, but is not allowed to cross. Essentially the object is not considered a collision threat, so the quadrotor does not give way to the object. For the black encounter, the object has appeared on the left in a non-collision near head on encounter. The avoidance decision has recognised that the object is not a collision threat, but takes precautionary action to increase separation. For the red encounter, the object approaches from the left and is then allowed to cross behind the quadrotor. Again, the avoidance decision has tried to increase separation by moving to the right, having recognised the object as non threatening. Again, in all cases the appropriate avoidance decision was adopted. A successful resolution decision occurs relatively early in each encounter (bar one), causing decreased deviation from the original flight path. The reason the quadrotor trajectory appears to diverge, is only because of the non-collision avoidance control forcing a zero altitude and heading rate. As such, all three encounters can be classified as **PACR** outcomes. The single unresolved encounter can be classified as a **PAIR** outcome, and there are no **IC** outcomes.

6.4.3 Avoidance Decision

The practical implementations in earlier chapters required a simplified avoidance decision to account for the limitations in implementing the entire automated system. The previous section used the original avoidance strategy, but only a small subset of the avoidance decision outcomes for a single avoidance decision threshold were analysed. Although the results are promising, it is unlikely that an appropriate avoidance decision will be adopted for an arbitrary encounter, given the diversity of object trajectories, image features uncertainty, ambient conditions and imperfect actuation. A full analysis of the avoidance decision strategy is then required. The analysis needs to be conducted by keeping all the simulation parameters, predictive control and resolution decision parameters fixed. The encounters and avoidance decision parameters or thresholds are then varied.

Recalling that the avoidance decision is based on comparing a measurement of image feature convergence or divergence $\check{\mathbf{s}}$, to a fixed avoidance threshold η where

$$\check{\mathbf{s}} = \mathrm{di}(\mathbf{s})\,\dot{\mathbf{s}}^T, \quad \eta = \mathrm{di}(\xi_{\mathbf{s}})\,\xi_{\dot{s}}^T \tag{6.14}$$

To simplify the analysis, consider that the uncertainty on the image feature position ξ_s and velocity $\xi_{\dot{s}}$ measurements is equal in both colatitude σ and azimuth γ. In this way, the avoidance threshold η for each image feature can be set to the same value. Consider first setting the avoidance threshold equal to the product of the variance on the image feature position and velocity as above. The threshold can then be scaled to be larger or smaller than this reference value using a scaling factor λ_η such that

$$\eta = \lambda_\eta(\mathrm{di}(\xi_{\mathbf{s}})\,\xi_{\dot{s}}^T) \tag{6.15}$$

where $\lambda_\eta \in \{1/8, 1/4, 1/2, 1, 2, 4, 8\}$. For each threshold, 1000 random collision and non-collision encounters are generated, resulting in 2000 encounters per threshold. The resolution decision was fixed such that $\epsilon = 0.15$, and a collision boundary r_c was set such that a collision is declared if $r(t) \leq 0.25$m. Scaling the collision boundary up results in a relatively conservative region, given the typical size of small unmanned aircraft. The resulting modified System Operating Curve is shown in Fig 6.9. The first shows the relative location of each threshold using the full axis. The second focuses on the region interest. Using the curve, two important observations can be made.

First, it shows that the avoidance decision strategy is generally effective at avoiding collision, irrespective of the avoidance threshold. This is shown by the the collection of operating points located toward the ideal operating point. Collision avoidance is approximately 90% effective, with a relatively low rate of induced collisions and missed avoidances.

Second, it highlights that it is indeed difficult to uniquely distinguish between a collision and non-collision object under realistic flight conditions. This is depicted by the small variations in operating point location, that lack a unique and distinguishable pattern with respect to decision threshold variations. For example, a scale factor of 4 shows improved

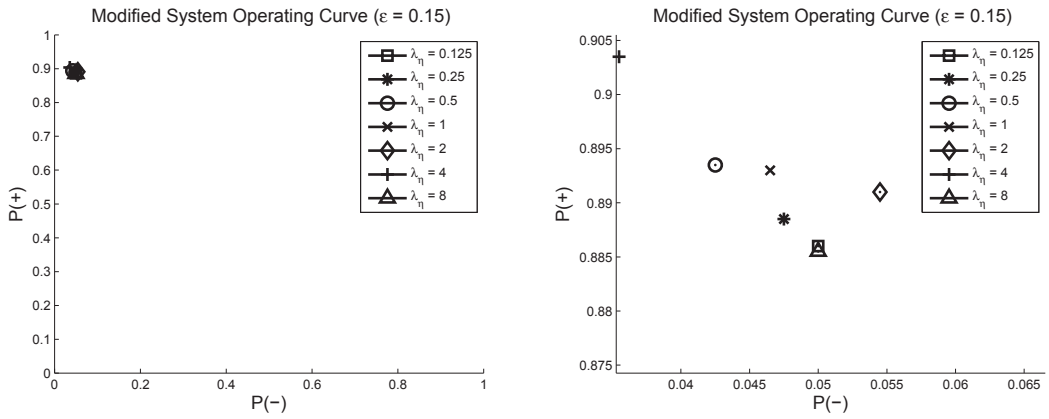

(a) Original

(b) Zoomed

Figure 6.9: Example modified System Operating Curve (SOC) for variable avoidance decision threshold $\lambda_\eta \in (0.125, 8)$. The resolution threshold and collision boundary are such that $\epsilon = 0.15$ and $r_c = 0.25$m respectively.

performance, but a scale factor of 8 shows degraded performance compared to a scale factor 2. Interestingly, the result compliments the experimental results obtained in [130] without the need for flight tests. However, having always acted, and using a different decision variable $\check{\mathbf{s}}$, the approach is more effective than using bearing rate alone.

The above results were verified for different resolution thresholds and simulation parameters, ensuring the observed effects can indeed be attributed to the avoidance decision strategy. As such, although the avoidance strategy is effective, it is still unclear exactly how to the select η. It may be such that further simulations are required to ensure the encounter sample size is sufficient enough to appropriately represent the true system behaviour. In this way, an optimal avoidance decision threshold $\eta^\star$ might be obtained. This is yet to be determined.

Remark It may be possible to define a region $\Omega_{\bar{\eta}}$ about an arbitrary nominal operating point $\bar{\eta}$ corresponding to $\lambda_\eta = 1$, such that the system performance will remain in that region for all avoidance decision thresholds $\eta \in (\bar{\eta}/8, 8\bar{\eta})$. This may be coarsely approximated by first finding the maximum 2-norm (Euclidean) distance d_η between the point where $\lambda_\eta = 1$ and all points corresponding to $\eta \in (\eta/8, 8\eta)$. A circular (or elliptical) region of radius d_η can then be defined about the nominal operating point. Of note, it is reiterated that this is only an approximation based on preliminary results, such there may exist some $\eta \in (\bar{\eta}/8, 8\bar{\eta})$ that lie outside $\Omega_{\bar{\eta}}$.

6.4.4 Resolution Decision

The practical implementations in earlier chapters allowed the full resolution strategy to be implemented. This was also the case in the previous section, but similar to the avoidance decision, only a small subset of the resolution decision outcomes for a single resolution threshold were analysed. Again, it is unlikely that an appropriate resolution decision will be adopted for an arbitrary encounter, given the diversity of system parameters and added uncertainty. Indeed, some examples in which no resolution decision was made have already been presented. A full analysis of the avoidance resolution strategy is then required. In this case the analysis needs to be conducted by keeping all the simulation parameters, predictive control and avoidance decision parameters fixed. The encounters and resolution decision parameters or thresholds are then varied.

Recalling that the resolution decision is based on comparing the value of an augmented objective function J_ψ, to a fixed resolution threshold ϵ where

$$J_\psi = J_s + \lambda_\psi \int_t^{t+T_p} (\psi(\tau) \ominus \psi^*)_{Q_s} \, d\tau, \quad \epsilon > 0 \tag{6.16}$$

The resolution threshold ϵ can be chosen somewhat arbitrary. However, the previous chapter suggests that the value should be kept relatively small and scaled according to the weighting matrix elements Q_s, λ_R and λ_P (see §5.5.2). Given the weighting matrix elements are fixed for all simulations, consider $\epsilon \in \{0.02, 0.05, 0.07, 0.09, 0.15, 0.30\}$. For each threshold, 1000 random collision and non-collision encounters are generated, resulting in 2000 encounters per threshold. The avoidance decision was fixed such that $\eta = \lambda_\eta(\mathrm{di}(\xi_\mathbf{s})\, \xi_{\dot{s}}^T)$ where $\lambda_\eta = 1$. Again, a collision boundary r_c was set such that a collision is declared if $r(t) \leq 0.25$m. The resulting modified System Operating Curve is shown in Fig 6.10. The same curve is shown in two different scales. The first shows the relative location of each threshold using the full axis. The second focuses on the region interest. Again, two important observations can be made.

First, it shows that the resolution decision strategy is generally effective at resolving avoidance and precautionary avoidance behaviour, irrespective of the resolution threshold. Specifically, the resolution decision helps to avoid inducing future collisions, indicated by the set of operating points located toward the ideal operating point between 65%-95%.

Second, system performance is significantly influenced by the placement of the resolution threshold compared to the avoidance decision curve presented earlier. This is depicted by the variation in operating point location that gives the plot the same characteristic shape as typically receiver operating curves. The operating points move toward the ideal operating point as the threshold increases, before diverging when the threshold is increased past 0.15. To this end, the plot is actually visualising the tradeoff between maintaining the reference spiral or ceasing avoidance behaviour. For small thresholds, a greater number of collisions are probable as the platform is reluctant to leave the spiral having never reached the resolution threshold. This causes the aircraft to continually attempt to spiral the object, which can be problematic for some cases. As the threshold

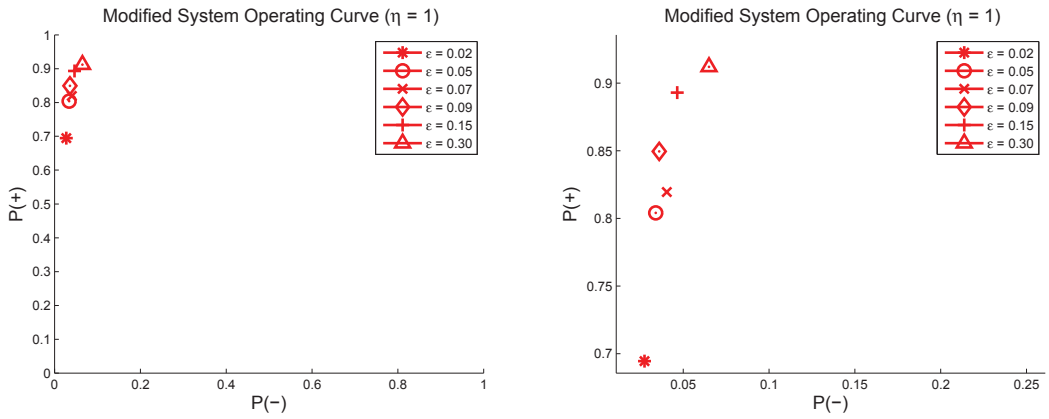

(a) Original (b) Zoomed

Figure 6.10: Example modified System Operating Curve (SOC) for variable resolution decision threshold $\epsilon \in (0.02, 0.30)$. The avoidance threshold and collision boundary are such that $\lambda_\eta = 1$ and $r_c = 0.25$m respectively.

increases, more encounters are resolved so fewer encounters offer the opportunity to induce future collisions. As the threshold is increased further, it becomes very conservative. As such, there is an increased potential for the avoidance action to be stopped at an earlier instance. This is because it is considerably easier to meet the resolution threshold. As a consequence, the aircraft may not have maneuvered enough to avoid a collision or near miss. This means that the amount of correct resolutions will increase, but at the expense or increased risk of collision. This can be seen as the curve starts to move away from the ideal operating point.

The above results were verified for different avoidance thresholds and simulation parameters, ensuring the observed effects can indeed be attributed to the resolution decision strategy. This implies that the modified system operating curve can then be used to optimise the resolution decision threshold to obtain $\epsilon^\star$. To this end, optimisation is with respect to the designer preferences. If collision was of utmost importance, then a lower then a larger, conservative threshold would be preferred over a lower value.

6.4.5 System Performance Example

Combining the above results, a series of system operating curves can be easily derived for different collision boundaries $r_c \in \{0, 0.125, 0.25, 0.5\}$ to help visualise the complete system performance and variability. The operating point for each resolution threshold $\epsilon \in \{0.02, 0.05, 0.07, 0.09, 0.15, 0.30\}$ can first be plotted. A curve can then be constructed by connecting each operating point, in order to visualise an approximate curve for all other resolution thresholds $0.02 < \epsilon < 0.30$. An approximate region of operating points Ω_η corresponding to the bounded set of avoidance thresholds where $\lambda_\eta \in (1/8, 8)$ can then be superimposed at each operating point. Taking this approximate approach, the results are depicted in Fig 6.11.

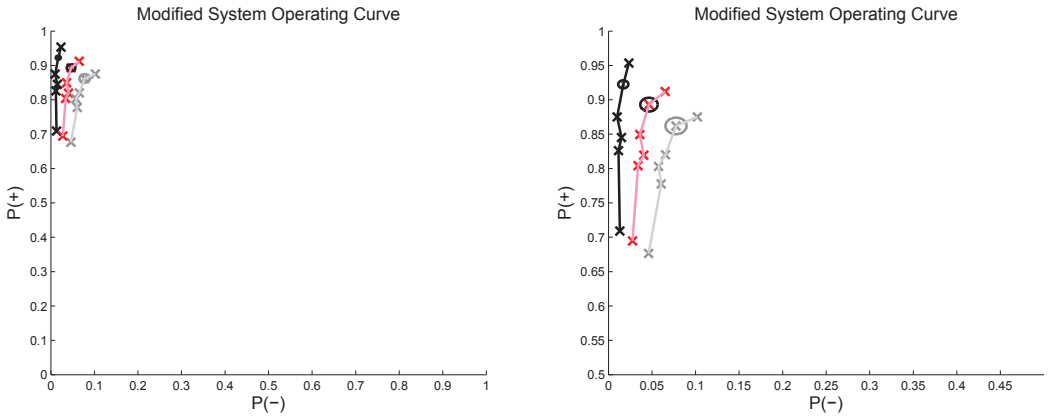

(a) Original (b) Zoomed

Figure 6.11: Example modified System Operating Curve (SOC) for variable resolution $\epsilon \in (0.02, 0.30)$ ($\times$) and avoidance $\lambda_\eta \in (0.125, 8)$ decision threshold. The set of avoidance thresholds is approximated using the region $\Omega_{\bar{\eta}}$ ($\circ$). Three curves are shown for a collision boundary r_c of 0.125 ($-$), 0.25 ($-$) and 0.5 ($-$).

As expected, a smaller collision boundary shows improved system performance provided all other parameters remain the same. However, the results are open to interpretation given the aircraft sizes need to be considered and the arbitrary selection of the collision boundary. For example, the results suggest that for $r_c = 0.125$ the system is 95% effective at avoiding collisions. This could mean that a collision is only just avoided (with no additional separation) for two small aircraft with an approximate radius ≤ 0.0625, and an actual collision boundary of zero. Conversely, for $r_c = 0.5$ the system is only 85% effective at avoiding collisions. This could mean that each collision is now avoided with non-zero separation if considered the same two small aircraft. Of course, the result could also be interpreted as collision avoidance of lager aircraft with zero separation. The results then provide an indication of the relative system performance with respect to aircraft size and separation standards.

When using system operating curves, it is also important to consider the notion of a feasible avoidance trajectory. In some instances, it may be such that given the control constraints, collision boundary, aircraft size and initial geometry, a collision is unavoidable. In these cases, no decision or control strategy can provide collision avoidance. Removing such encounters from the data set and then recalculating the operating curve would move each operating point slightly closer to the ideal operating point. However, if the number of occurrences is small compared to the total number of encounters, the shift will be barely visible. Although such cases are possible under the simulation parameters used to construct the system operating curves above, no such encounters were observed. This was verified by checking each case resulting in collision, and re-simulating the encounter using the control input limits (maximum and minimum) until the point of collision. If the maximum of the minimum range value for each of the two avoidance directions (left and right) remains less than the collision boundary, it is infeasible to avoid the encounter. Some example results are shown in Fig 6.12 for $\epsilon = 0.15$ and $\lambda_\eta = 1$.

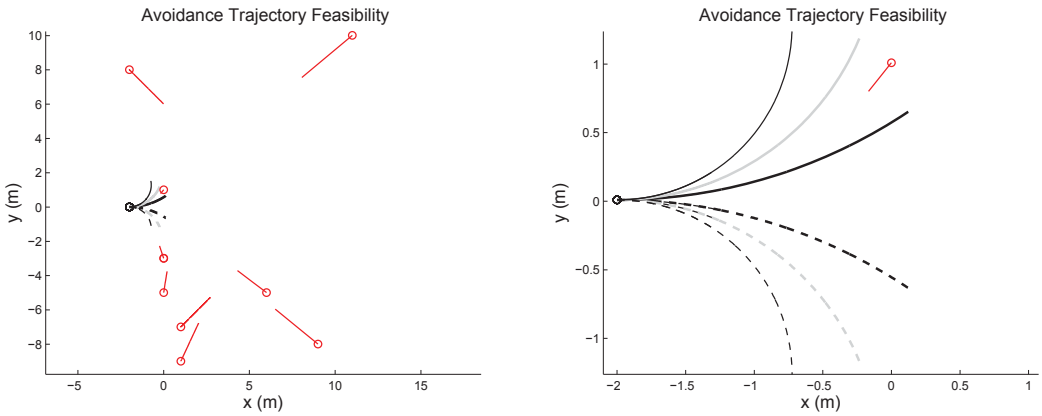

(a) Feasible Trajectories

(b) Potentially Infeasible Trajectory

Figure 6.12: Example illustration of feasible and infeasible collision encounters. Object trajectories (−) and initial positions (∘) are shown until the original collision time. The possible lateral avoidance trajectories are shown using rate 1 (−), rate 2 (−) and rate 3 (−) coordinated turns. Trajectories for both left (*solid*) and right (*dashed*) decisions are also shown.

Fig 6.12(a) shows a set of nine re-simulated collision encounters along with the possible avoidance trajectories that correspond to a maximum turn rate of 1 (3°/s), 2 (6°/s) and 3 (9°/s). Both avoidance decisions are shown such that six aircraft trajectories result. It can be seen that only one object trajectory may represent an infeasible avoidance encounter. This encounter is then isolated and depicted in Fig 6.12 along with the same six possible avoidance trajectories. Importantly, the encounter is only infeasible if the control is limited to a rate 1 turn, the avoidance decision is fixed to a left hand turn and $r_c = 0.5$. In this case, the corresponding minimum separation is less than 0.5m so the collision is unavoidable regardless of the controller. If a right hand turn is used instead, the encounter then becomes feasible. As the avoidance decision strategy allows for right and left hand turns, the encounter then remains in the feasible set, and is included in the results used to derived the system operating curves above. Of note, if an infeasible encounter had occurred and was retained, the resulting system operating curve is not incorrect. Instead, the curve then corresponds to a conservative representation of system performance, given infeasible encounters are actually possibly in a realistic See and Avoid environment.

As a final point, the proposed performance evaluation strategy can also be used to help identify the attributes of particular encounters which cause difficultly for the See and Avoid system. More precisely, the encounters that result in a collision can be analysed in more detail to help identify the set of collision geometries which cause the system to fail. This is done by re-simulating each collision encounter (missed avoidance and induced collision) up until the time of minimum separation, and inspecting the resulting set of object trajectories. Taking this approach, some example results are shown in Fig 6.13 for $\epsilon = 0.15$, $\lambda_\eta = 1$ and $r_c \in \{0.125, 0.5\}$ for 2000 encounters. The results clearly show that the avoidance strategy handles head-on type encounters very well, as only a single collision results from such geometry. Crossing encounters are more difficult to manage, with object trajectories originating near ±90 degrees causing the most difficulty. This can

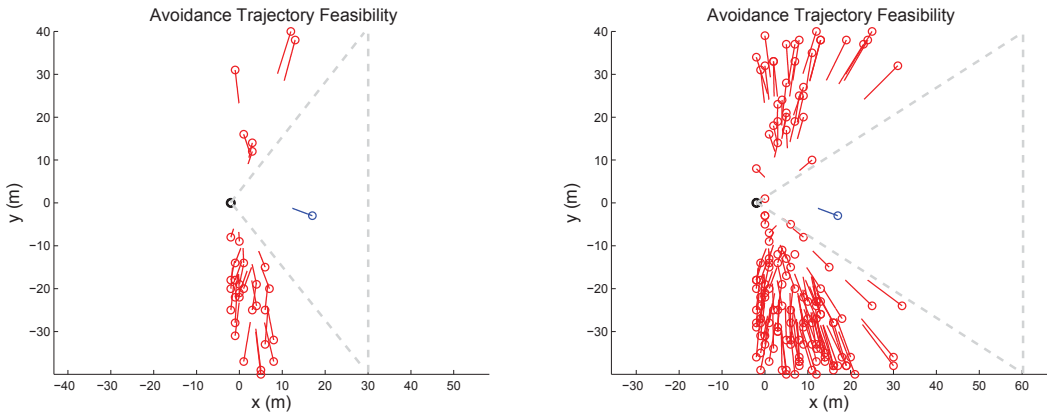

(a) $r_c = 0.125$ (b) $r_c = 0.5$

Figure 6.13: Example illustration of all collision trajectories for an encounters set of 2000. Object initial position (∘/∘) and trajectory for missed avoidance (−) and induced collision (−) are shown until the original collision time. The aircraft initial position (∘) without avoidance trajectory is also shown.

be inferred by considering that only the object trajectories originating very near to ±90 degrees are retained as the collision boundary decreases. Similar results are also obtained for different avoidance and resolution thresholds.

Remark In support of the preceding analysis, the resulting counts of each encounter outcome for some example sets of 2000 encounters with are given in table 6.1 and 6.2.

Encounter Type	r_c	**MA/IC**	**CACR/PACR**	**CAIR/PAIR**	**Total**
Collision	0.125	33	874	93	1000
Non-Collision	0.125	1	877	121	1000
Collision	0.5	154	755	91	1000
Non-Collision	0.5	1	877	121	1000

Table 6.1: Example Monte-Carlo simulation outcomes for $\lambda_\eta = 1$ and $\epsilon = 0.15$ for two different conflict boundaries $r_c \in \{0.125, 0.5\}$.

Encounter Type	r_c	**MA/IC**	**CACR/PACR**	**CAIR/PAIR**	**Total**
Collision	0.125	47	904	49	1000
Non-Collision	0.125	0	954	46	1000
Collision	0.5	204	747	49	1000
Non-Collision	0.5	0	954	46	1000

Table 6.2: Example Monte-Carlo simulation outcomes for $\lambda_\eta = 1$ and $\epsilon = 0.30$ for two different conflict boundaries $r_c \in \{0.125, 0.5\}$.

6.5 Summary

This chapter introduced a probabilistic analysis methodology for quantitatively evaluating the performance of reactive, vision-based collision avoidance systems. An additional set of collision avoidance outcomes were introduced to account for the reactive and completely autonomous nature of the proposed system. Leveraging System Operating Curves, the performance and design tradeoffs of the threshold-based avoidance and resolution decision strategies can then be visualised. A preliminary performance evaluation was then conducted using the avoidance decision strategy presented in Chapter 3, and the control and resolution decision strategy presented in Chapter 5. The key findings, benefits and drawbacks of using the evaluation framework and collision avoidance system include:

Performance Evaluation Framework

- The framework is flexible and extensible in the sense that it can be applied on a subsystem level, consistent with See and Avoid system partitioning, or combined to evaluate the closed loop system performance. For example, it can be used to evaluate avoidance and resolution decision strategies together or independently. Additionally, the general principles behind the framework can be extended to other automated systems (or subsystems) for avoidance and resolution. This is possible even if the decision thresholds or variables do not use the same sensor measurements.
- The framework is well aligned with existing evaluation techniques for object detection using Receiver Operating Curves. The framework therefore complements and extends the available tools for See and Avoid system performance evaluation.
- The framework provides a means to interpret, visualise and optimise the tradeoffs associated with the avoidance and resolution decision threshold placement. For a vision-based avoidance decision, the approach provides a novel mapping from the image space to the overall system performance attributes. As a consequence, and given the method in which the avoidance decisions are based, it could be extended to help evaluate the effectiveness of aviation right of way rules. This is possible if all system uncertainty is removed and the avoidance decision threshold is set to zero.
- Although the framework is intended to optimise system thresholds and parameters, it can also be used to analyse the sensitivity of the system to sensor (image feature) uncertainty, or unique encounter geometries. Each will influence the collision avoidance system performance, and alter the perceived utility of the system. To this end, it is acknowledged that the results presented in this work are preliminary, using a consistent and fixed set of auxiliary parameters and noise models. A real set of encounters would be required to guarantee the presented results are representative of the true system performance.

Actual System Performance

- The results suggest that the general approach to See and Avoid can be effective, even if the decision and resolution thresholds are not optimised. In most cases, the system is between 80-95% effective at avoiding collisions with a less than 10% chance of a negative outcome (collision). Of note, this is the first time a See and Avoid system has been analysed using such probabilistic approaches, consistent with aviation practise. This is an important step toward certification of such systems.
- System Operating Curves are leveraged in the design of the performance evaluation framework, and can be extended to include System Performance (SP) Curves [194]. Risk Ratio is represented as a function of specific relative velocities, system thresholds or other parameters. For the system presented here, and regardless of the threshold values, the risk ratio is consistently less than 0.5. This indicates a significant overall improvement in safety when using the collision avoidance system.
- The resolution threshold has significantly more influence on the overall system performance compared to the avoidance threshold. The result compliments the perception that it is difficult to discriminate between collision and non-collision objects using only visual observations, without range or intent information. Importantly, it does not mean that an effective avoidance strategy cannot be derived. The results presented here are indeed to the contrary, and suggest that always taking decisive action can be beneficial. The notion is also supported by trends in adaptive campaigning for military applications [362, 363].
- The avoidance and resolution decisions are not perfect, and a balance must be struck regarding the conservativeness of the associated threshold values. Interestingly, regardless of the threshold values, the system is shown to manage head-on encounters very well, with performance degrading as the object approaches from ± 90 degrees. The result has implications for See and Avoid system design concerning sensor choice and placement in the context of providing improved situational awareness.
- System performance is derived with respect to three collision boundaries. The results then provide an example of system performance, as it is dependent not only on initial encounter geometry, but the platform size, control limits (platform limitations) and choice of collision boundary. Of note, the notion of feasible avoidance should also be considered. In some cases, a conflict can be considered unavoidable due to difficult initial geometry and a short encounter duration. Such encounters could be removed before deriving the system operating curve. Of course, it is possible that such difficult encounters can actual occur, so including the results corresponds to a realistic and conservative representation of system performance.
- The performance metrics used to derive the system operating curves neglect the negative outcome concerning precautionary avoidance that remains unresolved. It was demonstrated that under optimised thresholds, encounters resulting in precautionary avoidance that remain unresolved were significantly less frequent (i.e. 46/1000) than precautionary avoidance that is resolved. Including the metrics in the negative outcomes will shift each curve to the right, whilst retaining similar avoidance performance.

Chapter 7

Conclusions

7.1 Research Summary

In order to realise the full commercial potential of unmanned aircraft, they must be granted access to civilian airspace. This means unmanned aircraft must demonstrate a range of safety critical tasks to a level that is deemed at least equivalent to that of manned aircraft. These include task such as forced landing, visual approach and more importantly, See and Avoid. This refers to a form of decentralized short term collision avoidance in which the pilot must visually detect and avoid any unplanned hazard, be it static or dynamic. Developing effective autonomous See and Avoid capability is considered as a major step toward unmanned aircraft integration into the national airspace.

The focus of this work was then an investigation into vision-based collision avoidance methods for practical aircraft See and Avoid applications. Assuming object detection and given the problem constraints, the research was directed toward successfully automating the decision and control components of a completely vision-based system, that does not leverage range or intent information. First, a novel avoidance decision strategy was proposed based on aviation rules of the air, properties of conical spiral motion and the expected uncertainty on image feature estimates. Second, a new spherical partitioned image-based control approach was presented and then extended using principles from nonlinear model predictive control. Further extensions included the first application of stability based design approaches using quasi-infinite horizon nonlinear model predictive control. Third, a simple set of resolution strategies were proposed, coupling the platform orientation and visual controllers to ensure the avoidance behaviour can be autonomously stopped. Lastly, the tradeoffs in design parameter selection and overall system performance was simultaneously visualised and assessed using a probabilistic approach. The evaluation technique was based on novel extension to existing approaches used in aviation to asses collision avoidance system performance.

Viewing the research elements in a collective manner, and considering the See and Avoid problem is currently unresolved, this work complements global efforts directed at solving this complex problem. A number of additional contributions were also made in the investigation, development and implementation of each of the system components. In some cases, the significant contributions pertain explicitly to the See and Avoid problem, whilst in other cases the advances constitute a more general scientific contribution.

7.2 Addressing Research Questions

This thesis posed a series of initial questions to help guide the research program. This section details how the results of the research investigation and associated contributions directly address each of these research questions (see §1.2.2 and §1.2.4).

Question 1 How can the state obtained from computer vision alone be used directly in decision and control for collision avoidance and resolution of both static and dynamic objects?

Considering the avoidance function, it was found that it is very difficult to estimate relative state (position, velocity, range etc.) for an arbitrary collision encounter using visual sensors in the Sense and Avoid environment. Highly accurate relative angular positions can however be obtained directly, at initial detection distances that leave between 10-15 seconds available for avoidance action (see §1.3). The useful state information is thus limited, which restricts many existing approaches used in aviation and robotics. However, by considering human navigation models, the notion of replicating pilot performance and the desire for predictable behaviour, a novel and effective avoidance decision strategy was derived (see §3). The approach relies on assessing visual cues (object divergence or convergence) to actively re-positioning the object on the image surface. The specific reference position is determined by considering the properties of spiral motion in light of existing aviation right-of-way rules (see §3.2). Importantly, the approach relies on tuning a single image-based decision threshold, that can be derived by considering the level of expected uncertainty on image feature measurements (see §3.3). As such, it can be extended to other camera systems, avoidance control solutions and platform types. Of note, the decision structure bears strong resemblance to existing aviation alerting system such as TCAS, yet does not rely on cooperative technology (sensors) or range estimation.

Considering the resolution function, it was found that many existing systems do not explicitly employ discrete resolution decisions. Instead, the notion of taking avoidance action is used to imply that the collision can also be resolved. This is a valid assumption for systems in which a human is in the loop or tactical design approaches such as path planning (see §1.3). The same cannot be said for a reactive vision-based system, so an explicit resolution decision strategy was required. The result was the development of a threshold-based resolution decision, that uses an augmented objective function to couple the platform heading and visual controller (see §5.5.2). Although achieving the threshold value cannot be guaranteed, tuning a single parameter reduces the development work required, as in the case for the avoidance decision strategy. Additionally, the approach remained decoupled from the avoidance decision, allowing the optimisation of the decision thresholds.

Using a series on Monte-Carlo simulations the approach to avoidance and resolution was shown to be over 90% effective at avoiding collision, with less than 10% of encounters resulting in negative outcomes (see §6.4.2 and §6.4.3). This is true for both liberal and conservative decision thresholds, suggesting the framework offers a potentially generic

approach to making avoidance and resolution decisions. The research outcomes are collectively described by **Contribution 4**, advancing technological solutions to the See and Avoid problem.

Question 2 How can visual-based control be used to safely re-position the unmanned aircraft (avoidance manuever) once an avoidance decision has been made?

Considering the avoidance control, it was found that despite the large number of proposed vision-based approaches, no existing approach was able to manage the See and Avoid problem constraints. Indeed, many approaches using See and Avoid as the target application often fall short of expectations, having not realistically considered the operational environment (see §1.3,§4.2 and §5.2). This may be in terms of available image features, range estimation, encounter scale or camera modelling. Research was then primarily directed toward the derivation and practical application of effective spherical image-based controllers, using only a single point feature for feedback control. Two general control approaches were used, including a classical partitioned image-based visual servoing scheme and a visual predictive control framework.

The classical approach provided a refined derivation and implementation of a partitioned spherical image-based control scheme, extending general image-based control concepts (see §4.3). The research constitutes the first successful practical implementation of spherical image-based control on an aerial platform. It is found that a simple controller was capable of satisfactory spiral tracking for collision avoidance, with modest tuning requirements. Adding a weighted integral term allowed the extension of the controller to dynamic objects, without explicitly considering the object motion in the control design. From a theoretical perspective, the effect of over and under estimating the reference range parameter required for the controller was also verified. Results suggested that if the parameter cannot be obtained, an overestimated value should be used for a conservative approach. In this case, the platform converges to the spiral quicker than if an underestimated value had been used. Indeed, the result might seem counter intuitive to many, as it is often assumed more aggressive control could be invoked by assuming a closer object. The result also suggests that accurate range estimation is not required for effective collision avoidance. The research outcome is described by **Contribution 1**, complementing both See and Avoid system development and general image-based control research.

The predictive approach provided a novel derivation and implementation of spherical visual predictive control, further extending general image-based control concepts (see §5.3). The research constitutes not only the first successful practical implementation of S-VPC, but the only existing practical implementation of image-based visual predictive control of any form for aerial platform (see §5.5). It was found that the predictive framework offered improved spiral tracking and collision avoidance control for both static and dynamic objects, compared to the classical image-based approach. Visibility, platform and state constraints can be considered explicitly in the control design, whilst remaining computationally tractable. Local stability can be assured and the framework showed satisfactory

robustness to added uncertainty, image processing delay and model mismatch. The research outcome is described by **Contribution 2**, complementing both See and Avoid system development and general image-based control research.

The visual predictive approach was then extended through the application of stability-based design concepts original derived for general model predictive control (see §5.4). The research constitutes the first application of Quasi-Infinite Horizon Nonlinear Model Predictive Control (QIH-NMPC) to visual predictive control (QIH-VPC) of any form. Apart from the inherent nominal stability guarantees, it was found that the control design implicitly enables better selection of reference image features for spiral tracking and collision avoidance, and decreases the control effort by 3% compared to visual predictive control. Although designed for the nominal (ideal) system, the new approach to visual control demonstrated satisfactory robustness to added uncertainty, image processing delay and model mismatch. Importantly, the calculation of the controller parameters (terminal region and terminal penalty term) was also presented in a manner amenable to fast and simple using modern computing tools (MATLAB). It could then be used for an arbitrary visual control application, provide the necessary conditions are met. The research outcome is described by **Contribution 3**, extending general image-based control research.

Question 3 What are the performance limitations of the proposed collision avoidance and resolution system, and what impact does this have on using such a concept as a general framework for designing automated vision-based systems?

Considering performance assessment, it was found that autonomous vision-based collision avoidance systems are typically analysed with respect to stability and feasibility, having assumed the correct avoidance decision has been made (§1.3). This is often accomplished using a set of simplifying assumption, such as known object intent or relative state, but is unrealistic for typical See and Avoid encounters. Given the system variability (sensor uncertainty, encounter geometries etc.), and considering the inclusion of two decision thresholds, probabilistic evaluation techniques are better suited to analyse the proposed system. System Operating Curves have been used to asses aircraft alerting system performance, so it then makes sense to extend such analysis tools to accommodate vision-based See and Avoid systems.

A novel probabilistic evaluation technique was then derived that explicitly consider the unique collision outcomes, and associated performance metrics, for a completely automated collision avoidance system (see §6.2 and §6.3). It was found that using the evaluation framework, the tradeoff in system performance and threshold parameter section could be visualised. This feature was used to interpret the effect of parameter placement on overall system performance, and subsequently optimise the decision thresholds. Although the predictive controller was used in each case, the framework could be used to visualise performance tradeoffs with respect to other control approaches, instead of decision threshold placement.

Using a set of Monte-Carlo trials, initial results suggest the resolution threshold has more effect on the system performance. The avoidance decision is less sensitive to the associated threshold placement, complimenting recent suggestions that it is difficult to discriminate between collision and non-collision object using only visual information. However, the avoidance strategy remained effective and suggested that always taking decisive action can be beneficial. Indeed, preliminary figures suggest that a tuned system is over 90% effective at avoiding collisions, producing a negative outcome for less than 10% of encounters. Considering a risk ratio of less than 0.5 is maintained in all cases, significant safety improvement when using the collision avoidance system can be guaranteed. The research outcomes are collectively described by **Contribution 5**, advancing technological solutions to the See and Avoid problem.

Having addressed each of the research questions, this work has presented

---

How a practical and effective automated vision-based decision and control strategy for collision avoidance and resolution can be derived, using only the available state information from monocular computer vision observations, that remains consistent with existing aviation practise and procedure in See and Avoid encounters.

7.3 Other Contributions

Although not detailed in this thesis, a simple image-based approach to spiral tracking and collision avoidance control was also derived for small fixed wing unnamed aircraft. The control structure allowed for lateral or vertical separation depending on the designers choice. The control commands were transformed to reference pitch and roll angles, having decoupled the image kinematics. Similar avoidance behaviour to that of the quadrotor model used in this work was observed, and helped to verify the flexibility of the approach to multiple platforms and large scale encounters. The results can be found in the authors publication list.

7.4 Considerations & Further Work

There are a few items that should be considered when applying or extending the proposed vision-based See and Avoid system framework. These include:

Practical Implications

- The avoidance and resolution decision strategies used in this thesis represent an effective avoidance approach, but not a perfect system. Having only considered simulation results, it is yet to be determined if the stated performance is indeed reflective of the system performance in a real environment. It is also hypothesised that further simulation studies using nominal (no uncertainty) conditions, could be used to quantitatively assess the effectiveness of existing aviation rules of the air.
- To provide a more comprehensive system analysis, it would be useful to determine the distribution of the initial image feature behaviour (position and velocity) associated with the set of collision encounters that cause difficultly for the avoidance decision strategy. To this end, the probabilistic analysis technique presented in this work could be further exploited, and the avoidance decision strategy refined.
- The See and Avoid system was implemented in a scaled environment using a quadrotor platform, in a proof-of-concept approach. Although novel and suitable for a preliminary analysis, complete validation of the utility of the approach would require a full scale implementation on other platforms. This could be done using optionally piloted aircraft, circumventing regulatory issues in the absence of flight approvals.
- Existing aircraft detection and tracking approaches would need to be modified to ensure the object can be tracked during the avoidance control, with minimal processing delay.

Quasi-Infinite Horizon Visual Predictive Control

- For circular motion, determining the nominal range value at which the terminal region shrinks to a point would provide insights into the global domain of operation. Considering the implications for spiral tracking and collision avoidance, it may then be possible to provide some added safety guarantees for at least a subset of encounters.
- For spiral motion, a nonlinear controller could be used to derive the terminal region and penalty matrix. The controller could derived using classical image-based approaches, similar to that derived in this work. It is expected a terminal region and not a terminal equality constraint may then result for such spirals. Of course, the nonlinear controller could also be used for circular motion, generalising the controller for an arbitrary spiral and offering a more flexibility framework.
- Care should be taken when applying the control framework to ensure the necessary conditions in which to derive suitable controller can be satisfied. To this end, it would be possible to extend the control approach beyond spiral tracking and to the common 4-point coplanar servoing problem if desired. A large number of visual control approaches are compared with respect to this problem, so the endeavour may yield some informative results. The same can be said for other camera models and image feature types.

Appendix A

Aviation Flight Rules

Excerpt from Civil Aviation Safety Authority (CASA) Civil Aviation Regulations (CAR) 1998 [29] regarding Rules-of-Air and Right-of-Way

161 Right of way

(a) *CAR 161 (1) - An aircraft that is required by the rules in this Division to keep out of the way of another aircraft shall avoid passing over or under the other, or crossing ahead of it, unless passing well clear.*

(b) *CAR 161 (2) - The pilot in command of an aircraft that has the right of way must maintain its heading and speed, but nothing in the rules in this Division shall relieve the pilot in command of an aircraft from the responsibility of taking such action as will best avert collision*

162 Rules for prevention of collision

(a) *CAR 162 (1) - When 2 aircraft are on converging headings at approximately the same height, the aircraft that has the other on its right shall give way (with additional caveats for airships, gliders and balloons*

(b) *CAR 162 (2) - When two aircraft are approaching head-on or approximately so and there is danger of collision, each shall alter its heading to the right*

(c) *CAR 162 (3) - An aircraft that is being overtaken has the right-of-way and the overtaking aircraft, whether climbing, descending, or in horizontal flight, shall keep out of the way of the other aircraft by altering its heading to the right, and no subsequent change in the relative positions of the two aircraft shall absolve the overtaking aircraft from this obligation until it is entirely past and clear.*

(d) *CAR 162 (4) - An overtaking aircraft shall not pass the aircraft that it is overtaking by diving or climbing*

163 Operating near other aircraft

(a) *CAR 163 (1) - The pilot in command of an aircraft must not fly the aircraft so close to another aircraft as to create a collision hazard.*

Appendix B

Coordinate Systems and Transformations

B.1 Local Geodetic Frame, Earth Fixed or World Frame

The world coordinate frame $\mathcal{F}_w$ is a fixed reference coordinate system with origin at a given point of the Earth ellipsoid. For a North-East-Down (NED) representation, the x-axis points North and the y-axis points East. The z-axis is normal to the Earth's surface pointing downwards and assumed to be aligned with the gravity vector.

B.2 Body Fixed Frame

The Body Fixed Frame coordinate frame $\mathcal{F}_b$ is a reference coordinate system with origin at the centre of the mass of the body/vehicle to which it is associated with. The x-axis is aligned with the longitudinal axis of the vehicle, pointing forwards. The z-axis is aligned with the normal axis of the vehicle, pointing downwards. The y-axis completes the right-handed system. Although rotary and fixed wing aircraft differ significantly, the x-axis can be used to denote the nominal direction of travel or forward velocity. A variable rotation and transform will exist between the body and world frame as the vehicle rotates and translates (moves).

B.3 Camera Fixed Frame

The Camera Fixed Frame coordinate frame $\mathcal{F}_c$ is a reference coordinate system attached to a camera. For a perfectly calibrated camera, the origin of the frame coincides with the camera focal point. The orientation of the camera frame is somewhat arbitrary and may depend on the camera type (or lens). For the spherical camera used in this work, the x-axis is the optical axis. The z-axis points downwards and is tangential to the imaging surface at the optical axis. The y-axis completes the right-handed system. A fixed rotation and transform will exist between the camera and body frame to which it is attached (eye-in-hand system).

Angular measurements taken with respect to a camera frame $\mathcal{F}_c$ (colatitude and azimuth) and body $\mathcal{F}_b$ (bearing and elevation) can be considered equivalent, if the observations are of an object positioned at least two orders of magnitude greater than the relative displacement between the body and camera frame origins. Consider a body fixed frame origin $\mathcal{F}_b^o$ positioned at the point $(0, 0, 0)$. Consider now a camera frame origin $\mathcal{F}_c^o$ normally distributed about the body frame origin such that $\mathcal{F}_c^o \sim \mathcal{N}(0, 0.05)$. The camera inherits the

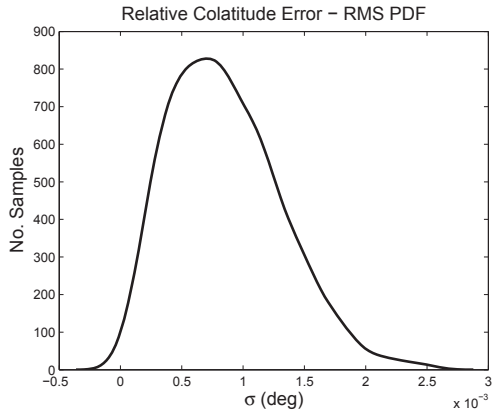

(a) Probability Distribution

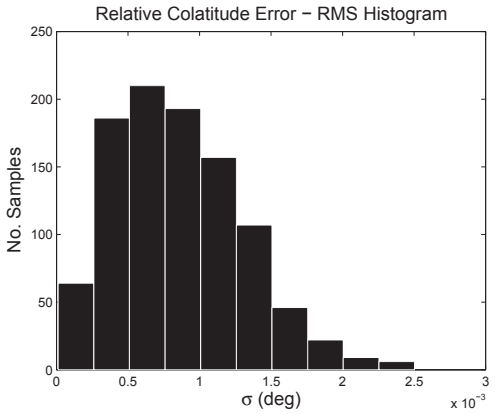

(b) Histogram

Figure B.1: Expected errors in colatitude angle under the assumption that $r \gg \left\| {}^{b}\mathbf{t}_c \right\|$

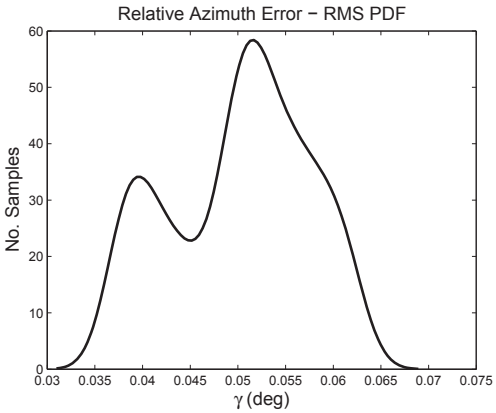

(a) Probability Distribution

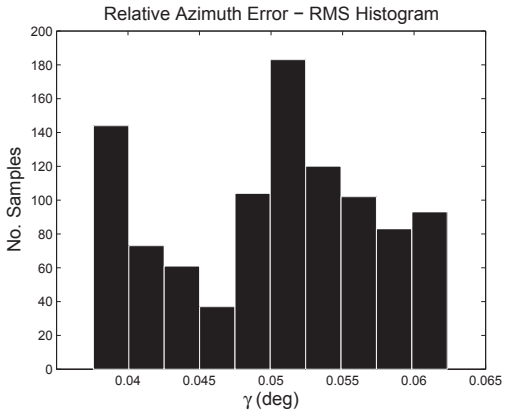

(b) Histogram

Figure B.2: Expected errors in azimuth angle under the assumption that $r \gg \left\| {}^{b}\mathbf{t}_c \right\|$

same orientation as the body frame and observes an object $\mathcal{O}$ at a distance of 10m. The camera and body frames are then successively rotated by 5 deg increments in clockwise direction in both pitch θ and yaw ψ. The results from a Monte-Carlo simulation used to estimate the distribution of the root-mean-square error of the angular measurements and range are given below.

B.4 Rotation Matrices

A rotation matrix describes the relationship between the orientation of one coordinate frame to another. The rotation can be parametrised using Euler angles ϕ, θ and ψ or constructed using quaternions. Using Euler angles, three successive rotations $R_z(\psi), R_y(\theta)$ and $R_x(\phi)$ from an arbitrary frame $\mathcal{F}_a$ to another frame $\mathcal{F}_b$ are performed. The rotation sequence is dented as a Z-Y-X sequence, due to the order of intrinsic rotations and the associated rotation angles are commonly known as yaw, pitch and roll respectively. The vector $\mathbf{x}_a$ defined in $\mathcal{F}_a$ can be defined in $\mathcal{F}_b$ such that

$$\mathbf{x}_b = {}^{a}R_b(\phi, \theta, \psi)\mathbf{x}_a \tag{B.1}$$

$$\mathbf{x}_b = R_x(\phi)R_y(\theta)R_z(\psi)\mathbf{x}_a \tag{B.2}$$

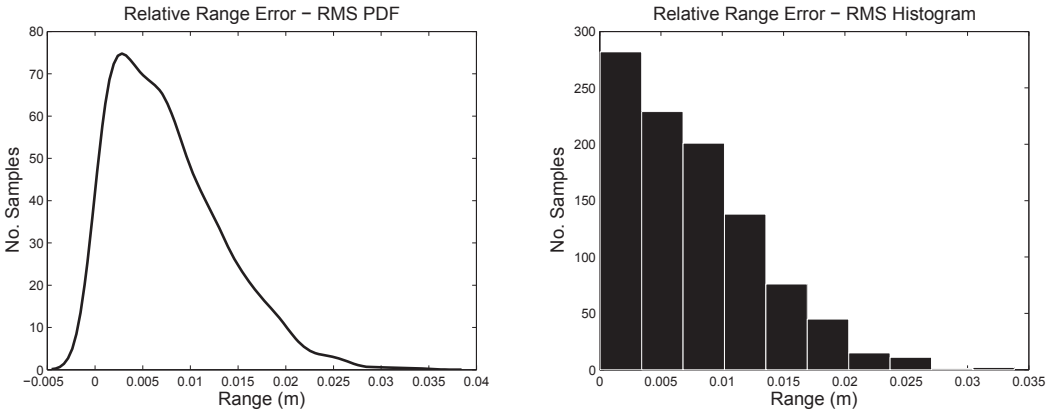

(a) Probability Distribution (b) Histogram

Figure B.3: Expected error in range under the assumption that $r \gg \left\|{}^{b}\mathbf{t}_{c}\right\|$

where

$$R_x(\phi) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\phi & \sin\phi \\ 0 & -\sin\phi & \cos\phi \end{pmatrix}, \; R_y(\theta) = \begin{pmatrix} \cos\theta & 0 & -\sin\theta \\ 0 & 1 & 0 \\ \sin\theta & 0 & \cos\theta \end{pmatrix}, \; R_z(\psi) = \begin{pmatrix} \cos\psi & \sin\psi & 0 \\ -\sin\psi & \cos\psi & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{(B.3)}$$

and

$${}^{a}R_b(\phi,\theta,\psi) = \begin{pmatrix} \cos\theta\cos\psi & \sin\phi\sin\theta\cos\psi - \cos\phi\sin\psi & \cos\phi\sin\theta\cos\psi + \sin\phi\sin\psi \\ \cos\theta\sin\psi & \sin\phi\sin\theta\sin\psi + \cos\phi\cos\psi & \cos\phi\sin\theta\sin\psi - \sin\phi\cos\psi \\ -sin\theta & \sin\phi\cos\theta & \cos\phi\cos\theta \end{pmatrix} \quad \text{(B.4)}$$

Similarly, by taking the matrix inverse, the vector $\mathbf{x}_b$ defined in $\mathcal{F}_b$ can be defined in $\mathcal{F}_a$ such that

$$\mathbf{x}_a \;=\; {}^{a}R_b(\phi,\theta,\psi)^{-1}\mathbf{x}_b \quad \text{(B.5)}$$

B.5 Transformation Matrices

A transformation matrix describes the relationship between the position of the origin of one coordinate frame to another. The transformation can be parametrised using simple vector addition/difference ${}^{a}\mathbf{t}_b$ or formulated into a diagonal matrix ${}^{a}T_b$ made up of this vector. Usually, to completely represent a vector in an arbitrary frame $\mathcal{F}_b$ and initial stated with respect to another arbitrary frame $\mathcal{F}_a$ both rotation and transformation is required such that

$$\mathbf{x}_b \;=\; {}^{a}R_b(\phi,\theta,\psi)\mathbf{x}_a + {}^{a}T_b \quad \text{(B.6)}$$

where ${}^{a}T_b = (T_x \; T_y \; T_z)^T$ in Cartesian notation.

Appendix C

Spirals

C.1 Optimal Spirals

This Appendix describes the setup of an optimal bang-bang control problem for the avoidance of static objects in two dimensions based on [220]. The aircraft is considered to move at a fixed velocity in the xy-plane, and permitted to turn left or right via a change in heading (angular velocity). Only two possible control actions are considered such that a single control action u is limited to the control constraint domain $\mathbb{U} \in \{-1, 1\}$. These actions can be considered as a proxy for left and right avoidance manuevers, and not necessarily the only possible choices.

First, consider the aircraft state $x(t)$ defined by

$$x(t) = [r(t)\ \alpha(t)\ \psi(t)]^T$$

where $r(t)$ is the range to the object, $\psi(t)$ is the aircraft heading and $\alpha(t)$ is the relative angular displacement to the object. The equations of motion for a static object encounter with an aircraft moving with constant velocity v are given by $\dot{x}(t) = f(x(t), u(t)$ where

$$\dot{r}(t) = -v\cos(\alpha(t)) \tag{C.1}$$

$$\dot{\alpha}(t) = v\sin(\alpha(t))/r - u(t) \tag{C.2}$$

$$\dot{\psi}(t) = -u(t) \tag{C.3}$$

Second, consider an arbitrary encounter where $r(0) \gg 0$ and $T > 0$ defines a maneuver or terminal time considered to be the time of closest approach. The objective of the optimal control problem is to then maximize the miss distance, or ensure $r(T)$ is a large as possible, using only the available controls defined by $\mathbb{U}$. To this end, consider a cost function J to be maximized such that

$$J(x(t), u(t), t) = r(T) \tag{C.4}$$

$$J^\star(x(t), u(t), t) = \max_{u\in\mathbb{U}}(r(T)) \tag{C.5}$$

The problem is a Mayer problem with free terminal end points, as T is unknown a priori. To solve the optimal control problem (C.5), consider the conditions for optimality and define the Hamiltonian

$$H(\lambda(t), x(t), u(t), t) = \lambda^T \cdot f(x(t), u(t), t) + \mathcal{L}(x(t),\ u(t)) \tag{C.6}$$
$$H(\lambda(t), x(t), u(t), t) = -\lambda_r v\cos(\alpha) + \lambda_\alpha(v\sin(\alpha)/r - u) - \lambda_\psi u \tag{C.7}$$

where $\lambda = [\lambda_r \lambda_\alpha \lambda_\psi]$ defines the Lagrangian multipliers (adjoint variables) and $\mathcal{L}(x(t), u(t)) = 0$ due to the structure of the cost function. The Lagrangian multipliers satisfy the equations

$$\dot{\lambda} = -\nabla H \tag{C.8}$$
$$\dot{\lambda} = -\left(\frac{\partial H}{\partial r}, \frac{\partial H}{\partial \alpha}, \frac{\partial H}{\partial \psi}\right) \tag{C.9}$$
$$\dot{\lambda} = \begin{pmatrix} \lambda_\alpha v\cos(\alpha)/r^2 \\ -\lambda_r v\sin(\alpha) - \lambda_\alpha v\cos(\alpha)/r \\ 0 \end{pmatrix} \tag{C.10}$$

with boundary conditions

$$\lambda(T) = \nabla J|_{t=T} = \left(\frac{\partial J}{\partial r}\ \frac{\partial J}{\partial \alpha}\ \frac{\partial J}{\partial \psi}\right)|_{t=T} = [1\ 0\ 0]^T \tag{C.11}$$

The transversatility condition is given by

$$\dot{J}\,|_{t=T} = \frac{\partial J}{\partial t}\,|_{t=T} + \nabla J \cdot \dot{x}\,|_{t=T} = 0 \tag{C.12}$$
$$\dot{J}\,|_{t=T} = [1\ 0\ 0] \cdot [\dot{r}\ \dot{\alpha}\ \dot{\psi}]^T|_{t=T} = 0 \tag{C.13}$$
$$\dot{J}\,|_{t=T} = \dot{r}|_{t=T} = 0 \tag{C.14}$$
$$\tag{C.15}$$

which implies that

$$\dot{r}(T) = 0 \tag{C.16}$$
$$-v\cos(\alpha(T)) = 0 \tag{C.17}$$
$$\cos(\alpha(T)) = 0 \tag{C.18}$$

Therefore, if the optimal control $u^\star$ that maximizes the minimum miss distance is followed, the terminal range condition means $\alpha(T) = \pm\pi/2$. This suggests that the optimal spiral to establish for collision avoidance is in fact a left or right circle about the object.

Continuing the analysis further, the optimal control in the vicinity of the terminal time can be found using Pontryagins maximum (minimum) principle. For the optimal control $u^\star$ to be a solution to the Mayer problem (C.5) subject to (C.1)-(C.3), it is necessary that there exists a non-zero continuous function $\lambda(t)$ that satisfies (C.10)-(C.11) so that for every $t \in [0, T]$

$$\begin{aligned}
H(\lambda(t), x(t), u^\star(t), t) &= \max_{\mathbb{U}}(H(\lambda(t), x(t), u(t), t)) && \text{(C.19)}\\
H(\lambda(t), x(t), u^\star(t), t) &= \max_{u\in[-1,1]}(-\lambda_r v\cos(\alpha) + \lambda_\alpha(v\sin(\alpha)/r - u) - \lambda_\psi u) && \text{(C.20)}\\
H(\lambda(t), x(t), u^\star(t), t) &= \max_{u\in[-1,1]}(u(-\lambda_\alpha - \lambda_\psi)) && \text{(C.21)}\\
&- \lambda_r v\cos(\alpha) + \lambda_\alpha v\sin(\alpha)/r && \text{(C.22)}
\end{aligned}$$

The non-singular control that maximises the Hamiltonian is therfore given by

$$u = -\operatorname{sign}(\lambda_\alpha + \lambda_\psi) \tag{C.23}$$

Introducing the retrograde or backward time $\tau = T - t$, and considering that the Lagrangian multipliers vary continuously and are subject to the terminal condition $\lambda(T) = [1\,0\,0]^T$, there exists $\delta > 0$ in the vicinity of the terminal time for $\tau \in (0, \delta)$ such that

$$\operatorname{sign}(\lambda_\alpha(\tau) + \lambda_\psi(\tau)) = \operatorname{sign}(\mathring{\lambda}_\alpha(0) + \mathring{\lambda}_\psi(0)) \tag{C.24}$$

where the retrograde time is given by the open circles. Writing the control in terms of the Lagrangian multipliers

$$\begin{aligned}
u(\tau) &= -\operatorname{sign}(\mathring{\lambda}_\alpha(0) + \mathring{\lambda}_\psi(0)) && \text{(C.25)}\\
u(\tau) &= -\operatorname{sign}(\mathring{\lambda}_\alpha(0)) && \text{(C.26)}\\
u(\tau) &= -\operatorname{sign}(\lambda_r v\sin\alpha(0)) && \text{(C.27)}\\
u(\tau) &= -\operatorname{sign}(\sin\alpha(0)) && \text{(C.28)}
\end{aligned}$$

Finally, substituting the optimal values for $\alpha(T)$ and recalling that $\alpha(0)$ in retrograde time corresponds to $\alpha(T)$ in normal time

$$\begin{aligned}
u^\star(\tau) &= -1 \qquad \alpha(T) = \pi/2 && \text{(C.29)}\\
u^\star(\tau) &= 1 \qquad \alpha(T) = -\pi/2 && \text{(C.30)}
\end{aligned}$$

As the control is applied for $0 \leq t \leq T$, the above result associates the required control direction with the corresponding optimal spiral direction (inferred by $\alpha(T)$). If the object persists to the right of the aircraft, the aircraft must be turning to the left and vice versa. As switching the control for $t \in [0, T]$ is not permissible, the object must have been initially positioned to the right of the aircraft as well. This can be inferred using geometric intuition.

Remarks

- The magnitude of the control value is considered to be less important than its sign (direction) in the above analysis, and could be altered to a different scale with the same result.
- The above analysis can be pursued further to obtain the set of initial ranges and angular displacements for specific control constraints.
- Applying the same approach to determine optimal conical spiral angles for dynamic objects is not possible if the object intent (heading) is unknown. If the object heading and speed is known, numerical methods can be used to solve the problem however [221]. Given the See and Avoid problem constraints, the analysed is not pursued.

Appendix D

Publications

D.1 Journals

IEEE Trans. Robotics, vol. 30, no. 6, pp. 1441-1454, 2014

Visual Predictive Control of Spiral Trajectories,

A. Mcfadyen, P. Corke and L. Mejias

This paper deals with constrained image-based visual servoing (IBVS) of circular and conical spiral motion about an unknown object approximating a single image point feature. Effective visual control of such trajectories has many applications for small unmanned aerial vehicles including surveillance and inspection, forced landing (homing) and collision avoidance. A spherical camera model is used to derive a novel visual predictive controller (VPC) using stability-based design methods for general nonlinear model predictive control. In particular, a quasi-infinite horizon visual predictive control (QIH-VPC) scheme is derived. A terminal region, used as a constraint in the controller structure, can be used to guide appropriate reference image features for spiral tracking with respect to nominal stability and feasibility. Robustness properties are also discussed with respect to parameter uncertainty and additive noise. A comparison to competing visual predictive control schemes is made and some experimental results using a small quad rotor platform are given.

D.2 Conference Papers

Proc. IEEE/RSJ Int. Conf. Robotics and Intelligent Systems (IROS'13), pp. 50-56, Nov. 2013

Aircraft collision avoidance using spherical visual predictive control and single point features

A. Mcfadyen L. Mejias P. Corke and C. Pradalier

This paper presents practical vision-based collision avoidance for objects approximating a single point feature. Using a spherical camera model, a visual predictive control scheme guides the aircraft around the object along a conical spiral trajectory. Visibility, state and control constraints are considered explicitly in the controller design by combining image and vehicle dynamics in the process model, and solving the nonlinear optimization problem over the resulting state space. Importantly, range is not required. Instead, the principles of conical spiral motion are used to design an objective function that simultaneously guides the aircraft along the avoidance trajectory, whilst providing an indication of the appropriate point to stop the spiral behaviour. Our approach is aimed at providing a potential solution to the *See and Avoid* problem for unmanned aircraft and is demonstrated through a series of experimental results using a small quadrotor platform.

Proc. IEEE/RSJ Int. Conf. Robotics and Intelligent Systems (IROS'12), pp. 1199-1205, Oct. 2012

Rotorcraft collision avoidance using spherical image-based visual servoing and single point features

A. Mcfadyen P. Corke and L. Mejias

This paper presents a reactive collision avoidance method for small unmanned rotorcraft using spherical image-based visual servoing. Only a single point feature is used to guide the aircraft in a safe spiral like trajectory around the target, whilst a spherical camera model ensures the target always remains visible. A decision strategy to stop the avoidance control is derived based on the properties of spiral like motion, and the effect of accurate range measurements on the control scheme is discussed. We show that using a poor range estimate does not significantly degrade the collision avoidance performance, thus relaxing the need for accurate range measurements. We present simulated and experimental results using a small quad rotor to validate the approach.

Proc. Int. Conf. Unmanned Aircraft Systems (ICUAS'14), pp. 715-725, May 2014

Decision strategies for automated visual collision avoidance

A. Mcfadyen A. Durand-Petiteville and L. Mejias

This paper provides a preliminary analysis of an autonomous uncooperative collision avoidance strategy for unmanned aircraft using image-based visual control. Assuming target detection, the approach consists of three parts. First, a novel decision strategy is used to determine appropriate reference image features to track for safe avoidance. This is achieved by considering the current rules of the air (regulations), the properties of spiral motion and the expected visual tracking errors. Second, a spherical visual predictive control (VPC) scheme is used to guide the aircraft along a safe spiral-like trajectory about the object. Lastly, a stopping decision based on thresholding a cost function is used to determine when to stop the avoidance behaviour. The approach does not require estimation of range or time to collision, and instead relies on tuning two mutually exclusive decision thresholds to ensure satisfactory performance.

Proc. 28th Int. Congress of the Aeronautical Sciences (ICAS'12), Sep. 2012

Visual servoing approach to collision avoidance for aircraft

A. Mcfadyen, L. Mejias and P. Corke

This paper presents a reactive Sense and Avoid approach using spherical image-based visual servoing. Avoidance of point targets in the lateral or vertical plane is achieved without requiring an estimate of range. Simulated results for static and dynamic targets are provided using a realistic model of a small fixed wing unmanned aircraft.

D.3 Technical Reports

Scoping Study for Remotely Piloted Aircraft Systems Integration into Civil Airspace

A. Mcfadyen, R. Clothier, D. Campbell and G. William

CRC5055 Evaluating Unmanned Aircraft Systems for Deployment in Plant Biosecurity

A. Mcfadyen, F. Gonzalez, D. Campbell and D. Eagling

Appendix E

Parameters

E.1 Variables & Constants

Variable	Description
$\mathbf{p}$	Image feature point measured in camera frame
$\mathbf{s}$	Spherical image feature vector consisting of σ and γ angles
$\check{\mathbf{s}}$	Image feature convergence/divergence vector
$\mathbf{e}$	Spherical image feature error vector consisting of e_σ and e_γ
e_σ	Colatitude image feature error
e_γ	Azimuth image feature error
σ	Colatitude angle as observed from a spherical camera
γ	Azimuth angle as observed from spherical camera
$\check{\sigma}$	Image feature convergence with respect to colatitude angle.
$\check{\gamma}$	Image feature convergence with respect to azimuth angle.
$\mathbf{L}_s$	Full spherical image Jacobian/Interaction matrix
$\mathbf{L}_z$	Spherical image Jacobian/Interaction matrix for z axis components
$\mathbf{L}_{xy}$	Spherical image Jacobian/Interaction matrix for x and y axis components
$\mathbf{c}$	Conical angle vector consisting of β and α angles
β	Elevation angle as observed from a body frame
α	Bearing angle as observed from body frame
$\mathbf{u}_1$	Control vector for x velocity control using LQR/PID control
$\mathbf{u}_2$	Control vector for y velocity control using LQR/PID control
$\mathbf{u}_3$	Control vector for z axis using visual control
$\mathbf{u}$	Full control vector consisting of all inputs, $\mathbf{u}_1$, $\mathbf{u}_2$ and $\mathbf{u}_3$
u	Arbitrary single control
$\mathbf{a}$	Lower control limits for $\mathbf{u}_3$
$\mathbf{b}$	Upper control limits for $\mathbf{u}_3$
$\mathbf{x}_1$, $\mathbf{x}_x$	State vectors for x velocity control
$\mathbf{x}_2$, $\mathbf{x}_y$	State vectors for y velocity control
$\mathbf{x}_3$, $\mathbf{x}_z$	State vectors for z axis using visual controllers
$\mathbf{x}$	Full state vector consisting of all states, $\mathbf{x}_1$, $\mathbf{x}_2$ and $\mathbf{x}_3$
$\mathbf{A}$,$\mathbf{B}$,$\mathbf{C}$,$\mathbf{D}$	Linearised system state space matrices for LQR/PID control
$\mathbf{A_s}$,$\mathbf{B}_s$	System state space matrices for visual predictive control
U_n	Quadrotor motor command for motor n
Υ	Quadrotor low-level control vector

v_t	Object speed
Φ	Object inertial heading
$\bar{\alpha}$	Object relative bearing
$\bar{\beta}$	Object relative elevation
$\bar{\psi}$	Platform/Camera relative heading
r	Relative aircraft to object range
r_m	Quadrotor motor thrust force displacement
r_p	Quadrotor motor propeller radius
v_x	Platform/Camera forward velocity
v_y	Platform/Camera lateral velocity
v_z	Platform/Camera vertical velocity
ω_x	Platform/Camera roll rate
ω_y	Platform/Camera pitch rate
ω_z	Platform/Camera yaw rate
ϕ	Platform/Camera roll angle rate
θ	Platform/Camera pitch angle
ψ	Platform/Camera yaw angle
$\mathbf{F}$	Translational force
F_T	Thrust force
F_D	Drag force
$\mathbf{J}$	Inertia matrix
J_{nm}	Inertia matrix components about axis n and m
J	Objective function value with only image feature states included for visual predictive control.
J_s	Objective function value with only image feature states weighted for visual predictive control
J_ψ	Augmented objective function value with image feature and yaw states for visual predictive control
τ	Torque / independent time derivative variable
a_ϕ	Roll dynamic constant (1^{st} order approximation)
a_θ	Pitch dynamic constant (1^{st} order approximation)
a_ψ	Yaw rate dynamic constant (1^{st} order approximation)
δt_ϕ	Roll dynamic time delay
δt_θ	Pitch dynamic time delay
δt_ψ	Yaw rate dynamic time delay
δt_{F_T}	Roll dynamic time delay
t	Time
t_d	Collision avoidance detection and avoidance instance.
t_s	Collision avoidance resolution instance.
T_i	Image processing sampling time
T_s	Simulation sampling time
T_p	Prediction horizon for visual predictive controllers
T_c	Control horizon for visual predictive controllers
T_p	Prediction horizon for visual predictive controllers
T_v	Controller sampling time

λ	Partitioned spherical visual controller gain
λ_i	Partitioned spherical visual controller integral gain
λ_ψ	Augmented objective function weighting for visual predictive control
λ_R	Control scale factor for visual predictive control
λ_P	Terminal state scale factor for visual predictive control
λ_η	Avoidance decision threshold scaling factor
Q_s	State weighting matrix element for the image feature states for visual predictive control
$\mathbf{Q}$	State weighting matrix for visual predictive control / LQR controller and Kalman filter weighting matrix
$\mathbf{R}$	Control weighting matrix for visual predictive control / LQR controller and Kalman filter weighting matrix
$\mathbf{P}$	Terminal state weighting matrix for visual predictive control
k_p	Proportional control gain
k_i	Integral control gain
k_d	Derivative control gain
$\mathbf{K}$	Linear Quadratic Regulator / Kalman filter gain
$\mathbf{K}^e$	Linear Quadratic Regulator Integral gain
$\mathbf{K}_v$	Linear controller gain (LQR type) for QIH-VPC control design
$\mathbf{K}_{\text{F}}$	Kalman filter gain matrix
η	Avoidance decision threshold
ϵ	Avoidance decision threshold
κ	Scaling factor for closed loop poles of the linear feedback controller used in the derivation of the terminal penalty matrix for QIH-VPC
ς	Terminal region constant for QIH-VPC
Ω_ς	Terminal region for QIH-VPC (image region)
Ω_η	Approximate avoidance decision performance region
Ω_n	Quadrotor motor speed for motor n
M	Approximate domain of attraction associated (image region)
ξ_s,$\xi_{\dot{s}}$	Variance on $\mathbf{s}$ and $\dot{\mathbf{s}}$ measurements
ξ_σ,ξ_γ	Variance on σ and γ measurements
$\xi_{\dot{\sigma}}$,$\xi_{\dot{\gamma}}$	Variance on $\dot{\sigma}$ and $\dot{\gamma}$ measurements
q	Additive noise to image feature measurements
$\mathbf{q}_c$	Additive noise to controls/actuators
$\mathbf{w}_g$	Additive noise to platform orientation (attitude) to model turbulence
$\mathbf{w}_a$	Additive noise to platform translational velocity to model ambient wind
k,n	Indices
${}^aR_b$	Rotation matrix from frame b to frame a
${}^aT_b$	Transformation matrix from frame b to frame a
${}^a\mathbf{t}_b$	Relative vector from from frame b to frame a
m	Quadrotor mass
g	Gravitational constant
$\mathcal{F}$	Coordinate frame

Table E.1: Complete list of variables and constants including notation and description

E.2 System Parameters - Classic Visual Control

Parameter	Value	Parameter	Value
τ_R	155	$\mathcal{R}_m$	255
τ_G	70	$\mathcal{G}_m$	255
τ_B	43	$\mathcal{B}_m$	255
$\mathcal{A}_{\min}$	45pix	$\mathcal{A}_{\max}$	2500pix
λ	0.2	v_x	0.1
$\mathbf{a}$	$(-0.5, -0.35)$	$\mathbf{b}$	$(0.5, 0.35)$
T_s	0.1	-	-

Table E.2: Common system, visual controller and object detection (blob detection) parameters for custom *Parrot ARDrone*

Controller	k_p	k_i	k_d
v_x	0.25	0.25	0.01
v_y	0.25	0.25	0.01

Table E.3: Proportional-Integral-Derivative (PID) controller parameters for forward v_x and lateral v_y velocity control of custom *Parrot ARDrone* [287]

E.3 System Parameters - Visual Predictive Control

Parameter	Simulated - Circular Motion	Simulated - Spiral Motion
m (kg)	0.67	0.67
T_p	15	15
T_s (s)	0.10	0.10
r^* (m)	2.00	2.00
v_x^* (ms^{-1})	0.10	0.10
v_z^* (ms^{-1})	0.00	-
ω_z^* (degs^{-1})	2.87	-
κ	0.2635	-
ς	9.3937	-
$\mathbf{Q}$	$0.5\mathbf{I}_2$	$\mathbf{I}_2$
$\mathbf{R}$	$\mathbf{I}_2$	$\mathbf{0}_2$
$\mathbf{P}$	$\begin{pmatrix} 17.4134 & 2.1785 \\ 2.1785 & 1.4304 \end{pmatrix}$	$2\mathbf{I}_2$
$\mathbf{K}_v$	$\begin{pmatrix} -0.7057 & 0.0079 \\ -0.0192 & 0.7078 \end{pmatrix}$	-
$q(t)$ (rad)	$\mathcal{N}(0, 0.02^2)$	$\mathcal{N}(0, 0.02^2)$

Table E.4: Visual predictive controller parameters for simulated implementations

Parameter	Experimental - Spiral	Experimental - Collision
m (kg)	0.67	0.67
T_p	10	10
T_s (s)	0.04	0.04
r^* (m)	2.00	2.00
v_x^* (ms^{-1})	0.20	0.20
Q_s	1.00	1.00
λ_ψ	0.10	0.10
λ_R	1×10^{-8}	1×10^{-8}
λ_P	2.00	2.00
ϵ	-	$\{0.6, 0.8, 1.8, 1.9\}$
$q(t)$ (rad)	-	$\mathcal{N}(0, 0.02^2)$

Table E.5: Visual predictive controller parameters for experimental implementations

Parameter	Nominal (*)	Min	Max
v_z (ms^{-1})	0.00	-0.50	0.50
ω_z (degs^{-1})	2.87	-10.00	10.00

Table E.6: Visual predictive control constraints for simulated implementation

Parameter	Nominal (*)	Min	Max
v_z (ms^{-1})	0.00	-0.50	0.50
ω_z (degs^{-1})	2.87	-10.00	10.00

Table E.7: Visual predictive control constraints for experimental implementation

Controller	K	$\mathbf{K^e}$
v_x	$\begin{pmatrix} -0.5021 & 0.0157 & 0.0087 \end{pmatrix}$	0.3331
v_y	$\begin{pmatrix} 0.6314 & 0.0287 & 0.0143 \end{pmatrix}$	-0.4729

Table E.8: Linear Quadratic Regulator (LQR) with integral action (LQRI) controller gains for forward v_x and lateral v_y velocity control of custom *Ascending Technologies (AscTec) Hummingbird*[269]

Parameter	Value	Parameter	Value
a_θ	-10.1041	δt_θ	0.0602
a_ϕ	10.5129	δt_ϕ	0.0529
a_ψ	-5.8997	δt_ϕ	0.1107
-	-	δt_{F_T}	0.0410

Table E.9: Linearised x and y axis system dynamic constants custom *Ascending Technologies (AscTec) Hummingbird*

State	A	B	C	D
$\begin{pmatrix}\dot{x}\\ \theta\\ \dot{\theta}\end{pmatrix}$	$\begin{pmatrix}1 & -0.1054 & -0.0011\\ 0 & 0.3018 & 0.0107\\ 0 & -5.2759 & 0.9434\end{pmatrix}$	$\begin{pmatrix}-0.0036\\ 0.0566\\ 0.6560\end{pmatrix}$	$\begin{pmatrix}1 & 0 & 0\\ 0 & 1 & 0\end{pmatrix}$	$\begin{pmatrix}0\\ 0\end{pmatrix}$
$\begin{pmatrix}\dot{y}\\ \phi\\ \dot{\phi}\end{pmatrix}$	$\begin{pmatrix}1 & 0.1097 & 0.0012\\ 0 & 0.3365 & 0.0112\\ 0 & -4.6866 & 0.9504\end{pmatrix}$	$\begin{pmatrix}0.0032\\ 0.0496\\ 7.4097\end{pmatrix}$	$\begin{pmatrix}1 & 0 & 0\\ 0 & 1 & 0\end{pmatrix}$	$\begin{pmatrix}0\\ 0\end{pmatrix}$
$\begin{pmatrix}\psi\\ \dot{\psi}\end{pmatrix}$	$\begin{pmatrix}1 & 0.0107\\ 0 & 0.8993\end{pmatrix}$	$\begin{pmatrix}9.22\times10^{-4}\\ 0.1007\end{pmatrix}$	$\begin{pmatrix}1 & 0\\ 0 & 1\end{pmatrix}$	$\begin{pmatrix}0\\ 0\end{pmatrix}$

Table E.10: Linearised x and y axis state space system matrices. The systems are used for the Kalman filters and Linear Quadratic Regulator with integral action (LQRI) controllers for forward v_x and lateral v_y velocity control of the custom *Ascending Technologies (AscTec) Hummingbird*

State	Q	R
$\begin{pmatrix}\dot{x}\\ \theta\\ \dot{\theta}\end{pmatrix}$	$\begin{pmatrix}0.0011 & -0.0063 & -0.0275\\ -0.0063 & 0.0805 & 0.8369\\ -0.0275 & 0.8369 & 174.9295\end{pmatrix}$	$\begin{pmatrix}0.0056 & 0\\ 0 & 5.556\times10^{-4}\end{pmatrix}$
$\begin{pmatrix}\dot{y}\\ \phi\\ \dot{\phi}\end{pmatrix}$	$\begin{pmatrix}0.0012 & 0.0068 & 0.0334\\ 0.0068 & 0.0856 & 0.8920\\ 0.0334 & 0.8920 & 175.3554\end{pmatrix}$	$\begin{pmatrix}0.0056 & 0\\ 0 & 5.556\times10^{-4}\end{pmatrix}$
$\begin{pmatrix}\psi\\ \dot{\psi}\end{pmatrix}$	$\begin{pmatrix}0.0002 & 1.4582\times10^{-5}\\ 1.4581\times10^{-5} & 0.0016\end{pmatrix}$	$\begin{pmatrix}5.5556\times10^{-4} & 0\\ 0 & 0.0056\end{pmatrix}$

Table E.11: Kalman filter parameters for forward v_x, lateral v_y and yaw ψ velocity estimation of the custom *Ascending Technologies (AscTec) Hummingbird*

E.4 QIH-NMPC Conditions

The QIH-NMPC framework requires the following conditions to hold for the process model $\mathbf{f}(\mathbf{x}(t),\ \mathbf{u}(t))$ and associated states $\mathbf{x}(t)$ and controls $\mathbf{u}(t)$

- A1 That $\mathbf{f}(\mathbf{x}(t), \mathbf{u}(t))$ is twice continuously differentiable and $\mathbf{f}(\mathbf{0}, \mathbf{0}) = \mathbf{0}$. For a general non-zero reference the process model must be shifted such that $\mathbf{f}(\mathbf{x}^*, \mathbf{u}^*) = \mathbf{0}$.
- A2 $\mathbb{U} \subset \mathbb{R}^m$, $\mathbf{0} \in \mathbb{U}$ and contained in the interior of $\mathbb{U}$.
- A3 $\mathbf{f}(\mathbf{x}(t),\ \mathbf{u}(t))$ has a unique solution for any initial state $\mathbf{x}_0$ and piecewise continuous $\mathbf{u}(\cdot) \in \mathbb{U}$

E.5 System Parameters - Monte-Carlo Simulations

Parameter	Nominal (*)	Min	Max
T_s (s)	1/40	-	-
T_v (s)	1/40	-	-
T_i (s)	1/10	-	-
t (s)	60	-	-
r_c (m)	$\{0.125, 0.25, 0.5\}$	-	-
T_p	10	-	-
r^* (m)	2.00	-	-
Q_s	0.01	-	-
λ_ψ	0.10	-	-
λ_R	1×10^{-8}	-	-
λ_P	2.00	-	-
ϵ	$(0.02, 0.30)$	-	-
η	$(0.125, 8)$	-	-
$\mathbf{x}_o$ (m)	(-2,0,-10)	-	-
v_x (ms^{-1})	0.20	-	-
v_y (ms^{-1})	0.00	-	-
v_z (ms^{-1})	0.00	-0.50	0.50
ω_x (degs^{-1})	0.00	-	-
ω_y (degs^{-1})	0.00	-	-
ω_z (degs^{-1})	2.87	-10.00	10.00
r_o (m)	-	5	50
v_t (ms^{-1})	-	-0.50	0.50
t_{cpa} (s)	-	5	55
$q(t)$ (deg)	$\mathcal{N}(0, 2)$	-	-
$\mathbf{q}_c(t)$ (deg/degs^{-1}/N)	$\mathcal{N}_4(\mathbf{0}_4, \text{cov}(1, 1, 0.012, 0.5))$	-	-
$\mathbf{w}_a(t)$ (m)	$\mathcal{N}_3(\mathbf{0}_3, \text{cov}(0.02, 0.02, 0.00)$	-	-
$\mathbf{w}_g(t)$ (deg)	$\mathcal{N}(\mathbf{0}_3, \mathbf{I}_3)$	-	-

Table E.12: Monte-Carlo simulation parameters for analysing avoidance and resolution decision performance

E.6 Collision Avoidance System Metrics

P(**MA**)	Probability of missed detection
P(**IC**)	Probability of induced collision
P(**CACR**)	Probability of correct avoidance (no collision) and correct resolution (return to path)
P(**CAIR**)	Probability of correct avoidance (no collision) and incorrect resolution (no return to path)
P(**PACR**)	Probability of precautionary avoidance (improved miss distance) and correct resolution (return to path)
P(**PAIR**)	Probability of precautionary avoidance (improved miss distance) and incorrect resolution (no return to path)
P(+) = P(**CACR**) + P(**CAIR**) P(**PACR**) (Probability Positive Result)	Probability that the system, including avoidance and precautionary avoidance, results in success from a collision avoidance perspective
P(-) = P(**MA**) + P(**IC**) (Probability Negative Result)	Probability that the system, including avoidance and precautionary avoidance, results in failure from a collision avoidance perspective

Table E.13: Performance metrics used to derive modified System Operating Curves (SOC)

References

[1] International Civil Aviation Organisation (ICAO), "Cir 328 AN/190 Unmanned Aircraft Systems (UAS)," *Tech. Report*, Montreal, Canada, 2011

[2] A. Watts, V Ambrosia and E. Hinkley "Unmanned aircraft systems in remote sensing and scientific research: classification and considerations of use," *Remote Sensing*, vol. 4, pp. 1671-1692, June 2012

[3] C. Korpela, M. Orsag, T. Danko, B. Kobe, C. McNeil, R. Pisch and P. Oh, "Flight stability in aerial redundant manipulators," *IEEE Int. Conf. Robotics and Automation (ICRA'12)*, pp. 3529-3530, May 2012

[4] M. Orsag, C. Korpela, S. Bogdan and P. Oh, "Hybrid adaptive control for aerial manipulation," *Journal of Intelligent & Robotic Systems*, vol. 73, no. 1-4, pp. 693-707, Jan. 2014

[5] C. Zhang and J. Kovacs, "The application of small unmanned aerial systems for precision agriculture: a review," *Precision Agriculture*, vol. 13, no. 6, pp. 693-712, Dec. 2012

[6] A. Mcfadyen, F. Gonzalez, D. Campbell and D. Eagling, "Evaluating unmanned aircraft systems for deployment in plant biosecurity," *Tech. Report*, Canberra, Australia, 2014

[7] B. Faicala, F. Costaa, G. Pessinb, J. Ueyamaa, H. Freitasa, A. Colomboa, P. Finia, L. Villasc, F. Osórioa, P. Vargasd, T. Braune, "The use of unmanned aerial vehicles and wireless sensor networks for spraying pesticides," *Journal of Systems Architecture*, vol. 60, no. 4, pp. 393-404, April 2014

[8] J. Rasmussen, J. Nielsen, F. Garcia-Ruiz, S. Christensen and J. Streibig, "Potential uses of small unmanned aircraft systems (UAS) in weed research," *Weed Research*, vol. 53, no. 4, pp. 242-248, Aug. 2013

[9] J. Torres-Sánchez, J. Peña-Barragán, D. Gómez-Candón, A. De Castro and F. López-Granados, "Imagery from unmanned aerial vehicles for early site specific weed management," *Precision Agriculture '13*, sec. 2, pp.193-199, 2013

[10] V. Gonzalez-Dugo, P. Zarco-Tejada, E. Nicolás, P. Nortes, J. Alarcón, D. Intrigliolo, E. Fereres, "Using high resolution UAV thermal imagery to assess the variability in the water status of five fruit tree species within a commercial orchard," *Precision Agriculture*, vol. 14, no. 6, pp. 660-678, Dec. 2013

[11] R. Dyer, "The business case for investment in development of precision livestock management technologies and applications," *Australian and New Zealand Spatially Enabled Livestock Management Symposium*, pp. 13-18, July 2010

[12] L. Techy, D. Schmale III and C. Woolsey, "Coordinated aerobiological sampling of a plant pathogen in the lower atmosphere using two autonomous unmanned aerial vehicles," *Journal of Field Robotics*, vol. 27, no. 3, pp. 335-343, May-June 2010

[13] I. Maza, F. Caballero, J. Capitán, J. Martínez-de-Dios and A. Ollero, "Experimental results in multi-UAV Ccoordination for disaster management and civil security applications," *Journal of Intelligent & Robotic Systems*, vol. 61, no. 1-4, pp. 563-585, Jan. 2011

[14] S. Montambault, J. Beaudry, K. Toussaint and N. Pouliot, "On the application of VTOL UAVs to the inspection of power utility assets," *Int. Conf. Applied Robotics for the Power Industry (CARPI.10)*, pp. 1-7, Oct. 2010

[15] T. Bruggemann, J. Ford and R. Walker, "Control of aircraft for inspection of linear infrastructure ," *IEEE Trans. Control Systems Technology*, vol. 19, no. 6, pp. 1397-1409, Nov. 2010

[16] B. Fung, "We are much closer than you think to a revolution in drone shipping," *Washington Post*, Dec. 9, 2013

[17] Civil Aviation Safety Authority, "List of UAS operator certificate holders," *http://www.casa.gov.au*, accessed 17 June 2014

[18] Teal Group Corporation, "World Unmanned Aerial Vehicle Systems, Market Profile and Forecast 2013," *Tech. Report Excerpt*, June 2013

[19] K. Dalamagkidis, K. Valavanis, and L. Piegl, "On unmanned aircraft systems issues, challenges and operational restrictions preventing integration into the national airspace system," *Progress in Aerospace Sciences*, vol. 44, no. 7-8, pp. 503-519, Oct-Nov. 2008

[20] K. Valavanis, "Unmanned aircraft systems: the current state-of-the-art," *Springer*, 2013

[21] International Civil Aviation Organisation (ICAO), "Doc 9854 AN/458 Global Air Traffic Management Operational Concept," *Tech. Report*, Montreal, Canada, 2005

[22] Radio Technical Commission for Aeronautics (RTCA), "DO-304 Guidance material and considerations for unmanned aircraft systems," *Tech. Report*, Washington, United States, 2007

[23] Civil Aviation Safety Authority (CASA), "RPAs (drones) in civil airspace and challenges for CASA," Canberra, Australia, 2013

[24] A. Mcfadyen, R. Clothier, D. Campbell and G. Williams, "Scoping study for remotely piloted aircraft systems integration into civil airspace," *Tech. Report*, Melbourne, Australia, 2014

[25] D. Vos, "Five steps to facilitating convergence of manned and unmanned aviation," *Rockwell Collins*, 2009

[26] Single European Sky Air Traffic Management Research (SESAR), "R & D Roadmap for RPAS integration into non-segregated airspace and the ATM environments," Brussels, Belgium, 2012

[27] Radio Technical Commission for Aeronautics (RTCA), "DO-344, Operational and Functional Requirements and Safety Objectives (OFRSO) for Unmanned Aircraft Systems (UAS) Standards," *Tech. Report*, Washington, United States, 2013

[28] Airservices Australia (AsA), Designated Airspace Handbook, Canberra, Australia, 2013

[29] Civil Aviation Safety Authority (CASA), Civil Aviation Regulations (CAR), Canberra, Australia, 1988

[30] Civil Aviation Safety Authority (CASA), Civil Aviation Safety Regulations (CASR), Canberra, Australia, 1988

[31] Civil Aviation Safety Authority (CASA), Civil Aviation Orders (CAO), Canberra, Australia, 1988

[32] Civil Aviation Safety Authority (CASA), Advisory Circular (AC), Canberra, Australia, 2001

[33] Civil Aviation Safety Authority (CASA), Manual of Standards, Canberra, Australia, 2011

[34] Airservices Australia (AsA), Manual of Air Traffic Services (MATS), Canberra, Australia, 2013

[35] Airservices Australia (AsA), Aeronautical Information Package (AIP), Canberra, Australia, 2014

[36] Civil Aviation Safety Authority (CASA), Civil Aviation Safety Regulations (CASR), Part 91, General Operating and Flight Rules, Canberra, Australia, 1988

[37] Civil Aviation Safety Authority (CASA), Civil Aviation Regulations (CAR), Part 12, Section 162-163, Canberra, Australia, 1988

[38] Civil Aviation Safety Authority (CASA), Advisory Circular (AC), Part 101-3(0), Unmanned Aircraft and Rockets - Model Aircraft, Canberra, Australia, 2002

[39] Australian Transport Safety Bureau (ATSB), "Limitations of the see-and-avoid principle," *Tech. Report*, Canberra, Australia, Nov. 2004

[40] F. Barfield, W. Res, D. Center, and A. Wright-Patterson, "Autonomous collision avoidance: the technical requirements," *IEEE National Aerospace and Electronics Conf. (NAECON'00)*, pp. 808-813, Oct. 2000

[41] S. Hutchings, and S. Farmer, "Architecting UAV sense & avoid systems," *Conf.Autonomous Systems*, Institution of Engineering and Technology, pp. 1-8, 2007

[42] C. Morris "Midair collisions: limitations of the see-and-avoid concept in civil aviation," *Aviation, Space, and Environmental Medicine*, vol. 76, no. 4, pp. 357-365, April 2005

[43] K. Colvin, R. Dodhia and R. Dismukes, "Is pilots' visual scanning adequate to avoid mid-air collisions?" *Int. Symposium on Aviation Psychology*, pp. 104-109, 2005

[44] R. Kephart, "Comparison of see-and-avoid performance in manned and remotely piloted aircraft," *PhD Dissertation*, Ohio University, 2008

[45] A. Zeitlin, "UAS See & Avoid Standards: Challenges & Progress," *2009/2010 UAS Yearbook - UAS: The Global Perspective*, June 2009

[46] A. Zeitlin, "Issues and tradeoffs in Sense & Avoid for unmanned aircraft," *IEEE Systems Conf.*, pp. 61-65, April 2010

[47] P. Angelov, "Sense and avoid in UAS: research and application," *John Wiley & Sons*, 2012

[48] X. Prats, L. Delgado, J. Ramírez, P. Royo and E. Pastor, "Requirements, issues, and challenges for sense and avoid in unmanned aircraft systems," *Journal or Aircraft*, vol. 49, no. 3, pp. 677-687, May-June 2012

[49] M. Ebdon and J. Regan, "OPR: HQ ACC/DRUAV SMO Sense-and-avoid requirement for remotely operated aircraft (roa)," *Tech. Report*, June 2004

[50] A. Lacher, D. Maroney, and A. Zeitlin, "Unmanned aircraft collision avoidance technology assessment and evaluation methods," *Air Traffic Management Research & Development Seminar*, July 2007

[51] North Atlantic Treaty Organization (NATO), "Sense and Avoid requirements for unmanned aerial vehicle systems operating in non-segregated airspace," *Tech. Report*, Brussels, Belgium, 2008

[52] C. Geyer, C. Singh, and S. Chamberlain, "CMU-RI-TR-08-03 Avoiding collisions between aircraft: state of the art and requirements for UAVs operating in civilian airspace," *Tech. Report*, March 2008

[53] J. McCalmont, J. Utt, M. Deschenes, M. Taylor, R. Sanderson, J. Montgomery, R. Johnson, and D. McDermott, "Sense and avoid technology for unmanned aircraft systems," *SPIE 6566, Automatic Target Recognition XVII*, pp. 65660P, May 2007

[54] J. McCalmont, J. Utt, and M. Deschene, "Detect and avoid technology demonstration," *AIAA Infotech@Aerospace Conf.*, May 2002

[55] S. Hottman, K. Hansen and M. Berry, "DOT/FAA/AR-08/41 Literature review on detect, sense and avoid technology for unmanned aircraft systems," *Tech. Report*, US Department of Transportation, Sep. 2009

[56] M. Owen, S. Duffy and M. Edwards, "Unmanned aircraft sense and avoid radar: Surrogate flight testing performance evaluation," *IEEE Radar Conference*, pp. 0548-0551, 2014

[57] Integrated Robotics Imaging Systems, "UAV Radar Research," *http://www.uav-alaska.com*, 2014

[58] B. Karhoff, J. Limb, S. Oravsky, and A. Shephard, "Eyes in the domestic skies: an assessment of sense and avoid technology for the army's warrior unmanned aircraft," *IEEE Systems and Information Engineering Design Synopsium*, pp. 36-42, April 2006

[59] B. Korn, and C. Edinger, "UAS in civil airspace: demonstrating sense and avoid capabilities in flight trials," *IEEE/AIAA Digital Avionics Systems Conf. (DASC'08)*, pp. 4.D.1 1-7, 2008

[60] G. Calhoun, C. Miller, T. Hughes and M. Draper, "UAS sense and avoid system interface design and evaluation." *Human Factors and Ergonomics Society Annual Meeting*, vol. 58, no. 1, pp. 21-25, 2014

[61] M. Warren, L. Mejias, X. Yang, B. Arain, F. Gonzalez and B. Upcroft, "Enabling aircraft emergency landings using active visual site detection," *Int Conf. Field and Service Robotics (FSR'13)*, Dec. 2013

[62] L. Mejias, J. Lai, J. Ford and P. O'Shea, "Demonstration of closed-loop airborne sense-and-avoid using machine vision," *IEEE Aerospace and Electronic Systems Magazine*, vol. 27, no. 4, pp. 4-7, April 2012

[63] R. Clothier, D. Frousheger and M. Wilson, "The smart skies project : enabling technologies for future airspace environments," *Int. Congress of the Aeronautical Sciences (ICAS'12)*, Sep. 2012

[64] M. Kochenderfer, L. Espindle, J. Kuchar, and J. Griffith, "A Comprehensive aircraft encounter model of the national airspace system," *Lincoln Laboratory Journal*, vol. 17, no. 2, pp. 41-53, 2008

[65] M. Kochenderfer, L. Espindle, J. Griffith, and J. Kuchar, "Encounter modelling for sense and avoid development," *Integrated Communications, Navigation and Surveillance Conf. (ICNS'08)*, pp. 1-10, May 2008

[66] M. Kochenderfer, M. Edwards, L. Espindle, J. Kuchar and J. Griffith, "Airspace encounter models for estimating collision risk," *AIAA Journal Guidance, Control, and Dynamics*, vol. 33, no. 2, pp. 487-499, March-April 2010

[67] M. Kochenderfer, L. Espindle, M. Edwards, J. Kuchar and J. Griffith, 'Airspace encounter models for conventional and unconventional aircraft," *USA/Europe Air Traffic Management Research and Development Seminar (ATM'09)*, 2009

[68] B. Carpenter and J. Kuchar, "Probability-based collision alerting logic for closely-spaced parallel approach," *AIAA Aerospace Sciences Meeting and Exhibit*, Jan. 1997

[69] J. Holland, M. Kochenderfer and W. Olson, "Optimizing the next generation collision Avoidance system for safe, suitable, and acceptable operational performance," *USA/Europe Air Traffic Management Research and Development Seminar (ATM'13)*, 2013

[70] MATLAB and Statistics Toolbox Release 2012b, The MathWorks, Inc., Natick, Massachusetts, United States, 2012

[71] B. Stroustrup, The C++ Programming Language (4th Edition), Addison-Wesley, May 2013

[72] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler and A. Ng "ROS: an open-source Robot Operating System," *ICRA workshop on Open Source Software*, vol. 3, no. 3.2, p. 5, May 2009

[73] G. Bradski, "The OpenCV library," *Dr. Dobb's Journal of Software Tools*, 2000

[74] B. Huska and H. Ferreau and M. Diehl, "ACADO Toolkit - an open source framework for automatic control and dynamic optimization," *Optimal Control Theory and Methods*, vol. 32, No. 3, pp. 298-312, May-June 2011

[75] L. Biberman, "Electro-optical imaging : system performance and modeling," *International Society for Optical Engineering (SPIE) and ONTAR Corp*, Jan. 2001

[76] G. Holst, "Electro-optical imaging system performance" *International Society for Optical Engineering (SPIE)*, Dec. 2008

[77] R. Gonzalez and R. Woods, "Digital Image Processing," *Prentice Hall*, Aug. 2007

[78] M. Lucassen, "An example of applied colour vision research: The conspicuity of airplane colour," *Conf. Colour in Graphics, Imaging, and Vision (CGIV'02)*, pp. 73-76, Jan. 2002

[79] J. Lai, J. Ford, L. Mejias, A. Wainwright, P. O'Shea and R.Walker, "Field-of-view, detection range, and false alarm trade-offs in vision-based aircraft detection," *Int. Congress on Aeronautical Sciences (ICAS'12)*, Sep. 2012

[80] J. Griffith, M. Kochenderfer, and J. Kuchar, "Electro-optical system analysis for sense and avoid," *AIAA Guidance, Navigation and Control Conference and Exhibit*, pp. 18-21, Aug. 2008

[81] A. Nussberger, H. Grabner and L. Van Gool, "Aerial object tracking from an airborne platform," *IEEE Int. Conf. Unmanned Aircraft Systems (ICUAS'14)*, pp. 1284-1293, May 2014

[82] C. Geyer and K. Daniilidis, "A unifying theory for central panoramic systems and practical implications," *Eur. Conf. Computer Vision (ECCV'00)*, pp. 445-461, June-July 2000

[83] Point Grey Research, *http://www.ptgrey.com/products/spherical.asp*

[84] J. McCandless, "Detection of aircraft in video sequences using a predictive optical flow algorithm," *Optical Engineering*, vol. 38, no. 3, pp. 523-530, March 1999

[85] G. Recchia, G. Fasano, D. Accardo and A. Moccia, "An optical flow based electro-optical see-and-avoid system for UAVs," *IEEE Aerospace Conference*, pp. 1-9, March 2007

[86] S. Petridis, C. Geyer, and S. Singh, "Learning to detect aircraft at low resolutions," *Int. Conf. Computer Vision Systems (ICVS'08)*, pp. 474-483, May 2008

[87] D. Dey, C. Geyer, S. Singh, and M. Digioia, "Passive, long-range detection of aircraft: towards a field deployable Sense and Avoid system," *Field & Service Robotics*, Springer Berlin Heidelberg, pp. 113-123, Jan. 2010

[88] D. Dey, C. Geyer, S. Singh and M. Digioia, "A cascaded method to detect aircraft in video imagery," *Int. Journal Robotics Research*, vol. 30 no. 12,pp. 1527-1540, Oct. 2011

[89] Y. Barniv, "Dynamic programming solution for detecting dim moving targets," *IEEE Trans. Aerospace and Electronic Systems*, vol. 21, no. 1, pp. 144-156, Jan. 1985

[90] T. Gandhi, Y. Mau-Tsuen, R. Kasturi, O. Camps, L. Coraor, and J. McCandless, "Detection of obstacles in the flight path of an aircraft," *IEEE Trans. Aerospace and Electronic Systems*, vol. 39, no. 1, pp. 176-191, Jan. 2003

[91] R. Carnie R. Walker and P. Corke, "Image processing algorithms for UAV sense and avoid," *IEEE Int. Conf. Robotics and Automation (ICRA'06)*, pp. 2848-2853, May. 2006

[92] J. Lai, J. Ford, P. O Shea, R. Walker, and M. Bosse, "A study of morphological preprocessing approaches for track-before-detect dim target detection," *Australasian Conference on Robotics & Automation (ACRA'08)*, Dec. 2008

[93] J. Lai, J. Ford, P. O Shea, and R. Walker, "Hidden markov model filter banks for dim target detection from image sequences," *IEEE Int. Conf. Digital Image Computing: Techniques and Applications (DICTA'08))*, pp. 312-319, Dec. 2008

[94] J. Lai and J. Ford, "Relative entropy rate based multiple hidden Markov Model Approximation," *IEEE Trans. Signal Processing*, vol. 58, no. 1, pp. 165-174, July 2009

[95] J. Lai, L. Mejias, and J. Ford, "Airborne vision-based collision-detection system," *Journal of Field Robotics*, vol. 28, no. 2, pp. 137-157, March 2011

[96] J. Lai, J. Ford, L. Mejias, P. O'Shea, and R. Walker, "Detection versus false alarm characterisation of a vision-based airborne dim-target collision detection system," *IEEE Int Conf. Digital Image Computing: Techniques and Applications (DICTA'11)*, pp. 448-455, Dec. 2011

[97] J. Lai, J. Ford, L. Mejias and P. O'Shea, "Characterization of sky-region morphological-temporal airborne collision detection," *Journal of Field Robotics*, pp. 171-193, vol. 30, No. 2, March-April 2013

[98] L. Mejias, S. McNamara, J. Lai and J. Ford, "Vision-based detection and tracking of aerial targets for uav collision avoidance," *IEEE/RSJ Int. Conf. Intelligent Robots and Systems, (IROS'10)*, pp. 87-92, Oct. 2010

[99] F. Gustafsson, F. Gunnarsson, N. Bergman, U. Forssell, J. Jansson, R. Karlsson, and P. Nordlund, "Particle filters for positioning, navigation, and tracking," *IEEE Trans. Signal Processing*, vol. 50, no. 2, pp. 425-437, Feb. 2002

[100] S. Davey, M. Rutten, and B. Cheung, "A comparison of detection performance for several track-before-detect algorithms," *IEEE Int. Conf Information Fusion*, pp. 1-8, June-July 2008

[101] M. Zaveri, S. Merchant, and U. Desai, "Wavelet-based detection and its application to tracking in an IR sequence," *IEEE Trans. Systems, Man, and Cybernetics, Part C: Applications and Reviews*, vol. 37, no. 6, pp. 1269-1286, Nov. 2007

[102] D. Grilley, "Resolution requirements for passive Sense & Avoid," *AUVSI Conference*, 2005

[103] E. Kamen and K. Su, "Introduction to optimal estimation," *Springer-Verlag*, 1999

[104] J. Lai, J. Ford, L. Mejias, and P. O'Shea, "Vision-based estimation of airborne target pseudobearing rate using hidden Markov model filters," *IEEE Trans. Aerospace and Electronic Systems*, vol. 49, no. 4, pp. 2129-2145, Oct. 2013

[105] L. Forlenza, G. Fasano, D. Accardo, and A. Moccia, "Flight performance analysis of an image processing algorithm for integrated sense-and-avoid systems," *Int. Journal Aerospace Engineering*, vol. 2012, April 2012

[106] G. Fasano, D. Accardo, A. Tirri and A. Moccia, "Morphological filtering and target tracking for vision-based UAS sense and avoid," *IEEE Int. Conf. Unmanned Aircraft Systems (ICUAS'14)*, pp. 430-440, May 2014

[107] T. Bruggemann and L. Mejias, "Accuracy of monocular vision-derived angles to a threat for see and avoid control," *Int. Congress of the Aeronautical Sciences (ICAS'12)*, Sep. 2012

[108] T. Bruggemann and L. Mejias, "Airborne collision scenario flight tests: impact of angle measurement errors on reactive vision-based avoidance control," *Australian Int. Aerospace Conf. (AIAC'13)*, Feb. 2013

[109] N. Ayache, "Artificial vision for mobile robots: stereo vision and multisensory perception," *MIT Press*, 1991

[110] M. Warren, D. McKinnon, H. He, A. Glover, M. Shiel and B. Upcroft, "Large scale monocular vision-only mapping from a fixed-wing UAS," *Field and Service Robotics*, vol. 92, pp. 495-509, 2014

[111] Y. Bar-Shalom, X. Li, and T. Kirubarajan, "Estimation with applications to tracking and navigation," *Wiley-Interscience*, 2001

[112] S. Nardone and V. Aidala, "Observability criteria for bearings-only target motion analysis," *IEEE Trans. Aerospace and Electronic Systems*, vol. 17, no. 2, pp. 162-166, March 1981

[113] A. Payne, "Observability problem for bearings-only tracking," *Int. Journal of Control*, vol. 49, no. 3, pp. 761-768, 1989

[114] S. Hepner and H. Geering, "Observability analysis for target maneuver estimation via bearing-only and bearing-rate-only measurements," *Journal of Guidance, Control, and Dynamics*, vol. 13, no. 6, pp. 977-983, Nov-Dec. 1990

[115] S. Hammel and V. Aidala, "Observability requirements for three-dimensional tracking via angle measurements," *IEEE Trans. Aerospace and Electronic Systems*, vol. 21, no. 2, pp. 200-207, Feb. 2007

[116] B. La Scala, M. Mallick, and S. Arulampalam, "Differential geometry measures of nonlinearity for filtering with nonlinear dynamic and linear measurement models," *SPIE 6699, Signal and Data Processing of Small Targets*, Aug. 2007

[117] S. Blackman and R. Popoli, "Design and analysis of modern tracking systems," *Artech House*, 1999

[118] M. Ghalia, "Observability problem in target tracking via angles-only measurements," *SPIE 1950, Acquisition, Tracking, and Pointing VII*, vol. 171, Oct. 1993

[119] D. Woffinden and D. Geller, "Observability criteria for angles-only navigation," *IEEE Trans. Aerospace and Electronic Systems*, vol. 45, no. 3, pp. 1194-1208, July 2009

[120] J. Passerieux and D. Van Cappel, "Optimal observer maneuver for bearings-only tracking," *IEEE Trans. Aerospace and Electronic Systems*, vol. 34 , no. 3, pp. 777-788, July 1998

[121] Y. Oshman and P. Davidson, "Optimization of observer trajectories for bearings-only target localization," *IEEE Trans. Aerospace and Electronic Systems*, vol. 35, no. 3, pp. 892-902, July 1999

[122] V. Aidela and S. Hammel, "Utilization of modified polar coordinates for bearings-only tracking," *IEEE Trans. Automatic Control*, vol. 28 , no. 3, pp. 283-294, March 1983

[123] R. Gray, "Spiral guidance law for passive range estimation," *SPIE 1950, Acquisition, Tracking, and Pointing VII*, Oct. 1993

[124] N. Peach, "Bearings-only tracking using a set of range-parameterised extended Kalman filters," *IEE Proceedings - Control Theory and Applications*, vol. 142, no. 1, pp. 73-80, Jan. 1995

[125] A. Farina, "Target tracking with bearings-only measurements," *Signal Processing*, vol. 78, no. 1, pp. 61 78, Oct. 1999

[126] T. Erlandsson, "Angle-only target tracking," *Masters Thesis*, Linköping University, Sweden, 2007

[127] B. La Scala and M. Morelande, "An analysis of the single sensor bearings-only tracking problem," *IEEE Int. Conf. Information Fusion (FUSION'08)*, pp. 1-6, July 2008

[128] D. Franken and A. Hupper, "Unified tracking and fusion for airborne collision avoidance using log-polar coordinates," *IEEE Int. Conf. Information Fusion (FUSION'12)*, pp. 1246-1253, July 2012

[129] O. Shakernia, M. Chen, and V. Raska, "Passive ranging for uav sense and avoid applications," *AIAA Infotech@Aerospace Conf. and Exhibit*, pp. 1-10, March 2005

[130] O. Shakernia, W. Chen, S. Graham, J. Zvanya, A. White, N. Weingarten, and V. Raska, "Sense and avoid (SAA) flight test and lessons learned," *AIAA Infotech@Aerospace Conf. and Exhibit*, pp. 1-12, May 2007

[131] X. Rong Li and V. Jilkov, "Survey of maneuvering target tracking part I: dynamic models," *IEEE Trans. Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1333-1364, Oct. 2003

[132] X. Rong Li and V. Jilkov, "Survey of maneuvering target tracking part III: measurement models," *SPIE 4473, Signal and Data Processing of Small Targets*, Nov. 2001

[133] T. Kirubarajan and Y. Bar-Shalom, "Kalman filter versus imm estimator: when do we need the latter?" *IEEE Trans. Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1452-1457, Oct. 2003

[134] B. Ristic and M. Arulampalam, "Tracking a manoeuvring target using angle-only measurements: algorithms and performance," *Signal Processing*, vol. 83, no. 6, pp. 1223-1238, March 2003

[135] X. Rong Li and V. Jilkov, "Survey of maneuvering target tracking part V: multiple-model methods," *IEEE Trans. Aerospace and Electronic Systems*, vol. 41, no. 4, pp. 1255-1321, Oct. 2005

[136] P. Nordlund, "Efficient estimation and detection methods for airborne applications," *PhD Thesis*, Linköping, Sweden, 2008

[137] J. Lan, X. Li, V. Jilkov and C. Mu, "Second-order markov chain based multiple-model algorithm for maneuvering target tracking," *IEEE Trans. Aerospace and Electronic Systems*, vol. 49, no. 1, pp. 3-19, Jan. 2013

[138] M. Morelande and S. Challa, "Manoeuvring target tracking in clutter using particle filters," *IEEE Trans. Aerospace and Electronic Systems*, vol. 41, no. 1, pp. 252-270, Jan. 2005

[139] M. Bugallo, S. Xu, and P. Djurić "Performance comparison of EKF and particle filtering methods for maneuvering targets," *Digital Signal Processing*, vol. 17, no. 4, pp. 774-786, July 2007

[140] S. Shim, D. Won, M. Tahk, K.Seong and E.Kim, "Vision-based long-range target detection using coarse-to-fine particle filter," *Journal Aerospace Engineering*, vol. 49, no. 1, pp. 3-19, Oct. 2013

[141] A. Elena Tirri, G. Fasano, D. Accardo, A. Moccia and E. De Lellis, "Advanced sensing issues for UAS collision avoidance," *Int. Conf. Application and Theory of Automation in Command and Control Systems (ATACCS'12))*, pp. 12-19, May 2012

[142] D. Regan and A. Vincent, Visual processing of looming and time to contact throughout the visual field," *Vision Research*, vol. 35, no. 13, pp. 1845-1857, July 1995

[143] R. Gray and D. Regan, "Accuracy of estimating time to collision using binocular and monocular information," *Vision Research*, vol. 38, no. 4, pp. 499-512, Feb. 1998

[144] J. Yan, B. Lorv, H. Li and H. Sun, "Visual processing of the impending collision of a looming object: Time to collision revisited," *Journal of Vision*, vol. 11, no. 12, Oct. 2011

[145] D. Lee, R. Beard, P. Merrell and P. Zhan, "See and avoidance behaviors for autonomous navigation," *SPIE 5609, Mobile Robots XVII*, Dec. 2004

[146] J. Byrne and C. Taylor, "Expansion segmentation for visual collision detection and estimation," *IEEE Int. Conf. Robotics and Automation (ICRA'09)*, pp. 875-882, May 2009

[147] R. Bates, K. Kubala and A. Greengard, "Foveated imager providing reduced time-to-threat detection for micro unmanned aerial system," *SPIE 8360, Airborne Intelligence, Surveillance, Reconnaissance (ISR) Systems and Applications IX*, May 2012

[148] F. Martel, R. Schultz,W. Semke, Z. Wang and M. Czarnomski, "Unmanned aircraft systems sense and avoid avionics utilizing ADS-B transceiver," *AIAA Infotech@ Aerospace Conf.*, pp. 6-9, April 2009

[149] F. Martel, M. Mullins, N. Kaabouch and W. Semke, "Flight testing of an ADS-B-based miniature 4D sense and avoid system for small UAS," *AIAA Infotech@ Aerospace Conf.*, March 2011

[150] K. Foerster, M. Mullins, N. Kaabouch and W. Semke, "Flight testing of a right-of-way compliant ADS-B-based miniature sense and avoid system," *AIAA Infotech@ Aerospace Conf.*, June 2012

[151] B. Stark, B. Stevenson and Y. Chen, "ADS-B for small unmanned aerial systems: Case study and regulatory practices," *IEEE Int. Conf. Unmanned Aircraft Systems (ICUAS'13)*, pp. 152-159, May 2013

[152] J. Cutting, P. Vishton and P. Braren, "How we avoid collisions with stationary and moving obstacles," *Psychological Review*, vol. 102, no. 4, pp. 627-651, Oct. 1995

[153] D. Le Brun, Nouveau manuel de marin (New Seamans Manual), 2002

[154] M. Lenoir ,E. Musch ,E. Thiery and G. Savelsbergh "Rate of change of angular bearing as the relevant property in a horizontal interception task during locomotion," *Journal Motor Behaviour*, vol. 34, no. 4, pp. 385-401, Dec. 2002

[155] R. Ni and G. Andersen, "Detection of collision events on curved trajectories: optical information from invariant rate-of-bearing change," *Perception & Psychophysics*, vol. 70, no. 7, pp. 1314-1324, Oct. 2008

[156] M Kochenderfer, J. Griffith, and J. Kuchar, "Hazard alerting using line-of-sight rate," *AIAA Guidance, Navigation and Control Conf. and Exhibit*, pp. 2003-2011, Aug. 2008

[157] B. Fajen, "Guiding locomotion in complex, dynamic environments," *Frontiers in Behavioral Neuroscience*, vol. 7, no. 85, pp. 1-15, July 2013

[158] I. Koglbauer, R. Braunstingl and T. Haberkorn, "Modeling human and animal collision avoidance strategies," *Int. Symp. Aviat. Psychol*, pp. 554-559, 2013

[159] T. Haberkorn, I. Koglbauer, R. Braunstingl and B. Prehofer, "Requirements for future collision avoidance systems in visual flight: a human-centered approach," *IEEE Trans. Human-Machine Systems*, vol. 43, no. 6, pp. 583-594, Oct. 2013

[160] B. Fajen and W. Warren, "Behavioral dynamics of steering, obstacle avoidance, and route selection," *Journal Experimental Psychology-Human Perception and Performance*, vol. 29, no. 2, pp. 343-361, April 2003

[161] B. Fajen and W. Warren, "Behavioural dynamics of intercepting a moving target," *Experimental Brain Research*, vol. 180, no. 2, pp. 303-319, June 2007

[162] M. Zago , J. McIntyre, P. Senot and F. Lacquaniti, "Visuo-motor coordination and internal models for object interception," *Experimental Brain Research*, vol. 192, no. 4, pp. 571-604, Feb. 2009

[163] G. Diaz, F. Phillips and B. Fajen, "Intercepting moving targets: a little foresight helps a lot," *Experimental Brain Research*, vol. 195, no. 3, pp. 345-360, May 2009

[164] J. Kuchar and L. Yang, "A review of conflict detection and resolution modeling methods," *IEEE Trans. Intelligent Transportation Systems*, vol. 1, no. 4, pp. 179-189, Dec. 2000

[165] B. Albaker and N. Rahim, "A survey of collision avoidance approaches for unmanned aerial vehicles," *Int. Conf. Technical Postgraduates (TECHPOS'09)*, pp. 1-7, Dec. 2009

[166] J. Kuchar and A. Drumm, "The traffic alert and collision avoidance system," *Lincoln Laboratory Journal* vol. 16, no. 2, pp. 277-296, 2007

[167] S. Lee and E. Mueller, "A systems-based approach to functional decomposition and allocation for developing UAS separation assurance concepts," *AIAA Conf. Aviation Technology, Integration, and Operations*, 2013

[168] L. Winder and J. Kuchar, "Evaluation of collision avoidance maneuvers for parallel approach," *Journal of Guidance, Control and Dynamics*, vol. 22, no. 6, pp. 801-807, Nov-Dec. 1999

[169] G. Bakker, H. Kremer, and H. Blom, "Geometric and probabilistic approaches towards conflict prediction," *USA/Europe Air Traffic Management Research and Development Seminar (ATM'00)*, 2000

[170] L. Winder and J. Kuchar, "ICAT-2000-5 Generalized philosophy of alerting with applications for parallel approach collision prevention," *Tech. Report*, Lincoln Laboratory, Massachusetts Institute of Technology (MIT), 2000

[171] M. Prandini, J. Lygeros, A. Nilim, and S. Sastry, "A probabilistic framework for aircraft conflict detection," *AIAA Conf. Guidance, Navigation and Control*, pp. 1-36, Aug. 1999

[172] Y. Matsuno, and T. Tsuchiya, "Probabilistic conflict detection in the presence of uncertainty," *Air Traffic Management and Systems*, pp. 17-33, Feb. 2014

[173] M. Prandini, J. Hu, J. Lygeros, and S. Sastry, "A probabilistic approach to aircraft conflict detection," *IEEE Trans. Intelligent Transportation Systems*, vol. 1, no. 4, pp. 199-220, Dec. 2000

[174] P. Nordlund and F. Gustafsson, "Probabilistic near mid-air collision avoidance," *Tech. Report*, Linköping University, 2008

[175] F. Lindsten, P. Nordlund, F. Gustafsson, "Conflict detection metrics for aircraft Sense and Avoid systems," *Tech. Report*, Linköping University, 2009

[176] C. Van Daalen and T. Jones, "Fast conflict detection using probability flow," *Automatica*, vol. 45, no. 8, pp. 1903-1909, Aug. 2009

[177] R. Paielli and H. Erzberger, "Conflict probability estimate for free flight," *Journal Guidance, Control and Dynamics*, vol. 20 no. 3, pp. 588-596, June 1997

[178] J. Hu and M. Prandini, "Aircraft conflict detection: a method for computing the probability of conflict based on Markov chain approximation," *European Control Conf. (ECC'03)*, Sep. 2003

[179] T. Jones, "Tractable conflict risk accumulation in quadratic space for autonomous vehicles," *Journal Guidance, Control and Dynamics*, vol. 29, no. 1, pp. 39-48, Jan-Feb. 2006

[180] M. Albasman and J. Hu, "Probability of conflict analysis of 3D aircraft flight based on two-Level markov chain approximation approach," *Int. Conf. Networking, Sensing and Control (ICNSC'10)*, April 2010

[181] L. Yang, J. Yang, J. Kuchar, and E. Feron, "A real-time Monte Carlo implementation for computing probability of conflict," *AIAA Guidance, Navigation, and Control Conf. and Exhibit*, pp. 617-631, Aug. 2004

[182] S. Ramasamy, R. Sabatini and A. Gardi, "Avionics sensor fusion for small size unmanned aircraft sense-and-avoid" *IEEE Metrology for Aerospace (MetroAeroSpace'14)*, pp. 271-276, May 2014

[183] B. Vanek, T. PÃľni, P. Bauer and J. Bokor, "Vision only sense and avoid: a probabilistic approach," *American Control Conf. (ACC'14)*, pp. 1204-1209, June 2014

[184] P. Angelov, C. Bocaniala, C. Xydeas, C. Pattchett, D. Ansell, M. Everett, and G. Leng, "A passive approach to autonomous collision detection and avoidance in uninhabited aerial systems," *IEEE Int. Conf. Computer Modelling and Simulation*, pp. 64-69, April 2008

[185] W. Love, "Preview of TCAS II Version 7," *Air Traffic Control Quarterly*, vol. 6, no. 4, pp. 231-247, 1998

[186] J. Kuchar and L. Yang, "Prototype conflict alerting system for free flight," *Tech. Report*, National Aeronautics and Space Administration (NASA), 1997

[187] R. Chamlou, "Design principles and algorithm development for two types of NextGen airborne conflict detection and collision avoidance," *IEEE Int. Conf. Integrated Communications Navigation and Surveillance (ICNS'10)*, pp. N7-1-N7-12, May 2010

[188] M. Kochenderfer, J. Chryssanthacopoulos, L. Kaelbling, and T. Lozano-Perez, "ATC-360 Model-based optimization of airborne collision avoidance logic," *Tech. Report*, Lincoln Laboratory, Massachusetts Institute of Technology (MIT), 2010

[189] J. Kuchar, "Methodology for alerting-system performance evaluation," *AIAA Journal of Guidance, Control, and Dynamics*, vol. 19, no. 2, pp. 438-444, March-April 1996

[190] C. Metz, "Basic principles of ROC analysis," *Seminars in Nuclear Medicine*, vol. 8, no. 4, pp. 283-298, Oct. 1978

[191] J. Kuchar and L. Winder, "Hazard avoidance alerting with Markov decision processes," *Tech. Report*, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology (MIT), 2004

[192] L. Kaelbling and T. Lozano-Perez, "MIT-CSAIL-TR-2009-043 Finding aircraft collision-avoidance strategies using policy search methods," *Tech. Report*, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology (MIT), 2009

[193] H. Bai, D. Hsu, M. Kochenderfer and W. Sun Lee. "Unmanned aircraft collision avoidance using continuous-state POMDPs," *Robotics: Science and Systems*, vol. 7, pp. 1-8, 2012

[194] S. Temizer, M. Kochenderfer, L. Kaelbling, T. Lozano-Pérez, and J. Kuchar, "Collision avoidance for unmanned aircraft using markov decision processes," *AIAA Guidance, Navigation, and Control Conf.*, Aug. 2010

[195] T. Wolf and M. Kochenderfer, "Aircraft collision avoidance using Monte Carlo real-time belief space search," *Journal of Intelligent & Robotic Systems* vol. 64, no. 2 pp. 277-298, Nov. 2011

[196] J. Chryssanthacopoulos and M. Kochenderfer, "Accounting for state uncertainty in collision avoidance," *Journal of Guidance, Control, and Dynamics*, vol. 34, no. 4 pp. 951-960, July-Aug. 2011

[197] M. Kochenderfer, J. Chryssanthacopoulos and R. Weibel, "A new approach for designing safer collision avoidance systems," *Air Traffic Control Quarterly*, vol. 20, no. 1, 2012

[198] T. Billingsley, M. Kochenderfer and J. Chryssanthacopoulos, "Collision avoidance for general aviation," *IEEE Aerospace & Electronic Systems Magazine*, vol. 27, no. 7, pp. 4-12, July 2012

[199] A. Zeitlin and M. McLaughlin, "Modeling for UAS collision avoidance," *AUVSI Conf.*, 2006

[200] K. Nonami, F. Kendoul, S. Suzuki, W. Wang and D. Nakazawa, "Autonomous flying robots: unmanned aerial vehicles and micro aerial vehicles," *Springer*, 2010

[201] A. Alexopoulos, A. Kandil, P. Orzechowski and E. Badreddin,, "A comparative study of collision avoidance techniques for unmanned aerial vehicles," *IEEE Int. Conf. Systems, Man, and Cybernetics (SMC'13)*, pp. 1969-1974, Oct. 2013

[202] M. Hoy, A. Matveev and A. Savkin, "Algorithms for collision-free navigation of mobile robots in complex cluttered environments: a survey," *Robotica*, pp. 1-35, March 2014

[203] B. Sinopoli, M. Micheli, G. Donato and T. Koo, "Vision based navigation for an unmanned aerial vehicle," *IEEE Int. Conf. Robotics and Automation (ICRA'01)*, pp. 1757-1764, 2001

[204] S. Zingg, D. Scaramuzza, S. Weiss and R. Siegwart, "MAV navigation through indoor corridors using optical flow," *IEEE Int. Conf. Robotics and Automation (ICRA'10)*, pp. 3361-3368, May 2010

[205] V. Lippiello, G. Loianno and B. Siciliano, "MAV indoor navigation based on a closed-form solution for absolute scale velocity estimation using optical flow and inertial data," *IEEE Int. Conf. Decision and Control and European Control Conference (CDC-ECC'11)*, pp. 3566-3571, Dec. 2011

[206] M. Blosch, S. Weiss, D. Scaramuzza and R. Siegwart, "Vision based MAV navigation in unknown and unstructured environments," *IEEE Int. Conf. Robotics and Automation (ICRA'10)*, pp. 21-28, May 2010

[207] F. Fraundorfer, L. Heng, D. Honegger, G. Hee Lee, L. Meier, P. Tanskanen and M. Pollefeys, "Vision-based autonomous mapping and exploration using a quadrotor MAV," *IEEE/RSJ Int. Conf. Robots and Systems (IROS'12)*, pp. 4557-4564, Oct. 2012

[208] B. Call, R. Beard, C. Taylor and B. Barber, "Obstacle avoidance for unmanned air vehicles using image feature tracking," *AIAA Guidance, Navigation, and Control conf.* pp. 3406-3414, Aug. 2006

[209] F. Wen, Z. Qu, C. Wang and B. Hu, "Study on real-time obstacle avoidance of mobile robot based on vision sensor," *IEEE Int. Conf. Automation and Logistics (ICAL'08)*, pp. 2438-2442, Sept. 2008

[210] T. Mori and S. Scherer, "First results in detecting and avoiding frontal obstacles from a monocular camera for micro unmanned aerial vehicles," *IEEE Int. Conf. Robotics and Automation (ICRA'13)*, pp. 1750-1757, May 2013

[211] G. Fasano, D. Accardo, A. Tirri, A. Moccia and E. De Lellis, "Flight performance assessment of vision-based detection and tracking for UAS sense and avoid," *Infotech@ Aerospace Conf.*, Aug. 2013

[212] N. Franceschini, J. Pichon, C. Blanes and J. Brady, "From insect vision to robot vision," *Philosophical Transactions of The Royal Society of London. Series B: Biological Sciences* vol. 337, no. 1281, pp. 283-294, Sep. 1992

[213] T. Neumann and H. Bülthoff, "Insect inspired visual control of translatory flight," *Advances in Artificial Life*, pp. 627-636, 2001

[214] R. Nelson and J. Aloimonos, "Obstacle avoidance using flow field divergence," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 11, no. 10, pp. 1102-1106, Oct. 1989

[215] H. Chao, Y. Gu and M. Napolitano, "A survey of optical flow techniques for UAV navigation applications," *IEEE Int. Conf. Unmanned Aircraft Systems (ICUAS'13)*, pp. 710-716, May 2013

[216] G. De Croon, E. De Weerdt, C. De Wagter, B. Remes and R. Ruijsink, "The appearance variation cue for obstacle avoidance," *IEEE Trans. Robotics*, vol. 28, no. 2, pp. 529-534, Oct. 2011

[217] S. Bermudez i Badia, P. Pyk, and P. Verschure, "A fly-locust based neuronal control system applied to an unmanned aerial vehicle: the invertebrate neuronal principles for course stabilization, altitude control and collision avoidance," *International Journal of Robotics Research*, vol. 26, pp. 759, July 2007

[218] A. Merz, "Optimal aircraft collision avoidance," *Joint Automatic Control Conf.*, pp. 449-454, 1973

[219] A. Merz, "Maximum-miss aircraft collision avoidance," *Dynamics and Control*, vol. 1, no. 1, pp. 25-34, march 1991

[220] T. Tarnopolskaya and N. Fulton, "Optimal cooperative collision avoidance strategy for coplanar encounter: Merz's solution revisited," *Journal Optimization Theory and Applications*, vol. 140, no. 2, pp. 355-375, Feb. 2009

[221] T. Tarnopolskaya, N. Fulton and H. Maurer, "Synthesis of optimal bang-bang control for cooperative collision avoidance for aircraft (ships) with unequal linear speeds," *Journal Optimization Theory and Applications*, vol. 155, no. 1, pp. 115-144, Oct. 2012

[222] A. Smith and F. Harmon, "UAS collision avoidance algorithm based on an aggregate collision cone approach," *Journal Aerospace Engineering*, vol. 24, no. 4, pp. 463-477, Oct. 2011

[223] L. Schmitt and W. Fichter, "Collision-avoidance framework for small fixed-wing unmanned aerial vehicles," *Journal Guidance, Control, and Dynamics*, vol. 37, no. 4, pp. 1323-1329, 2014

[224] P. Fiorini and Z. Shiller, "Motion planning in dynamic environments using velocity obstacles," *International Journal of Robotics Research*, vol. 17, no. 7, pp. 760-772, July 1998

[225] J. Van den Berg, J. Snape, S. Guy and D. Manocha, "Reciprocal collision avoidance with acceleration-velocity obstacles," *IEEE Int. Conf. Robotics and Automation (ICRA'11)*, pp. 3475-3482, May 2011

[226] J. Snape, J. Van den Berg, S. Guy and D. Manocha, "The hybrid reciprocal velocity obstacle," *IEEE Trans. Robotics*, vol. 27, no. 4, pp. 696-706, April 2011

[227] J. Alonso-Mora, M. Schoch, A. Breitenmoser, R. Siegwart and P. Beardsley, "Object and animation display with multiple aerial vehicles," *IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS'12)*, pp. 1078-1083, Oct. 2012

[228] M. Rufli, J. Alonso-Mora and R. Siegwart, "Reciprocal collision avoidance with motion continuity constraints," *IEEE Trans. Robotics* vol. 29, no. 4 pp. 899-912, May 2013

[229] D. Alejo, J. Cobano, G. Heredia and A. Ollero, "Optimal reciprocal collision avoidance with mobile and static obstacles for multi-UAV systems," *IEEE Int. Conf. Unmanned Aircraft Systems (ICUAS'14)*, pp. 1259-1266, May 2014

[230] Y. Jenie, E. Van Kampen, C. De Visser and Q. Chu, "Velocity obstacle method for non-cooperative autonomous collision avoidance system for UAVs," *AIAA Guidance, Navigation, and Control Conf.*, 2014

[231] C. Goerzen, Z. Kong and B. Mettler, "A survey of motion planning algorithms from the perspective of autonomous UAV guidance," *Journal of Intelligent and Robotic Systems*, vol. 57, no. 1-4, pp. 65-100, Jan. 2010

[232] J. Cobano, R. Conde, D. Alejo and A. Ollero, "Path planning based on genetic algorithms and the monte-carlo method to avoid aerial vehicle collisions under uncertainties," *IEEE Int. Conf. Robotics and Automation (ICRA'11)*, pp. 4429-4434, May 2011

[233] G. Wang and S. Sam Ge, "General fight rule-based trajectory planning for pairwise collision avoidance in a known environment," *Int. Journal of Control, Automation and Systems*, vol. 12, no. 4, pp. 813-822, Aug. 2014

[234] J. Chuang and N. Ahuja. "An analytically tractable potential field model of free space and its application in obstacle avoidance," *IEEE Trans. Systems, Man, and Cybernetics, Part B: Cybernetics*, vol. 28, no. 5, pp. 729-736, Oct. 1998

[235] J. Ruchti, R. Senkbeil, J. Carroll, J. Dickinson, J. Holt and S. Biaz, "Unmanned aerial system collision avoidance using artificial potential fields," *Journal of Aerospace Information Systems* vol. 11, no. 3, pp. 140-144, March 2014

[236] M. Nieuwenhuisen, M. Schadler and S. Behnke, "Predictive potential field-based collision avoidance for multicopters," *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences (ISPRS'13)*, vol. 1, no. 2, pp. 293-298, Sep. 2013

[237] L. Mejias, J. Ford, and J. Lai, "Towards the implementation of vision-based uas sense-and-avoid," *Int. Congress of the Aeronautical Sciences (ICAS'10)*, Sep. 2010

[238] A. Beyeler, J. Zufferey and D. Floreano, "Vision-based control of near-obstacle flight," *Autonomous Robots*, vol. 27, no. 3 pp. 201-219, Oct. 2009

[239] C. Bills, J. Chen and A. Saxena, "Autonomous MAV flight in indoor environments using single image perspective cues," *IEEE Int. Conf. Robotics and automation (ICRA'11)*, pp. 5776-5783, May 2011

[240] X. Yang, L. Mejias, and T. Bruggemann, "A 3D collision avoidance strategy for uavs in an non-cooperative environment," *Journal of Intelligent Robotic Systems*, vol. 70, no. 1-4, pp. 315-327, April 2013.

[241] J. Saunders and R. Beard, "Reactive vision based obstacle avoidance with camera field of view constraints," *AIAA Guidance, Navigation and Control Conf. and Exhibit*, 2008

[242] J. Saunders, R. Beard, and J. Byrne, "Vision-based reactive multiple obstacle avoidance for micro air vehicles," *American Control Conference (ACC'09)*, pp. 5253-5258, June 2009

[243] R. Sharma, J. Saunders, and R. Beard, "Reactive path planning for micro air vehicles using bearing only measurements," *Journal of Intelligent Robotic Systems*, vol. 65, no. 1-4, pp. 409-416, Jan. 2012.

[244] H. Yu, R. Sharma, R. Beard and C. Taylor, "Observability-based local path planning and collision avoidance for micro air vehicles using bearing-only measurements," *American Control Conference (ACC'11)*, pp. 4649-4654, June 2011

[245] H. Yu, R. Sharma, R. Beard and C. Taylor, "Observability-based local path planning and obstacle avoidance using bearing-only measurements," *Robotics and Autonomous Systems*, vol 61, no. 12, pp. 1392-1405, Dec. 2013

[246] H. Yu and R. Beard, "Vision-Based Local-Level Frame Mapping and Planning in Spherical Coordinates for Miniature Air Vehicles," *IEEE Trans. Control Systems Technology*, vol. 21, no. 3, pp. 695-703, May 2012

[247] H. Choi, Y. Kim and I. Hwang "Reactive collision avoidance of unmanned aerial vehicles using a single vision sensor," *Journal Guidance, Control and Dynamics* vol. 36, no. 4, pp. 1234-1240, July-Aug. 2013

[248] S. Hutchinson, G. Hager and P. Corke, "A tutorial on visual servoing," *IEEE Trans. Robotics and Automation*, vol. 12, no. 5, pp. 651-670, Oct. 1996

[249] S. Han, H. Bang and C. Yoo, "Proportional navigation-based collision avoidance for UAVs," *Int. Journal of Control, Automation and Systems*, vol 7, no. 4 , pp. 553-565, Aug. 2009

[250] B. White, H. Shin and A. Tsourdos, "UAV obstacle avoidance using differential geometry concepts," *IFAC World Congress*, pp. 6325-6330, Aug. 2011

[251] Y. Lin and S. Saripalli, "Path planning using 3D dubins curve for unmanned aerial vehicles," *IEEE Int. Conf. Unmanned Aircraft Systems (ICUAS'14)*, pp. 296-304, May 2014

[252] M. Olivares-Mendez, L. Mejias, P. Campoy and I. Mellado-Bataller. "Cross-Entropy optimization for scaling factors of a fuzzy controller: a see-and-avoid approach for unmanned aerial systems," *Journal of Intelligent & Robotic Systems* vol. 69, no. 1-4, pp. 189-205, Jan. 2013

[253] N. Smith, R. Cobb, S. Pierce and V. Raska, "Optimal collision avoidance trajectories for unmanned/remotely piloted aircraft," *AIAA Guidance, Navigation, and Control Conf.*, 2013

[254] Y. Watanabe, A. Calise, E. Johnson and J. Evers. "Minimum-effort guidance for vision-based collision avoidance," *AIAA Atmospheric Flight Mechanics Conf. and Exhibit*, pp. 21-24, Aug. 2006

[255] D. Shim and S. Sastry, "An evasive maneuvering algorithm for UAVs in see-and-avoid situations," *American Control Conference (ACC'07)*, pp. 3886-3891, July 2007

[256] M. Hoy, A. Matveev, M. Garratt, and A. Savkin, "Collision-free navigation of an autonomous unmanned helicopter in unknown urban environments: sliding mode and MPC approaches," *Robotica*, vol. 30, no. 4, pp. 537-550, July 2012

[257] L. Filippis and G. Guglieri, "NMPC and genetic algorithm-based approach for trajectory tracking and collision avoidance of UAVs," *Int. Journal of Innovative Computing and Applications*, vol. 5, no. 3, pp. 173-183, Aug. 2013

[258] H. Fukushima, K. Kon, and F. Matsuno, "Model predictive formation control using branch-and-bound compatible with collision avoidance problems," *IEEE Trans. Robotics*, vol. 29, no. 5, pp. 1308-1317, May 2013

[259] Y. Wang, X. Zhu, Z. Zhou, H. Zhang and M. Xu, "NMPC based path following for a small fixed-wing unmanned aerial vehicle in a dynamic environment," *Future Computer and Information Technology*, vol. 86, pp. 75-82, 2014

[260] W. Green and P. Oh, "Optic-flow-based collision avoidance," *IEEE Robotics & Automation Magazine*, vol. 15, no. 1, pp. 96-103, March 2008

[261] D. Floreano, J. Zufferey, M. Srinivasan, C. Ellington, N. Franceschini, F. Ruffier, and J. Serres, "Optic flow based autopilots: speed control and obstacle avoidance," *Flying Insects and Robots*, pp. 29-50, Aug. 2009

[262] J. Zufferey, A. Beyeler and D. Floreano, "Optic flow to steer and avoid collisions in 3D," *Flying Insects and Robots*, pp. 73-86, Aug. 2010

[263] L. Mejias I. Mondragon and P. Campoy, "Omnidirectional bearing-only see and avoid for small aerial robots," *IEEE Int. Conf. Automation, Robotics and Applications (ICARA'11)*, pp. 23-28, Dec. 2011

[264] S. Bouabdallah, "Design and control of quadrotors with application to autonomous flying," *PhD Thesis*, École Polytechnique Federale de Lausanne, 2007

[265] S. Bouabdallah amd R. Siegwart, "Full control of a quadrotor," *IEEE/RSJ Int. Conf. Intelligent Robotics and Systems (IROS'07)*, pp. 153-158, Oct-Nov. 2007

[266] F. Kendoul, "Survey of advances in guidance, navigation, and control of unmanned rotorcraft systems," *Journal of Field Robotics*, vol. 29, no. 2, pp. 315-378, March-April 2012

[267] P. Zapotezny-Anderson and J. Ford, "Optimal-stopping control for airborne collision avoidance and return-to-course Flight," *Australian Control Conference (AUCC'11)*, pp. 155-160, Nov. 2011

[268] M. Blosch, S. Weiss, D. Scaramuzza, and R. Siegwart, "Vision based MAV Navigation in unknown and unstructured environments," *IEEE Int. Conf. Robotics and Automation (ICRA'10)*, pp. 21âĂŞ28, May 2010

[269] M. Burri, Y. Nikolic, C. Hurzeler, G. Caprari and R. Siegwart "Aerial service robots for visual inspection of thermal power plant boiler systems," *Int. Conf. Applied Robotics for the Power Industry (CARPI'12)*, Sept. 2012

[270] T. Bresciani, "Modelling, identification and control of a quadrotor helicopter," *Masters Thesis*, Lund University, 2008

[271] R. Mahony, V. Kumar and P. Corke, "Multirotor aerial vehicles: Modelling, estimation, and control of quadrotor," *IEEE Robotics & Automation Magazine*, vol. 19, no. 3 pp. 20-32, Aug. 2012

[272] M. Vanin, "Modelling, identification and navigation of autonomous air vehicles," *Masters Thesis*, Universita degli Studi di Padova, 2013

[273] P. Bristeau, E. Dorveaux, D. Vissière and N. Petit, "Hardware and software architecture for state estimation on an experimental low-cost small-scaled helicopter," *Control Engineering Practice* vol. 18, no. 7, pp. 733-746, July 2010

[274] T. Krajník, V. Vonásek, D. Fišer and J. Faigl, "AR-drone as a platform for robotic research and education," *Research and Education in Robotics EUROBOT 2011*, pp. 172-186, 2011

[275] L. Meier, P. Tanskanen, L. Heng, G. Hee Lee, F. Fraundorfer and M. Pollefeys, "PIXHAWK: A micro aerial vehicle design for autonomous flight using onboard computer vision," *Autonomous Robots*, vol. 33, no. 1-2, pp. 21-39, Aug. 2012

[276] P. Bristeau, F. Callou, D. Vissière and N. Petit, "The navigation and control technology inside the ar. drone micro uav," *IFAC World Congress*, vol. 18, no. 1, pp. 1477-1484, Aug. 2011

[277] Ascending Technologies, *http://www.asctec.de*, Germany

[278] H. Lim, J. Park, D. Lee and H. Jin Kim, "Build your own quadrotor: open-source projects on unmanned aerial vehicles," *IEEE Robotics & Automation Magazine*, vol. 19, no. 3, pp. 33-45, Sept. 2012

[279] K. Ogata, Modern Control Engineering, *Prentice-Hall*, 1995

[280] B. Stevens and F. Lewis, Aircraft control and simulation. *John Wiley & Sons*, 2003

[281] "PID vs LQ control techniques applied to an indoor micro quadrotor," *IEEE/RSJ Int. Conf. Intelligent Robotics and Systems (IROS'04)*, pp. 2451-2456, Sept. 2004

[282] M. Pena, E. Vivas and C. Rodriguez "Simulation of the quadrotor controller with LQR with integral effect," *ABCM Symposium Series in Mechatronics*, vol. 5, pp. 390-399, 2012

[283] P. Bouffard, A. Aswani and C. Tomlin, "Learning-based model predictive control on a quadrotor: onboard implementation and experimental results," *IEEE Int. Conf. Robotics and Automation (ICRA'12)*, pp. 279-284, May 2012

[284] P. Bouffard, "On-board model predictive control of a quadrotor helicopter: design, implementation and experiments," *Tech. Report*, Electrical Engineering and Computer Science, University of California, 2012

[285] M. Abdolhosseini, Y. Zhang and C. Rabbath, "An efficient model predictive control scheme for an unmanned quadrotor helicopter," *Journal of Intelligent Robotic Systems*, vol. 70, no. 1-4, pp. 27-38, April 2013

[286] W. Wahyudie, T. Bagus Susilo and H. Noura, "Robust PID controller for quadrotors," *Journal of Unmanned System Technology*, vol. 1, no. 1, 2013

[287] I. Mellado-Bataller, J. Pestana, M. Olivares-Mendez, P. Campoy, Pascual and L. Mejias, "MAVwork : a framework for unified interfacing between micro aerial vehicles and visual controllers," *Frontiers of Intelligent Autonomous Systems*, vol. 466, pp.165-179, 2013

[288] P. Corke, Visual Control of Robots: high-performance visual servoing, Research Studies Press, 1996

[289] K. Boyadzhiev, "Spiral and conchospirals in the flight of insects," *The college Mathematics Journal*, vol. 30, no. 1, pp. 22-31, Jan. 1999

[290] J. Lorimer, "Curved paths in raptor flight: deterministic model," *Journal of Theoretical Biology*, vol. 242, no. 4, pp. 880-889, Oct. 2006

[291] P. Salaris, L. Pallottino and A. Bicchi "Shortest path for finned, winged, legged and wheeled vehicles with side looking sensors," *Int. Journal of Robotics Research*, vol. 31, no. 8, pp. 997-1017, July 2012

[292] P. Chaumette and S. Hutchinson, "Visual servo control part II: advanced approaches," *IEEE Robotics and Automation Magazine*, vol. 14, no. 1, pp. 109-118, 2007

[293] P. Chaumette and S. Hutchinson, "Visual servo control part I: basic approaches," *IEEE Robotics and Automation Magazine*, vol. 13, no. 4, pp. 82-90, 2006

[294] P. Corke and S. Hutchinson, "A new partitioned approach to image-based visual servo control," *IEEE Trans. Robotics and Automation*, vol. 17, no. 4 pp. 507-515, Aug. 2001

[295] G. Mariottini and D. Prattichizzo, "Image-based visual servoing with central catadioptric cameras," *Int. Journal Robotics Research* vol. 27, no. 1, pp. 41-56, Jan. 2008

[296] F. Chaumette, "Potential problems of stability and convergence in image-based and position-based visual servoing," *Confluence of Vision and Control*, Springer London, pp. 66-78, 1998

[297] H. Hadj-Abdelkader, Y. Mezouar, N. Andreff and P. Martinet, "Omnidirectional visual servoing from polar lines," *IEEE Int. Conf. Robotics and Automation (ICRA'06)*, pp. 2385-2390, May 2006

[298] I. Sa and P. Corke, "Close-quarters quadrotor flying for a pole inspection with position based visual servoing and high-speed vision," *IEEE Int. Conf Unmanned Aircraft Systems (ICUAS'14)*, pp. 623-631, May 2014

[299] O. Tahri and F. Chaumette, "Point-based and region-based image moments for visual servoing of planar objects," *IEEE Trans. Robotics*, vol 21, no. 6, pp. 1116-1127, Dec. 2005

[300] R. Fomena and F. Chaumette, "Visual servoing from spheres using a spherical projection model," *IEEE Int. Conf. Robotics and Automation (ICRA'07)*, pp. 2080-2085, May 2007

[301] R. Fomena and F. Chaumette, "Improvements on visual servoing from spherical targets using a spherical projection model," *IEEE Trans. Robotics*, vol 25, no. 4, pp. 874-886, June 2009

[302] R. Fomena, O. Tahri and F. Chaumette, "Distance-based and orientation-based visual servoing from three points," *IEEE Trans. Robotics*, vol 27, no. 2, pp. 256-267, Feb. 2011

[303] M. Iwatsuki and N. Okiyama, "A new formulation of visual servoing based on cylindrical coordinate system," *IEEE Trans. Robotics*, vol. 21, no. 2, pp. 266-273, April 2005

[304] P. Corke, F. Spindler, and F. Chaumette, "Combining Cartesian and polar coordinates in IBVS," *IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS'09)*, pp. 5962-5967, Oct. 2009

[305] P. Corke, "Spherical image-based visual servoing and structure estimation," *IEEE Int. Conf. Robotics and Automation (ICRA'10)*, pp. 5550-5555, May 2010

[306] P. Corke and R. Mahony, "Sensing and control on the sphere," *Robotics Research*, vol. 70, pp. 71-85, 2011

[307] R. Mahony and T. Hamel, "Image-based visual servo control of aerial robotic systems using linear image features," *IEEE Trans. Robotics*, vol. 21, no. 2, pp. 227-239, April 2005

[308] N. Guenard, T. Hamel and R. Mahony, "A practical visual servo control for an unmanned aerial vehicle," *IEEE Trans. Robotics*, vol. 24, no. 2, pp. 331-340, April 2008

[309] O. Bourquardez, R. Mahony, N. Guenard, F. Chaumette, T. Hamel and L. Eck, "Image-based visual servo control of the translation kinematics of a quadrotor aerial vehicle," *IEEE Trans. Robotics*, vol. 25, no. 3 pp. 743-749, June 2009

[310] D. Lee, H. Lim, H. Jin Kim, Y. Kim and K. Jeong Seong, "Adaptive Image-Based visual servoing for an underactuated quadrotor system," *AIAA Journal of Guidance, Control, and Dynamics*, vol. 35, no. 4, pp. 1335-1353, 2012

[311] F. Le Bras, T. Hamel, C. Barat and R. Mahony, "Nonlinear image-based visual servo controller for automatic landing guidance of a fixed-wing aircraft," *European Control Conference (ECC'09)*, pp. 1836-1841, Aug. 2009

[312] L. Coutard, F. Chaumette and J. Pflimlin, "Automatic landing on aircraft carrier by visual servoing," *IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS'11)*, pp. 2843-2848, Sept. 2011

[313] D. Lee, T. Ryan and H. Jin Kim, "Autonomous landing of a VTOL UAV on a moving platform using image-based visual servoing," *IEEE Int. Conf. Robotics and Automation (ICRA'12)*, pp. 971-976, May 2012

[314] L. Mejias I. Mondragon and P. Campoy, "Omnidirectional bearing-only see and avoid for small aerial robots," *IEEE Int. Conf. Automation, Robotics and Applications (ICARA'11)*, pp. 23-28, Dec. 2011

[315] A. Mcfadyen and L. Mejias, "Visual servoing approach to collision avoidance for aircraft," *Int. Congress of the Aeronautical Sciences (ICAS'12)*, Sept. 2012

[316] E. Malis and P. Rives, "Robustness of image-based visual servoing with respect to depth distribution errors," *IEEE Int. Conf. Robotics and Automation (ICRA'03)*, vol. 1, pp. 1056-1061, May 2003

[317] C. Cheah and Y. Sun, "Region reaching control for robots with uncertain kinematics and dynamics," *IEEE Int. Conf. Robotics and Automation (ICRA'06)*, pp. 2577-2582, May 2006

[318] A. Remazeilles, N. Mansard and F. Chaumette, "A qualitative visual servoing to ensure the visibility constraint," *IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS'06)*, pp. 4297-4303, Oct 2006

[319] C. Cheah, D. Qun Wang and Y. Sun, "Region-reaching control of robots," *IEEE Trans. Robotics*, vol. 23, no. 6, pp. 1260-1264, ec.2007

[320] F. Schramm and G. Morel, "Ensuring visibility in calibration-free path planning for image-based visual servoing," *IEEE Trans. Robotics*, vol. 22, no. 4, pp. 848-854, Aug. 2006

[321] G. Chesi and Y. Hung, "Global path-planning for constrained and optimal visual servoing," *IEEE Trans. Robotics*, vol. 23, no. 5, pp. 22-31, Oct. 2007

[322] M. Kazemi, K. Gupta and M. Mehrandezh, "Path-planning for visual servoing: a review and issues," *Visual Servoing via Advanced Numerical Methods*, vol. 401, pp. 189-207, 2010

[323] P. Danes and D. Bellot, "Towards an LMI approach to multi criteria visual servoing in robotics," *European Journal of Control*, vol. 12, no. 1, pp. 86-110, 2006.

[324] T. Wang,W. Xie,G. Liu and Y. Zhao, "Quasi-min-max model predictive control for image-based visual servoing," *IEEE/ASME Int. Conf. Advanced Intelligent Mechatronics (ASMA'12)*, pp. 98-103, July 2012

[325] R. Findeisen and F. Allgower, "An introduction to nonlinear model predictive control," *21st Benelux Meeting on Systems and Control*, March 2002

[326] Y. Wang, H. Lang and C. De Silva, "An autonomous grasping system using visual servoing and nonlinear model predictive control," *IEEE Int. Conf. Robotics and Automation (ICRA'10)*, pp. 86-91, May 2010

[327] D. Lee, H. Lim and H. Kim, "Obstacle avoidance using image-based visual servoing integrated with nonlinear model predictive control," *IEEE Conf. Decision and Control (CDC-ECC'11)*, pp. 5689-5694, Dec. 2011

[328] A. Mcfadyen L. Mejias P. Corke and C. Pradalier, "Aircraft collision avoidance using spherical visual predictive control and single point features," *IEEE/RSJ Int. Conf. Robotics and Intelligent Systems (IROS'13)*, pp. 50-56, Nov. 2013

[329] G. Allibert, E. Courtial and Y. Touré, "Visual predictive control for manipulators with catadioptric camera," *IEEE Int. Conf. Robotics and Automation (ICRA'08)*, pp. 510-515, May 2008

[330] G. Allibert, E. Courtial and F. Chaumette, "Visual servoing via nonlinear predictive control," *Visual Servoing via Advanced Numerical Methods*, vol. 401, pp. 375-393, 2010

[331] H. Chen and F. Allgower, "Nonlinear model predictive control schemes with guaranteed stability," *Nonlinear Model Based Process Control*, vol. 353, pp. 465-494, 1998

[332] M. Sauvee, P. Poignet and E. Dombre, "Ultrasound image-based visual servoing of a surgical instrument through nonlinear model predictive control," *Int. Journal Robotics Research*, vol. 27, no. 1, pp. 25-40, Jan. 2008

[333] A. Burlacu, C. Copot, E. Vercera and C. Lazar, "Real time visual predictive control of manipulation systems," *IEEE Int. Conf. Advanced Robotics (ICAR'11)*, pp. 383-388, June 2011

[334] C. Copot, C. Lazar and A. Burlacu, "Predictive control of nonlinear visual servoing systems using image moments," *IET Control Theory & Applications*, vol. 6, no. 10, pp. 1486-1496, July 2012

[335] A. Burlacu and C. Lazar, "Reference trajectory-based visual predictive control," *Advanced Robotics*, vol. 26, no. 8-9, pp. 1035-1054, July 2012

[336] R. Coagienard, P. Grieder, E. Kerrigan and M. Morari, "Move blocking strategies in receding horizon control," *Journal of Process Control*, vol. 17, no. 6, pp. 563-570, July 2007

[337] G. Allibert E. Courtial and F. Chaumette, "Predictive control for constrained image-based visual servoing," *IEEE Trans. Automatic Control*, vol. 25, no. 5, pp. 933-939, Sept. 2010

[338] E. Ali and E. Zafiriou, "On the tuning of nonlinear model predictive control algorithms," *American Control Conf. (ACC'93)*, pp. 786-790, June 1993

[339] J.Garriga and M. Soroush, "Model predictive control tuning methods: a review," *Industrial & Engineering Chemistry Research*, vol. 49, no. 8, pp. 3505-3515, March 2010

[340] D. Mayne, and H. Michalska, "Receding horizon control of nonlinear systems," *IEEE Trans. Automatic Control*, vol. 35, no. 7, pp. 814-824, July 1990

[341] D. Mayne J. Rawling C. Rao and P. Scokaert, "Constrained model predictive control: stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789-814, June 2000

[342] M. Sauvée, P. Poignet, E. Dombre and E. Courtial, "Image based visual servoing through nonlinear model predictive control," *IEEE Int. Conf. Decision and Control*, pp. 1776-1781, Dec. 2006

[343] T. Murao, T. Yamada and M. Fujita, "Predictive visual feedback control with eye-in-hand system via stabilizing receding horizon approach," *IEEE Int. Conf. Decision and Control*, pp. 1758-1763, Dec. 2006

[344] H. Michalska and D. Mayne, "Robust receding horizon control of constrained nonlinear systems," *IEEE Trans. Automatic Control*, vol. 38, no. 11, pp. 1623-1633, Nov. 1993

[345] H. Chen and F. Allgower, "A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability," *Automatica*, vol. 34, no. 10, pp. 1205-1217, Oct. 1998

[346] J. Lofberg, "YALMIP: a toolbox for modelling and optimization in MATLAB," *IEEE Int. Symposium on Computer Aided Control System Design*, pp. 284-289, Sept. 2004

[347] B. Huska and H. Ferreau and M. Diehl, "ACADO Toolkit - an open source framework for automatic control and dynamic optimization," *Optimal Control Theory and Methods*, vol. 32, no. 3, pp. 298-312, May-June 2011

[348] P. Gill, W. Murray and M. Wright, Practical Optimization, *London, Academic Press*, 1981

[349] P. Boggs and J. Tolle, "Sequential quadratic programming," *Acta Numerica*, vol. 4, pp. 1-51, Jan. 1995

[350] J. Nocedal and S. Wright, Numerical Optimization, *Springer Verlag*, 2006

[351] P. Gill, W. Murray, M. Saunders and M. Wright, "Procedures for optimization problems with a mixture of bounds and general linear constraints," *ACM Trans. Math. Software*, vol. 10, no. 3, pp. 282-298, Sept. 1984

[352] P. Gill, W. Murray and M. Wright, Numerical Linear Algebra and Optimization, *Addison Wesley*, 1991

[353] S. Han, "A Globally Convergent Method for nonlinear programming," *Journal Optimization Theory and Applications*, vol. 22, no. 3, pp. 297, July 1977

[354] M. Powell, "A fast algorithm for nonlinearly constrained optimization calculations," *Numerical Analysis*, vol. 630, pp.144-157, Aug. 2006

[355] H.Ferreau, C. Kirches, A. Potschka, H. Bock and M. Diehl, "qpOASES: A parametric active-set algorithm for quadratic programming," *Mathematical Programming Computation*, vol. 6, no. 4, pp. 327-363, Dec. 2014

[356] B. Huska and H. Ferreau et al, ACADO Toolkit Users Manual, Version 1.0.2613beta, June 2011

[357] E. Ali and E. Zafiriou, "Optimization-based tuning of nonlinear model predictive control with state estimation," *Journal of Process Control*, vol. 3, no. 2, pp. 97-107, May 1993

[358] P. Viola and M. Jones, "Rapid object detection using a boosted cascade of simple features," *IEEE Int. Conf. Computer Vision and Pattern Recognition (CVPR'01)*, pp. 511-518, 2001

[359] M. Achtelik, S. Weiss and R. Siegwart, "Onboard IMU and monocular vision based control for MAVs in unknown in-and outdoor environments," *IEEE Int. Conf Robotics and Automation (ICRA'11)*, pp. 3056-3063, May 2011

[360] M. Barkat, Signal Detection and Estimation, *Artech House*, 2005

[361] L. Yang, J. Yang, E. Feron and V. Kulkarni, "Development of a performance-based approach for a rear-end collision warning and avoidance system for automobiles," *IEEE Intelligent Vehicles Symposium*, pp. 316-321, June 2003

[362] Australian Army, Adaptive Campaigning - Army's future land operating concept, *Tech. Report*, Canberra, Australia, 2009

[363] J. Kelly and M. Brennan, "OODA versus ASDA: Metaphors at war," *Australian Army Journal*, vol. 6, no. 3, pp.39-51, 2009