2016

Machine Vision Identification of Airport Runways With Visible and Infrared Videos

Andrew J. Moore

Matthew Schubert

Chester Dolph
Old Dominion University

Glenn Woodell

Original Publication Citation

Moore, A. J., Schubert, M., Dolph, C., & Woodell, G. (2016). Machine vision identification of airport runways with visible and infrared videos. *Journal of Aerospace Information Systems, 13*(7), 266-277. doi:10.2514/1.i010405

JOURNAL OF AEROSPACE INFORMATION SYSTEMS
Vol. 13, No. 7, July 2016

Machine Vision Identification of Airport Runways with Visible and Infrared Videos

Andrew J. Moore*
NASA Langley Research Center, Hampton, Virginia 23681
Matthew Schubert†
National Institute of Aerospace, Hampton, Virginia 23666
Chester Dolph‡
Old Dominion University, Norfolk, Virginia 23529
and
Glenn Woodell§
NASA Langley Research Center, Hampton, Virginia 23681

DOI: 10.2514/1.I010405

A widely used machine vision pipeline based on the Speeded-Up Robust Features feature detector was applied to the problem of identifying a runway from a universe of known runways, which was constructed using video records of 19 straight-in glidepath approaches to nine runways. The recordings studied included visible, short-wave infrared, and long-wave infrared videos in clear conditions, rain, and fog. Both daytime and nighttime runway approaches were used. High detection specificity (identification of the runway approached and rejection of the other runways in the universe) was observed in all conditions (greater than 90% Bayesian posterior probability). In the visible band, repeatability (identification of a given runway across multiple videos of it) was observed only if illumination (day versus night) was the same and approach visibility was good. Some repeatability was found across visible and shortwave sensor bands. Camera-based geolocation during aircraft landing was compared to the standard Charted Visual Approach Procedure.

Nomenclature

a,b	=	count of features in two images under comparison
e	=	time epoch; a sequential set of video frames in which a majority of reference images originating at a particular waypoint are matched
f_e	=	frames of epoch e
f_V	=	frames of video record V
M	=	match between a reference image and a video frame
n_e	=	number of epochs used to construct the universe of waypoint locations
n_V	=	number of videos used to construct the universe of waypoint locations
$P(e\|R)$	=	epoch-match probability; probability of waypoint location corresponding to epoch e given matches on reference image R
$P(R)$	=	probability of matching a particular reference image R
$P(R\|e)$	=	conditional match-epoch probability
$P_t(e)$	=	probability of location at waypoint corresponding to epoch e at time t
R	=	reference image
S	=	subframe of the waypoint image used to create a reference image
t	=	time, s
W	=	waypoint
V	=	specific video record

I. Introduction

MODERN commercial aircraft are routinely equipped with video sensors; both visible and infrared cameras may be employed, as the infrared bands provide extended visibility through atmospheric fog and rain. Modern machine vision algorithms are capable of rapid image identification, within limits. It is conceivable that runway descent assisted by machine vision, analogous to currently practiced human visual descent procedures, may be possible by processing visible and infrared cockpit videos with machine vision algorithms.

In this study, we examined the efficacy of a contemporary open-source image matching pipeline to determine which runway was in the field of view using Bayesian probability methods. For visible short-wave infrared (SWIR) and long-wave infrared (LWIR) videos, cockpit imagery was compared to a universe of stored runway images using the speeded-up robust features (SURF) feature detector, brute force matching, and geometric constraint checking using random sample consensus (RANSAC). Two metrics were used in this examination: detection specificity (identification of the runway approached and rejection of the other runways in the universe) and repeatability (identification of a specific runway across multiple videos of that runway, and rejection of the other runways in the universe). High detection specificity was found in all conditions

Received 16 June 2015; revision received 7 April 2016; accepted for publication 11 April 2016; published online Open Access 20 June 2016.

*Aerospace Technologist, Electromagnetics and Sensors Branch, 8 North Dryden Street.
†Research Scientist, 100 Exploration Way.
‡Doctoral Student, Vision Laboratory, 2107 Engineering Systems Building.
§Aerospace Technologist, Electromagnetics and Sensors Branch, 8 North Dryden Street.

(greater than 90% Bayesian posterior probability). In the visible band, repeatability was observed only if illumination (day vs. night) was the same and approach visibility was good. Measurable repeatability was observed between the visible and SWIR bands.

To test detection specificity, flight video from a single descent at each runway is all that is needed, but to test repeatability, flight video from multiple descents is required. This study benefits from a large dataset compared to previous published tests of runway identification (19 flight video recordings of final approach to nine runways), but it is still sorely limited: of the nine runways, repeated flights with the same (visible) sensor are available only at the same time of day for two runway descent glidepaths. We have enough data to arrive at confident conclusions regarding runway identification specificity, but our conclusions regarding repeatability must be considered tentative. In light of these results, comparisons are made between the new waypoint-based method for geolocation via machine vision and a standard FAA (Federal Aviation Administration) landmark-based visual landing method, the Charted Visual Flight Procedure (CVFP).

As machine vision techniques have developed in recent decades, those techniques have been applied in turn to improve the detection of runways and/or identification of specific runways from onboard video. Although not exhaustive, this progressive advancement in runway image processing is reviewed briefly in the following. Huertas et al. used localized edge detection and thresholding in 1990 to outline runways and create location-specific image templates for use in an expert system [1]; see also [2]. Fleming and his collaborators reviewed the literature through 2004 [3] (including work that harnessed a key innovation: the Hough transform) and applied the Hough transform to runway imagery, using stereo ranging to estimate the airplane viewpoint and landing distance [4]. Independently and contemporaneously, Shang and Shi [5] took a similar approach, using monocular perspective analysis instead of stereo analysis to estimate the landing geometry. In his thesis, Zongur [6] added a machine learning layer to previously applied techniques to recognize airports from orbital imagery. Pai et al. at USC and Honeywell, Inc. [7] used a new class of robust feature detector [8] and homographic perspective transformation to track runways in flight video; with image stabilization and image differencing, they could determine whether a runway was free of hazardous objects during the landing approach. Their application of a scale-invariant front end (SIFT) in 2009 represented a qualitative improvement in robustness in the application of machine vision to the processing of onboard video of a runway approach (we use a performance-optimized variant, SURF [9], in this study). Lastly, to address a problem related to runway detection (namely, the detection of safe landing areas), the machine vision methods of texture segmentation [10] and component trees [11] have been studied.

II. Methods

A. Runway Approach Video

For nine runways under various flight conditions, we obtained visible- and infrared-band videos of aircraft landing approaches along a straight-in glidepath under various flight conditions (resolution: visible, 1000×700 pixels or better; infrared: 640×480). The video frame rate was at least 30 frames per second; video was either uncompressed or compressed via Moving Picture Experts Group (MPEG) or equivalent at the time of capture at a low enough compression level to avoid obvious artifacts. Each video was annotated with seven variables to capture the unique aspects of the flight: 1) airport, using a standard designation as described below; 2) runway number, using the two-digit FAA heading designation; 3) sensor type (visible, short wave, or long wave); 4) time of day (day or night); 5) season (summer, winter, fall, spring, snow cover); 6) weather condition (clear, foggy, rain, snow, shadow); and 7) direction of approach (one- or two-letter compass direction, e.g., NW for approach from the northwest). The airports used in the study and their ICAO (International Civil Aviation Organization) designations are: 4G7 (Fairmont WV, USA), EDDF (Frankfurt, DE), KAVC (South Hill VA, USA), KCID (Cedar Rapids, IA, USA), KCKB (Clarksburg WV, USA), KLAS (Las Vegas NV, USA), KLFI (Hampton VA, USA), and KPHF (Newport News VA, USA).

Table 1 summarizes the dataset used in this study. Although this dataset includes several repeated approaches under varying conditions, there is not a single runway for which we have recordings of all of the tested conditions (e.g., day vs night, summer foliage vs snow, clear vs foggy). The abbreviated names in the approach video column include (in addition to the runway designator) annotations sufficient to uniquely specify each record. In the band column, "SW" and "LW" denote short-wave and long-wave infrared if the video was not captured in the visible band.

Superfluous footage at the beginning and end of the approach was removed from the raw approach recordings to produce clips that, at most, captured a continuous cockpit view beginning from 3–5 nmi distant from the runway through to approximately 1000 ft from the runway threshold. The starting distance of 3–5 nmi was chosen because it was the furthest distance at which the runway portion of the video image extended across several pixels, and therefore could be processed for gradient-based image feature detection. The ending distance of 1000 ft from the runway edge was chosen so that the video contained the "go/no-go" distance, or minimum descent altitude for a precision approach. The go/no-go distance was defined by a decision height of 200 ft [12] from the threshold of a runway that, at a typical glideslope descent angle of 3 deg, equated to approximately a 2400 ft distance from threshold (typically 3800 ft distance from the touchdown point). Since the touchdown zone was within 2800 ft of the runway edge, a video ending distance of 1000 ft from the runway edge was guaranteed to include the go/no-go point.

The starting distance of 3–5 nmi is within or equal to the final approach segment distance of each of the runways in this study (approximately 5 nmi, except for Frankfurt at approximately 11 nmi), and so geolocation via machine vision processing during final approach can be realistically studied.

B. Waypoint and Reference Image Selection

1. Waypoint Separation Considering Out-of-Plane Perspective Change

Selected frames ("waypoints") of the edited video were chosen as the origin of the reference images (Fig. 1, left). Waypoints were spaced along the landing descent glidepath so that the change in visual perspective from one waypoint to the next was within the out-of-plane rotation limitations of the machine vision employed. The matching ability of the SIFT/SURF class of image matching algorithms was known to degrade significantly with rotation out of the image plane due to geometric distortion (e.g., foreshortening); both Lowe [8] and Bay et al. [9] noted a marked drop in match repeatability for out-of-plane rotation angles greater than 30 deg. The viewpoint perspective of the runway and its surroundings, as seen from the cockpit, varied gradually during descent; correspondingly, the out-of-plane rotation angle between a reference image viewpoint and the instantaneous viewpoint changed smoothly as the airplane descended through the waypoint (e.g., the images at right in Fig. 1). Past a certain angle, matching degraded, and a new waypoint and reference image were needed for a successful match.

This constraint dictated our waypoint selection method. Referring to the left of Fig. 1, a frame was selected as waypoint 1 shortly after the start of the landing video. As the video progressed, the distance to the runway decreased, and the (positive) out-of-plane rotation angle between the viewpoint of reference image 1 and the instantaneous viewpoint increased until matching began to degrade. A frame for waypoint 2 was selected thereafter, such that the (negative) out-of-plane rotation angle between the viewpoint of reference image 2 and the instantaneous viewpoint was small enough for good matching to reference image 2 before matching to reference image 1 degraded. Matching to reference image 2 was then stable until descent progressed to the point that the (positive) out-of-plane rotation angle between the viewpoint of reference image 2 began to degrade. A frame for waypoint 3 was selected thereafter such that it was stably detectable (at a negative out-of-plane rotation angle) before reference image 2 detection degraded, and so on. Although this waypoint assignment technique was applied manually in this study, it was systematic and may have been amenable to automation.

Table 1 19 runway approach video records used in this study [a]

Approach video	Runway	Band	Time	Weather	Waypoints	References/ waypoint	Epochs (no filter)	References/ epoch (no filter)	Epochs (high pass)	Δ references/ epoch (high pass)
4G7_23_night	1	— —	Night	Clear	3	4,4,4	3	2,3,4	3	— —
EDDF_25L_night	2	— —	Night	Clear	3	4,4,4	3	4,4,4	3	— —
KAVC_01_fog_S	3[b]	— —	Day	Fog/ rain	3	2,2,2	3	1,1,2	3	2,2,2
KAVC_01_fog_SE	3[b]	— —	Day	Fog/ rain	2	1,1	2	1,1	2	— —
KCID_27_LW	4[b]	LW	Day	Fog/ snow	6	4,4,4,4,4,4	6	4,4,4,4,4,4	6	— —
KCID_27_SW	4[b]	SW	Day	Fog/ snow	6	4,4,4,4,4,4	6	4,4,4,4,4,4	6	— —
KCID_27_VIS	4[b]	— —	Day	Fog/ snow	8	4,4,4,4,4,4,4,4	8	1,3,4,4,4,4,4,4	8	4,4,4,4,4,4,4,4
KCID_27_LW_night	4[b]	LW	Night	Clear	5	4,4,4,4,4	5	4,4,4,4,4	5	— —
KCID_27_SW_night	4[b]	SW	Night	Clear	4	4,4,4,4	4	4,4,4,4	4	— —
KCID_27_VIS_night	4[b]	— —	Night	Clear	5	4,4,4,4,4	5	3,3,4,4,4	5	4,4,4,4,4
KCKB_21_night	5	— —	Night	Clear	5	4,4,4,4,4	5	1,1,4,4,4	2	1,1
KLAS_25L_LW_night	6[b]	LW	Night	Clear	4	4,4,4,4	3	4,4,4	4	4,4,4,1
KLAS_25L_SW_night	6[b]	SW	Night	Clear	7	4,4,4,4,4,4,4	7	4,4,4,4,4,4,4	7	— —
KLAS_25L_VIS_night	6[b]	— —	Night	Clear	6	4,4,4,4,4,4	6	4,4,4,4,4,4	6	— —
KLFI_08_S	7[b]	— —	Day	Clear	3	4,4,4	3	4,4,4	3	— —
KLFI_08_W	7[b]	— —	Day	Clear	4	4,4,4,4	4	4,4,4,4	4	— —
KLFI_08_fog_NW	7[b]	— —	Day	Fog/ rain	3	1,1,3	1	3	3	1,1,3
KLFI_26_rain	8	— —	Day	Fog/ rain	5	4,4,4,4,4	4	2,3,4,4	5	2,4,4,4,4
KPHF_25	9	— —	Day	Clear	3	4,4,4	3	3,4,3	3	4,4,4
Total	9	— —	— —	— —	85	321	81	289	82	295

[a]The six columns to the right show the number of waypoint assignments, the number of epochs (video frame blocks) defined by imagery from the waypoints, and the reference image count. Image enhancement increases the number of usable reference images in some cases; for those cases, the rightmost two columns restate the number of epochs and the increased reference image count.

[b]For four runways, there are multiple videos.

2. Constructing a Universe of Waypoint-Based Locations

For the collection of 19 videos, the number of waypoints W ranged from two to eight (Table 1). Up to four reference images R were selected by outlining up to four subframes S at each waypoint: one centered on the runway, one centered on the landscape to the left of the runway, one centered on the landscape to the right of the runway, and one approximately included the three other reference images. (For some waypoints, fewer than four subframes were selectable: in some cases, the runway was not centered in the frame; in other cases, a substantial part of the frame was obscured by fog, e.g., the KAVC_01 videos). After reference image selection at all waypoints of all videos, the universe of visual references images could be represented as a matrix indexed by the video of origin v, the waypoint w, and the subframe s:

$$\boldsymbol{R}(\boldsymbol{v},\boldsymbol{w},\boldsymbol{s}) = \boldsymbol{R}(\boldsymbol{V},1:\boldsymbol{W}_{\boldsymbol{V}},1:\boldsymbol{S}_{\boldsymbol{W}}); \qquad 1 \le \boldsymbol{V} \le 19; \qquad 2 \le \boldsymbol{W}_{\boldsymbol{V}} \le 8; \qquad 1 \le \boldsymbol{S}_{\boldsymbol{W}} \le 4 \tag{1}$$

Although 321 images were potential candidates for membership in the universe reference image set, image quality was so poor in some subframes at some waypoints, due to foggy conditions, that matching was not possible. For unenhanced video, 289 of the 321 candidates could be matched; high-pass filtering increased the usable image set to 295.

An example set of reference images at one waypoint is shown at the bottom right of Fig. 1. Each reference image file name was annotated to track the indices v, w, and s: the image name was appended with the name of its video file of origin V, the number of the frame in video V selected for its waypoint of origin W_V, and the subframe position S_W of the reference bounding box center relative to the runway (center, left, right, and full).

C. Image Matching

A Linux-based computer with 16 CPUs (32 logical cores; Xeon E5; 3.1 GHz clock speed), four GPUs (graphics processing units; Nvidia Tesla K20X), and 128 GB of RAM was used for all processing. The OpenCV open-source computer vision bundle was used for most of the image processing and recognition modules [13]. Many modules were accelerated using Nvidia CUDA bindings in OpenCV.

After waypoints were defined for all videos, image matching was attempted for each frame of each landing approach, using the entire set of reference images as candidate matches. Matching result statistics for each video file included a binary vector for each frame, indicating whether a match was found in the frame to one or more of the 321 reference images in the universe.

Reference image matching to each frame proceeded in four steps: 1) color transformation of the original video color space to a YUV (luminance-chrominance) space to extract the approximately grayscale Y luma channel if the video was not grayscale; 2) identification of image features of each frame and each reference image using the multiscale SURF feature detector [9] on the gray (or luma) channel; 3) feature matching

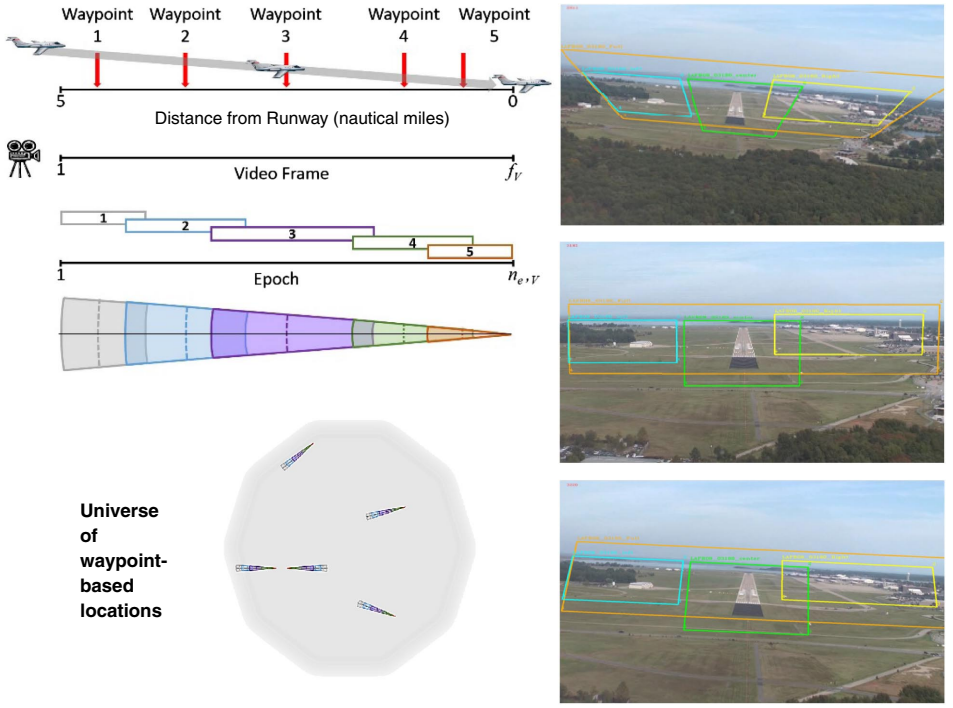

Fig. 1 On the left are shown a selection of waypoints and the definition of epochs. Waypoints are spaced to avoid extreme perspective change; epochs, defined from reference imagery at the waypoint location, are overlapping ensemble blocks of frames to avoid bias arising from temporal continuity. At the bottom is a notional depiction of a universe of waypoint-based locations along the glidepaths to known runways. On the right are shown example frames from KLFI_08_W a few frames before (top), at (middle), and after (bottom) the waypoint frame, showing the change in perspective through an epoch.

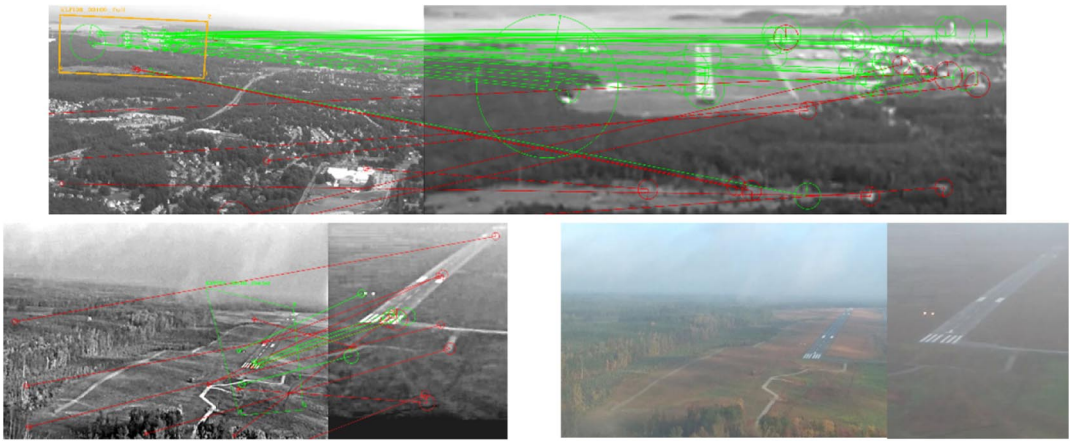

Fig. 2 Matching examples. Matching features (circles) are connected by lines; indicated are matches classified as inliers and outliers by the RANSAC algorithm. At the top is a video frame (left) from approach KLFI_08_S that is compared to a stored reference image from the same approach (right). A portion of the video frame (upper-left bounding box projection) is strongly matched, demonstrating waypoint-based location specificity. No image enhancement is applied. At the bottom left is cross matching between two landing videos in conditions of light fog, with image enhancement on both the video frame and reference image. A portion of the video frame (bounding box projection) from landing approach KAVC_01_fog_S is matched to a reference taken from a different approach recording (KAVC_01_fog_SE) at the same runway, demonstrating repeatability. At the bottom right is shown the same imagery as the bottom left without image enhancement. Cross matching is not observed due to low contrast. Reference images are scaled up by four times in this depiction.

between the frame and reference, as described in the following; and 4) geometric validation of matched features by attempting to project the bounding box of the reference image onto a bounding box on the video frame and rejecting unlikely perspective changes [14]. After matching, the frames were reassembled into a video with the frame number and the colored matching box superimposed (e.g., images at right of Fig. 1) for results presentation and inspection.

Color transformation (step 1) from the red green blue (RGB) to the Y luma channel was computed as $Y = 0.299 * \mathrm{R} + 0.587 * \mathrm{G} + 0.114 * \mathrm{B}$. Optionally, image enhancement was applied via high-pass filtering (unsharp masking, retinex) [15–17] by subtracting three low-pass versions of the image: short range, midrange, and long range (5-, 20-, and 240-pixel Gaussian sigma, respectively). The three subtracted low-pass images were each weighted with a value of 0.333, and the resulting high-pass image was renormalized to an eight-bit dynamic range. Images at the bottom of Fig. 2 show examples of matching a frame to a reference image before (right) and after (left) high-pass filtering.

Feature detection (step 2) was performed using the OpenCV 2.4 GPU implementation of SURF using its default parameters (64-bit descriptor, four octaves of multiscale processing) and a midrange Hessian threshold of 400.

Feature matching (step 3) was accomplished by comparing b features on the video frame with a features on the reference image ("$b \times a$ brute force matching"), using the L2 norm of the b feature descriptor vectors to find and rank the two best matches in the pool of a features. As a crosscheck, a second "reverse" $a \times b$ comparison was performed to find, for each of the a reference image features, the closest single match among

the b features on the video frame. (This asymmetry exists because the reference image is almost always a subset of the full video frame.) Highest-ranking matches from the first check were used only if they were validated by the second check. For example, if frame feature $b1$ had two matching features $a1$ and $a2$ in the first comparison, the $< b1, a1 >$ match pair was validated only if reference image feature $a1$ had matching frame feature $b1$ in the second comparison. Additionally, the $< b1, a1 >$ match pair was rejected if its L2 norm was not sufficiently unique. In particular, if the L2 norm of $<b1, a2>$ was within 30% of the L2 norm of $< b1, a1 >$, then the match was discarded; this greatly reduced spurious matches (e.g., to textured areas such as foliage). The 30% heuristic was originally determined by Lowe [8] and was carefully validated in our study.

Finally, the matches were validated geometrically (step 4). The matches were used to find a homography from the reference image to the video frame using RANSAC with the default OpenCV reprojection threshold distance of three pixels [18]. The homography was used to project the corners of the reference image onto the video frame, creating a box bounding the reference imagery in the video frame. This bounding box was tested against logical geometric constraints; flipping, mirroring, extreme-angled parallelogram projections, and "bowtie" projections were rejected. Figure 2 shows two examples of observed feature matching.

D. Determining Waypoint-Based Location Using Statistical Ensembles of Video Frames

1. *Motivations for Decomposition into Ensemble Blocks of Frames*

To determine airplane location via comparison of real-time cockpit video with prerecorded "ground truth" video, one can match each incoming frame with each prerecorded frame and compute the most likely location from the most likely match. This brute force approach has two disadvantages. First, it relies on the naïve assumption that each video frame is an independent experiment, which overestimates probability due to frame-to-frame similarity (temporal continuity). Second, it is compute intensive, scaling badly with the number and length of ground truth videos. In this section, we describe a method that addresses both of these disadvantages.

To avoid the statistical bias of the naïve frame independence assumption, we used an ensemble statistical method that grouped blocks of contiguous frames of each video V into time epochs e. Each epoch within a video was defined by a sequential set of frames in which a majority of reference images originating at a particular waypoint W_V were matched. In this ensemble statistical method, the assumption was that an epoch, rather than a frame, was an independent experiment [19]. Also, by decomposing each video into waypoint-based epochs, computation could be performed up front (preflight) to produce an epoch-based ensemble probability model. During flight, a relatively compact conditional Bayesian probability model could be used, in conjunction with the reference images used to create the model, to efficiently estimate the probable waypoint-based location of the airplane. The cost of improved scalability (besides the upfront computation) was that the precision of the universe of possible geographical locations was limited by the number and geographical distribution of waypoints.

Figure 1 illustrates epoch decomposition for a runway approach video V with f_V frames. Five frames are chosen as waypoints $w \in W_V = 1, 2, \ldots 5$ along the flight path, and reference images from up to four subframes S_W are chosen at each of the five waypoints. Due to temporal continuity, if the video's frames are compared with all its reference images $R(v, w, s) = R(1, 1:5, 1:S_W)$, a set of five epochs corresponding to the five waypoints is found, with each reference matched for a range of frames, starting before the frame of its waypoint of origin and continuing past it. The change in the shape of the subframe match bounding boxes (Fig. 1, right) is expected from the change in perspective during egomotion [20] and can be predicted with formal optical flow analysis [21].

2. *Epoch Definition*

Similarly, each of the 19 flight videos can be decomposed into epochs (time-contiguous ensembles of frames) by matching all of the video's frames only to reference images from the video's waypoints, i.e., matching f_V frames to $R(v, w, s) = R(V, W, S_W)$ for $W = 1, 2, \ldots W_V$. To algorithmically determine the first and last frames of the epoch corresponding to waypoint W, we compute a binary vector

$$\boldsymbol{M}_f = \begin{cases} 1, \text{if there is a match to reference image } R(\boldsymbol{v}, \boldsymbol{w}, \boldsymbol{s}) = R(\boldsymbol{V}, \boldsymbol{W}, 1{:}\boldsymbol{S}_W) \text{ in frame } f \\ 0, \text{otherwise} \end{cases} \tag{2}$$

The median of the frame range of nonzero matches was found, and then medians of the upper and lower halves were bisected, yielding four quartiles. The upper and lower quartiles were modified by standard Tukey outlier rejection techniques [22]; this resulted in an epoch that did not include many sparse matches while reasonably containing the bulk of the matches. Figure 3 illustrates decomposition of the SWIR landing video, KCID_27_SW_night, into four epochs.

Some waypoint frames have poorly matched reference images due to low-visibility weather conditions. As a result, not all of the 85 defined waypoints (unfiltered video: 81 epochs; enhanced video: 82 epochs) and 321 reference images (unfiltered video: 289 images; enhanced video: 295 images) were usable to define the universe of epoch-based locations (bottom row of Table 1).

3. *Statistical Modeling and Bayesian Waypoint-Based Location Probability*

After video decomposition into epochs, an epoch-based ensemble probability model is constructed by matching, per epoch, the f_e frames of each epoch e to all reference images $R(v, w, s) = R(V = 1{:}19, W_V, S_W)$ from all waypoints and all videos. Some of the f_e frames match (by definition) reference images from the epoch's video and waypoint of origin; they may also match reference images from other videos and waypoints. For example, epoch 2 of Fig. 1 will contain matches mostly to waypoint 2 images but also some matches to waypoints 1 and 3 images. Although only matches to reference images of a given waypoint are used to define an epoch [Eq. (2)], matches of the epoch's frames to other reference images in the universe capture the coincidental ability of the ensemble frame block to estimate the probability of location at other waypoints in the universe. Again, precomputing can reduce the computational burden in real-time image stream machine vision image processing.

The conditional match-epoch probability $P(R|e)$ of a match M from a frame f in the block of f_e frames to any reference image in the universe is the match count weighted as a fraction of ensemble size f_e:

$$\boldsymbol{P}(\boldsymbol{R}|\boldsymbol{e}) = \frac{\sum_{f=1}^{f_e} \boldsymbol{M}_i}{\boldsymbol{f}_e}; \qquad \boldsymbol{M}_i = \begin{cases} 1, \text{if there is a match to } R = R(\boldsymbol{V}, \boldsymbol{W_V}, 1{:}\boldsymbol{S}_W) \text{ in epoch frame } i \\ 0, \text{otherwise} \end{cases} \tag{3}$$

This is the most compute-intense step in the method, as all frames of all epochs are compared to all reference images from all waypoints of all videos. Assuming that all n_e epochs are equally probable (see the following), the probability $P(R)$ of matching a particular reference image R across all epochs of all videos in the universe can be precomputed as

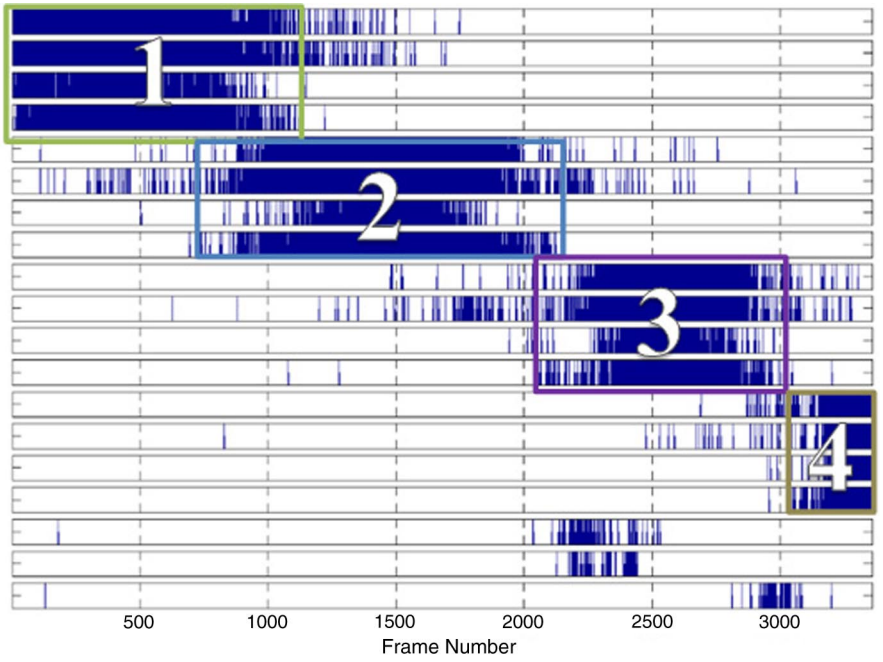

Fig. 3 Quartile detection of epochs for the short-wave infrared night approach video KCID_27_SW_night. All 16 reference images from four waypoints match imagery in the 3359 frames, as indicated by a vertical bar in the top 16 subplots. The bottom three subplots show matches against KCID_27_VIS_night, which is a visible-band video of the same runway approach.

$$\boldsymbol{P(R)} = \frac{\sum_{e=1}^{n_e} \boldsymbol{P(R|e)}}{\boldsymbol{n}_e} \tag{4}$$

Since the match-epoch probability $P(R|e)$ is known, the unconstrained match probability $P(R)$ is known, and the epoch probability $P(e)$ is assumed to equal $1/n_e$, it is possible to precompute the epoch-match probability $P(e|R)$ that a match to image R determines the waypoint-based location corresponding to epoch e as

$$\boldsymbol{P(e|R)} = \frac{\boldsymbol{P(R|e)P(e)}}{\boldsymbol{P(R)}} \tag{5}$$

We wish to estimate the airplane location given the reference image matches M at any time on any of the n_V videos. Since each epoch $e = 1, 2, \ldots n_e$ is defined from imagery at a specific waypoint location

$$w = \{1, 2, \ldots \sum_{v=1}^{n_V} W_v\}$$

estimation of the epoch probability via matching to imagery used to define the epochs is an estimation of the waypoint-based location probability with a confidence greater than chance ($1/n_e$). From images visible out the cockpit window at time t, the waypoint-based location probability $P_t(e)$ can be determined by matching to the universe of k reference images $R_1, R_2, \ldots, R_k$ and weighting each match by the epoch-match probabilities $P(e|R)$:

$$\boldsymbol{P}_t\boldsymbol{(e)} = \left[1 - \prod_{i=1}^{k} (1 - \boldsymbol{P(e|R}_i) \cdot \boldsymbol{M}_t)\right] \tag{6}$$

where

$$\boldsymbol{M}_t = \begin{cases} 1, \text{if there is a match at time } t \text{ to the } i\text{th reference image } R \\ 0, \text{otherwise} \end{cases} \tag{7}$$

It is important to note that the assumption of a uniform prior probability assignment $P(e) = 1/n_e$ is a simplification, based on approximately equally-spaced waypoints and the approximately equal number of "ground truth flights" per runway in the universe. The effect of airplane velocity is removed by the normalization to the number of frames in Eq. (3), but care must be taken in this method to ensure approximate waypoint equidistance and an approximately equal number of ground truth flights per sensor per runway. To illustrate the latter with an extreme example, if the universe was constructed with 1000 ground truth flights to runway A but just a single ground truth flight to runway B, the uniform prior estimate for $P(e)$ would be severely distorted. Regarding the assumption of equidistant waypoints, we took care to enforce adherence to this assumption in our (manual) creation of the universe. Regarding the assumption of an approximately equal number of ground truth flights per runway, approximate adherence was enforced by the input data distribution. On a sensor-by-sensor basis, our data (19 videos, three sensor types, and nine runways) were not heavily skewed toward any particular runway. Within the accuracy of these approximations, Eq. (6) gives an estimate of airplane geographic location within a universe comprising

$$w = \left\{1, 2, \ldots \sum_{v=1}^{n_V} W_v\right\}$$

waypoint locations.

III. Results

A. Histogram Representation of Framewise Probability

A matrix plot of the classification of flight records across the universe of runways requires a compact representation of match results for all of the frames of each flight record. In this section, we describe a histogram representation of framewise probability with two examples.

Figure 4 shows frame-by-frame plots (left) of the most probable epoch (waypoint-based location) for two flight records and the runway-based location probability histograms (right) derived from them. The frame-by-frame plots show the most probable epoch at each instant in the flight record: for each frame [time t in Eq. (6)], all matches between the frame image and the universe of reference images $P(R|e)$ are found, epoch probabilities $P_t(e)$ for all n_e epochs are computed, and the epoch with maximum probability is found. (In the rare case of multiple maxima, the last epoch in our alphabetically sorted list of epochs is used.) The waypoint-based location corresponding to the epoch with maximum probability is the best posterior estimate of the geographic location for each instant of time (video frame) in the frame-by-frame plots of Fig. 4; it is used as an intermediate result for specificity plotting in Fig. 5 and as an end result for repeatability plotting in Fig. 6.

To capture both the confidence and relative duration of location estimation, the frame-by-frame epoch results are sorted and binned into runway-based location (Fig. 4, right; Fig. 5) or video-specific waypoint-based location (Fig. 6) probability deciles as follows:

1) For each frame, the name of the most probable epoch is converted to its waypoint of origin and then to its runway of origin (Fig. 5) or video record of origin (Fig. 6) by table lookup.

2) The frame probability is tallied separately, by decile, for each lookup result. Frames with no matches or weak matches [$P(e_t) <= 0.1$] are converted to a null lookup result called "no signal" and tallied separately.

This compact representation allows for a rendition of the results across the entire universe of flight records in Figs. 5 and 6. Due to the limitations of graphical rendering, the histogram bar width cannot be drawn faithfully for low frame counts in Figs. 5 and 6. The histogram axis color is used to mitigate this: the axis is drawn in black if the percentage of frames in any bin exceeds 1.5% (e.g., the top right EDDF_25L results histogram in Fig. 4) and is colored gray otherwise (e.g., the empty EDDF_25L no signal histogram in Fig. 4).

There are three reasons that a flight video frame may have no or low probability of location (no signal) in this method. By far, the most common reason is that the frame is captured in poor atmospheric visibility conditions and has low contrast and few image features. The second most common reason is that a rapid yaw/pitch/roll change as the aircraft is buffeted by wind causes image smearing in the video sensor, transiently distorting image features. These first two reasons are discussed further in the following. Third, such a frame may arise due to a methodological issue: the location in view has no assigned epoch. For example, during the first approximately 60 frames of flight video KPHF_25 (lower plots of Fig. 4), the aircraft is banking such that the runway is not in the camera field of view and no epoch is defined: there is literally no defined location for the portion of the universe in view. We assign epochs to minimize this methodological reason for a no-signal result.

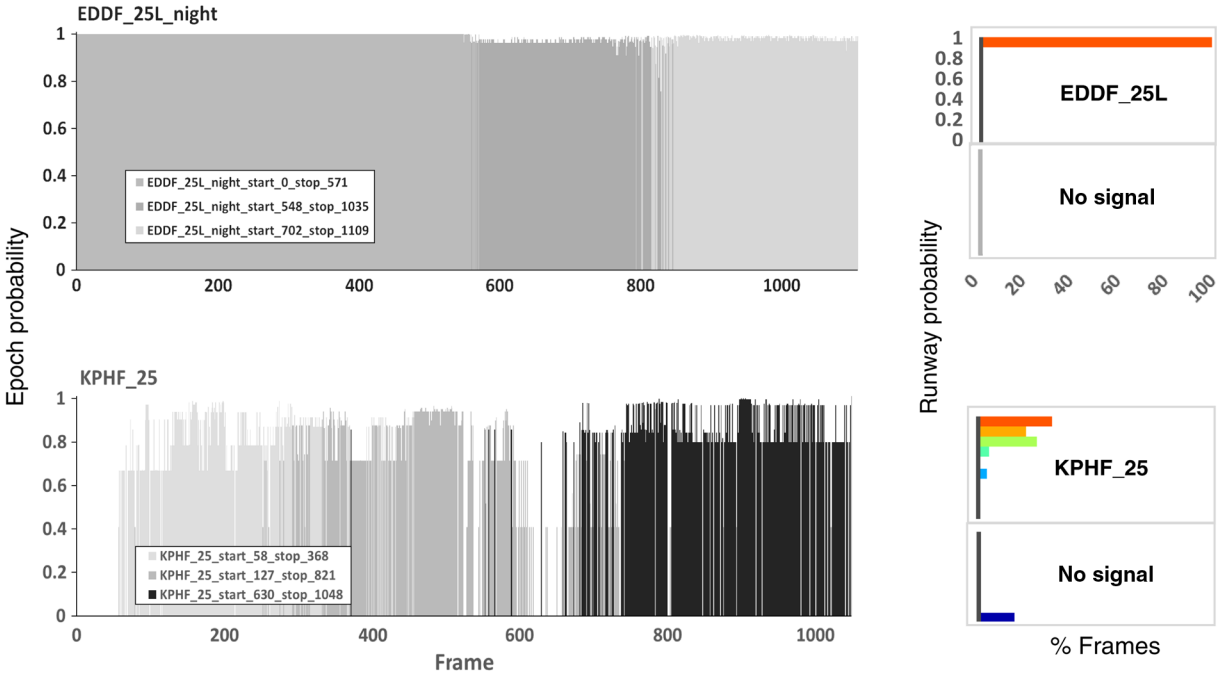

Fig. 4 Histogram summary of waypoint-based location probability. The maximum epoch probability per frame is plotted at left, and runway location probabilities are tallied in the histograms at right. Frames with low/no (≤ 0.1) probability are tallied in the no-signal histogram, whereas all others are assigned to runway location probability bins. The top left shows the framewise results for video EDDF_25L_night. Across all reference images in the universe, matches are found only to three epochs (inset) that correspond to waypoints in this video record. Very strong (greater than 90%) posterior epoch probability is found for nearly all frames. The top right shows the histogram of the framewise probability locating the aircraft in the EDDF_25L glidepath with a 0.9–1.0 probability bin for nearly 100% of the video frames in the record. The bottom left shows the framewise results for video KPHF_25. Matches are found only to epochs that correspond to waypoints in this video record for most frames, with probability ranging from 0.4 to 1.0. The airplane is not in the glidepath of this runway during the first part of the record, and its pitch, yaw, and roll shift during the landing approach. The bottom right shows, accordingly, that the histogram of runway-based location probability locates the aircraft in the KPHF_25 descent glidepath with a range of probabilities for the majority of frames, and it locates it in the null location of no-signal for the remaining 15% of the frames.

Fig. 5 Specificity of runway glidepath identification. Probability histograms are shown without image enhancement; if image enhancement notably changes the results, the result is highlighted with a gray box and the corresponding result from the image enhancement experiment is overlaid to its lower left in a gray dashed box. For each approach video, the probable runway glidepath location of the aircraft is aggregated for each of the nine runways according to the legend at top left. In the left columns, the histogram bar width varies as the percentage of frames in which the most probable ($p > 0.1$) runway glidepath location is detected, and the histogram bar color corresponds to the probability decile, as per Fig. 4. The right column shows the percentage of frames with no ($p \leq 0.1$) significant location probability. The axis shade is black if the percentage of frames in any bin exceeds 1.5%, and it is gray otherwise.

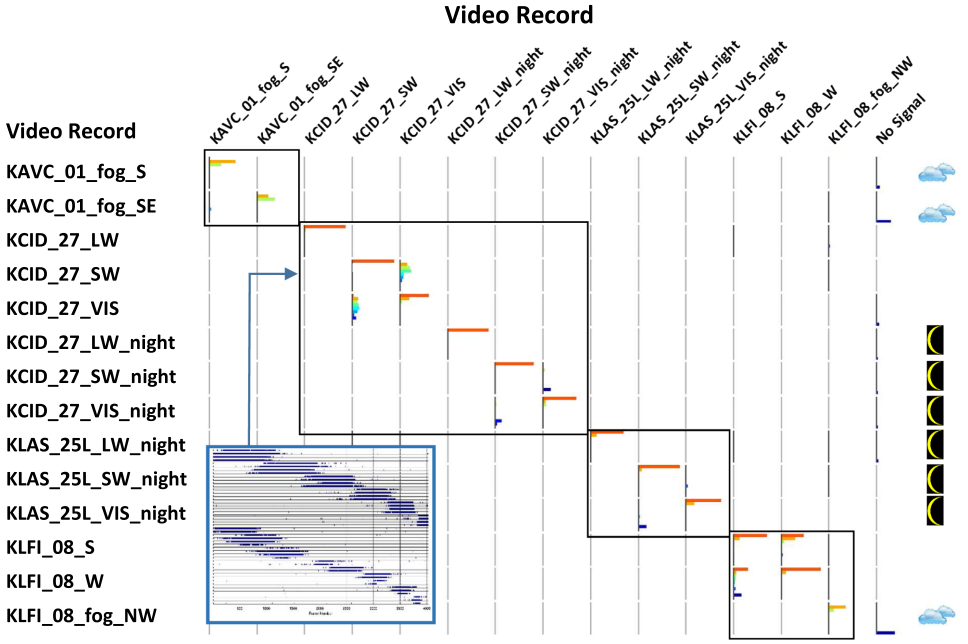

Fig. 6 Repeatability of runway glidepath localization. Epoch match probabilities are sorted by video of origin to determine which runway descent video records in the universe of known geographic locations are effective for waypoint-based localization. Along four runway glidepaths (grouped by black outlines), repeatability is observed for the same time of day, clear weather conditions, and same sensor type. Cross-sensor repeatability is also observed between visible- and SWIR-band records. The inset shows raw match data for video record KCID_27_SW to the SWIR imagery (top 20 traces) and to visible-band imagery from record KCID_27_VIS (bottom traces).

B. Specificity of Runway Glidepath Identification

For each of the 19 runway approach video records, we computed the percentage of frames of the record for which a significant ($p > 0.1$) probability was found that the airplane was at one of the nine known runway-based locations. The matrix of histograms for all videos and all possible runway-based locations is plotted in Fig. 5. This may be interpreted as a detection specificity plot: if the only probable runway-based location is identical to the runway of origin, then there are no false positives and the result is a true positive detection. For example, the histograms for the Frankfurt (EDDF) approach video are empty except in the EDDF_25L column: this means that the method locates the airplane to the glidepath of the final approach to Frankfurt runway 25L with perfect specificity. As a counterexample, there is a small false-positive count for the

4G7_23_night video; although a great majority of the frames are correctly reported as located in the glidepath of the final approach to 4G7_23, there is a small number of frames incorrectly reported as located in the glidepath of the final approach to KLAS_25L. (The percentage is too small to render well as a histogram bar in this figure but is indicated by the black axis color.)

The method used in this study displays excellent runway identification specificity for all flight records and strong confidence for clear weather flights. Fog or rain, however, degrade image contrast, reduce image feature matching, and therefore reduce location confidence. For example, a large percentage of frames of the two KAVC_01 runway approach videos' records, captured on a foggy day, have no match to any of the reference images in this universe of runway glidepath waypoint locations. Image enhancement improved frame contrast, image matching, and therefore glidepath location assignment in these (KAVC_01_fog_SE, KAVC_01_fog_S) and other (KLFI_08_fog_NW, KLFI_26_rain) poor weather recordings. This is shown in Fig. 5 by overlaying (dashed boxes) histograms of matching after image enhancement to the lower left of the unenhanced probability histograms.

Image enhancement is not always beneficial in improving specificity of runway identification. In one case in this study (KCKB_21_night), the video stream of a nighttime clear weather flight is compressed by the video electronics before digital storage on the camera. The MPEG-style compression artifacts, although benign without enhancement, are amplified in contrast by enhancement [23] and introduce spurious image features. Degradation of this type is not seen in the KLAS nighttime records, which are captured without compression.

C. Repeatability of Runway Glidepath Identification

To test the stability of waypoint-based localization across variations in runway descent records (sensor type, time of day, weather conditions), we sorted $P_t(e)$ of Eq. (6) by video of origin before tallying the results into histograms. Figure 6 shows the results for the four runway glidepaths for which we had multiple video records; the waypoint-based localization probability histograms for 14 records are grouped by runway. Repeatability (i.e., identification of a given runway across multiple videos) is demonstrated if there are offdiagonal histograms in the array.

Measureable repeatability is found between runway descent recordings if the time of day and sensor type are the same and the weather is clear. (Image enhancement is used for all results plotted in Fig. 6.) In the first experiment (airport KAVC), weak repeatability is observed between a pair of daytime visible-band records taken in foggy conditions (but is not observed without image enhancement). In the second experiment (airport KCID), strong repeatability ($p > 0.1$ for greater than 90% of frames) is observed between daytime visible-band and SWIR records, even though the visible sensor is constructed with filters that block SWIR. The inset of Fig. 6 shows the corresponding raw image match data for short-wave infrared recoding of the daytime flight (KCID_27_SW). Much weaker cross-sensor repeatability is observed between nighttime visible-band and SWIR records. In the third experiment (airport KLAS), weak cross-sensor repeatability is observed between nighttime visible-band and SWIR records. Finally, in the fourth experiment (airport KLFI), strong repeatability is observed among the two daytime clear weather flights into the runway, but there is no cross matching between either of them and the record taken in foggy conditions. This last result indicates that image enhancement, at least as we have applied it, is not generally sufficient to counter the image feature degradation that results from atmospheric turbidity.

D. Granularity of Waypoint-Based Location

With final approach speeds of about 100 kt, a final approach segment distance of about 5 nmi, and waypoints defined at about each nautical mile of distance from touchdown, the airplane will transition from one epoch to another within about 1 min. As a backup to gyros and GPS this is a level of spatiotemporal granularity that is helpful and useful, since an analogous alert, the aural Traffic Collision Avoidance System (TCAS) alert, is known to be highly effective, even during the attention-limit conditions of airplane landing [24]. As a standalone navigation method, finer granularity is desirable.

Finer granularity may be possible by interpolating across waypoints. Subsequent to matching and probability determination, we reassembled the processed frames into a video with a "fuel gauge" representation of probability for all nonzero waypoint locations displayed next to the video imagery. We observed that, as the airplane descended, the waypoint probability moved smoothly from one waypoint to the next. Referring to the sketch at the left of Fig. 1, for example, as the probability of location at waypoint 1 declined, that of waypoint 2 increased; as the probability of location at waypoint 2 declined, that of waypoint 3 increased; and so on through the set of waypoints. Quantitative tests of inter-waypoint interpolation were not conducted in this study but might be a fruitful area of future research.

IV. Discussion: Machine Vision Constraints and Cockpit Video

We applied a widely used machine vision pipeline to the problem of runway identification using videos from 19 runway approaches along the final segment glidepath of nine runways. Our hypothesis was that, for rigid objects and fixed scenes, current machine vision technology was capable of identifying imagery rapidly and with specificity over a modest range of camera viewpoints and scene illumination. It contained four constraints (rigid objects, fixed scenes, moderate viewpoint change, and moderate illumination change) and predicted two positive results (rapid and specific identification). In this section, we discuss how this study proves or disproves the hypothesis, given its adherence to the hypothesis constraints.

A. Adherence to Constraints

By design, the adherence to the constraints is as follows.

1) The visual subject (an airport runway and its surrounding terrain) is rigid, and so this constraint is met without exception.

2) The fixed scene constraint varies slightly. By using approach video from just one seasonal variation (full foliage, no snow cover), we artificially remove one source of scene variation. We allow one type of minor change within the scene (the movement of vehicles on the roads within the camera view), and it has no discernible effect on the results.

3) Camera viewpoint is constrained by separating the approach video into time epochs that are treated as independent observations of a fixed scene, with a moderate variation of viewpoints.

4) Since no approach video is recorded at dawn or dusk, we limit scene illumination to two variants (day and night).

B. Prediction Success/Failure

Rapid identification was observed for daytime videos, both with and without image enhancement. We were able to identify specific runways rapidly (about 10 frames per second), even when matching against a universe of approximately 300 reference images. For night videos, matching was appreciably slower, especially with high-pass filtering (slowdown factor: three times at a minimum; five times is typical; 20 times at a maximum). This was because the spurious features were "detected" in the noise of the supposedly featureless parts of the scene.

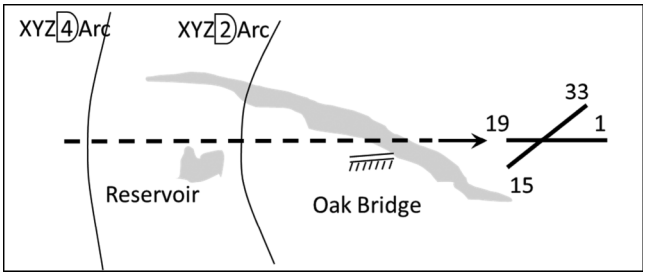

Fig. 7 Simplified example of a charted visual flight procedure, adapted from [28].

Specific identification was superb, as predicted in the hypothesis, with one important exception: image smear on light aircraft in windy conditions. In this and the previous [25] study, we uncovered an exception to the notion that runway identification could be achieved as long as the viewpoint change was moderate. We observed (Fig. 4) that the continuity of waypoint-based location was more variable for a small aircraft (Cessna, at bottom) than a large aircraft (Boeing 747, at top). Review of the KPHF_25 video revealed that the camera viewpoint in recordings was prone to sudden changes in pitch, yaw, or roll, presumably due to wind shifts buffeting the Cessna during the flight. The KPHF_25 video (and other Cessna recordings) implicated the image sensor: it did not respond instantly to a sudden viewpoint change. Streaking, smearing, and other frame readout artifacts were evident with each perturbation. We concluded that a sensor with a higher frame rate, a faster pixel readout response, or frame-shielded design (see, for example, [26]) was needed to eliminate this effect. This sensor constraint should especially be observed in general aviation and in small unmanned aerial vehicle applications.

In our prior study [25], we concluded that specific identification failed in poor weather conditions. By introducing image enhancement, we mitigated that limitation, somewhat, for this dataset. Four runway approaches were recorded in foggy or rainy conditions. A change in air turbidity was, in essence, a scene change. We found that specificity was superb, even with this violation of the fixed scene constraint. The constancy of specificity throughout the flight, however, was degraded. As a striking illustration of this, consider the second and third rows of Fig. 5 (videos EDDF_25L_night and KAVC_01_fog_S). The runway location in the former case was reported strongly for nearly 100% of the approach, whereas the runway location of the latter was reported strongly for only about 20% of the descent glidepath approach (and was unknown for the remaining 80%). Image enhancement improved detection to about 70% of the frames in this case (top, dashed overlay box of Fig. 5).

C. Temporal Continuity to Ameliorate Loss of Identification

The probability model used in this study has no temporal component: indeed, the waypoint/epoch method is used to prevent the statistical bias that arises from frame-to-frame continuity. No-signal frames, caused by poor imaging conditions (fog, camera smearing) are an artificial result of the instantaneous probability calculation and independence assumptions inherent in Eq. (6). In practice, a human observer readily compensates for these dropouts while viewing the results videos, using temporal continuity to ameliorate temporary loss of identification. This raises the prospect that a future study that employs a probabilistic model with a temporal component can infer runway identification based not just on the current frame but also on the prior series of frames.

D. Extension to Repeatability, and Applicability to Geolocation

Repeatable identification across two or more recordings of the same runway glidepath was observed to be strong within the same sensor type in clear weather. Waypoint-based image matching can repeatedly locate a runway in clear daylight or night conditions (e.g., approaches KLFI_08_S and KLFI_08_W, or rows 12 and 13 of Fig. 6), but, as expected from intuition, no repeatability is found between day and night recordings at the same runway location (KCID 27).

Repeatability degrades in poor weather, as seen in several rows of Fig. 6. The KAVC data, taken in foggy conditions, shows only weak repeatability (top two rows) and only with image enhancement (see also Fig. 2), and the KLFI_08 data show strong repeatability between records taken in the clear conditions but none at all between clear and foggy conditions (bottom rows of Fig. 6), regardless of image enhancement. We find repeatability between visible and short-wave infrared recordings of a daytime pair of records (rows 4 and 5 of Fig. 6). This cross-sensor result is weaker but detectable for night recordings.

With these observations, the machine vision method can be compared to an analogous visual landing method, the Charted Visual Flight Procedure (CVFP) [27,28], which uses prominent landmarks as waypoints during descent to guide navigation along a predetermined path. A simplified CVFP is illustrated in Fig. 7. There are two notable restrictions inherent in the CVFP descent method: visibility minimums are specified for each waypoint (e.g., [29,30]), and it is not authorized at night. The latter restriction is perhaps due to occasional pilot misidentification of landscapes at night (Hart [31] emphasizes human error in nighttime runway identification and cites several examples, including confusion between two runways used in this study, 4G7_23 and KCKB_21).

V. Conclusions

With these specificity and repeatability results, it can be said that current machine vision technology is applicable to airborne geolocation, within limitations. By comparing imagery from an onboard sensor to a database of reference images, an aircraft can geolocate as long as visibility is clear and the database imagery is collected within the fixed scene and fixed illumination constraints cited previously. For example, the database should include night and day imagery, seasonal variations such as snow cover, wet vs dry pavement, and full vs sparse foliage; and it should be reasonably current. Minor scene changes such as vehicular traffic do not degrade results, and we expect that similarly minor additions (or deletions) such as traffic cones, commercial signage, and cell towers will not do so either, as their contribution to the total set of image features is small. Major changes, such as new roadways and building additions/renovations, may modify the set of image features significantly, and so may violate the fixed scene assumption. The clear visibility limitation may be relaxed if image enhancement is applied, but we cannot quantify a visibility cutoff from the data used in this study.

In commercial and general aviation airplanes, geolocation is determined visually and with many assistive technologies (e.g., inertial guidance systems, radiolocation, Global Positioning System). Waypoint-based geolocation via machine vision image matching may supplement these methods or provide the same capability in the final approach segment for some clear conditions at a lower cost. With only visible and infrared sensor inputs, however, the limitations of machine vision geolocation in poor weather ensure that it cannot be the sole backup navigational technology.

An analogous visual landing method, called a charted visual flight procedure [26,27] uses prominent landmarks visible at a safe altitude as waypoints during descent to guide navigation along a predetermined path. A simplified charted visual flight procedure is illustrated in Fig. 7. Like the method used in this study, visual recognition of landmarks along the landing approach is used to navigate to the runway. Both human and machine waypoint-based geolocations fail in foggy conditions; visibility minimums at each waypoint are documented in charted visual flight procedure charts (e.g., [28,29]). Waypoint-based geolocation via machine vision has one important advantage over its humancentric counterpart: we find excellent specificity at night but charted visual flight procedure is not authorized at night. The Charted Visual Flight Procedure resembles the method used in this study, as it employs visual recognition of landmarks along a prescribed landing approach to navigate to the runway. The regulatory visibility minimums at each waypoint of a CVFP final descent codify limitations of human vision in low visibility conditions akin to the limitations of the machine vision method that we found in rain and fog. Potential pilot misidentification of landscapes at night is mitigated in the CVFP regulations, which disallow the method at night, in all visibility conditions. In contrast, we found that for clear weather conditions, waypoint-based geolocation via machine vision shows excellent location specificity at night. This advantage of artificial over natural vision may be of operational use. Further studies using imagery outside the visible band (in the infrared, as in this study, or in the radar band) may demonstrate, more conclusively than this study, that geolocation via machine vision is reliable over a greater range of atmospheric visibilities and landing times than the regulatory ranges specified for CVFP landings.

In conclusion, an aircraft with an onboard image sensor and computing power that includes GPUs, a database of reference images, and a probability table corresponding to that database can geolocate during the final approach segment as long as visibility is clear and the database 1) includes imagery sampled at the same season and time of day, 2) covers the geographical areas likely to be encountered in course of the flight, and 3) contains imagery of the same sensor type as the onboard sensor. By matching across the visible and infrared bands, the sensor requirement can be broadened to some degree: some matching is possible between visible and short-wave infrared bands.

Acknowledgments

This work was supported in part by the NASA Atmospheric and Environmental Safety Technologies and the NASA Safe Autonomous Systems Operations programs. We are grateful to Robert Neece, Taumi Daniels, and Dan Jobson for comments on experimental design, and to Tim Dugan for discussions on experimental design and operational feasibility. An invaluable flight video was provided by Rockwell Collins, and we are grateful to Carlo Tiana, Tim Etherington, and Weston Lahr for helpful discussions.

References

[1] Huertas, A., Cole, W., and Nevatia, R., "Detecting Runways in Complex Airport Scenes," *Computer Vision, Graphics, and Image Processing*, Vol. 51, No. 2, 1990, pp. 107–145
doi:10.1016/0734-189X(90)90027-S

[2] Tandra, S., and Rahman, Z., "Robust Edge-Detection Algorithm for Runway Edge Detection," *SPIE Image Processing: Machine Vision Applications*, SPIE Paper 68130L, Bellingham, WA, 2008.

[3] Gong, X., Abbot, A. L., and Fleming, G. A., "A Survey of Techniques for Detection and Tracking of Airport Runways," *4th AIAA Aerospace Sciences Meeting and Exhibit*, AIAA Paper 2006-1436, 2006, pp. 9–12.

[4] Trisiripisal, P., Parks, M. R., Abbott, A. L., Liu, T., and Fleming, G. A., "Stereo Analysis for Vision-Based Guidance and Control of Aircraft Landing," *44th AIAA Aerospace Science Meeting and Exhibit*, AIAA Paper 2006-1438, 2006, pp. 1–17.

[5] Shang, J., and Shi, Z., "Vision-Based Runway Recognition for UAV Autonomous Landing," *International Journal of Computer Science Network Security*, Vol. 7, No. 3, 2007, pp. 112–117.

[6] Zongur, U., "Detection of Airport Runways in Optical Satellite Images," M.S. Thesis, Middle East Technical Univ., Ankara, Turkey, 2009.

[7] Pai, C. H., Lin, Y. P., Medioni, G. G., and Hamza, R. R., "Moving Object Detection on a Runway Prior to Landing Using an Onboard Infrared Camera," *IEEE Computer Vision and Pattern Recognition*, IEEE, Piscataway, NJ, pp. 1–8, 2007.

[8] Lowe, D., "Distinctive Image Features from Scale-Invariant Keypoints," *International Journal of Computer Vision*, Vol. 60, No. 2, 2004, pp. 91–110.
doi:10.1023/B:VISI.0000029664.99615.94

[9] Bay, H., Tuytelaars, T., and Van Gool, L., "SURF: Speeded-Up Robust Features," *Computer Vision and Image Understanding (CVIU)*, Vol. 110, No. 3, 2008, pp. 346–359.
doi:10.1016/j.cviu.2007.09.014

[10] Shen, Y. F., Rahman, Z., Krusienski, D., and Li, J., "A Vision-Based Automatic Safe Landing-Site Detection System," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 49, No. 1, 2013, pp. 294–311.
doi:10.1109/TAES.2013.6404104

[11] Sun, X., Christoudias, C. M., Lepetit, V., and Fua, P., "Real-Time Landing Place Assessment in Man-Made Environments," *Machine Vision and Applications*, Vol. 25, No. 1, 2014, pp. 211–227.
doi:10.1007/s00138-013-0560-7

[12] "1.1.9 Instrument Landing System (ILS)," *Aeronautical Information Manual: Official Guide to Basic Flight Information and ATC Procedures*, Federal Aviation Administration, 2015, pp. 1-1-7–1-1-11, www.faa.gov/air_traffic/publications/media/AIM.pdf [retrieved 10 Dec. 2015].

[13] Bradski, G., "The OpenCV Library," *Doctor Dobbs Journal*, Vol. 25, No. 11, 2000, pp. 120–126.

[14] Zhang, Z., Deriche, R., Faugeras, O., and Luong, Q. T., "A Robust Technique for Matching Two Uncalibrated Images Through the Recovery of the Unknown Epipolar Geometry," *Artificial Intelligence*, Vol. 78, No. 1, 1995, pp. 87–119.
doi:10.1016/0004-3702(95)00022-4

[15] Oppenheim, A. V., Shafer, R. W., and Stockham, T. G., Jr., "Non-Linear Filtering of Multiplied and Convolved Signals," *Proceedings of the IEEE*, Vol. 56, No. 8, IEEE Publ., Piscataway, NJ, 1968, pp. 1264–1269.

[16] Moore, A., Allman, J., and Goodman, R. M., "A Real-Time Neural System for Color Constancy," *IEEE Transactions on Neural Networks*, Vol. 2, No. 2, 1991, pp. 237–247.
doi:10.1109/72.80334

[17] Rahman, Z., Jobson, D. J., and Woodell, G. A., "Retinex Processing for Automatic Image Enhancement," *Journal of Electronic Imaging*, Vol. 13, No. 1, 2004, pp. 100–110.
doi:10.1117/1.1636183

[18] Torr, P. H., and Murray, D. W., "The Development and Comparison of Robust Methods for Estimating the Fundamental Matrix." *International Journal of Computer Vision*, Vol. 24, No. 3, 1997, pp. 271–300.
doi:10.1023/A:1007927408552

[19] Tekalp, M., "Video Segmentation," *Handbook of Image and Video Processing*, Academic Press, Orlando, FL, 2000, pp. 383–399.

[20] Lee, D. N., and Reddish, P. E., "Plummeting Gannets: A Paradigm of Ecological Optics," *Nature*, Vol. 293, No. 5830, 1981, pp. 293–294.
doi:10.1038/293293a0

[21] Horn, B. K. P., and Schunck, B. G., "Determining Optical Flow," *Artificial Intelligence*, Vol. 59, No. 1, 1993, pp. 81–87.
doi:10.1016/0004-3702(81)90024-2

[22] Tukey, J. W., *Exploratory Data Analysis*, Addison-Wesley, Reading, MA, 1977, pp. 43–44.
[23] Rahman, Z., Jobson, D. J., and Woodell, G. A., "Investigating the Relationship Between Image Enhancement and Image Compression in the Context of the Multi-Scale Retinex," *Journal of Visual Communication and Image Representation*, Vol. 22, No. 3, 2011, pp. 237–250. doi:10.1016/j.jvcir.2010.12.006
[24] Young, S. D., Daniels, T. S., Evans, E. T., Uijt de Haag, M., and Duan, P., "Understanding Crew Decision-Making in the Presence of Complexity: A Flight Simulation Experiment," *AIAA InfoTech Aerospace Conference*, AIAA Paper 2013-4894, 2013, pp. 1–22. doi:10.2514/6.2013-4894
[25] Schubert, M., Moore, A., Dolph, C., and Woodell, G., "Machine Vision for Airport Runway Identification," *Proceedings of SPIE*, 9477, Optical Pattern Recognition XXVI, 2015; also SPIE Defense+ Security, International Soc. for Optics and Photonics, 2015, Paper 94770G. doi:10.1117/12.2177320
[26] Mendis, S. K., Kemeny, S. E., Gee, R. C., Pain, B., Staller, C. O., Kim, Q., and Fossum, E. R., "CMOS Active Pixel Image Sensors for Highly Integrated Imaging Systems," *IEEE Journal of Solid-State Circuits*, Vol. 32, No. 2, 1997, pp. 187–197. doi:10.1109/4.551910
[27] "5.4.24 Charted Visual Flight Procedure (CVFP)," *Aeronautical Information Manual: Official Guide to Basic Flight Information and ATC Procedures*, Federal Aviation Administration, 2015, pp. 5-4-61–5-4-62, www.faa.gov/air_traffic/publications/media/AIM.pdf [cited 10 Dec. 2015].
[28] "Information for Operators 11003: Pilot's Roles and Responsibilities During Visual Approaches," Federal Aviation Administration, Jan. 2011, p. 2, www.faa.gov/other_visit/aviation_industry/airline_operators/airline_safety/info/all_infos/media/2011/InFO11003.pdf [retrieved 25 Jan. 2011].
[29] "Amdt 1 09SEP99 SMOKE STACK VISUAL RWY 02," Federal Aviation Administration, Sept. 1999, http://aeronav.faa.gov/d-tpp/1605/00762smokestack_vis2.pdf [retrieved 10 Dec. 2015].
[30] "Orig 12MAY05 LA RIVER VISUAL RWY 12," Federal Aviation Administration, May 2005, http://aeronav.faa.gov/d-tpp/1605/00236lariver_vis12.pdf [retrieved 10 Dec. 2015].
[31] Hart, C. A., "Safety Recommendation A-15-009 and A-15-010," National Transportation Safety Board, Washington, D.C., May 2015, www.ntsb.gov/safety/safety-recs/recletters/A-15-009-010.pdf [retrieved 4 May 2015].

J. P. How
Associate Editor