Experiments in Leader Classification and Following with an Autonomous Wheelchair

Procópio Stein, Anne Spalanzani, Vitor Santos, Christian Laugier

▶ To cite this version:

Experiments in Leader Classification and Following with an Autonomous Wheelchair

Procópio Stein[1]*, Anne Spalanzani1[12], Vítor Santos[3], and Christian Laugier[1]

[1] INRIA Rhône-Alpes, 655 Avenue de l'Europe, 38334 Saint Ismier cedex, France
[2] Université Grenoble Alpes, Lab. LIG, Grenoble, France
[3] Universidade de Aveiro, Campus Universitário de Santiago 3800 Aveiro, Portugal
procopiostein@gmail.com,anne.spalanzani@inria.fr,
vitor@ua.pt,christian.laugier@inria.fr

Abstract. With decreasing costs in robotic platforms, mobile robots that provide assistance to humans are becoming a reality. A key requirement for these types of robots is the ability to efficiently and safely navigate in populated environments. This work proposes to address this issue by studying how robots can select and follow human leaders, to take advantage of their motion in complex situations. To accomplish this, a machine learning framework is proposed, comprising data acquisition with a real robot, data labeling, feature extraction and the training of a leader classifier. Preliminary experiments combined the classification system with a multi-mode navigation algorithm, to validate this approach using an autonomous wheelchair.

Keywords: leader selection, leader following, assistance robotics, human-aware navigation

1 Motivation, problem statement and related framework

The development of assistance and service robots requires special attention to the problem of navigation in populated environments, as how to interact with humans, respect social conventions and deal with unpredicted situations. The capacity of a robot to address these requirements directly impact its acceptance by humans.

Several successful techniques of motion planning address the problem of navigation in populated environments using probabilistic and predictive approaches [1], [2] and on models of social interactions [3]. The idea is that, by estimating areas that ought to be occupied in the future, these algorithms can create collision-free motion plans that may also respect social conventions.

Unfortunately, these algorithms usually do not incorporate the notion that humans will also adapt their motion to the environment, and it is usually assumed that the robot is the sole responsible for collision avoidance maneuvers.

One consequence of these assumptions, is that, due to the uncertainty in predictions allied with situations of high human density, these algorithms may fail to find a path. This problem has been described in [4] and is known as the Freezing Robot Problem (FRP).

The motivation for this work derives from the observation that humans can seamlessly move in highly populated environments, while addressing complex situations and interactions with other humans. Due to this complexity, the mechanisms that take place in this behavior are not yet fully understood. This is an active topic of research [5], [6] with some findings incorporated into motion planning algorithms [7], [8].

According to these observations, we study the novel hypothesis that a robot can that properly select persons as leaders in a complex environment, will benefit from their advanced navigation and interaction skills, by following them [9]. Figure 1 illustrates the advantages of this approach in comparison with a state of the art motion planning algorithm (RiskRRT), using a pedestrian simulator in a narrow corridor.
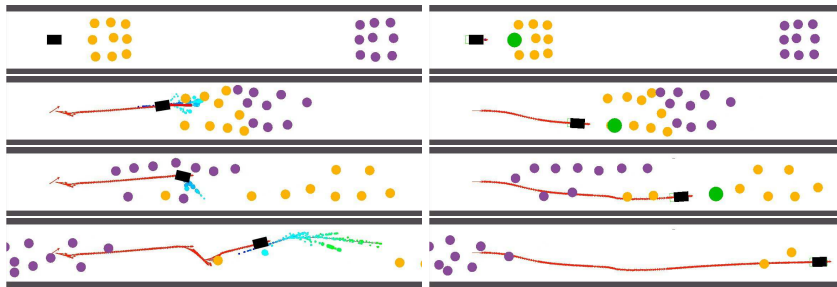

Fig. 1. Qualitative comparison of two navigation techniques in a narrow corridor. Left column: RiskRRT, right: leader following. Each row depicts the same instant during tests. Agents are represented by circles, the robot by the black rectangle and its path by a sequence of red arrows, showing its orientation in each instant.

2   Technical Approach

Our experiments investigated how humans perform leader selection in populated environments. The approach was to associate LIDAR measurements with leader behaviors, in order to train a machine learning algorithm to classify persons as good or bad leaders. The general framework was built using ROS [10], while for machine learning we used the OpenCV library.

2.1   Data acquisition

Data was acquired using a small car-like robot. Only two sensors were used: a laser scanning range finder and a camera fitted with a wide-angle lens. In

this way, videos could be recorded and associated with laser scans during tests. Figure 2 shows the robotic platform, and a sample of the data collected from both sensors at a time instant. The purpose of this setup was to obtain information relative to the behavior and reactions of persons being followed by the robot.

Fig. 2. Left: platform used for data acquisition; middle: sample image from a wide-angle camera; right: sample of laser measurements showing three tracked subjects and their orientation as the dark vectors. The robot is represented as a black rectangle.

2.2 Data Labeling and Feature Extraction

In many datasets, the labeling process is clear and objective, as marking faces in images, or pedestrians in laser scans. However, in this work the objective is to learn how humans decide when to start or stop following someone, or in other words, when someone is a good or bad leader. In order to do so, participants create labels based on their *feeling* about someone being a good leader or not.

A binary labeling system was used, with candidates identified either as good or bad leader. The process is depicted in . Volunteers should press a button whenever they *felt* a transition from good to bad leader occurred, while watching a video of persons being followed by the robot. This results in a mark that is stored along with all the original data recorded from experiments.

Videos of tests were shown to participants, who created labels based on their *feeling* about someone being a good leader or not. In the labeling process, the volunteers watched a video of persons being followed by the robot, and pressed a button whenever they *felt* a transition from good to bad leader occurred, as shown in Figure 3.

As a result, a mark is stored along with all the original data recorded from experiments. The final label was computed as an average of individual inputs and then, LIDAR measurements were used to obtain the descriptors of each followed person. Figure 4 depicts the following features, which were calculated based on the robot's and subject's current state:

- subject's absolute velocity v_{c_i};
- relative heading between the robot and the subject α_{rc_i};
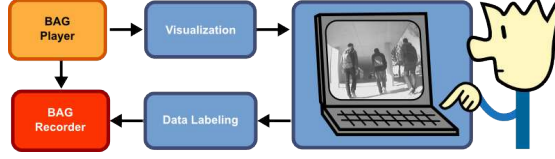- subject's relative velocity w.r.t. the robot reference frame vx_{rc_i} and vy_{rc_i};

Fig. 3. Diagram of the data labeling process.

$v_r = \|AB\|$
$v_{c1} = \|CD\|$
$\alpha_{rc1} = \angle BEC$
$vx_{rc1} = CG\text{-}v_r$
$vy_{rc1} = GD$
$\beta_{rc1} = \angle BAC$
$d_{rc1} = \|AC\|$
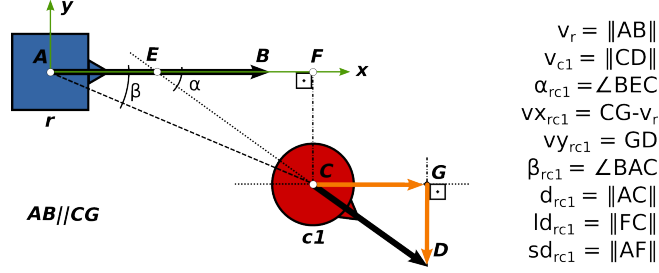$ld_{rc1} = \|FC\|$
$sd_{rc1} = \|AF\|$

Fig. 4. Graphical representation of the features extracted while following. The robot is represented as a blue square while the human being followed as a red circle.

- angle between the robot heading and the subject's position β_{c_i};
- distance between the subject and the robot d_{c_i};
- lateral and sagittal displacement of the subject ld_{c_i} and sd_{c_i} ;

All the laser measurements were transformed to a fixed frame, to compensate for the robot's motion. Then, a motion tracker [11] was used to manage the identification and tracking of moving targets. The standard deviation and first derivative of each feature were also used, totaling 24 features.

2.3   Classifier Training and Features Evaluation

The chosen classifier is the adaptive Boosting (or AdaBoost) [12]. Among the advantages of this algorithm is the exponential convergence of the training error to zero and good generalization properties. But most important to the experiments conducted here, is the capacity of the algorithm to inherently select the features that contribute the most to the classification process.

This allows us to study of feature contribution, which is an important aspect of this work, enabling us to draw conclusions about why some features are more important than others, giving insights on what are the criteria used by humans when engaging or disengaging in following behaviors.

Learning and understanding what matters in following-leading engagements will be essential to the development of a more refined theory about human-following behaviors for robots.

Fig. 5. Good leader behavior, although the robot is noticed, the leaders kept moving without changing their behaviors

Fig. 6. Bad leader: the person notices the robot and moves aside, giving room for the robot to pass.

3   Results

3.1   Data Acquisiton

In total 47 runs of the robot following persons or group of persons were recorded, with a mean duration of about 20 seconds each. Tests were conducted in an open corridor, about $3m$ wide. The robot operator's goal was to move along this corridor back and forth, covering a distance of about $20m$ in each pass.

Although each test had its own peculiarities, after the acquisition process, the data could be organized in four different classes, according to the type of transition from good to bad leader.

Before the labeling began, examples of different classes of good and bad leader situations were shown to the volunteers. Later they were asked to tell which of these situations occurred in each experiment. Although each test has its own peculiarities, the experiments can roughly fit one of the following classes:

1. **good leader (gd)** - leader(s) maintained their speed and orientation, withouth changing their behavior while being followed (Figure 5);
2. **bad leader, moved aside (as)** - leaders gave room for the robot to pass, generally moving aside, while keeping their original motion direction (Figure 6);
3. **bad leader, far or fast (fr)** - the distance between the leader and the robot grew to a point where it was not advantageous to keep following them (Figure 7);
4. **bad leader, stopped (st)** - when the person being followed stops moving (Figure 8).

Fig. 7. Bad leader: the subject marked with a white circle moves to fast compared to the robot's speed and gets too far from it.

Fig. 8. Bad leader: the person stops moving.

Besides the four typical situations mentioned before, two other situations were recorded, which did not involve leader following, but could enhance the classifier training. These are the case when candidates were not moving, but rather standing close to the robot's path (nm), or when they were moving towards the opposite direction (od).

3.2   Classifier Training

After labeling, each resulting dataset was given a name to identify the situation they represented, together with a number to differentiate them. This organization allowed training and test datasets to encompass all the cases encountered. In total, 12911 samples were obtained, with the proportion of 37% of bad leader labels and 63% of good leader labels. This is equivalent to 451 seconds of tests, divided in 47 experiments. Two datasets were created, one for the training comprising 8504 samples with 34% cases of bad leader, and a test dataset with 2715 samples and 39% cases of bad leader. The training dataset is composed of the following situations:

An iterative process was used in the training of AdaBoost structures. In each iteration, the most important features were kept and then the training data was reprocessed and a new AdaBoost was trained. The objective of this procedure was to evaluate the influence played by different sets of features in the classifier performance and on the importance of features.

3.3   Feature Contribution

Figure 9 shows the contribution ratio of each feature used, after the first training the classifier using 24 features. The most important features are the lateral displacement and the distance between the robot and the leader. Besides those,

Table 1. Composition of training dataset

number of datasets used	dataset situation
9	good leader
7	bad leader, leader moved aside
5	bad leader, leader too far or fast
9	bad leader, leader stopped
7	bad leader, candidate not moving
10	bad leader, candidate coming from opposite direction

Table 2. Contribution Ratio of Most Important Features

feature	size of feature set			
	24	16	12	8
σ distance	0.077	**0.085**	**0.104**	**0.181**
lateral displacement	**0.120**	**0.130**	**0.118**	**0.169**
distance	**0.120**	**0.131**	**0.125**	**0.122**
angle to robot	**0.087**	0.080	0.090	0.119
sagittal displacement	0.052	0.053	0.078	0.118
relative heading	0.078	0.083	0.081	0.113
target velocity	0.068	0.068	0.082	0.093
σ relative velocity in y	0.057	0.062	0.075	0.085

the standard deviation of the distance, of the relative heading and of the target velocity have significant impact on the classification process. Finally, the derivatives features have almost no role at all in classification.

After an iterative process of picking the most important features and retraining the AdaBoost algorithm, results pointed to the **standard deviation of the distance between the robot and the leader** and the **lateral displacement of the leader** as the two most important features, as shown in Table 2. Together they contribution ratio was approximately 35% when training the classifier using only the eight most important features.

3.4   Classifier Performance

To evaluate classification, two metrics were created. The first is the **false good leader**, where the classifier labeled a sample as *good leader* but the ground truth has a *bad leader* label. This is the most critical error, because it means that the algorithm would potentially select a bad leader to be followed, possibly disturbing the person. The second metrics is the **false bad leader**, that occurred when the classifier output a *bad leader* and the ground truth is labeled as *good*
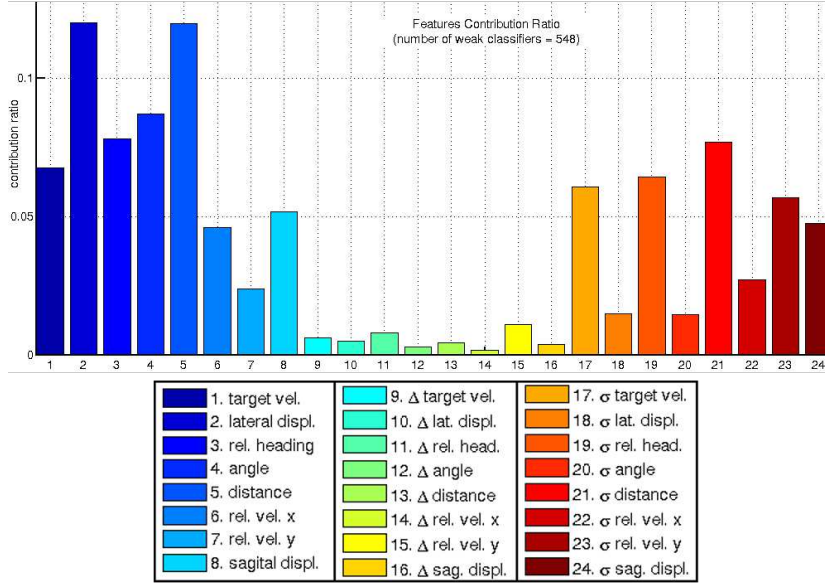
Fig. 9. Feature contribution ratio using 24 features, number of weak learners: 548.

leader. This error is less critical because it translates only in a lost opportunity to classify someone as a good leader.

Regarding the performance of the classifiers, applied to the registered data, the *false good leader* relative error was particularly large at the *move aside* situations. However, the error was much smaller in the remaining situations, and the overall performance of the classifier was remarkable, with the total relative error across all the situations around 3%.

4   Experiments

4.1   Platform and Setup

Although the aforementioned framework provided promising results, the classifier only operates over instantaneous measurements, and the ground truth was based on subjective impressions of volunteers. Therefore, to properly validate the proposed approach, interactive experiments are required.

Such experiments were conducted with an autonomous robotic wheelchair, equipped with two LIDAR sensors, the bottom one at $0.1m$ and the top other at aproximately $0.55m$ in height. The platform also incorporated a RGBD camera, an on-board computer that takes care of low-level hardware control tasks and also a notebook computer that hosts the high-level algorithms, as the localization, planning and leader selection and following.

Figure 11 illustrates the developed setup:

Fig. 10. Experimental platform, showing the two LIDARs and a RGBD camera.
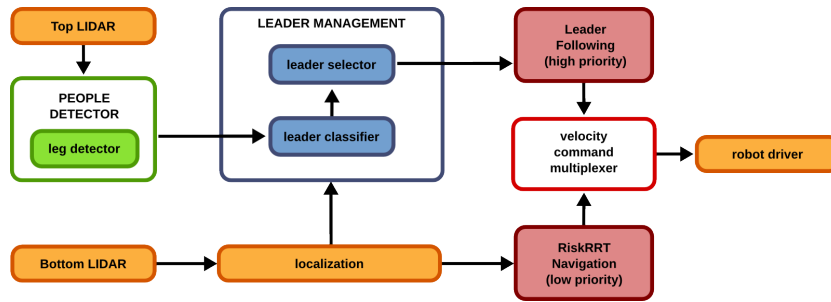

Fig. 11. Experimental architecture.

1. *People tracking*: the technique used for detecting and tracking persons uses the top LIDAR measurements in an AdaBoost classifier, which was previously trained with examples of legs patterns. If two legs are found and satisfy some requirements, it is assumed that a person has been detected.
2. *Leader classifier*: the information about tracked persons is passed to a module that implements the leader classification system presented in Section 2.3. This module provides a binary output (good/bad leader) based on a set of features extracted based on a tracked subject.
3. *Leader selector*: the previous module provides instantaneous classification of subjects based on the state of the robot and on their state. But to select a leader, a history of classifications must be considered. This module keeps a list of tracked subjects and associates a score to each one of them, which is a result of votes cast by different classifications. A good leader classification equals a vote of 0.01 points, while a bad leader vote is −0.1 points. Only
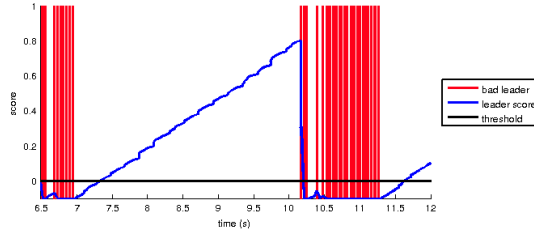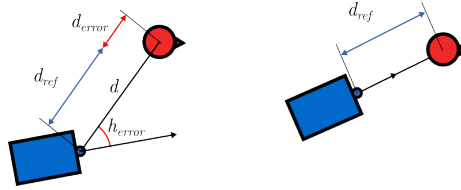
Fig. 12. Leader score example.

Fig. 13. Leader follower controllers.

after the score of a subject passes a threshold, that person can be considered a good leader to be followed. An example of this system is shown in Figure 12.

Such a difference of vote values means that someone needs several classifications of good leader to be considered to be followed, while only a couple of bad leader classifications is enough for a leader to be abandoned. Another advantage of this method is that a list with candidates scores allows for an easy switching of leaders. Once a leader is chosen, this module uses two proportional controllers, one for the distance and other for the heading, to compute velocity commands that are sent to the robot driver module, as shown in Figure 13.

4. *RiskRRT Navigation*: if no feasible leader was found, the robot still needs to be able to navigate in dynamic environments. To accomplish this, this module uses the RiskRRT [1] algorithm to compute paths that avoid risky situations, based on the motion prediction of persons. This module is always active and computing possible paths, even when the robot is engaged in a leader following behavior. In this way it can keep navigating when a leader is abandoned.

5. *Velocity command multiplexer*: in this setup, there are two modules that generate velocity commands to the robot. To manage which command will be sent to the robot, a simple multiplexer is used, which forwards the command with higher priority. The leader selector has the preference, but once a leader is lost or abandoned, it stops sending commands and them the RiskRRT module takes over.
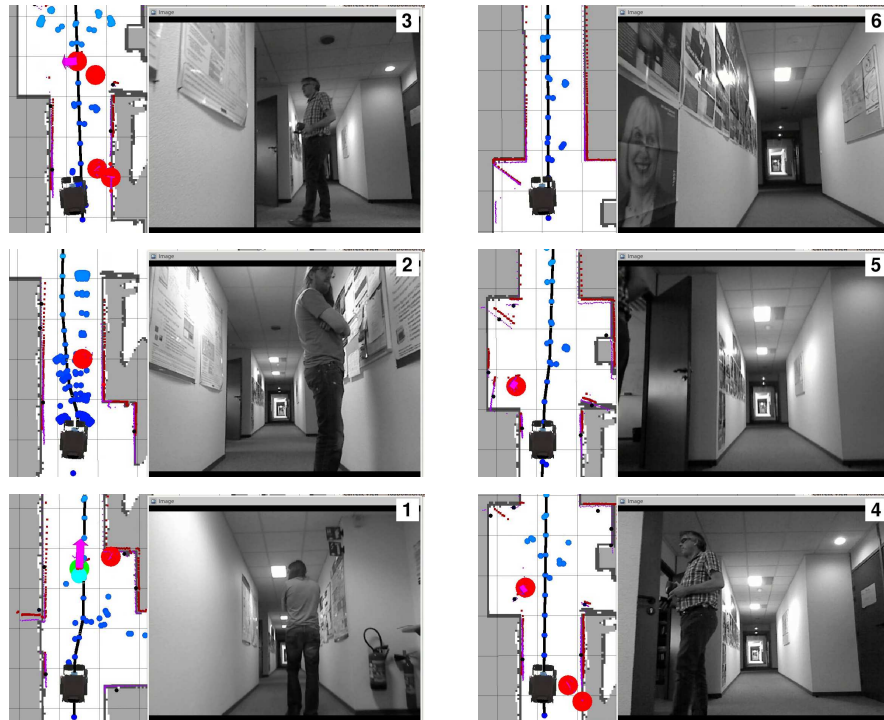
Fig. 14. Switching navigation method between leader following and independent navigation.

4.2   Tests on Switching Navigation Method

The first experiment tests the performance of the proposed setup in switching between the leader following mode and the independent navigation, using the RiskRRT navigation algorithm. Figure 14 illustrates this experiment, where detected persons are represented by either a red circle (bad leaders) or a green circle (goo leaders). The chosen leader is represented by a light blue circle overlapping a green one. The small blue nodes are the exploration nodes of RiskRRT and the black line is the best path to navigate.

The first image of the sequence (1) shows the robot following a leader while, at the same time, the independent algorithm computes possible navigation routes. Once the leader stops moving (2), a sequence of bad leader classifications makes his score fall below an acceptable threshold, and the robot abandons that leader. As soon as this happens, the robot starts to follow the path computed by the independent navigation algorithm, avoiding a collision with the former leader and also with another person crossing its path (3 - 5), which is also classified as a bad leader. The robot later continues to navigate using only the RiskRRT algorithm (6).

4.3   Tests on Leader Switching

This experiment tests the setup capacity of maintaining a list of feasible leaders and switching among them when the score of one surpass the score from another. Figure 15 shows two experiments (one in each column), and as in previous tests, a green circle means an instantaneous good leader classification, while a red one means the opposite. The blue circle marks the person with the highest leader score.

On the first examples (left column), the robot is behind two persons that are classified as good leaders and following the one on the right. As this person slows down to enter an office, he looses score points and the algorithm promptly switches to the other person, who continues to move along the corridor.

The second example (right column), shows the robot following a leader along a corridor. At a given moment, the leader slows down and another persons, which is moving faster passes the robot. According to the voting setup, the person that passes by the robot receives a higher score and then the robot changes the leader and starts to follow the new subject.

4.4   Tests on Leader Following Among Crowds

On the final experiment the robot must be able to select and follow leader in densely populated environments. The test occurs in a narrow corridor, where even small groups of people create difficult conditions for robot navigation. This test is shown in Figure 16, and follow the same nomenclature from previous examples.

The experiment starts with the robot behind two persons moving along a corridor and following the one in the right side. Soon, a group of persons appear coming from the opposite direction and at the same time, the current leader stops to enter an office. The proposed setup manages to correctly identify a suitable leader among several subjects and successfully navigates through the crowd, taking advantage of the natural interactions among persons.

5   Conclusions

This study and experiments provided insightful findings about the limitations and advantages of the proposed leader selection and following approach.

- *Autonomous wheelchairs are very interesting for experiments among persons.* This is mainly because people are used to see wheelchairs, and normally assume there is a human on the control of it. As a result, their reactions and behavior are more natural, which allows studies with minimal interference from the robot.
- *A robust person tracker is essential for a leader selection system.* When the robot followed someone, even short interruptions on the tracker resulted in erroneous leader classification and, therefore, in poor leader scoring, which hindered the navigation performance. Therefore, the identification of persons

Fig. 15. Two tests on leader switching. On the left column the leader stops and the robot switches to follow another person. On the right a better leader appears on the scene.

Fig. 16. Leader score example.

can be enhanced by using different sensors and detection methods, which must combined to provide a more robust detection and tracking system. Examples can be the fusion of images and laser scans to increase the reliability of people detection.

- *A balance must be found between taking advantage of a person's motion and disturbing that person.* Sometimes if the robotic wheelchair was too close to a leader, such person would stop to give room for the robot to pass, invalidating the advantages of the proposed method.
- *Distance-related features play an important role in leader selection.* Tests shown that the total and lateral distance to a candidate, together with their variance are crucial measurements to engage or disengage in a following behavior.
- *The feature space should be enriched with additional sensors and descriptors.* In some situations the selected features were not sufficient to properly classify a candidate. Different sensors and features must be used to improve the performance of machine learning classifiers. For example, a camera could be used detect gaze direction, and a RGBD camera could detect torso, head and arm patterns. Another interesting descriptor would be a distance from a subject to a precomputed path for the robot, which could avoid the situation where the robot drifts away from an objective.
- *Leader following has a great potential for enhancing robot navigations.* The presented framework and experiments validated the initial assumption that leader following can improve robotic navigation in dynamic environments, taking advantage of complex interactions among persons.

References

1. Fulgenzi, C., Tay, C., Spalanzani, A., Laugier, C.: Probabilistic navigation in dynamic environment using rapidly-exploring random trees and gaussian processes. In: 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). (September 2008) 1056–1062
2. Bennewitz, M., Burgard, W., Cielniak, G., Thrun, S.: Learning motion patterns of people for compliant robot motion. The International Journal of Robotics Research **24**(1) (2005) 31
3. Rios-Martinez, J., Spalanzani, A., Laugier, C.: Understanding human interaction for probabilistic autonomous navigation using Risk-RRT approach. In: 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). (September 2011) 2014–2019
4. Trautman, P., Krause, A.: Unfreezing the robot: Navigation in dense, interacting crowds. In: 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). (October 2010) 797–803
5. Mombaur, K., Truong, A., Laumond, J.: From human to humanoid locomotion - an inverse optimal control approach. Autonomous Robots **28** (December 2009) 369–383
6. Dyer, J.R., Johansson, A., Helbing, D., Couzin, I.D., Krause, J.: Leadership, consensus decision making and collective behaviour in humans. Philosophical Transactions of the Royal Society B: Biological Sciences **364** (March 2009) 781–789

7. Althoff, D., Wollherr, D., Buss, M.: Safety assessment of trajectories for navigation in uncertain and dynamic environments. In: 2011 IEEE International Conference on Robotics and Automation (ICRA). (May 2011) 5407–5412

8. Henry, P., Vollmer, C., Ferris, B., Fox, D.: Learning to navigate through crowded environments. In: 2010 IEEE International Conference on Robotics and Automation (ICRA). (May 2010) 981–986

9. Stein, P., Santos, V., Spalanzani, A., Laugier, C.: Navigating in populated environments by following a leader. In: International Symposium on Robot and Human Interactive Communication (RO-MAN). (2013)

10. Quigley, M., Conley, K., Gerkey, B., Faust, J., Foote, T., Leibs, J., Wheeler, R., Ng, A.Y.: ROS: an open-source robot operating system. In: ICRA workshop on open source software. Volume 3. (2009)

11. Almeida, J.: Target tracking using laser range finder with occlusion. Master's thesis, Universidade de Aveiro, Aveiro, Portugal (2010)

12. Freund, Y., Schapire, R.E.: A decision-theoretic generalization of on-line learning and an application to boosting. Journal of Computer and System Sciences **55**(1) (August 1997) 119–139