Online Learning of Human Navigational Intentions

Mahmoud Hamandi[1] and Pooyan Fazli[2(✉)]

[1] Electrical Engineering and Computer Science Department, Cleveland State University, Cleveland, OH 44115, USA
m.hamandi@csuohio.edu

[2] Department of Computer Science, San Francisco State University, San Francisco, CA 94132, USA
pooyan@sfsu.edu

Abstract. We present a novel approach for online learning of human intentions in the context of navigation and show its advantage in human tracking. The proposed approach assumes humans to be motivated to navigate with a set of imaginary social forces and continuously learns the preferences of each human to follow these forces. We conduct experiments both in simulation and real-world environments to demonstrate the feasibility of the approach and the benefit of employing it to track humans. The results show the correlation between the learned intentions and the actions taken by a human subject in controlled environments in the context of human-robot interaction.

Keywords: Navigational intentions · Human tracking
Human-robot interaction

1 Introduction

With recent developments in artificial intelligence and robotics, robots are increasingly being assigned tasks where they have to navigate in crowded areas [1–4]. While humans learn over the years to understand each other's intentions and plan their paths accordingly, robots still have difficulty understanding human intentions, forcing them to navigate in an over conservative way in human-populated environments.

Previous work on robot navigation within crowds mostly rely on the Social Force Model (SFM) [5] to understand humans, where each is assumed to navigate with a set of known imaginary social forces. Luber *et al.* [6] assumed fixed weights for the social forces based on average human weight and dimensions and track humans with the corresponding motion model. This approach might fail in the real world where humans have different characteristics and might change their intentions over time.

Ferrer *et al.* [7] proposed to control a robot with the Social Force Model to navigate similar to humans by learning a fixed weight for each force from a

S. S. Ge et al. (Eds.): ICSR 2018, LNAI 11357, pp. 1–10, 2018.
https://doi.org/10.1007/978-3-030-05204-1_1

dataset on human navigation. Although learning a fixed weight for each force works for controlling a robot to navigate similar to humans, it might fail to track multiple humans in the real world where each has different preferences for these forces.

On the other hand, Vasquez *et al.* [8] used the social forces and other human features to learn to navigate around humans using an Inverse Reinforcement Learning framework [9] without the need to track humans.

Recently, Alahi *et al.* [10] suggested to use an LSTM-based neural network to track humans with the network learning the connection between one human's position and another. While this approach is very promising, it is not obvious how to extract human intentions from the end-to-end neural network.

In this work, we present Human Intention Tracking (HIT). HIT learns the intentions of each human in the scene to reach a fixed target point or to interact with a robot directly from the Social Force Model. The assumption made here is that intentions are valid for the current time span and change over time. In addition, we assume that instantaneous navigational intentions can be fully understood from observing the human navigation. We do acknowledge that the incorporation of other cues such as gaze or incorporating more information about the environment, such as a semantic map, can allow a better understanding of human intentions. However, we assume this information is not available for the robot, which relies solemnly on the humans' positions in an occupancy grid map to learn their intentions. The proposed approach integrates a Kalman Filter with a motion model based on SFM to track humans and learns their intentions in the environment. While reducing the difference between the predicted and observed human positions, we learn SFM weights specific to each human. Our experiments show the advantage of this method in tracking humans as well as the direct connection between the learned intentions and the actual human motions in the environment.

2 Human Tracking and Intention Learning

We rely on the Social Force Model [5] to track humans and learn their intentions. For a human moving to a fixed target with robots and other humans in the environment as shown in Fig. 1, the resultant social force can be expressed as:

$$\mathbf{F} = \alpha_3 F_{robot} + \alpha_2 F_{human} + \alpha_1 F_{obstacle} + \alpha_0 F_{target}, \tag{1}$$

where $\mathbf{F}$ is the resulting force driving the human, F_{robot} is the force pushing the human toward or away from the robot, F_{human} is the force pushing the human toward or away from other humans, $F_{obstacle}$ is the force driving the human away from obstacles, and F_{target} is the force pushing the human to the target.

Each of the forces is exponentially related to the distance between the two objects enforcing it, with the exception of the last force which is linearly related to the human speed. α_0, α_1, α_2, and α_3 represent the weight of each force, and it can be considered as the intention of the human to consider the corresponding force while navigating. For example, if the human ignores the robot's existence

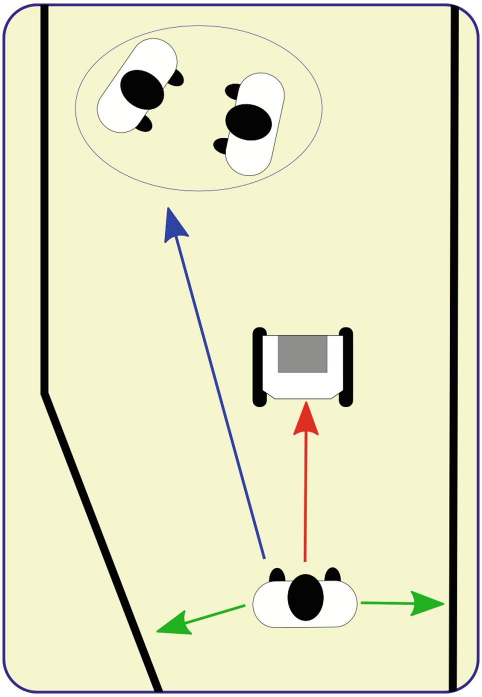

Fig. 1. Example social forces in the environment, showing interaction force to a robot, interaction force to other humans, and interaction force to obstacles. (Color figure online)

completely, the corresponding α should be zero. If the human interacts with the robot, the corresponding α should be positive, while if he runs away from the robot, the corresponding α should be negative. Mathematically, the forces are represented as follows:

$$F_o = A_o \times e^{(\delta_o - \|\mathbf{d_o}\|)/B_o} \times \frac{\mathbf{d_o}}{\|\mathbf{d_o}\|}, \tag{2}$$

where o is a member of the set $O = \{robot, human, obstacle\}$, A_o, δ_o, and B_o are fixed parameters specific to each member of the set, $\mathbf{d_o}$ is the distance vector between the human and the corresponding object in O, and $\|\mathbf{d_o}\|$ is its norm. Ferrer *et al.* [7] show how to learn A_o, δ_o, and B_o from a human dataset and provide typical values for each. On the other hand, we model the force to a fixed target as:

$$F_{target} = \kappa \frac{\mathbf{v}}{\|\mathbf{v}\|}(1 - cos\theta), \tag{3}$$

where κ is a fixed parameter, θ is the angle between the human trajectory and the target direction, and $\|\mathbf{v}\|$ is the norm of the human velocity $\mathbf{v}$. This equation emphasizes the difference in direction between the actual trajectory and the one leading to the target, which helps the robot learn the intention of the human to reach the corresponding position.

Consequently, for a set of learned weights, the social force $\mathbf{F}$ can be calculated based on the observed environment, and we can model the human motion as presented in [6]:

$$\begin{bmatrix}\mathbf{x_t}\\ \mathbf{v_t}\end{bmatrix} = \begin{bmatrix}\mathbf{x}_{t-1} + \mathbf{v}_{t-1}\Delta\mathrm{t} + \frac{\mathbf{F}}{2}\Delta\mathrm{t}^2\\ \mathbf{v}_{t-1} + \mathbf{F}\Delta\mathrm{t}\end{bmatrix}, \tag{4}$$

where $\mathbf{x_t}$ is the position of the human, $\mathbf{v_t}$ is the velocity of the human at time step t, and Δt is the time difference between the two time frames. The motion model of each human can be used to track the human using a Kalman Filter, which predicts their future positions after each observation. Our framework learns the underlying weights that could lead to the observed position by reducing the error between the observed and the predicted positions. As such, the tracking of each human starts with an assumption for each α and updates them for each human as the robot receives more observations. Specifically, the algorithm updates the parameters to reduce the difference between the predicted and the observed human position. This can be achieved as the observed position presents the real social force driving the human, while the predicted position presents the estimated one. As such, the difference between the two is linearly related to the error in the estimate of the Social Force Model. Mathematically, we denote the difference between the two positions as $\mathit{diff}(\mathbf{F})$ and learn each α as:

$$\alpha_{i,t} = \alpha_{i,t-1} + \mathit{diff}(\mathbf{F}) \times F_i \times \gamma, \tag{5}$$

where F_i is the interaction force corresponding to α_i as presented in Eq. 1, and γ is the learning rate.

In this work, we are mainly concerned with the two interaction forces that show the intention of the human to interact with the robot and the intention to reach a fixed target point in the environment.

3 Experiments and Results

We have proposed a method to learn human intentions while observing their navigation paths. Due to the complexity of human intentions, it is difficult to define a single test that can prove the viability of the proposed algorithm. Instead, we split our experiments into three parts:

1. First, we investigate the tracking ability of our algorithm on the ETH walking pedestrians dataset [11]. The dataset provides annotated trajectories of 650 humans recorded over 25 min of time on two different maps referred to as *ETH-Univ* and *ETH-Hotel*.
2. Second, we choose scenes from the dataset with an obvious change in the human direction and study the change in the weight of reaching the human's final goal. This test shows the ability of the algorithm to learn the intention of the human to reach a fixed target.

3. Finally, we test the system on a real robot with humans in the scene. As the humans navigate around the robot, we study their intentions to interact with it.

Table 1. Comparison of average displacement error (m)

	ETH-Univ	*ETH-Hotel*
HIT	0.11	0.036
Target [6]	0.16	0.085
Social-LSTM [10]	0.008	0.15

3.1 Human Tracking

To assess the tracking ability, we compare the average displacement error between the predicted and the observed human position of our algorithm against the one achieved by the methods in [6,10] based on a one-step look-ahead analysis.

Luber *et al.* [6] presented *Target*, a tracking algorithm that combines the Social Force Model with a Kalman filter to predict humans' future positions. Their approach assumes fixed intentions for each human and learns their targets online. While this method allows the tracker to adapt to the target location, it does not adapt to the changes or preferences in intentions toward other humans and obstacles.

On the other hand, Alahi *et al.* [10] presented *Social-LSTM*, a deep learning algorithm for human tracking. Their approach employs an LSTM based network to predict future positions based on previous ones. In addition, they introduce the social pooling layer where the network predicting a human's position shares a hidden layer with other humans' networks. This approach allows the network to learn the interaction among humans in a scene and predict the future positions accordingly.

Our results presented in Table 1 show that our algorithm outperforms *Target* in both datasets and outperforms *Social-LSTM* on the *ETH-Hotel* dataset. While *Social-LSTM* outperforms our algorithm on the *ETH-Univ* dataset, its performance drops drastically on the *ETH-Hotel* dataset, where obstacles are closer than the former and human crowds are denser. This can be related to the network not being able to generalize to a dataset with settings different than the social aspects it was trained on. However, the increased human proximity improved the performance of our algorithm and *Target*'s due to the importance of social forces in such scenes.

3.2 Intention Learning

We mapped the ETH dataset into a 2-dimensional simulator as explained in our previous work [12]. In this simulator, we searched manually for scenarios where the human intends to reach a final goal that is changing over time, represented by a sudden or gradual change of motion direction and plotted the learned intention to reach that goal.

We show two samples of these scenarios in Fig. 2. In Fig. 2(a), we can see the human is traversing in a direction that might not lead to the target for the first few frames and then changing his direction toward the target. It can be observed in Fig. 2(b) that the change in direction is directly related to the stabilization of α_0, namely the intention weight for reaching the target, after decreasing for the first few frames. Figure 2(c) shows an opposite scenario where the human moves toward a target in the first few frames, after which he changes his direction away from the target. Consequently, it can be observed in Fig. 2(d) that the intention weight decreases substantially after the change in the direction.

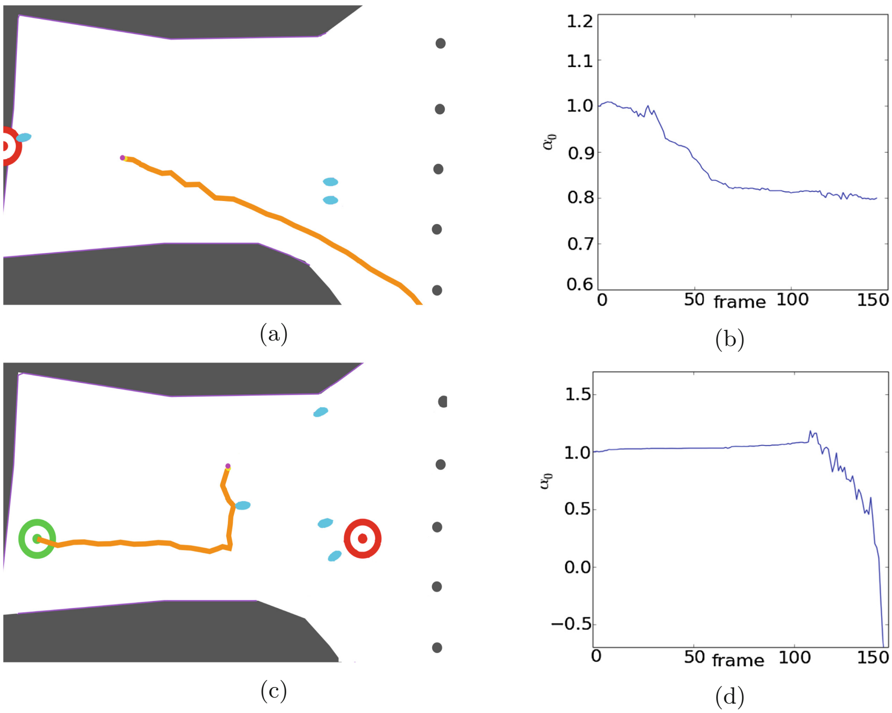

Fig. 2. Two sample trajectories from the dataset mapped into the simulator. The environments in (a) and (c) show static obstacles in dark gray and humans as blue ellipses. The start point is shown in green and the target region is shown in red. The start point in (a) is in the lower-right corner outside the view frame. The orange line depicts the human's trajectory. (b) shows the intention to reach the target for the trajectory in (a), and (d) shows the intention to reach the target for the trajectory in (c). (Color figure online)

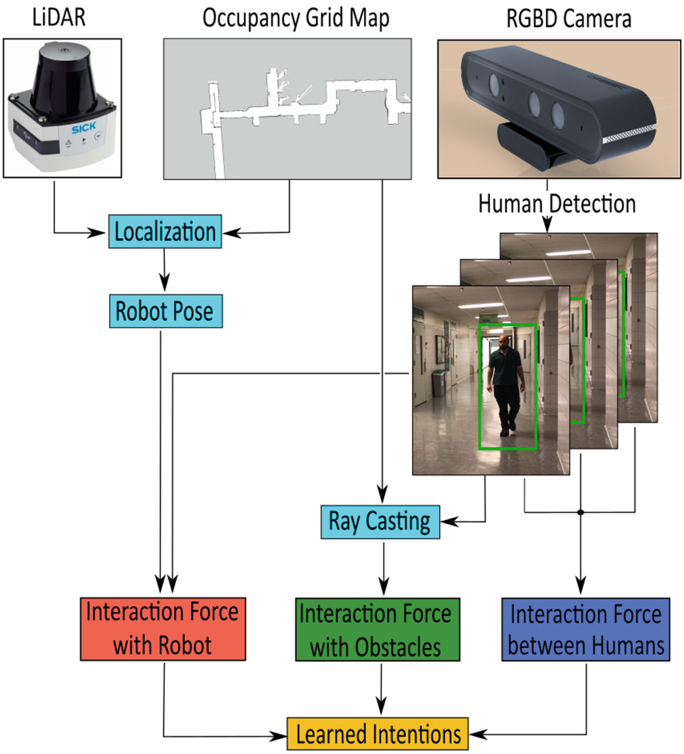

Fig. 3. System implementation on the real robot.

These two scenarios show that our algorithm is able to learn the intention of the human to reach the target and update its belief about the intentions as they change.

3.3 Robot Experiments

Our system implementation on the robot is outlined in Fig. 3. The experiments were conducted on a Segway RMP110 based robot [13] equipped with a SICK TiM LiDAR scanner for localization and an Orbbec Astra Pro RGBD camera for human detection and tracking. To detect humans, we rely on the human detection open-source code presented by the Spencer project [14], which provides a variety of algorithms to detect humans using an RGBD camera.

The robot continuously localizes itself in an occupancy grid map with the aid of the LiDAR. At the same time, its position and velocity as well as the human's are employed to calculate the interaction force between the two entities. To calculate the interaction forces between the human and nearby obstacles, we apply Eq. 2 between his location and the closest obstacle to that location in the occupancy grid map. Finally, the relative positions of the detected humans allow the calculation of the interaction forces between them.

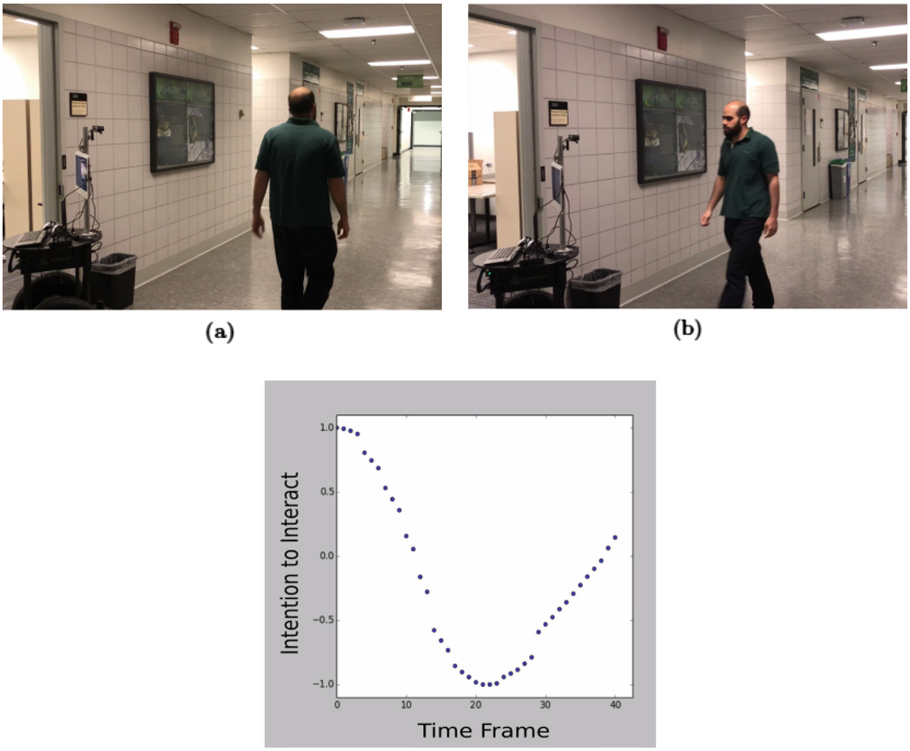

(c)

Fig. 4. (a) Human walking away from the robot. (b) Human approaching the robot to interact with it after he was walking away from it in the previous frame. (c) Plot of the learned intention to interact decreasing when the human was not reaching for the robot and then increasing gradually to the value corresponding to the scene in (b).

During the experiments, the robot was either static or navigating in the environment. In both cases, the robot continuously detected humans around it and learned their intentions. When the robot is navigating, it stops just before the human when it detects an intention to interact. For the sake of clarity and brevity, we only show the intention analysis of the human when the robot is static in the environment, as this analysis is not affected by the robot's movement.

Figure 4 shows an example scenario where the human started its path by moving away from the static robot to come back later and interact with it. The learned intention to interact shows a decrease while the human was moving away from the robot and then increases while the human moves toward the robot. This shows the algorithm was able to adapt to the change in intentions and correct its parameters as soon as the human changed their intentions.

These experiments show the viability of the algorithm when applied on a real robot, where the robot was able to learn the human intentions to interact despite the short range of the camera.

4 Conclusion and Future Work

We presented HIT, a novel approach to track humans while learning their navigational intentions. The proposed method was tested in simulation and real-world scenarios, where in the former we observed the change in the learned intentions as the human changed their direction of motion, and in the latter, we observed the learned intentions of a human to interact with the robot in controlled test scenarios. These experiments proved the ability of the algorithm to learn human navigational intentions and adapt to changes quickly. In addition, we tested the effect of the learned intentions on human tracking and showed its advantage over other tracking algorithms from the literature.

In the future, our approach can be implemented into a hierarchical system where the locally learned intentions can be modeled to infer global human intentions. In such a system, the local intentions can be treated as the observations of a Hidden Markov Model used to learn the latent global intentions similar to [15]. We would expect the implementation of such a system to be around a semantic map representing the function of each object and location in the environment and the connections among them. We also intend to use the proposed approach to improve the legibility and social navigation of service robots in human-populated environments.

References

1. Veloso, M., et al.: CoBots: collaborative robots servicing multi-floor buildings. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS, pp. 5446–5447 (2012)
2. Hawes, N., et al.: The STRANDS project: long-term autonomy in everyday environments. IEEE Robot. Autom. Mag. **24**(3), 146–156 (2017)
3. Chen, X., Ji, J., Jiang, J., Jin, G., Wang, F., Xie, J.: Developing high-level cognitive functions for service robots. In: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems, AAMAS, pp. 989–996 (2010)
4. Patel, U., Hatay, E., D'Arcy, M., Zand, G., Fazli, P.: Beam: a collaborative autonomous mobile service robot. In: Proceedings of the AAAI Fall Symposium Series, pp. 126–128 (2017)
5. Helbing, D., Molnar, P.: Social force model for pedestrian dynamics. Phys. Rev. E **51**(5), 4282–4286 (1995)
6. Luber, M., Stork, J.A., Tipaldi, G.D., Arras, K.O.: People tracking with human motion predictions from social forces. In: Proceedings of the IEEE International Conference on Robotics and Automation, ICRA, pp. 464–469 (2010)
7. Ferrer, G., Zulueta, A.G., Cotarelo, F.H., Sanfeliu, A.: Robot social-aware navigation framework to accompany people walking side-by-side. Auton. Robot. **41**(4), 775–793 (2017)

8. Vasquez, D., Okal, B., Arras, K.O.: Inverse reinforcement learning algorithms and features for robot navigation in crowds: an experimental comparison. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS, pp. 1341–1346 (2014)
9. Ziebart, B.D., Maas, A.L., Bagnell, J.A., Dey, A.K.: Maximum entropy inverse reinforcement learning. In: Proceedings of the AAAI Conference on Artificial Intelligence, AAAI, pp. 1433–1438 (2008)
10. Alahi, A., Goel, K., Ramanathan, V., Robicquet, A., Fei-Fei, L., Savarese, S.: Social LSTM: human trajectory prediction in crowded spaces. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR, pp. 961–971 (2016)
11. Pellegrini, S., Ess, A., Schindler, K., Van Gool, L.: You'll never walk alone: Modeling social behavior for multi-target tracking. In: Proceedings of the 12th IEEE International Conference on Computer Vision, ICCV, pp. 261–268 (2009)
12. Hamandi, M., D'Arcy, M., Fazli, P.: DeepMoTIon: learning to navigate like humans. arXiv preprint arXiv:1803.03719 (2018)
13. Khandelwal, P., et al.: BWIBots: a platform for bridging the gap between AI and human-robot interaction research. Int. J. Robot. Res. **36**(5–7), 635–659 (2017)
14. Linder, T., Arras, K.O.: Multi-model hypothesis tracking of groups of people in RGB-D data. In: Proceedings of the International Conference on Information Fusion, FUSION, pp. 1–7 (2014)
15. Kelley, R., Tavakkoli, A., King, C., Nicolescu, M., Nicolescu, M., Bebis, G.: Understanding human intentions via hidden markov models in autonomous mobile robots. In: Proceedings of the ACM/IEEE International Conference on Human Robot Interaction, HRI, pp. 367–374 (2008)