Learning Human Navigational Intentions

Mahmoud Hamandi and Pooyan Fazli

***Abstract*—We present a novel approach for intention learning in the context of human navigation. 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 show the correlation between the learned preferences and the intentions of the human subject and study these intentions in the context of human-robot interaction and human tracking. We conduct experiments both in simulation and real-world environment to demonstrate the feasibility of the approach and the benefit of employing it to track humans.**

I. Introduction

With the recent developments in artificial intelligence and robotics, robots are increasingly being assigned tasks where they have to navigate in crowded areas. While humans learn over the years to understand one another and plan their paths accordingly, robots still cannot understand 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) [1] to understand humans. Luber *et al.* [2] assume 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.* [3] proposed to control a robot with the social force model to navigate similar to humans by learning the weight of each force from a dataset on human navigation. Although learning one 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.* [4] used the social forces and other human features to learn to navigate directly around humans using an Inverse Reinforcement Learning framework [5], without the need to track humans.

Recently, Alahi *et al.* [6] suggested to use an LSTM-based neural network to learn 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 intention from the end-to-end neural network.

In this work, we suggest to learn the intentions of each human in the scene to follow a moving robot or reach a fixed target directly from the social force model. The proposed approach integrates a Kalman Filter with a motion model based on SFM to track humans in the environment and learn the aforementioned intentions. Observing 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 those weights and the intention of the humans.

II. Human Tracking and Intention Learning

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

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

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. In the case where the target is the actual robot, F_{target} will be included in F_{robot}. Each of the forces is exponentially related to the distance between the two objects enforcing it, with the exception of the last force that is linearly related to the human speed. αs 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 completely, the corresponding α should be zero, if he follows the robot, it should be positive, while if he runs away from the robot, the corresponding α should be negative.

Mathematically the forces are represented as follows:

$$F_k = A_k e^{(\delta_k - \|\mathbf{d_k}\|)/B_k} \frac{\mathbf{d_k}}{\|\mathbf{d_k}\|}, \quad (2)$$

where k is a member of the set $O = \{robot, human, obstacle\}$, A_k, δ_k, and B_k are fixed parameters specific to each member of the set, $\mathbf{d_k}$ is the distance vector between the human and the corresponding object in O, and $\|\mathbf{d_k}\|$ is its norm. Ferrer *et al.* [3] present a framework to learn A_k, δ_k, and B_k. On the other hand:

$$F_{target} = \kappa(\mathbf{v}^* - \mathbf{v}), \quad (3)$$

where κ is a fixed parameter, $\mathbf{v}^*$ is the desired velocity (which we assume to be the velocity of the moving target), and $\mathbf{v}$ is the actual human velocity. This force emphasizes the difference between the desired and actual velocity making it suitable for a moving target, where the human has to match their velocity with the target's. However, this force does not adapt to the velocity preference of each human when following a fixed target, where the desired velocity can be considered as a personal preference. Instead, we model the force to a fixed target as:

$$F_{target} = \kappa \frac{\mathbf{v}}{\|\mathbf{v}\|}(1 - cos\theta), \tag{4}$$

where 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 location.

As such, for a set of learned weights the social force can be calculated based on the observed environment, and we can model the human motion as presented in [2]:

$$\begin{bmatrix} \mathbf{x_t} \\ \mathbf{v_t} \end{bmatrix} = \begin{bmatrix} \mathbf{x}_{t-1} + \mathbf{v}_{t-1}\Delta \mathrm{t} + \frac{\mathbf{F}}{2}\Delta \mathrm{t}^2 \\ \mathbf{v}_{t-1} + \mathbf{F}\Delta \mathrm{t} \end{bmatrix}, \tag{5}$$

where $\mathbf{x_i}$ is the position of the human, $\mathbf{v_i}$ is the velocity of the human at time step i, 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 position after each observation. In addition, our framework learns the underlying weights that could lead to the observed position and reduce its error with the predicted. As such, the tracking of each human starts with an approximate value of each α and updates them for each human as we receive more observations. Specifically, we update the parameters to reduce the difference between the predicted and the observed human location. This can be achieved as the observed location 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 note the difference between the two positions as $diff(F)$ and learn each α as :

$$\alpha_{i,t} = \alpha_{i,t-1} + diff(F).\mathbf{f}_i.\gamma, \tag{6}$$

where $\mathbf{f}_i$ 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 in the environment. These forces can show us if the human is trying to reach the designated target or trying to follow (or escort) a moving robot.

III. Experiments

In this paper, 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 [7].
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 a fixed point in the environment. 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 following a human. As the human stops or no longer matches its velocity to the robot's, the robot has to detect the change in the human intention and stop.

Thus far, we have completed the first two parts and successfully coded the third part. In the first two parts, we rely on the ETH dataset mapped to a 2-dimensional simulator as explained in our previous work [8]. For the real robot experiments, we implemented our system in Robot Operating System [9]. The system controls the robot to follow the closest human in the environment. To detect humans, we rely on the human detection open-source code presented by the Spencer project [10], which provides a variety of algorithms to detect humans in an RGB-D camera.

IV. Results

Figure 1 shows two sample trajectories of humans as they change their intentions to reach the target and the corresponding changes in the learned weights. In Figure 1(a) we can see the human is traversing in a direction that does not lead to the target for the first few frames and then changing his direction toward the target. It can be observed in Figure 1(b) that the change in direction is directly related to a stabilization of α_0 -the intention weight for reaching the target- after decreasing for the first few frames. Figure 1(c) shows an opposite scenario where the human moves toward the target in the first few frames, after which he changes his direction away from the target. Consequently, it can be observed in Figure 1(d) that the intention weight decreased substantially after the change in the direction.

Our experiments tracked the humans and compared the error between each predicted and observed position. The results show an average error of 0.10 m between the two positions. This metric shows that the instant prediction of the future position is accurate. In the final submission, we will compare our results against other algorithms in the literature ([2], [6]), where the first employs a method similar to ours with fixed weights, and the second trains an LSTM based neural network to track humans in the environment. It should be noted that our tracking tests are conducted only to prove the ability of the system to learn the SFM weights and not to show a long-term tracking ability. In fact, we cannot confirm without testing if our algorithm will be able to track humans accurately in a long horizon or not, as it is not designed to do so.

V. Conclusion and Future Work

We demonstrated an algorithm able to track humans while learning their intentions online and conducted experiments to assess both the tracking and intention learning capability of the method. The experiments showed the proposed method is able to cope with changes in human intentions and track humans

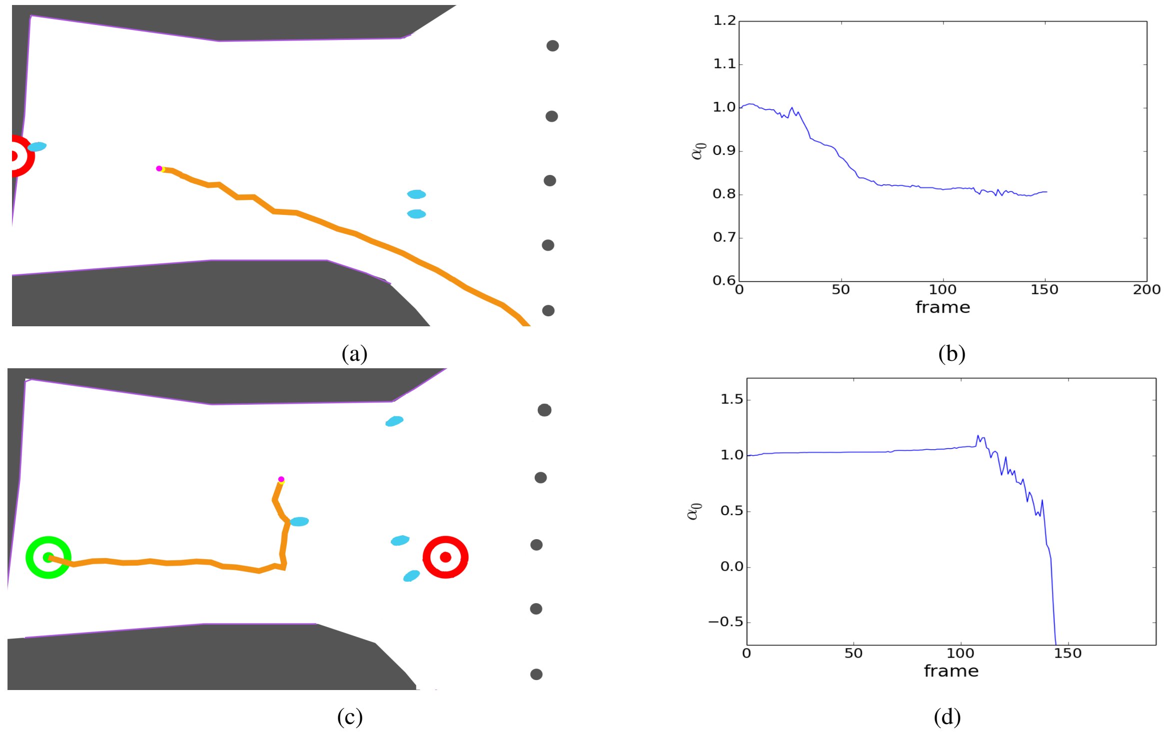

Fig. 1: Two sample trajectories from the dataset mapped into the simulator. The simulated environments in (a) and (c) show static obstacles in dark grey and humans as blue ellipses. The start point is shown in green and the target region is shown in red. The orange line depicts the humans trajectory. (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).

even when their intentions changed. In addition, we suggested a scenario where we can investigate the algorithm on a real robot. We expect the robot to learn the intention of the human to interact with it and stop once it detects that the human is no longer trying to match his velocity to the robot's.

References

[1] D. Helbing and P. Molnar, "Social force model for pedestrian dynamics," *Physical review E*, vol. 51, no. 5, p. 4282, 1995.

[2] M. Luber, J. A. Stork, G. D. Tipaldi, and K. O. Arras, "People tracking with human motion predictions from social forces," in *Proceedings of the IEEE International Conference on Robotics and Automation, ICRA*, 2010, pp. 464–469.

[3] G. Ferrer, A. G. Zulueta, F. H. Cotarelo, and A. Sanfeliu, "Robot social-aware navigation framework to accompany people walking side-by-side," *Autonomous robots*, vol. 41, no. 4, pp. 775–793, 2017.

[4] D. Vasquez, B. Okal, and K. O. Arras, "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*, 2014, pp. 1341–1346.

[5] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning." in *Proceedings of the AAAI Conference on Artificial Intelligence, AAAI*, 2008, pp. 1433–1438.

[6] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social lstm: Human trajectory prediction in crowded spaces," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2016, pp. 961–971.

[7] S. Pellegrini, A. Ess, K. Schindler, and L. Van Gool, "You'll never walk alone: Modeling social behavior for multi-target tracking," in *Proceedings of the 12th IEEE International Conference on Computer Vision, ICCV*, 2009, pp. 261–268.

[8] M. Hamandi, M. D'Arcy, and P. Fazli, "Deepmotion: Learning to navigate like humans," *arXiv preprint arXiv:1803.03719*, 2018.

[9] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "Ros: an open-source robot operating system," in *ICRA workshop on open source software*, vol. 3, no. 3.2. Kobe, Japan, 2009, p. 5.

[10] T. Linder and K. O. Arras, "Multi-model hypothesis tracking of groups of people in rgb-d data," in *Proceedings of the 17th International Conference on Information Fusion, FUSION*, 2014, pp. 1–7.