

Predicting Pedestrian Trajectories for Robot Navigation (Extended Abstract)

Sujeong Kim¹ Stephen J. Guy² Wenxi Liu³ David Wilkie¹ Rynson W. H. Lau³ Ming C. Lin¹
Dinesh Manocha¹

Abstract—We introduce a novel online method to predict pedestrian trajectories using agent-based velocity-space reasoning for improved human-robot interaction and collision-free navigation. Our formulation uses velocity obstacles to model the trajectory of each moving pedestrian in a robot’s environment and improves the motion model by adaptively learning relevant parameters based on sensor data. The resulting motion model for each agent is computed using statistical inferencing techniques, including a combination of Ensemble Kalman filters and a maximum-likelihood estimation algorithm. This allows a robot to learn individual motion parameters for every agent in the scene at interactive rates. We highlight the performance of our approach for collision-free robot navigation among pedestrians based on noisy, sparsely-sampled data and highlight the results in our simulator.

I. INTRODUCTION

Robots are becoming increasingly common in everyday life. As more robots are introduced into human surroundings, it becomes increasingly important to develop safe and reliable techniques for human-robot interaction. Robots working around humans must be able to successfully navigate to their goal positions in dynamic environments with multiple people moving around them. A robot in a dynamic environment thus needs the ability to sense, track, and predict the position of all people moving in its workspace to navigate complicated environments without collisions.

Sensing and tracking the position of moving humans has been studied in robotics and computer vision, e.g. [1], [2], [3]. These methods often depend upon an a priori motion model fitted for the scenario in question. However, these motion priors, which are typically generalized motions rather than motions specific to an individual’s movements, usually do not accurately capture or predict the trajectory of each pedestrian. For example, we frequently observe uncommon pedestrian motions, such as moving against the flow of other agents in a crowd, or quick velocity changes to avoid collisions. In order to address these issues, many pedestrian tracking algorithms use multi-agent or crowd motion models based on local collision avoidance [4], [5]. These multi-agent interaction models effectively capture short-term deviations from goal-directed paths, but in order to do so, they must

already know each pedestrian’s goal positions; they often use handpicked destination information, or other heuristics that require prior knowledge about the environment. As a result, these techniques have important limitations: they are unable to account for unknown environments with multiple destinations, or times when pedestrians take long detours or make unexpected stops. In general, the assumption that the destination information remains constant can often result in large errors in predicted trajectories.

In this work, we seek to overcome these limitations by presenting a new motion model (BRVO) that is built on agent-based, velocity-space reasoning combined with Bayesian statistical inference; BRVO can provide an individualized motion model for each agent in a robot’s environment. We apply Ensemble Kalman Filtering (EnKF) to estimate the parameters for a human motion model based on *Reciprocal Velocity Obstacles* (RVO) [6]. We infer the most likely state for each observed person: its position, velocity, and goal velocity. More over, our formulation is capable of dynamically adjusting the parameters for each individual in the presence of sensor noise and model uncertainty. We address the problems associated with fixed goal positions by integrating learning into our predictive framework and by adjusting short-term steering information.

We use BRVO to compute collision-free trajectories for robots, which takes into account kinematic constraints; our approach can compute these paths at interactive rates in scenarios with dozens of pedestrians. Our experiments with real-world pedestrian datasets demonstrate that BRVO can improve planning in uncertain environments: environments with dynamic obstacles or with sensor limitations, including noisy or sparse sensor inputs.

II. BAYESIAN-RVO

Given an agent’s state \mathbf{x}_k at timestep k , we use the RVO collision-avoidance motion model, denoted here as f , to predict the agent’s next state \mathbf{x}_{k+1} . We denote the state prediction error of f at each time step as \mathbf{q} . This leads to our motion model of:

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k) + \mathbf{q}. \quad (1)$$

Additionally, we assume that the sensing of the robot can be represented by a function h that provides an observed state \mathbf{z}_k , which is a function of the person’s true state $\hat{\mathbf{x}}_k$ plus some noise from the sensing processing, which is denoted

¹Sujeong Kim, David Wilkie, Ming C. Lin and Dinesh Manocha are with the Department of Computer Science, University of North Carolina at Chapel Hill.

²Stephen J. Guy is with the Department of Computer Science, University of Minnesota.

³Wenxi Liu and Rynson W. H. Lau are with the Department of Computer Science, City University of Hong Kong.

as \mathbf{r} . That is:

$$\mathbf{z}_k = h(\hat{\mathbf{x}}_k) + \mathbf{r}. \quad (2)$$

We assume that the error terms q and r are independent at each timestep, and that they follow a zero-mean Gaussian distribution with covariances Q and R , respectively.

We represent each agent's state, \mathbf{x} , as the six-dimensional set of RVO parameters:

$$\mathbf{x} = \begin{bmatrix} \mathbf{p} \\ \mathbf{v} \\ \mathbf{v}_{pref} \end{bmatrix}, \quad (3)$$

where \mathbf{p} is the agent's position, \mathbf{v} the velocity, and \mathbf{v}_{pref} the preferred velocity. The crowd dynamics model f is then:

$$f\left(\begin{bmatrix} \mathbf{p} \\ \mathbf{v} \\ \mathbf{v}_{pref} \end{bmatrix}\right) = \begin{bmatrix} \mathbf{p} + \mathbf{v}\Delta t \\ \underset{\mathbf{v}_{pref}}{\operatorname{argmin}_{\mathbf{v} \in ORCA}} \|\mathbf{v} - \mathbf{v}_{pref}\| \end{bmatrix}. \quad (4)$$

We assume that the robot has the ability to sense the relative positions of the pedestrians. This assumption leads to a simple h function of the form

$$h\left(\begin{bmatrix} \mathbf{p} \\ \mathbf{v} \\ \mathbf{v}_{pref} \end{bmatrix}\right) = \mathbf{p}. \quad (5)$$

We use EnKF, a sampling-based extension of Kalman Filtering, to estimate the corresponding agent state for each pedestrian. EnKF takes as input an estimate of the prediction error Q and observations $\mathbf{z}_0 \dots \mathbf{z}_k$ and produces an estimate of the true pedestrian states \mathbf{x}_k as a distribution of likely states \mathcal{X}_k . EnKF provides an estimate of the true distribution of likely pedestrian states by representing it with these m samples. We use the EM-algorithm [7] to improve our estimate of Q , which will in turn improve the quality of the learning and the predictiveness of the method. More details can be found in [8].

III. ROBOT NAVIGATION WITH BRVO

One potential application of BRVO is for safer navigation for autonomous robot vehicles through areas of dense pedestrian traffic. Here we describe a method to integrate BRVO with the GVO (Generalized Velocity Obstacle) motion-planning algorithm proposed in [9] to achieve more effective navigation of a simulated car-like robot through a busy walkway. The robot uses BRVO technique to predict the motion of each pedestrian as it navigates through the crowded walkway to its goal position. We assume a non-cooperative environment, where the pedestrians may not actively avoid collisions with the robot and the robot must assume 100% responsibility for collision avoidance (i.e., asymmetric behavior).

The robot's configuration is represented as its position (x, y) and the orientation ϕ . The robot has controls u_s and u_ϕ ,

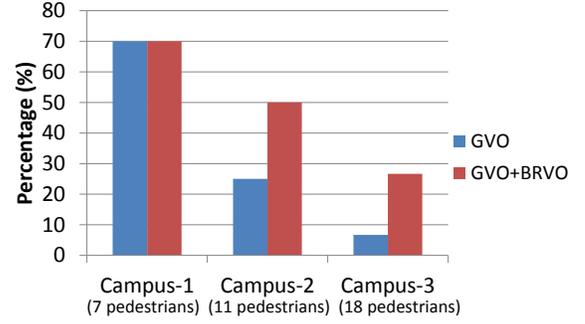


Fig. 1: Performance of the robot navigation We measured the percentage of the trajectories in which the robot reached farther than halfway to the goal position without any collision during the data sequence, using only GVO (blue bars) and using GVO with BRVO (red bars), both with 15cm sensor noise. In many of these scenarios, using GVO for navigation often caused the robot to stop moving and to avoid passing through the crowd to avoid collisions, (the freezing robot problem). However, GVO-BRVO algorithm improves the navigation especially for more challenging scenarios with more pedestrians.

which are the speed and steering angle of the robot, respectively. Its constraints are defined as follows:

$$\begin{aligned} x'(t) &= u_s \cos \theta(t), \\ y'(t) &= u_s \sin \theta(t), \\ \theta'(t) &= u_s \frac{\tan u_\phi}{L}, \end{aligned} \quad (6)$$

where L is the wheelbase of the robot. We assume $L = 1\text{m}$.

Assuming that the control remains constant for the time interval, the robot's position $R(t, u)$ at time t given the control u can be derived as follows:

$$R(t, u) = \begin{pmatrix} \frac{1}{\tan(u_\phi)} \sin(u_s \tan(u_\phi)t) \\ -\frac{1}{\tan(u_\phi)} \cos(u_s \tan(u_\phi)t) + \frac{1}{\tan(u_\phi)} \end{pmatrix}. \quad (7)$$

For further details, please refer to [9].

We use following datasets to measure the performance of the robot navigation.

Campus This video was recorded from the top of the ETH main building in Zurich [4]. We extracted three sequences from this data, each containing about 10 seconds of pedestrian interaction: Campus-1 (7 pedestrians), Campus-2 (11 pedestrians), Campus-3 (18 pedestrians).

Students This video was recorded from a street view [10]. The dataset contains the motion of 434 students over a period of roughly 3 minutes, tracked manually.

The robot is given an initial position at one side of the walkway and is asked to move through the pedestrians to the opposite side. Given the importance of safety in pedestrian settings, if at any point the robot fails to find collision free trajectory which moves forward along the path, the robot will stop or turn back towards the start. Given this setup, we

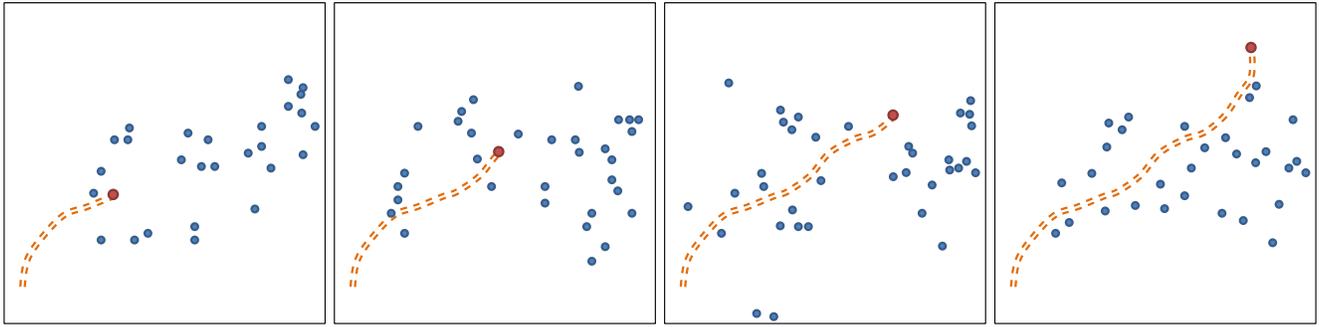


Fig. 2: Example robot trajectory navigating through the crowd in *Students* dataset. Blue circles represent current pedestrian positions, red circles are the current position of the robot, and orange dotted lines are the previous positions of the robot.

evaluate the percentage of times the robot is able to make it more than halfway through the pedestrian crossing without colliding with a pedestrian, or needed to stop or turn back. We run experiments on all the three sequences, assuming a sensing error of $\pm 15\text{cm}$ noise in positional estimates, and a sampling rate of 2.5Hz to allow adequate time for any visual processing needed to detect pedestrians in the sensed area. We collect the mean of 30 runs for each sequence; the robot's initial position and goal position are randomly chosen for each run.

Fig. 1 shows the result using only GVO (blue bars) and the result using GVO with BRVO (red bars). As can be seen from the graph, as the scenarios got more dense, robots navigating using just GVO tends to avoid collisions by staying still, displaying the same freezing robot problem discussed by other researchers (see for example [11]). However, the combined GVO-BRVO algorithm improves the navigation, especially for more challenging scenarios with more pedestrians where it more than doubled the task completion rate.

This series of experiments indicates that BRVO can improve planning in uncertain environments: environments with dynamic obstacles or with sensor limitations, including noisy or sparse sensor inputs. Though better prediction does not guarantee better navigation, and the freezing robot problem can still occur [11], we believe that better prediction algorithms can improve the robotic navigation.

We also believe that BRVO can be combined with other robot navigation methods in a cooperative setup, since it improves the performance of RVO-based motion models in noisy-data situations, as discussed in [12]. More importantly, BRVO does not need any prior knowledge about the scene, like destination or goal positions and motion priors, which can be a considerable benefit for navigating robots or autonomous wheelchairs in real-world scenes.

ACKNOWLEDGMENT

We would like to thank Jur van den Berg for his help. This research is supported in part by ARO Contract W911NF-10-1-0506, NSF awards 0917040, 0904990, 100057 and 1117127, and Intel.

REFERENCES

- [1] M. Luber, J. Stork, G. Tipaldi, and K. Arras, "People tracking with human motion predictions from social forces," in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2010, pp. 464–469.
- [2] M. Rodriguez, S. Ali, and T. Kanade, "Tracking in unstructured crowded scenes," in *Proc. of the IEEE 12th International Conference on Computer Vision*, 2009, pp. 1389–1396.
- [3] L. Kratz and K. Nishino, "Tracking pedestrians using local spatio-temporal motion patterns in extremely crowded scenes," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, no. 99, p. 11, 2011.
- [4] S. Pellegrini, A. Ess, K. Schindler, and L. Van Gool, "You'll never walk alone: Modeling social behavior for multi-target tracking," in *Proc. of the IEEE 12th International Conference on Computer Vision*, 2009, p. 261268.
- [5] K. Yamaguchi, A. Berg, L. Ortiz, and T. Berg, "Who are you with and where are you going?" in *Proc. of the 2011 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2011, pp. 1345–1352.
- [6] J. Van Den Berg, S. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," *Robotics Research*, p. 319, 2011.
- [7] G. J. McLachlan and T. Krishnan, *The EM Algorithm and Extensions (Wiley Series in Probability and Statistics)*, 2nd ed. Wiley-Interscience, Mar. 2008.
- [8] S. Kim, S. Guy, W. Liu, R. Lau, M. Lin, and D. Manocha, "Predicting pedestrian trajectories using velocity-space reasoning," in *The Tenth International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2012.
- [9] D. Wilkie, J. van den Berg, and D. Manocha, "Generalized velocity obstacles," in *Proceedings of the 2009 IEEE/RSJ international conference on Intelligent robots and systems*, ser. IROS'09. Piscataway, NJ, USA: IEEE Press, 2009, pp. 5573–5578. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1732643.1732952>
- [10] A. Lerner, E. Fitusi, Y. Chrysanthou, and D. Cohen-Or, "Fitting behaviors to pedestrian simulations," in *Symp. on Computer Animation*, 2009, p. 199208.
- [11] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, 2010, pp. 797–803.
- [12] P. Trautman, J. Ma, R. M. Murray, and A. Krause, "Robot navigation in dense human crowds: the case for cooperation," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, 2013, pp. 2153–2160.