

Minimally Invasive Social Navigation

Stefan H. Kiss, K. Y. Cadmus To, Chanyeol Yoo, Robert Fitch and Alen Alempijevic

University of Technology Sydney, NSW Australia

{stefan.h.kiss,cadmus.to}@student.uts.edu.au, {chanyeol.yoo,robert.fitch,alen.alempijevic}@uts.edu.au

Abstract

Integrating mobile robots into human society involves the fundamental problem of navigation in crowds. This problem has been studied by considering the behaviour of humans at the level of individuals, but this representation limits the computational efficiency of motion planning algorithms. We explore the idea of representing a crowd as a flow field, and propose a formal definition of path quality based on the concept of invasiveness; a robot should attempt to navigate in a way that is minimally invasive to humans in its environment. We develop an algorithmic framework for path planning based on this definition and present experimental results that indicate its effectiveness. These results open new algorithmic questions motivated by the flow field representation of crowds and are a necessary step on the path to end-to-end implementations.

1 Introduction

Humans and robots co-existing in an environment require an understanding of each other’s motion to perform safe interactions. For humans, the ability to predict the motion of others within an environment enables implicit path planning to reach intended goals in a self-centred manner. For robots, the capability to anticipate human motion can facilitate more fluent interaction [Hoffman and Breazeal, 2007]. We are interested in this fundamental problem in the case of dense human crowds, where a mobile robot must consider the motion of many humans in order to navigate. This problem is important to any application of mobile robots in crowded spaces, such as public indoor and outdoor areas. Our main focus is on how to use observed human motion to develop a path planning framework that seeks to find a path to a goal region while minimising interference with humans.

When individuals move in a crowd, the decision-making process is more than merely reactive be-

haviour [Trautman *et al.*, 2015]; humans incorporate reasoning about the possible actions of others. Ignoring inter-dependencies in human motion may lead to overly conservative motion planning [Turnwald and Wollherr, 2019]. This issue can be avoided if the robot can anticipate humans’ cooperative collision avoidance and can take into account the goal-driven nature of human decision making.

Encoding human cooperative interaction in motion planning can be coarsely divided into approaches that (1) build upon sets of social rules, or (2) consider reactive behaviours. Methods that encode social rules have evolved from Helbing’s seminal work on social forces [Helbing and Molnár, 1995] to incorporate relative velocity and intended direction of travel [Moussaïd *et al.*, 2009]. Deep learning strategies have recently come into focus [Alahi *et al.*, 2016]. Alternatively, [Van den Berg *et al.*, 2008] propose to take into account reactive behaviour, assuming humans can be modelled as agents that make similar collision-avoidance reasoning.

Integrating either of these approaches into a motion planning framework requires a means of evaluating pairwise interactions [Rudenko *et al.*, 2018] and of inferring the goal of each agent in order to predict its future trajectory [Vemula *et al.*, 2017; Bevilacqua *et al.*, 2018]. The computation required to implement these operations depends on the number of pedestrians in the scene, and can become prohibitive in large crowd densities [Trautman *et al.*, 2015]; the computational complexity of modelling pairwise interactions is linear in the number of pedestrians in the best case. Additionally, modelling accuracy is typically also dependent on time, contributing poor predictions at larger time scales.

Bootstrapping planning by exploiting social rules of spatial occupancy [Alempijevic *et al.*, 2013] or relative velocity [Bera *et al.*, 2019] has proved to be effective in socially-aware robot navigation in low to medium-density environments. However, to cater towards higher densities we argue that looking macroscopically at flow (both relative and average) will enable more scalable so-

lutions to be developed. Recent analysis of human-robot interaction around emergency egress points indicates interplay between the direction of pedestrian flow [Chen *et al.*, 2018] and the relative velocity between robot and pedestrian. In addition, the energy cost in the navigation process has been identified as being of vital importance [Chen *et al.*, 2018]. We therefore posit that the objective function of a motion planner should be inherently goal driven (such as the goal-directed motion of humans) while social compliance should be related to minimal invasiveness (minimal disruption to flow of others).

To operationalise these ideas, we encode human motion macroscopically as a flow field and propose a corresponding formal definition of invasiveness given the macroscopic properties of the crowd. This formulation is defined with respect to both deterministic and stochastic flows, and matches our intuition of invasiveness, which is a subjective concept that has no well-established mathematical representation. We then develop an algorithmic framework for non-myopic minimally-invasive social navigation using this definition. The non-myopic property, in this context, means that the algorithm aims to find a globally optimal solution as opposed to one that is locally optimal or reactive. Similar work has been done to navigate through wind fields or ocean currents [Lee *et al.*, 2017; Yoo *et al.*, 2016; Lee *et al.*, 2019b; To *et al.*, 2019b; Yoo *et al.*, 2019; To *et al.*, 2019a]. To the best of our knowledge, our approach is the first to formulate social navigation problem using flow fields. Our framework is presented using the well-known *optimal Probabilistic RoadMap* (PRM*) algorithm for sampling-based motion planning, but other variants could be easily substituted.

We illustrate the behaviour of our algorithms in simulation of artificial scenarios that allow us to systematically examine a set of macroscopic (high-level) features of crowd-based flow fields. Further, we consider an indoor scenario that shows how these features would appear in an application of a robot moving through a crowded building. Results show that the paths generated by our algorithm are less-invasive than those of comparison methods.

The main contributions of this paper are a novel formal definition of invasiveness given the macroscopic properties of a crowd, and a novel framework for non-myopic, minimally-invasive social navigation. The significance of this work is that it contributes a necessary first step towards the larger goal of finding computationally efficient solutions that are suitable for real-world applications. As a result of this work, a number of new problems arise that are motivated by the flow-field representation. We present an overview of these problems in the concluding section of the paper.

2 Problem Formulation

The robot’s environment consists of obstacles and a pedestrian crowd in 2D space $\mathbb{X} \subset \mathbb{R}^2$. Our robot can manoeuvre through it independent of the crowd’s behaviour. Hence the first-order dynamics of this system is:

$$\dot{\mathbf{x}}_r = \mathbf{v}_r, \quad (1)$$

where \mathbf{x}_r is the position of the robot in 2D space free from obstacles, and \mathbf{v}_r is its velocity in some limited velocity space $\mathbb{V} \subseteq \mathbb{R}^2$. The trajectory of the robot over time t is denoted as $\mathbf{x}_r(t)$.

The interaction of the robot’s presence in the crowd is measured as *social invasiveness* \mathcal{I}_r . This invasiveness arises from the instantaneous amount of influence of the robot amongst the pedestrians. This quantity is a function of robot and environment states. Existing methods to measure invasiveness include social force [Helbing and Molnár, 1995; Moussaïd *et al.*, 2009] and proximity-based discomfort [Hoogendoorn and Bovy, 2003; Treuille *et al.*, 2006].

In the interest of minimising the overall intrusion of the robot traversal through crowds, the path planning problem is formally defined as a minimisation of total interference along the path:

Problem 1 (Minimally invasive path planning in crowds). *Given a measure of invasiveness with the crowd \mathcal{I}_r , the robot’s dynamics (1), its initial position \mathbf{x}_{init} , and goal position \mathbf{x}_{goal} , find the optimal trajectory*

$$\mathbf{x}_r^* = \arg \min_{\mathbf{x}_r} \int_0^T \mathcal{I}_r dt \quad \text{for } T \in \mathbb{R}_{\geq 0}, \quad (2)$$

such that

$$\begin{aligned} \mathbf{x}_r(0) &= \mathbf{x}_{init}, \\ \mathbf{x}_r(T) &= \mathbf{x}_{goal}. \end{aligned}$$

To address this problem we must first formulate an analytical measure of social invasiveness that alleviates the computational bottleneck during planning. The optimum trajectory is then found through the application of sampling-based motion planning minimising the total invasiveness along the trajectory.

3 Social Invasiveness

Measures of social invasiveness can be formulated by treating pedestrians as point masses, however the minimum time complexity of these approaches are linear. This implies that their computation becomes prohibitively expensive in dense crowds. Additionally, modelling accuracy is typically also dependent on the prediction horizon time, contributing poor predictions at larger time scales.

However, we propose that there exists properties of pedestrian crowds that remain fairly consistent over time [Treuille *et al.*, 2006; Karamouzas *et al.*, 2014] which describe their macroscopic behaviour. We present our formulation of social invasiveness based on these ideas.

3.1 Invasiveness in Crowd Flows

Before we present our analytical definition of social invasiveness, we begin by introducing the continuous representation of crowds. Consistent crowd behaviours can be modelled as time-invariant flow fields with density. A *crowd flow* is defined as the 2D density and velocity field pair (ρ, \mathbf{V}) , where

$$\begin{aligned} \rho : \mathbb{X} &\rightarrow \mathbb{R}_{\geq 0}, \\ \mathbf{V} : \mathbb{X} &\rightarrow \mathbb{R}^2. \end{aligned}$$

The density of the crowd $\rho(\mathbf{x})$ is a scalar field that describes the expected number of pedestrians per unit of area at \mathbf{x} . The crowd velocity $\mathbf{V}(\mathbf{x})$ is a vector field corresponding to the speed and direction of the pedestrians at \mathbf{x} .

We can now define the invasiveness of the robot as

$$\mathcal{I}_r = \rho_r \|\Delta \mathbf{V}_r\|^2, \quad (3)$$

where

$$\begin{aligned} \rho_r &= \rho(\mathbf{x}_r), \\ \mathbf{V}_r &= \mathbf{V}(\mathbf{x}_r), \text{ and} \\ \Delta \mathbf{V}_r &= \mathbf{V}_r - \mathbf{v}_r. \end{aligned}$$

This formulation consists of two factors which account for different aspects of invasiveness. The first factor $\rho \|\Delta \mathbf{V}\|$ is proportional to the expected number of pedestrians per second that cross a *region of interaction* around the robot. The constant of proportionality can be thought of as related to the cross-sectional width of this region. A second factor $\|\Delta \mathbf{V}\|$ assumes that the interference caused upon interaction is proportional to the velocity difference between the two parties, and is therefore inversely proportional to the available time to act.

Equation (3) matches our intuition of what we expect to be invasive. The robot will be less invasive if it travels through areas of low crowd density or if it travels at the same velocity as the crowd.

Crowd flows start to become poor models of pedestrian behaviour as multiple flows overlap [Jodoin *et al.*, 2013]. The velocity at each point can no longer be described by a single vector as pedestrian intents are not dictated solely by their position in space.

3.2 Invasiveness in Stochastic Crowd Flows

To account for multiple overlapping pedestrian flows, we consider *stochastic crowd flows*. More specifically, we

augment the concept of crowd flows with a variance of velocity. This allows us to rigorously reinterpret the crowd flow in a probabilistic sense, and redefine social invasiveness in expectation:

$$\mathcal{I}_r = \mathbb{E} \left[\rho_r \|\Delta \mathbf{V}_r\|^2 \right],$$

from which we can derive

$$\mathcal{I}_r = \rho_r \left(\|\boldsymbol{\mu}_{\mathbf{V}_r} - \mathbf{v}_r\|^2 + \sigma_{\mathbf{V}_r}^2 \right), \quad (4)$$

where

$$\begin{aligned} \boldsymbol{\mu}_{\mathbf{V}} : \mathbb{X} &\rightarrow \mathbb{R}^2, & \boldsymbol{\mu}_{\mathbf{V}} : \mathbf{x} &\mapsto \mathbb{E} [\mathbf{V}(\mathbf{x})], \\ \sigma_{\mathbf{V}}^2 : \mathbb{X} &\rightarrow \mathbb{R}_{\geq 0}, & \sigma_{\mathbf{V}}^2 : \mathbf{x} &\mapsto \text{Var} [\mathbf{V}(\mathbf{x})], \\ \boldsymbol{\mu}_{\mathbf{V}_r} &= \boldsymbol{\mu}_{\mathbf{V}}(\mathbf{x}_r), \text{ and} \\ \sigma_{\mathbf{V}_r}^2 &= \sigma_{\mathbf{V}}^2(\mathbf{x}_r). \end{aligned}$$

Note that $\text{Var}[\cdot]$ here indicates the *scalar* variance, defined as the expected value of the squared euclidean distance from the mean: $\text{Var}[\mathbf{x}] = \mathbb{E} [\|\mathbf{x} - \mathbb{E}[\mathbf{x}]\|^2]$.

These additional macroscopic properties of the crowd can be understood intuitively. The mean velocity $\boldsymbol{\mu}_{\mathbf{V}}(\mathbf{x})$ describes the typical flow of pedestrians around point \mathbf{x} . The variance of velocity $\sigma_{\mathbf{V}}^2(\mathbf{x})$ describes the irregularity of the flow around that point. A low variance indicates a coherent flow of pedestrians in the mean direction, whereas a large velocity variance with a mean close to zero indicates that pedestrians walk through the region from multiple directions with no dominant flow. In practice, variance is always expected to be greater than zero as crowd flows approximate multiple individuals with different destinations.

4 Minimally Invasive Path Planning in Crowd Flows

To find a minimally invasive trajectory with general path planning algorithms, we can express invasiveness between two points as a cost function. Assuming that the robot travels in a straight line between points, we need to determine its speed along the path.

4.1 Minimally Invasive Speed

As the robot travels along a straight line segment, the invasiveness along an infinitesimal step of length ds can be expressed as

$$\mathcal{I}_r dt = \mathcal{I}_r \frac{ds}{v_r}, \quad (5)$$

where v_r is the robot's speed towards the goal. We can now determine the minimally invasive speed for the infinitesimal step:

$$\begin{aligned} v_r^* &= \arg \min_{v_r} \frac{\mathcal{I}_r}{v_r} \\ v_r^* &= \sqrt{\|\boldsymbol{\mu}_{\mathbf{V}_r}\|^2 + \sigma_{\mathbf{V}_r}^2} \end{aligned} \quad (6)$$

Table 1: Parameters and results for comparing macroscopic properties

Scenario	macroscopic properties			interference of trajectory	
	ρ 1/m ²	$\mu\mathbf{v}$ m/s	$\sigma^2\mathbf{v}$ m ² /s ²	social	naïve
density	0.5 – 1.5	0	1	17.1	27.7
velocity	1	$\ \cdot\ = 1$	0.25	5.5	22.3
variance	1	0	0.25 – 1.25	23.1	24.5

Interestingly, this result implies that the optimal instantaneous speed does not depend on the direction of motion itself.

This analytical value of minimally invasive speed allows us to calculate the invasiveness along a path by integrating with the appropriate discretisation.

4.2 Integration with Optimal Path Planner

We use PRM* to generate a bidirectional graph using samples randomly drawn from a bounded 2D configuration space of position, including the initial and goal positions. The weights of the graph edges correspond to the invasiveness between two nodes. In this context, the straight line assumption for invasiveness along these edges becomes negligible as distance between considered samples approaches zero as the number of sample points increases [Karaman and Frazzoli, 2011].

Dijkstra’s algorithm is then applied on the resulting graph with respect to the starting position to find a minimally invasive tree to every other node in the graph. Trajectories generated this way simultaneously avoid obstacles while approaching the optimal trajectory as the number of samples increases.

5 Experiments

We test our algorithm on different simple scenarios demonstrating its behaviour when one macroscopic property is varied and a complex scenario where a combination of them is varied. The macroscopic parameters we consider is the crowd density, velocity, and the variance of the velocities. The path that our algorithm generates is compared against a planner that only aims to reduce path length and ignores the crowd flow. Understandably, the trajectories generated from our planner have lower invasiveness than the baseline in all cases.

5.1 Independent Comparisons

Each of the macroscopic properties, density, velocity and variance, are varied independently while others stay uniform across the space. Table 1 shows the 3 different scenarios named after the parameter that varies in them.

The density scenario is analogous to a stationary crowd gathered around a point of interest. Figure 1a shows our algorithm avoiding the dense crowd, as it is more invasive to wade through the crowd than it is to move around it through less dense regions. In the less

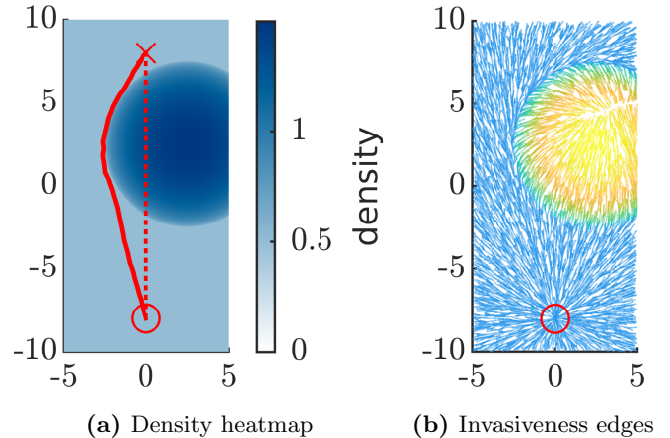


Figure 1: Simple case varying crowd density. The left compares the trajectory from the naïve planner and our social planner, shown as dotted and solid lines respectively, from an initial position (circle) to a goal (cross). PRM*’s minimal invasive tree from the initial position is shown on the right, with edges coloured by the invasiveness per unit distance.

dense regions, the algorithm prefers straight paths to minimise the time spent around the pedestrians.

It is interesting to see critical points in the minimally invasive tree in figure 1b in the upper right of the crowd which we will refer to as *breakpoints*. These breakpoints indicate regions where small changes in goal specification drastically changes the optimal path. In this case, the breakpoint is not centred at the dense crowd region, but is more towards the far side of the crowd. This is intuitive since if we need to plan a path to the middle of the crowd, we can only minimise the invasiveness of a path to the boundary of this crowd, which corresponds to a straight line in this scenario.

The abstract scenario varying velocity corresponds to a crowd of people orbiting a point. In this situation, we can travel with minimal invasiveness by following the crowd’s movement. By specifying a particular destination, our algorithm generates an interesting trajectory illustrated in figure 2a. Following our intuition, it travels in the direction of flow. On close inspection we also note that the trajectory bulges over time. Since the destination is not directly “downstream” of the starting location, the robot must cross the flow of the crowd. The rate at which the robot crosses the flow appears consistent through the trajectory as well since this would minimise the sum square of relative velocities. This effect can be seen across the minimal invasive tree in figure 2b. The breakpoints in this scenario highlight the regions where it is equally invasive to move a short distance against the flow as to move a much longer distance with the flow.

The final simple scenario varies the variance of crowd

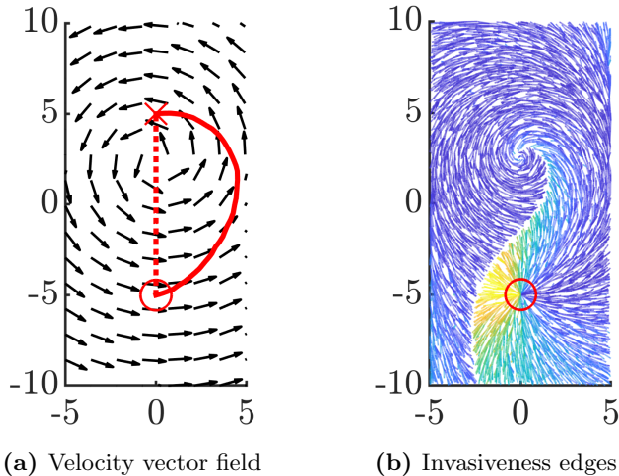


Figure 2: Simple case varying crowd velocity. The left compares the trajectory from the naïve planner and our social planner, shown as dotted and solid lines respectively, from an initial position (circle) to a goal (cross). PRM*’s minimal invasive tree from the initial position is shown on the right, with edges coloured by the invasiveness per unit distance.

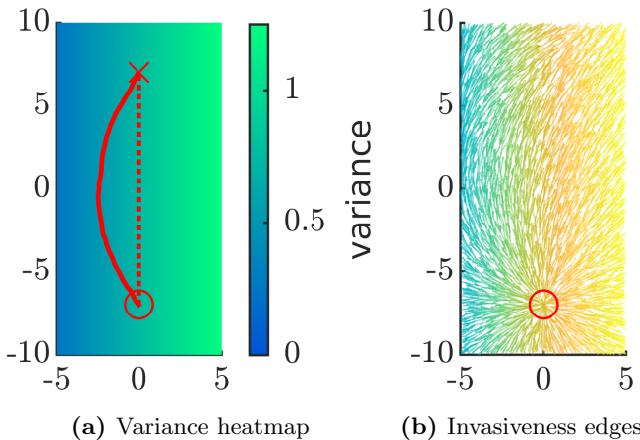


Figure 3: Simple case varying the variance of the crowd velocity. The left compares the trajectory from the naïve planner and our social planner, shown as dotted and solid lines respectively, from an initial position (circle) to a goal (cross). PRM*’s minimal invasive tree from the initial position is shown on the right, with edges coloured by the invasiveness per unit distance.

velocities, as a result of summing 3 pedestrian flows. Two flows travel in opposite directions in the vertical axis in figure 3a with equal speed, but with more pedestrians on the right side. A third flow is added with zero velocity to ensure that the density of the crowd across the space is even.

In this case, our planner finds a trajectory that is the compromise between travelling longer in regions with lower variance and travelling shorter in regions with

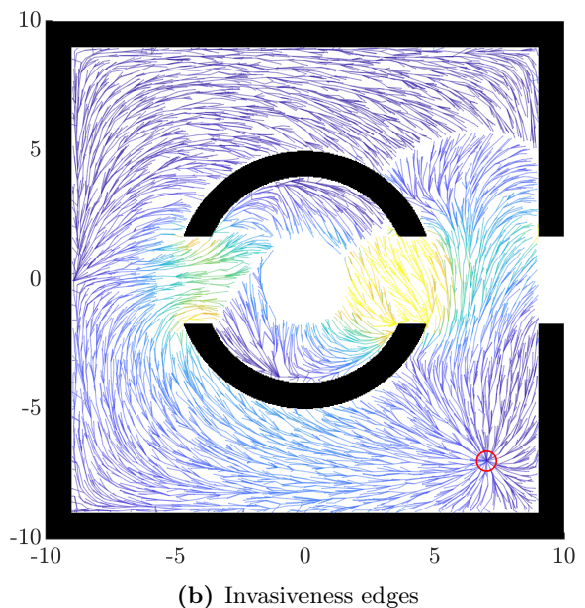
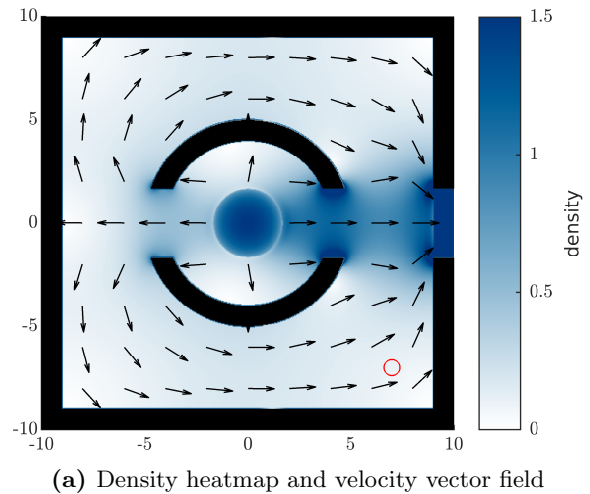


Figure 4: A hypothetical concert hall scene, showing the application of our work to a larger and more complex environment. PRM*’s minimal invasive tree from an initial position (circle) is shown below, with edges coloured by the invasiveness per unit distance.

higher variance. It is intuitive to avoid regions towards the right of the environment despite some of the crowd going in the robot’s direction since there are also others with the opposite velocity travelling against the robot. In practice, humans avoid causing regions of high variance of velocity as they are associated with higher rates of collision which leads to social phenomena such as lane forming.

5.2 Concert Hall

We construct a hypothetical scene corresponding to audience members vacating a concert hall shown in figure 4a. The audience leaves the inner room through the

two exits and then leaves the outer room through one main exit. Logically, more people form around the right inner door and the outer door.

In figure 4b we visualise the minimally invasive trajectories from an initial position in the bottom-right to every position in the hall. As the colour of each edge indicates the invasiveness per unit of distance, it can be seen that it is very costly to enter the inner room through the right door, or to cross between the inner and outer doors. Additionally, the breakpoint in the top right indicates the region where it is less invasive to move around the inner room than to push through the dominant crowd flow towards the main exit.

This example illustrates that the algorithm is capable of planning over long time horizon to avoid myopic decisions in the presence of non-uniform velocity and density fields.

6 Discussion and Future Work

We have presented a new perspective on the problem of mobile robot navigation in dense crowds that is based on a flow field representation. This representation is advantageous in that it enables the development of computationally efficient planning algorithms whose running time is independent of the total number of individual pedestrians. Our results provide initial evidence to support this idea, but there are a number of areas of future work that would need to be addressed in order to develop an end-to-end implemented system.

One important question is how to create a flow field estimate given observations that could be readily acquired by a perception system. One approach would be to adapt recent results for estimating incompressible flows using specialised Gaussian process regression [Lee *et al.*, 2019b]. Another important avenue to explore is the case of dynamic (time-dependent) flows. This is an instance of a time-dependent shortest path problem, which remains open in general. However, promising efficient solutions have recently been developed for relevant special cases including in our previous work [Lee *et al.*, 2019a]. Finally, an interesting question for real-world systems is how to validate the fidelity of a flow field representation of crowds. This could be approached by comparing predicted behaviour of individuals, drawn from a distribution induced by the flow field, with observations of individuals in actual crowds.

Acknowledgments

This work is supported in part by Australian Government Research Training Program (RTP) Scholarships and the University of Technology Sydney.

References

- [Alahi *et al.*, 2016] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese. Social LSTM: human trajectory prediction in crowded spaces. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition*, pages 961–971, 2016.
- [Alempijevic *et al.*, 2013] A. Alempijevic, R. Fitch, and N. Kirchner. Bootstrapping navigation and path planning using human positional traces. In *Proc. of IEEE ICRA*, pages 1242–1247, 2013.
- [Bera *et al.*, 2019] Aniket Bera, Tanmay Randhavane, and Dinesh Manocha. The emotionally intelligent robot: Improving socially-aware human prediction in crowded environments. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition Workshops*, June 2019.
- [Bevilacqua *et al.*, 2018] P. Bevilacqua, M. Frego, D. Fontanelli, and L. Palopoli. Reactive planning for assistive robots. *IEEE Robotics and Automation Letters*, 3(2):1276–1283, 2018.
- [Chen *et al.*, 2018] Z. Chen, C. Jiang, and Y. Guo. Pedestrian-robot interaction experiments in an exit corridor. In *Proc. of International Conference on Ubiquitous Robots*, pages 29–34, 2018.
- [Helbing and Molnár, 1995] Dirk Helbing and Péter Molnár. Social force model for pedestrian dynamics. *Physical Review E*, 51:4282–4286, 1995.
- [Hoffman and Breazeal, 2007] Guy Hoffman and Cynthia Breazeal. Effects of anticipatory action on human-robot teamwork efficiency, fluency, and perception of team. In *Proc. of ACM/IEEE International Conference on Human-Robot Interaction*, pages 1–8, 2007.
- [Hoogendoorn and Bovy, 2003] Serge Hoogendoorn and Piet H.L. Bovy. Simulation of pedestrian flows by optimal control and differential games. *Optimal Control Applications and Methods*, 24(3):153–172, 2003.
- [Jodoin *et al.*, 2013] Pierre-Marc Jodoin, Yannick Benezeth, and Yi Wang. Meta-tracking for video scene understanding. In *2013 10th IEEE International Conference on Advanced Video and Signal Based Surveillance*, pages 1–6. IEEE, 2013.
- [Karaman and Frazzoli, 2011] Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *The International Journal of Robotics Research*, 30(7):846–894, 2011.
- [Karamouzas *et al.*, 2014] Ioannis Karamouzas, Brian Skinner, and Stephen J Guy. Universal power law governing pedestrian interactions. *Physical Review Letters*, 113(23):238701:1–238701:5, 2014.

- [Lee *et al.*, 2017] James Ju Heon Lee, Chanyeol Yoo, Raewyn Hall, Stuart Anstee, and Robert Fitch. Energy-optimal kinodynamic planning for underwater gliders in flow fields. In *Proc. of ARAA ACRA*, 2017.
- [Lee *et al.*, 2019a] James Ju Heon Lee, Chanyeol Yoo, Stuart Anstee, and Robert Fitch. Efficient optimal planning in non-FIFO time-dependent flow fields. In *arXiv:1909.02198 [cs.RO]*, pages 1–10, 2019.
- [Lee *et al.*, 2019b] Ki Myung Brian Lee, Chanyeol Yoo, Ben Hollings, Stuart Anstee, Shoudong Huang, and Robert Fitch. Online estimation of ocean current from sparse GPS data for underwater vehicles. In *Proc. of IEEE ICRA*, pages 3443–3449, 2019.
- [Moussaïd *et al.*, 2009] Mehdi Moussaïd, Dirk Helbing, Simon Garnier, Anders Johansson, Maud Combe, and Guy Theraulaz. Experimental study of the behavioural mechanisms underlying self-organization in human crowds. *Proceedings of the Royal Society B: Biological Sciences*, 276(1668):2755–2762, 2009.
- [Rudenko *et al.*, 2018] A. Rudenko, L. Palmieri, and K. O. Arras. Joint long-term prediction of human motion using a planning-based social force approach. In *Proc. of IEEE ICRA*, pages 4571–4577, 2018.
- [To *et al.*, 2019a] Kwun Yiu Cadmus To, James Ju Heon Lee, Chanyeol Yoo, Stuart Anstee, and Robert Fitch. Streamline-based control of underwater gliders in 3d environments. In *Proc. of IEEE CDC*, 2019.
- [To *et al.*, 2019b] Kwun Yiu Cadmus To, Ki Myung Brian Lee Lee, Chanyeol Yoo, Stuart Anstee, and Robert Fitch. Streamlines for motion planning in underwater currents. *Proc. of IEEE ICRA*, 2019.
- [Trautman *et al.*, 2015] Pete Trautman, Jeremy Ma, Richard M Murray, and Andreas Krause. Robot navigation in dense human crowds: Statistical models and experimental studies of human–robot cooperation. *The International Journal of Robotics Research*, 34(3):335–356, 2015.
- [Treuille *et al.*, 2006] Adrien Treuille, Seth Cooper, and Zoran Popović. Continuum crowds. In *ACM Transactions on Graphics (TOG)*, volume 25, pages 1160–1168. ACM, 2006.
- [Turnwald and Wollherr, 2019] Annemarie Turnwald and Dirk Wollherr. Human-like motion planning based on game theoretic decision making. *International Journal of Social Robotics*, 11(1):151–170, 2019.
- [Van den Berg *et al.*, 2008] J. Van den Berg, Ming Lin, and D. Manocha. Reciprocal velocity obstacles for real-time multi-agent navigation. In *Proc. of IEEE ICRA*, pages 1928–1935, 2008.
- [Vemula *et al.*, 2017] A. Vemula, K. Muelling, and J. Oh. Modeling cooperative navigation in dense human crowds. In *Proc. of IEEE ICRA*, pages 1685–1692, 2017.
- [Yoo *et al.*, 2016] Chanyeol Yoo, Robert Fitch, and Salah Sukkarieh. Online task planning and control for fuel-constrained aerial robots in wind fields. *Int. J. Robot. Res.*, 35(5):438–453, 2016.
- [Yoo *et al.*, 2019] Chanyeol Yoo, Stuart Anstee, and Robert Fitch. Stochastic path planning for autonomous underwater gliders with safety constraints. In *Proc. of IEEE/RSJ IROS*, 2019.