

Uncertainty Models for TTC-Based Collision Avoidance

Zahra Forootaninia*, Ioannis Karamouzas[†], and Rahul Narain*

*Department of Computer Science and Engineering, University of Minnesota

[†]School of Computing, Clemson University

Supplementary material available at <http://www-users.cs.umn.edu/~foro0012/UTTC/>

Abstract—We address the problem of uncertainty-aware local collision avoidance within the context of time-to-collision based navigation of multiple agents. We consider two specific models that account for uncertainty in the future trajectories of interacting agents: an *isotropic* model which conservatively considers all possible errors, and an *adversarial* model that assumes the error is towards a head-on collision. We compare the two models experimentally via a number of simulation scenarios, and also provide theoretical guarantees about the collision avoidance behavior of the agents.

I. INTRODUCTION

Local collision avoidance plays an important role in multi-agent navigation and planning. Whether there are Roombas cleaning the floor, autonomous vehicles driving in a crowded city center, or animated characters walking through a virtual world, the agents should be able to sense their surroundings and react accordingly in order to avoid collisions while successfully completing their tasks. As such, numerous models for local collision avoidance have been proposed in robotics, traffic engineering, and computer animation.

Many well-known collision avoidance models rely on the concept of *velocity obstacles* (VO) introduced by Fiorini and Shiller for planar, disc-shaped agents [9]. The VO is the set of all relative velocities between two agents that can lead to a potential collision in the future, assuming that the agents are moving at constant linear velocity. Hence, agents can safely navigate by selecting new velocities that are outside of any VOs induced by their neighbors. The VO provides a tractable alternative to the notion of *inevitable collision states*, which is inherently expensive to compute for most real world problems [10, 5]. Since the work of Fiorini and Shiller [9], many extensions of velocity obstacles have been proposed that account for reciprocity among the agents, non-linear motion, controllers, kinematically constrained robots, rectangular-shaped agents, and group formations [31, 11, 35, 33, 12, 17, 22].

Much of the appeal of VO-based approaches is due to the ORCA framework proposed by van den Berg et al. [32]. ORCA provides an efficient way for computing a collision-free velocity outside the union of all VOs by conservatively approximating each VO as a half-plane and using linear programming to quickly find a feasible solution. As ORCA can also provide formal guarantees about the collision-free behavior of the agents, it has become very popular and many variants have been proposed based on modified VO formulations, including

approaches for non-holonomic and car-like robots [2, 3], as well as elliptical agents [8] to name just a few.

Overall, velocity-based approaches provide an intuitive framework to reason about collision avoidance, and different models have been introduced to account for uncertainty in the sensing data allowing implementation on physical robots. The PVO [11] method, for example, considers uncertainty in the movement and sensing of the robots but does not assume reciprocity among the interacting robots. The HRVO [27] formulation addresses this issue, whereas the approach in [15] extends ORCA to bound the error introduced by localization.

However, despite their stability and applicability to robots, VO-based and ORCA-like approaches are very conservative in nature. In an attempt to guarantee collision-free behavior, the agents tend to throw away too many admissible velocities exhibiting inefficient behavior. In many interaction scenarios, for example, the agents end up stop moving, focusing on not colliding and forgetting about their tasks in hand. In contrast, *force-based* models can allow more flexibility in the behavior of the agents (see, e.g., Fig. 2).

Force-based models traditionally resolve collisions by using repulsive artificial potential fields [4, 20], or a mixture of distance-based physical and social forces [14, 24]. In all these models, though, the agents are prone to collisions and unrealistic behaviors such as oscillatory and backwards motions, since they only react when they get too close to each other and do not account for the velocities of their neighbors. To address these issues, *predictive* force-based models have been recently introduced, in which two agents only experience an interaction force if they are approaching each other. The resulting force is typically based on the time that it takes for the agents to collide [25, 26], or the time at which the distance between the interacting agents becomes minimal [37].

A representative predictive force-based model is the *time-to-collision* (TTC) approach proposed recently by Karamouzas et al. [18]. The authors analyzed a large corpus of crowd data, including interactions of pedestrians in commercial streets and college campuses, as well as motion capture experiments where participants had to navigate through dense and narrow bottlenecks. Their analysis showed that, independent of the task in hand, the interaction force between two pedestrians follows an inverse *power-law* relationship as a function of their estimated time to collision. The resulting TTC model allows agents to emulate better how humans resolve collisions in real

life, leading to more efficient behavior as compared to VO-based approaches. Another benefit that we show here is that unlike other previously mentioned force-based approaches, the TTC model can guarantee collision-free navigation in generic scenarios. This property has never been established in previous work, and we provide a proof in Sec. IV.

However, despite providing better flexibility in the velocities that the agents can adopt, the TTC model, as well as other predictive force-based approaches, assume that the agents are equipped with perfect sensors. In the real world, such an assumption is unrealistic since there is typically some noise when a robot senses the positions and velocities of its neighbors. Similarly, in interactive virtual environments, some uncertainty exists when the user is interacting with AI agents (e.g., a user-controlled avatar walking amidst NPCs).

The problem of motion planning under sensing and motion uncertainty can be generically modeled as a partially observable Markov decision process (POMDP) [1], where a control policy needs to be computed over the space of belief states. POMDPs lead to optimal solutions, but their computational complexity increases exponentially with the dimensions of the state space. To address this issue, approximate models have also been proposed that have a polynomial running time in the state space, such as approaches that locally optimize an input feasible trajectory assuming Gaussian beliefs [30, 34]. However, such approaches are still intractable to the problem of navigating multiple agents in real time without collisions.

As such, in this paper, we focus on local multi-agent navigation and generalize the TTC model to account for uncertainty in the future trajectories of interacting agents. Our aim is to maintain the structure and simplicity of the TTC model, while extending its collision-free guarantees to the uncertain case. Overall, this paper makes the following contributions:

- We extend the time-to-collision approach for local navigation to incorporate a model of uncertainty in the sensed positions and velocities of other agents. We call this approach *uncertainty-TTC*, or UTTC for short.
- We present two uncertainty models for collision avoidance in the presence of bounded sensing error: an *isotropic* model (UTTC-I) which conservatively considers error in all possible directions, and an *adversarial* model (UTTC-A) where the error is assumed to be directly towards a head-on collision.
- We provide a theoretical proof of collision-free navigation of multiple agents under both the TTC and UTTC-I models in the absence of sensing error.
- We experimentally demonstrate collision-free behavior of both UTTC-I and UTTC-A in a number of simulation scenarios, including with nonzero sensing error.

II. BACKGROUND

Here, we provide a brief overview of the TTC model along with the necessary equations needed in the following section to extend this model by incorporating uncertainty. We refer the reader to Karamouzas et al. [18] and the supplementary material for more details.

A. Time-to-Collision Avoidance Model

Consider two agents or robots, A_i and A_j , which have current positions \mathbf{x}_i and \mathbf{x}_j , and current velocities \mathbf{v}_i and \mathbf{v}_j , respectively. For simplicity we assume that the physical workspace of A_i and A_j is in \mathbb{R}^2 , and the agents can be modeled as discs of radii r_i and r_j , respectively, moving on the 2D plane. The interaction between them depends only on their relative displacement $\mathbf{x}_{ij} = \mathbf{x}_i - \mathbf{x}_j$, their relative velocity $\mathbf{v}_{ij} = \mathbf{v}_i - \mathbf{v}_j$, and their combined radii $r_{ij} = r_i + r_j$. To avoid notational clutter, we will drop the subscripts ij and write \mathbf{x} , \mathbf{v} , and r for \mathbf{x}_{ij} , \mathbf{v}_{ij} , and r_{ij} respectively.

Given the current positions and velocities of the agents, we predict collisions by assuming a linear extrapolation of their trajectories. The agents' predicted relative position at any time $t \geq 0$ is therefore $\mathbf{x} + \mathbf{v}t$. Their expected *time to collision* τ is the earliest time at which the predicted positions are colliding,

$$\tau = \tau(\mathbf{x}, \mathbf{v}, r) = \min \{t : t \geq 0, \|\mathbf{x} + \mathbf{v}t\| \leq r\}. \quad (1)$$

If the set of colliding times is empty, we take $\tau = \infty$.

In the TTC model, the collision avoidance behavior of the two agents follows a *power-law* relationship with respect to the time to collision τ . In particular, we define an interaction energy $U = U(\mathbf{x}, \mathbf{v}, r)$ as a function of τ alone,

$$U = f(\tau) = k\tau^{-m}e^{-\tau/\tau_0}, \quad (2)$$

where m denotes the exponent of the power law, τ_0 is the truncation point of the energy (similar to the notion of time horizon), and k is a scaling constant. The interaction energy U is nonnegative and decreasing with increasing τ . It is zero when the agents are not on a collision course ($\tau = \infty$), and rises to infinity as a collision becomes imminent ($\tau \rightarrow 0$). The agents A_i and A_j avoid the collision by following the negative gradient of the interaction energy with respect to \mathbf{x}_i and \mathbf{x}_j respectively. That is, agent A_i experiences a repulsive force \mathbf{f}_{ij} that pushes it towards a lower potential configuration,

$$\mathbf{f}_{ij} = -\frac{\partial U}{\partial \mathbf{x}_i} = -f'(\tau) \frac{\partial \tau}{\partial \mathbf{x}_i}, \quad (3)$$

and similarly agent A_j experiences a force $\mathbf{f}_{ji} = -\partial U / \partial \mathbf{x}_j$.

To compute \mathbf{f}_{ij} and \mathbf{f}_{ji} , we note the following facts:

- 1) The derivative of the power law is

$$f'(\tau) = -\frac{ke^{-\tau/\tau_0}}{\tau^{m+1}} \left(m + \frac{\tau}{\tau_0} \right). \quad (4)$$

- 2) As τ depends only on $\mathbf{x} = \mathbf{x}_i - \mathbf{x}_j$, we have

$$\frac{\partial \tau}{\partial \mathbf{x}_i} = \frac{\partial \tau}{\partial \mathbf{x}} = -\frac{\partial \tau}{\partial \mathbf{x}_j}. \quad (5)$$

This also implies that the forces \mathbf{f}_{ij} and \mathbf{f}_{ji} are equal and opposite.

- 3) The time to collision can be computed explicitly as the smallest positive root of $\|\mathbf{x} + \mathbf{v}\tau\|^2 = r^2$, which is a quadratic equation in τ . As the leading coefficient is

independent of \mathbf{x} , the derivative of τ with respect to \mathbf{x} can therefore be shown to be

$$\frac{\partial \tau}{\partial \mathbf{x}} = \frac{\mathbf{x} + \mathbf{v}\tau}{\sqrt{D}}, \quad (6)$$

where $D = (\mathbf{x} \cdot \mathbf{v})^2 - \|\mathbf{v}\|^2(\|\mathbf{x}\|^2 - r^2)$ is the discriminant of the quadratic.

Therefore, we have

$$\mathbf{f}_{ij} = \frac{ke^{-\tau/\tau_0}}{\tau^{m+1}} \left(m + \frac{\tau}{\tau_0} \right) \left(\frac{\mathbf{x} + \mathbf{v}\tau}{\sqrt{D}} \right), \quad (7)$$

$$\mathbf{f}_{ji} = -\mathbf{f}_{ij}. \quad (8)$$

We refer to the force defined in Equation (7) as the TTC repulsion force $\mathbf{f}_{\text{TTC}}(\mathbf{x}, \mathbf{v}, r)$. The magnitude of this force is inversely proportional to the $(m+1)$ th power of the time to collision τ , and its direction is along the relative displacement of the two agents at the moment of impact.

B. Time to Collision and Velocity Obstacles

The TTC model is closely related to the concept of *velocity obstacles* (VO), i.e., the set of all relative velocities between two agents that will lead to a collision at some moment in time. The VO can be geometrically interpreted as a cone with its apex at the origin in the velocity space, and its legs tangent to the disc centered at $\mathbf{x}_j - \mathbf{x}_i$ having radius r . In the TTC model, the interaction energy U is nonzero if and only if the agents' relative velocity \mathbf{v} lies inside the VO.

Figure 1 shows the gradient of the τ -based repulsive potential for a simple 2-agent scenario, and compares it to the instantaneous change in relative velocity that a typical VO-based model assumes for the same scenario. As can be seen, the VO approach forbids any relative velocities inside the cone, forcing the agents to choose a new relative velocity on the boundary of the VO to avert a collision. On the other hand, in TTC, the gradient guides the agent either to the left or right of the VO, or slows it down depending on the position of the relative velocity inside the VO.

In state-of-the-art VO-based approaches such as ORCA, a collision between two agents is typically resolved by projecting the relative velocity to the closest point on the VO boundary. Furthermore, in an attempt to guarantee fast computations of collision-free velocities, ORCA approximates each VO with a line and uses linear programming to find an optimal new velocity on the boundary of the convex hull formed by the intersection of all lines. However, due to such a conservative approximation, the set of feasible velocities of an agent becomes significantly restricted, compromising the agent's goal-directed behavior. In contrast, the TTC model allows much more flexibility in the velocities that the agents can adopt, leading to more human-like avoidance behavior.

As an example, consider Fig. 2 where two agents going from left to right have to interact with another agent coming from the opposite direction. ORCA forces the single agent to wait until the other two agents have walked around it. In contrast, in the TTC model, the acceleration of the agents gradually increases with decreasing τ values; the standalone




Fig. 1: (a) Two agents i and j and the associated velocity obstacle (light gray). Any relative velocity \mathbf{v}_{ij} lying in the VO will lead to a collision in the future. (b) Change in velocity as a function of \mathbf{v}_{ij} in a VO model. (c) Interaction forces as a function of \mathbf{v}_{ij} in the TTC model.




Fig. 2: Comparison between ORCA and the TTC model in a small scenario involving three agents. At each discrete time, the traces of the agents are shown as colored disks which are light at their initial positions and dark at their current positions.

agent is willing to “take a step” towards a collision, assuming that such a step will help the agent solve the collision more efficiently in the near future.

C. Multi-Agent Navigation

The TTC formulation can be used for independent navigation of multiple holonomic agents A_1, A_2, \dots, A_n sharing a common 2D workspace as follows. Each agent A_i independently performs a continuous cycle of sensing and acting with time step Δt . A global planning routine computes a preferred velocity $\mathbf{v}_i^{\text{pref}}$ that will lead the agent towards its goal; in this work, we assume a simple $\mathbf{v}_i^{\text{pref}}$ that is directed towards A_i 's goal with a magnitude equal to its preferred speed. This is translated into a goal force $\mathbf{f}_i^{\text{goal}} = k_g(\mathbf{v}_i^{\text{pref}} - \mathbf{v}_i)$ which adapts the agent's current velocity to the preferred one. The agent also senses the position, velocity, and radius of each nearby agent A_j , and computes an avoidance force \mathbf{f}_{ij} as defined in (7). The new velocity and position of the agent is then obtained through numerical integration using the total force, and the sensing-acting cycle repeats.

III. ADDING UNCERTAINTY

The TTC model assumes that robots have perfect sensors and are able to sense other nearby agents' positions and velocities exactly. In the real world, though, such an assumption is unrealistic, and other agents' positions and velocities can only be measured with some amount of error. If a small sensing

error changes the predicted trajectories from colliding to non-colliding, the TTC force will be zero and the agents will not be able to avoid the collision. Thus, if one wished to apply TTC-style collision avoidance to robots in practice, the TTC model in its current form would not be adequate. In this section, we remove this drawback by extending the TTC model to account for sensor uncertainty.

For agent A_i to compute its avoidance force with respect to another agent A_j , it needs to observe the other agent's position and velocity. Let us denote by $\hat{\mathbf{x}}$ and $\hat{\mathbf{v}}$ the relative displacement and velocity, respectively, between A_i and the sensed position and velocity of A_j . We assume that the sensor error is bounded by known constants,

$$\|\hat{\mathbf{x}} - \mathbf{x}\| \leq \delta, \quad (9)$$

$$\|\hat{\mathbf{v}} - \mathbf{v}\| \leq \epsilon. \quad (10)$$

(δ and ϵ need not be the same for different pairs of agents.) Thus, given $\hat{\mathbf{x}}$ and $\hat{\mathbf{v}}$, the true relative displacement \mathbf{x} could be any point in the ball $B(\hat{\mathbf{x}}, \delta)$ centered at $\hat{\mathbf{x}}$ with radius δ , and similarly $\mathbf{v} \in B(\hat{\mathbf{v}}, \epsilon)$. Our goal is to modify the TTC forces to avoid any possible collision within these bounds.

Accounting for the position uncertainty is straightforward. The initial conditions (\mathbf{x}, \mathbf{v}) lead to a collision if at some time $t \geq 0$ we have $\|\mathbf{x} + \mathbf{v}t\| \leq r_i + r_j$. Applying the triangle inequality, this implies that $\|\hat{\mathbf{x}} + \mathbf{v}t\| \leq r_i + r_j + \delta$, that is, the trajectories starting from the sensed relative position $\hat{\mathbf{x}}$ collide under a larger combined radius $r_i + r_j + \delta$. Thus, if we assume that $\mathbf{x} = \hat{\mathbf{x}}$ and $r = r_i + r_j + \delta$, in effect enlarging the radius of the other agent by δ , we are certain to avoid collisions even if the true relative position differs from $\hat{\mathbf{x}}$ by up to δ . From here on, we assume that the position uncertainty has been accounted for in this way, and restrict attention only to the velocity uncertainty ϵ . Note that enlarging the radius alone is not sufficient to prevent collisions in the presence of velocity uncertainty: $\hat{\mathbf{v}}$ may be outside the VO cone even if the agents are about to collide, in which case the TTC force will be zero.

We propose two different generalizations of the TTC model to account for error in the sensed velocities. Our *isotropic* model considers collisions with all possible velocities \mathbf{v} within the bounds of sensor error, and defines the interaction energy based on a conservative time to collision $\hat{\tau}_{\text{iso}}$. Our *adversarial* model is derived from a simplifying assumption that the true velocity differs from the sensed velocity by an error directly in the direction of $-\mathbf{x}$, as though chosen by an adversary seeking to drive the agents directly towards each other. The isotropic model is easier to analyze theoretically, and we provide theoretical guarantees of collision avoidance. On the other hand, the adversarial model performs equally well at preventing collisions in practice, and we observe that it tends to choose more efficient paths.

A. The isotropic model

In the isotropic model, we choose the estimated time to collision $\hat{\tau}_{\text{iso}}$ to be the time of the earliest collision among all

possible values of \mathbf{v} ,

$$\hat{\tau}_{\text{iso}} = \min_{\tilde{\mathbf{v}} \in B(\hat{\mathbf{v}}, \epsilon)} \tau(\mathbf{x}, \tilde{\mathbf{v}}, r). \quad (11)$$

Equivalently, we seek the earliest time $t \geq 0$ at which there exists *some* possible velocity $\tilde{\mathbf{v}} \in B(\hat{\mathbf{v}}, \epsilon)$ for which the predicted positions are colliding. The set of predicted relative positions forms a ball,

$$\{\mathbf{x} + \tilde{\mathbf{v}}t : \tilde{\mathbf{v}} \in B(\hat{\mathbf{v}}, \epsilon)\} = B(\mathbf{x} + \hat{\mathbf{v}}t, \epsilon t) \quad (12)$$

which has a nonempty intersection with the set of colliding relative positions $B(\mathbf{0}, r)$ if and only if $\|\mathbf{x} + \hat{\mathbf{v}}t\| \leq r + \epsilon t$. Therefore, the time to collision under isotropic uncertainty can again be found as the smallest positive root of a quadratic equation in $\hat{\tau}_{\text{iso}}$,

$$\|\mathbf{x} + \hat{\mathbf{v}}\hat{\tau}_{\text{iso}}\|^2 = (r + \epsilon\hat{\tau}_{\text{iso}})^2. \quad (13)$$

Finally, we have

$$\frac{\partial \hat{\tau}_{\text{iso}}}{\partial \mathbf{x}} = \frac{\mathbf{x} + \hat{\mathbf{v}}\hat{\tau}_{\text{iso}}}{\sqrt{D}} \quad (14)$$

where the discriminant is

$$D = (\mathbf{x} \cdot \hat{\mathbf{v}} - r\epsilon)^2 - (\|\hat{\mathbf{v}}\|^2 - \epsilon^2)(\|\mathbf{x}\|^2 - r^2). \quad (15)$$

We denote the associated interaction energy of the isotropic model by $U_{\text{iso}} = f(\hat{\tau}_{\text{iso}})$, and define the repulsion force

$$\mathbf{f}_{\text{iso}}(\mathbf{x}, \hat{\mathbf{v}}, r) = -f'(\hat{\tau}_{\text{iso}}) \frac{\partial \hat{\tau}_{\text{iso}}}{\partial \mathbf{x}} \quad (16)$$

as before.

B. The adversarial model

We motivate the adversarial model by considering the behavior of the isotropic model in the limit as the agents approach a collision. In this case, a first-order expansion of Equation (13) about $\|\mathbf{x}\| \rightarrow r$ gives

$$\hat{\tau}_{\text{iso}} \approx \frac{\|\mathbf{x}\| - r}{\frac{\mathbf{x}}{\|\mathbf{x}\|} \cdot (\hat{\mathbf{v}} - \epsilon \frac{\mathbf{x}}{\|\mathbf{x}\|})}. \quad (17)$$

For the TTC model $\|\mathbf{x} + \mathbf{v}\tau\|^2 = r^2$, the corresponding approximation is

$$\tau \approx \frac{\|\mathbf{x}\| - r}{\frac{\mathbf{x}}{\|\mathbf{x}\|} \cdot \mathbf{v}}. \quad (18)$$

Comparing the two, we see that when the agents are close together, the isotropic model behaves like the TTC model applied to a modified velocity $\hat{\mathbf{v}} - \epsilon \frac{\mathbf{x}}{\|\mathbf{x}\|}$.

The adversarial model arises by adopting this simple approximation for all \mathbf{x} and $\hat{\mathbf{v}}$, not just the ones that are close to a collision. That is, we define the repulsive force in the adversarial model as

$$\mathbf{f}_{\text{adv}}(\mathbf{x}, \hat{\mathbf{v}}, r) = \mathbf{f}_{\text{TTC}}(\mathbf{x}, \tilde{\mathbf{v}}, r), \quad (19)$$

where $\tilde{\mathbf{v}} = \hat{\mathbf{v}} - \epsilon \frac{\mathbf{x}}{\|\mathbf{x}\|}$.

We have also considered defining an interaction potential $U_{\text{adv}} = f(\tau(\mathbf{x}, \tilde{\mathbf{v}}, r))$ and deriving the force $\mathbf{f}_{\text{adv}} = -\partial U_{\text{adv}} / \partial \mathbf{x}$, analogous to the isotropic model. This




Fig. 3: Comparison of the support sets of the TTC, UTTC-I and UTTC-A energies. The uncertainty-aware models enlarge the set of relative velocities which are affected.

model is not identical to using $\mathbf{f}_{\text{TTC}}(\mathbf{x}, \tilde{\mathbf{v}}, r)$ because it involves an additional term proportional to $\partial\tilde{\mathbf{v}}/\partial\mathbf{x}$, but it behaves very similarly in practice. Indeed, both formulations lead to nonzero forces on the same set of inputs, namely those in which the velocity $\tilde{\mathbf{v}}$ leads to a collision. As the model using $\mathbf{f}_{\text{TTC}}(\mathbf{x}, \tilde{\mathbf{v}}, r)$ is much simpler to implement, we use it in the following, and discuss the U_{adv} variant further in the supplementary material.

C. Relationship with velocity obstacles

It is instructive to consider how our two uncertainty models affect the set of relative velocities $\hat{\mathbf{v}}$ which yield a nonzero force \mathbf{f}_{ij} . Equivalently, these are the velocities for which a potential collision is predicted. For the TTC model, let us denote this set $S_{\text{TTC}} = \{\hat{\mathbf{v}} : \mathbf{f}_{\text{TTC}}(\mathbf{x}, \hat{\mathbf{v}}, r) \neq \mathbf{0}\}$; as described in Section II-B, this set is precisely the VO cone. Both the isotropic and the adversarial model enlarge this set, but in slightly different ways. The isotropic model predicts a potential collision for $\hat{\mathbf{v}}$ if there exists a colliding velocity $\tilde{\mathbf{v}}$ within ϵ of $\hat{\mathbf{v}}$. Thus S_{iso} is the Minkowski sum of S_{TTC} with a disc $B(\mathbf{0}, \epsilon)$. This is exactly analogous to the uncertainty model used in HRVO [27], which accounts for uncertainty in the sensed velocity of another agent by taking the Minkowski sum of the corresponding reciprocal velocity obstacle with a disc of radius ϵ .

On the other hand, in the adversarial model, we only check whether a collision exists with the modified velocity $\tilde{\mathbf{v}} = \hat{\mathbf{v}} - \epsilon \frac{\mathbf{x}}{\|\mathbf{x}\|}$. Therefore, S_{adv} is the translation of S_{TTC} by $\epsilon \frac{\mathbf{x}}{\|\mathbf{x}\|}$. As S_{TTC} is a cone that contains the vector $-\frac{\mathbf{x}}{\|\mathbf{x}\|}$, we have $S_{\text{TTC}} \subset S_{\text{adv}}$, that is, the support of the adversarial model contains that of the standard TTC model. In fact, because $\epsilon \frac{\mathbf{x}}{\|\mathbf{x}\|} \in B(\mathbf{0}, \epsilon)$, we further have the containment $S_{\text{TTC}} \subset S_{\text{adv}} \subset S_{\text{iso}}$. This relationship is shown in Figure 3.

IV. ANALYSIS

In this section, we first prove the collision-free guarantee for the TTC model in the absence of uncertainty, which has not been proved before. We then extend the proof to provide collision-free guarantees in the isotropic model with nonzero uncertainty.

A. Collision-free guarantees without sensor error

Our proof takes the form of an energy dissipation argument. We show, first, that the TTC interaction forces are *dissipative*, i.e. they do not increase the total energy of the system over time; and second, that for a collision to occur, the system must attain arbitrarily large amounts of energy. Therefore, as long as the system starts in a state of finite energy and all the other forces acting on it are either conservative or dissipative, it cannot undergo a collision.

In this section, we collect the positions and velocities of the n agents into the vectors $\mathbf{q}, \dot{\mathbf{q}} \in \mathbb{R}^{2n}$, respectively. The system evolves under the influence of generalized forces $\mathbf{f}_1, \mathbf{f}_2, \dots \in \mathbb{R}^{2n}$ according to the second-order ODE

$$\ddot{\mathbf{q}} = \sum_j \mathbf{f}_j. \quad (20)$$

The total kinetic energy of the system is $T = \frac{1}{2}\|\dot{\mathbf{q}}\|^2$, and the rate of work done on the system by a force \mathbf{f}_j is $\dot{W}_j = \mathbf{f}_j \cdot \dot{\mathbf{q}}$. The evolution equation implies that $\dot{T} = \sum_j \dot{W}_j$.

We assume that the forces \mathbf{f}_j acting on the system are of two types: conservative forces, and dissipative forces. A force \mathbf{f}_j is *conservative* if there exists a potential function $U_j(\mathbf{q})$ such that $\mathbf{f}_j = -\partial U_j / \partial \mathbf{q}$. In this case, the work done by the force is always equal to the decrease in the associated potential, $\dot{W}_j = -\dot{U}_j$. A force \mathbf{f}_j is *dissipative* if it always performs negative work, $\dot{W}_j = \mathbf{f}_j \cdot \dot{\mathbf{q}} \leq 0$. We can define the total energy of the system as $E = T + U$, where U is the sum of the potentials of all *conservative* forces. It is readily shown that when all the forces are conservative or dissipative, the total energy is a nonincreasing function of time.

Lemma 1. *The TTC force \mathbf{f}_{ij} defined by (7) is dissipative, and performs work at the rate $\dot{W} = f'(\tau)$.*

Proof: First, we compute the quantity $\partial\tau/\partial\mathbf{x}_{ij}$. Recall that τ is the smallest positive root of a function of the form $\phi(\mathbf{x}_{ij} + \mathbf{v}_{ij}t)$, namely $\phi(\mathbf{x}) = \|\mathbf{x}\|^2 - r_{ij}^2$. Performing implicit differentiation with respect to \mathbf{x}_{ij} , we have

$$\frac{\partial\phi}{\partial\mathbf{x}_{ij}} + \frac{\partial\phi}{\partial\tau} \frac{\partial\tau}{\partial\mathbf{x}_{ij}} = 0, \quad (21)$$

i.e.

$$\nabla\phi + (\nabla\phi \cdot \mathbf{v}_{ij}) \frac{\partial\tau}{\partial\mathbf{x}_{ij}} = 0. \quad (22)$$

Therefore, $\partial\tau/\partial\mathbf{x}_{ij} = -(\nabla\phi \cdot \mathbf{v}_{ij})^{-1} \nabla\phi$, which immediately implies that

$$\dot{W} = -f'(\tau) \frac{\partial\tau}{\partial\mathbf{x}_{ij}} \cdot \mathbf{v}_{ij} = f'(\tau). \quad (23)$$

Note that f is a decreasing function, and so the work done by the TTC force is always negative. ■

Lemma 2. *The force $\mathbf{f}_{ij}^{\text{iso}}$ in the isotropic model is dissipative when $(\mathbf{x}_{ij} + \mathbf{v}_{ij}\hat{\tau}_{\text{iso}}) \cdot \mathbf{v}_{ij} \leq 0$.*

Proof: We follow the same argument as above, except that $\hat{\tau}_{\text{iso}}$ is a root of $\phi_{\text{iso}}(\mathbf{x}_{ij} + \mathbf{v}_{ij}t, t)$ with $\phi_{\text{iso}}(\mathbf{x}, t) = \|\mathbf{x}\|^2 -$

$(r + \epsilon t)^2$. Denoting $\partial\phi_{\text{iso}}/\partial\mathbf{x}$ and $\partial\phi_{\text{iso}}/\partial t$ by $\nabla\phi_{\text{iso}}$ and ϕ'_{iso} respectively, implicit differentiation now yields

$$\frac{\partial\hat{\tau}_{\text{iso}}}{\partial\mathbf{x}_{ij}} = \frac{-\nabla\phi}{\nabla\phi \cdot \mathbf{v}_{ij} + \phi'} \quad (24)$$

$$= \frac{-(\mathbf{x}_{ij} + \mathbf{v}_{ij}\hat{\tau}_{\text{iso}})}{(\mathbf{x}_{ij} + \mathbf{v}_{ij}\hat{\tau}_{\text{iso}}) \cdot \mathbf{v}_{ij} - (r + \epsilon\hat{\tau}_{\text{iso}})\epsilon}. \quad (25)$$

When $(\mathbf{x}_{ij} + \mathbf{v}_{ij}\hat{\tau}_{\text{iso}}) \cdot \mathbf{v}_{ij} \leq 0$, this quantity is negative, and consequently the rate of work done by the force,

$$\dot{W} = f'(\hat{\tau}_{\text{iso}}) \frac{(\mathbf{x}_{ij} + \mathbf{v}_{ij}\hat{\tau}_{\text{iso}}) \cdot \mathbf{v}_{ij}}{(\mathbf{x}_{ij} + \mathbf{v}_{ij}\hat{\tau}_{\text{iso}}) \cdot \mathbf{v}_{ij} - (r + \epsilon\hat{\tau}_{\text{iso}})\epsilon} \quad (26)$$

is negative too. ■

So far we have considered the forces at a single instant in time. To analyze the possibility of collisions, we now consider the trajectories of agents as a function of time. Let us restrict attention to two agents in the system, A_i and A_j , and consider their separation $s(t) = \|\mathbf{x}_{ij}(t)\| - r_{ij}$. In the TTC model it is common for agents to graze tangentially past each other without colliding; at the moment of tangency the separation $s(t)$ falls to zero but does not cross it, and we have $s'(t) = 0$. As we do not wish to forbid such trajectories, in this section we only consider a collision to occur when $s(t)$ reaches zero with a nonzero time derivative, $s'(t) < 0$. The following lemma shows that in arriving at such a state requires an unbounded amount of energy.

Lemma 3. *If two agents collide at a time t^* , and all forces other than the TTC force \mathbf{f}_{ij} are bounded, there exists an interval $[t_0, t^*]$ in which the TTC force between the agents performs an unbounded amount of negative work.*

We refer the reader to the supplementary material for the proof as it is more involved than the others. Lemma 3 implies that in the generic case, when one pair of agents approaches each other but all other pairs have a finite separation, a collision cannot occur. However, we cannot yet rule out simultaneous collisions where multiple pairs of agents collide at the same time t^* , for example when an agent gets squeezed between other agents approaching from different directions at the same time. In such cases, multiple interaction forces would become unbounded. We leave the analysis of such situations for future work.

Theorem 1. *In a multi-agent system that begins with finite energy and evolves under the influence of (i) conservative forces with nonnegative potentials, and (ii) dissipative forces, there cannot be a time instant where exactly one pair of agents is colliding.*

Proof: The total energy of the system is nonincreasing and nonnegative. If two agents were to approach a collision at time t^* with no other collisions in a neighborhood of t^* , there would exist a time interval $[t_0, t^*]$ within that neighborhood in which all other forces were bounded, and in which the TTC force removes an unbounded amount of energy from the system. Therefore, $E(t_0)$ is finite and $E(t^*) \rightarrow -\infty$,




Fig. 4: (left to right) The 3-agent, 8-agent, crossing, and hallway scenarios used to evaluate our proposed uncertainty models. The depicted simulations were obtained using the UTTC-I model.

which is impossible by the assumption that the potentials are nonnegative. ■

A similar argument may be made for the isotropic model in the absence of sensor error, except for two modifications. First, the isotropic force is not always dissipative, so the argument of nonincreasing energy cannot immediately be employed. However, as it is dissipative for agents that are approaching each other, the argument still holds in the pre-collision interval $[t_0, t^*]$. Second, the expression for the rate of work done by the force is more complicated. As $s \rightarrow 0, \tau \rightarrow 0$, we approximate $(\mathbf{x}_{ij} + \mathbf{v}_{ij}\hat{\tau}_{\text{iso}}) \cdot \mathbf{v}_{ij} \approx \mathbf{x}_{ij} \cdot \mathbf{v}_{ij} \approx r\dot{s}$ and find that the integrated work done is

$$W = \int f'(\hat{\tau}_{\text{iso}}(t)) \frac{r\dot{s}}{r\dot{s} - r\epsilon} dt. \quad (27)$$

Because our assumption is that \dot{s} is bounded away from zero as $\hat{\tau}_{\text{iso}} \rightarrow 0$, so is the fraction in the integrand, and the integral remains infinite.

B. Collision-free motion with sensor error

The isotropic model guarantees that if two agents were to approach a collision in finite time, we would have $\hat{\tau} \rightarrow 0$ even in the presence of bounded sensor error. This is because $\hat{\tau}$ is defined by taking the minimum time to collision over all relative velocities in $B(\hat{\mathbf{v}}, \epsilon)$, and by the bounded error assumption we know that the true velocity \mathbf{v} is in this set. This property is encouraging because it means the repulsion force would rise to infinity at least as rapidly as in the error-free case, suggesting that the same collision avoidance guarantees should apply. However, the interaction forces are no longer reciprocal ($\mathbf{f}_{ij} \neq \mathbf{f}_{ji}$), because the error in $\hat{\mathbf{v}}_{ij}$ may not equal the error in $\hat{\mathbf{v}}_{ji}$. This complicates the energy dissipation argument, and we do not yet have full theoretical guarantees of the absence of collisions. Nevertheless, the results of our simulation experiments are consistent with the encouraging theoretical properties mentioned above, and indicate that both the isotropic and adversarial models prevent all collisions even with nonzero sensing error.

V. EVALUATION

We have evaluated the performance of our method via simulation experiments on four navigation benchmarks, shown

in Fig. 4. These include:

- 3-agents: A group of two agents interacts with a single agent.
- 8-agents: Eight agents are placed along the circumference of a disc and have to walk to their antipodal positions.
- Hallway: Two groups of 75 agents each cross paths in a hallway while coming from opposite directions. The width of the hallway can accommodate 7 to 10 agents abreast.
- Crossing: Four groups of 30 agents each enter from four perpendicular directions and cross each other.

The simulation parameters used in our experiments are the same as in Karamouzas et al. [18], except the time step $\Delta t = 5$ ms.

We model sensing error as follows. At each time t , the sensed relative velocity $\hat{\mathbf{v}}_{ij}(t)$ for a pair of interacting agents is randomly perturbed from the true velocity $\mathbf{v}_{ij}(t)$ via

$$\hat{\mathbf{v}}_{ij}(t) = \mathbf{v}_{ij}(t) + \boldsymbol{\eta}_{ij}(t), \quad (28)$$

where $\boldsymbol{\eta}_{ij}$ is a random process $\mathbb{R} \rightarrow \mathbb{R}^2$. In our experiments, we have considered a number of different models for the distribution and temporal behavior of $\boldsymbol{\eta}_{ij}(t)$.

For simplicity, we assume that at any time t , the errors $\boldsymbol{\eta}_{ij}(t)$ for different pairs of agents i, j are independent. For the temporal behavior of the error, we consider two cases:

- 1) White noise: For a given pair of agents i, j , the errors at different time instants $\boldsymbol{\eta}_{ij}(t_1)$ and $\boldsymbol{\eta}_{ij}(t_2)$ are independent and identically distributed.
- 2) Systematic error: For a given pair of agents i, j , the error is constant over time, $\boldsymbol{\eta}_{ij}(t) = \boldsymbol{\eta}_{ij}(0)$.


Finally, for the distribution of the error $\boldsymbol{\eta}_{ij}(t)$ over \mathbb{R}^2 , we consider two classes of distributions, parametrized by noise magnitude ν :

- 1) The uniform distribution on a disc centered at $\mathbf{0}$ with radius ν , which we denote $\mathcal{D}(\mathbf{0}, \nu)$; this satisfies our assumption that the error is bounded by ϵ if $\nu \leq \epsilon$.
- 2) The bivariate normal distribution $\mathcal{N}(\mathbf{0}, \frac{1}{4}\nu^2\mathbf{I})$, which has the same mean and covariance as $\mathcal{D}(\mathbf{0}, \nu)$.

The combination of error distribution and temporal model give rise to four error models. For each error model and each navigation scenario, we ran $N = 100$ simulations with different random seeds. For each run, we checked whether any collisions between two agents occurred over the course of the simulation, and computed the mean travel time over all agents.

A. Results

a) Collision-free behavior: To experimentally verify the collision-free guarantees of our models, we ran 100 simulations for the 8-agents scenario for each interaction model (TTC, UTTC-I, UTTC-A) and error model. We counted the number of runs in which collisions occurred and plotted the empirical probability of collision as a function of error magnitude ν . For white noise, we observed no collisions in all cases, even for the TTC model; this is not surprising as the effect of randomly time-varying error is likely to cancel



(a) disc distribution error (b) normal distribution error

Fig. 5: Fraction of runs that result in collisions as a function of sensor error. The dashed line denotes the uncertainty bound. Both UTTC-I and UTTC-A are collision-free as long as the noise does not exceed the uncertainty bound.

out over time. For systematic error, the results are shown in Figure 5.

As can be seen in the plots, the standard TTC model cannot prevent collisions for even a small amount of error. Our models guarantee collision-free interactions as long as the sensing errors of agents do not exceed the velocity uncertainty in the model. However, if the errors increase past the uncertainty bound, the assumptions of the model are violated and collisions may begin to occur. Interestingly, while the normal distribution is unbounded and therefore never satisfies our bounded error assumption, our method is still able to prevent most collisions as long as the magnitude of the error is not too large. This suggests that our approach may be useful in practice even when the error does not follow a bounded distribution.

b) Effect on motion efficiency: The addition of uncertainty in our model causes agents to move more cautiously and maintain a greater separation from each other. Therefore, one may expect that the collective navigation behavior may become less time efficient than in the standard TTC model. To quantify this effect, we ran simulations comparing the mean travel time over all agents under the TTC model with the mean travel time under UTTC-I and UTTC-A with $\epsilon = 0.2$. In the UTTC cases, we tested two settings of error magnitude, $\nu = 0$ and $\nu = 0.2$. This gives five different conditions for each of the four scenarios. We ran each condition 100 times and report the mean and standard deviation of the mean travel time in Table I.

The results illustrate that adding uncertainty causes only a modest increase in travel time. The TTC model is equivalent to UTTC with $\epsilon = 0$, so the error-free case ($\nu = 0$) quantifies the effect of uncertainty alone on the collective motion of the agents. It can be seen that the mean travel time under the UTTC-I model is almost always larger than that under UTTC-A, which is in turn larger than that under TTC. This difference increases in congested scenarios such as the hallway, and under large uncertainty UTTC may fail to generate self-organized phenomena such as lane formation, which allow interactions to be solved efficiently at a macroscopic level. This is likely related to the fact that UTTC agents are more conservative and

	TTC	UTTC-I		UTTC-A	
		$\nu = 0$	$\nu = 0.2$	$\nu = 0$	$\nu = 0.2$
3-agents	13.04	13.13	13.30± 0.26	13.03	13.25 ± 0.36
8-agents	14.81±0.31	15.07 ± 0.42	14.49 ± 0.74	15.35 ± 0.56	14.10 ± 0.59
hallway	56.17	67.23	70.98± 8.08	58.97	63.06 ± 3.30
crossing	55.83	58.17	59.04± 0.51	55.85	57.08 ± 0.39

TABLE I: Mean travel time for all scenarios and interactions models. Reported numbers are the averages over 100 simulation runs. UTTC-I and UTTC-A use an uncertainty of $\epsilon = 0.2$. For the 8-agent scenario, we perturbed the initial and goal positions of all agents randomly to break the symmetry, thus we have a nonzero standard deviation even for the $\nu = 0$ cases.

respond to a larger range of potentially colliding velocities ($S_{\text{TTC}} \subset S_{\text{adv}} \subset S_{\text{iso}}$), as shown in Section III-C), leading them to keep a greater distance from each other. Mean travel times increase further in the presence of nonzero error ($\nu = 0.2$), but in this case the TTC model cannot prevent collisions. The difference between the mean travel times of UTTC-I and UTTC-A continues to hold.

c) Computational cost: We measured the performance of TTC and the UTTC models with our single-threaded implementation running on a 3.5GHz Intel Xeon processor. The computational cost depends on the number of agents in the scenario, but all our examples ran in real time. Similar to TTC, UTTC takes $O(k)$ time to compute a new velocity for an agent at a given time step, where k denotes the number of neighbors that are inside the sensing range of the agent. In scenarios where the crowd density remains constant, UTTC runs asymptotically in $O(n)$ time per simulation cycle for n agents. In the hallway scenario, for example, UTTC took between 1.67 ms/frame (for UTTC-A) to 1.96 ms/frame (for UTTC-I) to compute new velocities for the agents, an increase of only 11% and 36% over TTC respectively.

VI. DISCUSSION

Our paper builds on prior work on time-to-collision based local navigation of multiple agents. We have focused on the problem of collision avoidance while accounting for uncertainty in the sensor data of the agents. Two specific sensor error models were introduced, an isotropic model which conservatively considers all possible errors, and an adversarial model that assumes the error is towards a head-on collision. We have provided theoretical guarantees on the collision avoidance behavior of the agents and demonstrated the applicability of both of our models via simulations. In the future, we would like to test the applicability of our models on physical robots in real environments, as well as run simulation and real-world comparisons with existing VO-based uncertainty models.

Our method relies on the assumption that the sensor error is bounded by known constants. In practice, these error bounds could be estimated a priori using a calibration procedure [29]. The uncertainty could also be estimated on the fly using methods such as LTA [23] and BRVO [21] which estimate the motion of other agents from noisy data, using a pre-trained pedestrian model or ensemble Kalman filtering respectively. These methods are thus complementary to our work.

Our work has a number of limitations that we would like to address in the future. Firstly, we have provided theoretical guarantees only for TTC and UTTC-I in the error-free case. In the future, we would like to provide a more thorough analysis establishing similar guarantees in all cases. Second, neither the TTC model nor our UTTC models respect kinodynamics and differential constraints. As system dynamics can significantly affect the complexity of multi-agent navigation problems [6, 16], we would like to extend our formulations to account for different kinematic systems. Such an extension will allow us to generalize our uncertainty-aware interaction models to a wide range of agents (the approach of [28], for example, can be adapted to enable UTTC on differential-drive agents). Third, to guarantee numerical stability, a very small time step Δt must be chosen during each sensing-acting cycle. In recent work [19], we have shown that collision-free motion can be guaranteed with arbitrary time steps in the context of crowd simulation, albeit using a centralized solver. It would be interesting to see if such techniques can be adapted to robot navigation.

As shown in our experiments, incorporating uncertainty in the TTC formulation leads to agents that exhibit more conservative behavior than vanilla TTC agents. For example, a pair of agents approaching each other in a hallway will keep some extra distance between them due to the uncertainty in sensing each other’s velocities. In future work, we would like to quantify this effect and further investigate the properties of UTTC-A, which appears to be less conservative than UTTC-I in all of our simulations.

Another limitation of this and other VO-based work, is that collisions are detected by linearly extrapolating the agents’ trajectories based on their current velocities. In the future, we would like to extend our time-to-collision formulation to account for nonlinear motion. The recent work in [36] that uses the notion of continuous collision probabilities can provide some interesting ideas in this direction. Finally, we note that due to the decentralized nature of each agent’s decision, our UTTC models cannot provide any formal guarantees about the global behavior of the agents. Some local approaches have been proposed focusing on the adaptation of the goal velocity to alleviate congestion and deadlock situations [7, 13], which would be interesting to combine with our method. The work in [21] is also highly relevant as it can allow us to infer goal velocities from captured real crowd data.

REFERENCES

- [1] Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1):99 – 134, 1998.
- [2] Javier Alonso-Mora, Andreas Breitenmoser, Paul Beardsley, and Roland Siegwart. Reciprocal collision avoidance for multiple car-like robots. In *IEEE International Conference on Robotics and Automation*, pages 360–366, 2012.
- [3] Javier Alonso-Mora, Andreas Breitenmoser, Martin Ruffli, Paul Beardsley, and Roland Siegwart. Optimal reciprocal collision avoidance for multiple non-holonomic robots. In *Distributed Autonomous Robotic Systems*, pages 203–216. Springer, 2013.
- [4] Tucker Balch and Maria Hybinette. Social potentials for scalable multi-robot formations. In *IEEE International Conference on Robotics and Automation*, volume 1, pages 73–80, 2000.
- [5] Kostas E Bekris and Lydia E Kavraki. Greedy but safe replanning under kinodynamic constraints. In *IEEE International Conference on Robotics and Automation*, pages 704–710, 2007.
- [6] Kostas E Bekris, Devin K Grady, Mark Moll, and Lydia E Kavraki. Safe distributed motion coordination for second-order systems with different planning cycles. *The International Journal of Robotics Research*, 31(2):129–150, 2012.
- [7] Andrew Best, Sahil Narang, Sean Curtis, and Dinesh Manocha. Densesense: Interactive crowd simulation using density-dependent filters. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 97–102, 2014.
- [8] Andrew Best, Sahil Narang, and Dinesh Manocha. Real-time reciprocal collision avoidance with elliptical agents. In *IEEE International Conference on Robotics and Automation*, pages 298–305, 2016.
- [9] Paolo Fiorini and Zvi Shiller. Motion planning in dynamic environments using velocity obstacles. *The International Journal of Robotics Research*, 17:760–772, 1998.
- [10] Thierry Fraichard and Hajime Asama. Inevitable collision states: a step towards safer robots? *Advanced Robotics*, 18(10):1001–1024, 2004.
- [11] Chiara Fulgenzi, Anne Spalanzani, and Christian Laugier. Dynamic obstacle avoidance in uncertain environment combining PVOs and occupancy grid. In *IEEE International Conference on Robotics and Automation*, pages 1610–1616, 2007.
- [12] Andrew Giese, Daniel Latypov, and Nancy M Amato. Reciprocally-rotating velocity obstacles. In *IEEE International Conference on Robotics and Automation*, pages 3234–3241, 2014.
- [13] Julio E Godoy, Ioannis Karamouzas, Stephen J Guy, and Maria Gini. Adaptive learning for multi-agent navigation. In *International Conference on Autonomous Agents and Multiagent Systems*, pages 1577–1585. International Foundation for Autonomous Agents and Multiagent Systems, 2015.
- [14] Dirk Helbing and Peter Molnar. Social force model for pedestrian dynamics. *Physical Review E*, 51(5):4282, 1995.
- [15] Daniel Hennes, Daniel Claes, Wim Meeussen, and Karl Tuyls. Multi-robot collision avoidance with localization uncertainty. In *International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, pages 147–154. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- [16] Jeffrey Kane Johnson. A novel relationship between dynamics and complexity in multi-agent collision avoidance. In *Proceedings of Robotics: Science and Systems*, 2016.
- [17] Ioannis Karamouzas and Stephen J. Guy. Prioritized group navigation with formation velocity obstacles. In *IEEE International Conference on Robotics and Automation*, pages 5983–5989, 2015.
- [18] Ioannis Karamouzas, Brian Skinner, and Stephen J. Guy. Universal power law governing pedestrian interactions. *Physical Review Letters*, 113:238701, 2014.
- [19] Ioannis Karamouzas, Nick Sohre, Rahul Narain, and Stephen J. Guy. Implicit crowds: Optimization integrator for robust crowd simulation. *ACM Transactions on Graphics*, 36(4), 2017.
- [20] O. Khatib. Real-time obstacle avoidance for manipulators and mobile robots. *The International Journal of Robotics Research*, 5(1):90–98, 1986.
- [21] Sujeong Kim, Stephen J. Guy, Wenxi Liu, David Wilkie, Rynson W.H. Lau, Ming C. Lin, and Dinesh Manocha. BRVO: Predicting pedestrian trajectories using velocity-space reasoning. *The International Journal of Robotics Research*, 34(2):201–217, 2015.
- [22] Andrew Kimmel, Andrew Dobson, and Kostas Bekris. Maintaining team coherence under the velocity obstacle framework. In *International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, pages 247–256, 2012.
- [23] S. Pellegrini, A. Ess, K. Schindler, and L. van Gool. You’ll never walk alone: Modeling social behavior for multi-target tracking. In *IEEE 12th International Conference on Computer Vision*, pages 261–268, 2009.
- [24] C. W. Reynolds. Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4):24–34, 1987.
- [25] C. W. Reynolds. Steering behaviors for autonomous characters. In *Game Developers Conference*, pages 763–782, 1999.
- [26] Shawn Singh, Mubbasir Kapadia, Billy Hewlett, Glenn Reinman, and Petros Faloutsos. A modular framework for adaptive agent-based steering. In *ACM Symposium on Interactive 3D Graphics and Games*, pages 141–150, 2011.
- [27] J. Snape, J. v. d. Berg, S. J. Guy, and D. Manocha. The hybrid reciprocal velocity obstacle. *IEEE Transactions*

on Robotics, 27(4):696–706, 2011.

- [28] Jamie Snape, Jur Van Den Berg, Stephen J Guy, and Dinesh Manocha. Smooth and collision-free navigation for multiple robots under differential-drive constraints. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4584–4589, 2010.
- [29] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics*. MIT press, 2005.
- [30] J. van den Berg, S. Patil, and R. Alterovitz. Efficient approximate value iteration for continuous gaussian POMDPs. volume 3, pages 1832–1838, 2012.
- [31] Jur van den Berg, Ming Lin, and Dinesh Manocha. Reciprocal velocity obstacles for real-time multi-agent navigation. In *IEEE International Conference on Robotics and Automation*, pages 1928–1935, 2008.
- [32] Jur van den Berg, Stephen J. Guy, Ming Lin, and Dinesh Manocha. Reciprocal n-body collision avoidance. In *Robotics Research: The 14th International Symposium ISRR*, volume 70 of *Springer Tracts in Advanced Robotics*, pages 3–19. Springer, 2011.
- [33] Jur van den Berg, Jamie Snape, Stephen J Guy, and Dinesh Manocha. Reciprocal collision avoidance with acceleration-velocity obstacles. In *IEEE International Conference on Robotics and Automation*, pages 3475–3482, 2011.
- [34] Jur van den Berg, Sachin Patil, and Ron Alterovitz. Motion planning under uncertainty using iterative local optimization in belief space. *The International Journal of Robotics Research*, 31(11):1263–1278, 2012.
- [35] David Wilkie, Jur van den Berg, and Dinesh Manocha. Generalized Velocity Obstacles. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5573–5578, 2009.
- [36] David Wolinski, Ming C. Lin, and Julien Pettré. Warp-driver: Context-aware probabilistic motion prediction for crowd simulation. *ACM Transactions on Graphics*, 35(6):164:1–164:11, 2016. ISSN 0730-0301.
- [37] F Zanlungo, T Ikeda, and T Kanda. Social force model with explicit collision prediction. *EPL (Europhysics Letters)*, 93(6):68005, 2011.