

Pedestrian Simulation using Geometric Reasoning in Velocity Space

Sean Curtis and Dinesh Manocha

University of North Carolina at Chapel Hill,
Chapel Hill, NC, USA
{seanc, dm}@cs.unc.edu
<http://gamma.cs.unc.edu/PEDS>

Abstract. We present a novel pedestrian representation based on a new model of pedestrian motion coupled with a geometric optimization method. The model of pedestrian motion seeks to capture the underlying physiological and psychological factors which give rise to the *fundamental diagram* — the phenomenon that pedestrian speed reduces as density increases. The optimization method computes collision-free velocities directly in velocity space. The resultant method exhibits the same types of self-organizing behaviors shown by previous models, is both computationally efficient and numerically stable, can be intuitively tuned to model cross-cultural variation, and is sufficiently robust that a single set of simulation parameters produces viable results in multiple scenarios.

1 Introduction

The ultimate goal of pedestrian dynamics is to be able to evaluate pedestrian responses to various environments and conditions *in vitro*. Given the specification of a novel space, we seek to apply a simulation to predict how real humans will act in that space. As with all simulation domains, the challenge of simulating human pedestrians is one of balancing the need for a model which is comprehensible and computationally tractable against the requirement that the simulation results sufficiently reproduce real-world phenomena; the credibility of predictions based on simulations depend on the belief that the model accurately extrapolates behavior into new spaces. Unlike other simulation domains (e.g. fluid simulation), pedestrian dynamics possesses no governing equations against which pedestrian models can be compared. As such, it is impossible to definitively assert if a single “correct” model exists or what it may look like. Scientists actively search for models which exhibit human-like behaviors while simultaneously possessing desirable computational characteristics.

To that end, we present a novel microscopic model for simulating pedestrians. Like other pedestrian models, our model computes a collision-free velocity for each virtual pedestrian based on a *preferred* velocity and its local neighborhood of static and dynamic obstacles (e.g. other pedestrians.) Our model uses a geometric optimization technique, called ORCA[1], to directly compute, in velocity space, the collision-free velocity that most closely matches the preferred velocity.

The ORCA model has been shown to exhibit self-organizing behavior [20] and can reproduce small-scale pedestrian-pedestrian interactions [21]. It has also been shown to be efficient [22] and stable [23]. Together, efficiency and stability dictate how practical the approach is. Efficiency describes the model’s computation cost per simulation time step. The stability indicates how large a time step can be taken. High stability and efficiency means simulations can run very quickly, allowing multiple iterations to be taken.

Despite its computational strengths, pedestrians simulated with ORCA do not adhere to the fundamental diagram; their speed does not decrease as density increases (see Fig. 2.) We provide a unique model of pedestrian motion which seeks to capture the underlying factors which cause groups of pedestrians to exhibit the fundamental diagram. Pedestrian models typically function by adapting a preferred velocity to local dynamic conditions. Usually, the preferred velocity represents the speed and direction the pedestrian would take if he were alone in the space. We apply our motion model directly to the preferred velocity, adapting it to density conditions, and then pass the result to the pedestrian model. This density-dependent behavior model introduces behaviors consistent with the fundamental diagram into ORCA while preserving ORCA’s computational benefits. Finally, for the experiments given, the model is robust; a single set of simulation parameters produces good results across different scenarios.

The remainder of the paper is organized as follows: We give a brief overview of related work in section 2. In section 3 we present the basis of our geometric optimization algorithm, *velocity obstacles* and ORCA. Section 4 gives the motivation, details and impact of the density-dependent pedestrian behavior model. In Section 5 we generalize the model for two dimensions and show the results. Finally, in Section 6 we analyze the model’s strengths and weaknesses.

2 Related Work

There has been a great deal of work in the pedestrian dynamics community in devising models for pedestrian motion and interaction. The research has largely focused on two approaches: cellular automata (CA) and social forces (SF). Recently, motion-planning techniques in velocity space have been applied to pedestrian simulation with promising results.

Cellular automata approaches work by discretizing the simulation domain into a grid of “cells”. An agent may occupy at most one cell and a cell may contain no more than a single agent. Pedestrians move from cell to cell based on various techniques. These approaches were quite common in simulating vehicular traffic. Blue and Adler first applied these techniques to pedestrians using a set of crafted rules for determining pedestrian movement and showed that they were able to reproduce observable pedestrian phenomena in both uni-directional and bi-directional flows [3, 4]. Other approaches applied probabilistic techniques to determine agent movement and augmented the functionality with scalar and vector fields to reproduce emergent phenomena [5]. CA-based approaches tend to exhibit some artifacts due to the discretization of the space. There has been

extensive research to understand and mitigate these effects. Kirchner et al. examined the impact of the size of the cells [6]. Maniccam investigated the impact of exchanging square cells for hexagonal cells [7]. Yamamoto et al. introduced the “real-coded cellular automata” to minimize the aliasing artifacts which arise from traversing a rectilinear grid [8]. CA-based approaches remain an active branch of research; these methods have been used as a basis to simulate complex scenarios with more elaborate pedestrian behaviors. [9, 10].

Social-force-based models treat pedestrians like mass-particles and applies Newtonian-like physics to evolve the simulation. Helbing introduced the first such model in which each agent was motivated by a driving force which served to lead the agent towards a goal [11]. The agent would interact with obstacles and agents in the environment through the superpositioning of repulsive forces. Later SF-based models explored alternate formulations of the repulsive forces or novel forces including, compression and friction forces for evacuation [12], relative-velocity-dependent forces [13, 14], and group formations [15]. A more complete overview can be found in [16]. SF-based models have been shown to exhibit self-organizing behaviors [17].

Recently, techniques from robotics have been applied to pedestrian simulation with promising results. Agent interactions are based on *velocity obstacles* (VO). One agent defines a set of velocities for another agent which would lead to an inevitable collision. Agents then select a velocity outside of this set for a collision-free path. Originally, each agent assumed that all other agents were non-responsive [2]. Later VO-based models introduced alternative VO formulations which assumed that agents were aware that those around them would also respond [18, 19, 1]. Pedestrian simulation using VO-based models has been shown to exhibit self-organizing behaviors and varying flow through a bottleneck [20] as well as accurate microscopic interactions between individual pedestrians [21].

3 Velocity Obstacles

Before we discuss the details, we’ll define some notation and terminology. In discussing real and virtual humans, we need to be able to make a clear distinction between the two. To that end, we make the arbitrary choice to use the term “pedestrian” or “subject” to refer to real humans and “agents” for virtual humans. Each agent is modeled as a two-dimensional circle modeled with the state vector: $[r, \mathbf{p}, \mathbf{v}, \mathbf{v}^0]^T \in \mathbb{R}^7$, where r is the radius of the disk, \mathbf{p} , \mathbf{v} , and \mathbf{v}^0 are two-dimensional vectors representing the current position, current velocity and preferred velocity, respectively. By convention v^0 and v are the preferred speed and current speed (i.e. the magnitudes of the corresponding vectors.) When properties of agent i are discussed, the elements of the state will be subscripted with the agent’s index (e.g. r_i and \mathbf{v}_i^0 .)

Velocity obstacles exploit an intuitive principle. As two agents move through a shared space, each agent observes the position *and* velocity of the other agents and predicts whether a conflict will happen in the future. For example, agent

i may occupy a position on agent j 's intended path of travel, but if i is gone before j arrives, there is no conflict. This principle is formalized with the *velocity obstacle* [2].




Fig. 1. Velocity obstacles. (a) The velocity obstacle formed by agent j on i . (b) The truncated velocity obstacle for time window τ . (c) Agent i assumes j takes constant velocity \mathbf{v}_j ; VO_{ji} is displaced by \mathbf{v}_j . (d) The formulation of the half plane $ORCA_{ji}$, including the minimum change in relative velocity, \mathbf{u} , the direction of minimum change, \mathbf{n} . We also show the symmetric half plane $ORCA_{ij}$.

The velocity obstacle is, as the name implies, an obstacle, but rather than lying in workspace or configuration space, it lies in velocity space. Velocity space, for pedestrians, is a two-dimensional space in which each point represents a velocity. In Cartesian space, a position is a point and a velocity is a vector. Each velocity vector in Cartesian space maps to a point in velocity space. So, a workspace obstacle represents a set of positions which the agent cannot occupy; a velocity obstacle represents a set of velocities an agent cannot assume.

For agents i and j , agent j induces a velocity obstacle on i , VO_{ji} , and i induces a symmetric velocity obstacle on j , VO_{ij} . The velocity obstacle VO_{ji} is a cone, originating at \mathbf{p}_i , which tightly bounds the Minkowski sum of agent i 's geometry with j 's (Fig. 1(a)). If the *relative* velocity between agents i and j remains within this cone, there will be an inevitable collision. In practice, we are typically concerned with collisions that can occur within the next τ seconds. Including this term truncates the cone (Fig. 1(b)); very small relative velocities will not lead to collisions within τ seconds, so these velocities are not included in the obstacle.

This velocity obstacle consists of the set of *relative* velocities that lead to collision. A single agent cannot exert unilateral control over the relative velocity. If agent i assumes that j will not change velocity, i must take full responsibility for avoiding the collision. This behavior is modeled by translating VO_{ji} by j 's velocity (Fig. 1(c).) Given this assumption, agent i can avoid collision with agent j by selecting a *feasible* velocity — a velocity that lies outside the velocity obstacle. This is the original velocity obstacle formulation, in which each agent assumes that every other agent is a non-responsive, dynamic obstacle [2].

However, this model is a poor representation of pedestrians because pedestrians do, in fact, react to each other. This leads to two significant issues for agents.

First, if agent i predicts agent j 's velocity incorrectly, the velocity obstacle will have error (the amount of error depends on how bad the prediction is.) The error may be sufficiently large, that even if both agents select velocities lying outside the velocity obstacle, they may still collide within τ seconds. Second, both agents overreact to their neighbors (because they falsely assume that the other will make no effort to avoid collision.) This can easily lead to oscillatory motion as the agents overreact in successive steps.

Van den Berg et al. proposed an alternate formulation to VO which addresses these issues: Optimal Reciprocal Collision Avoidance (ORCA) [1]. The truncated cone, VO_{ji} , is replaced with a half plane, $ORCA_{ji}$. The half plane can be considered a union of many truncated cones. In effect, rather than assuming the agent j will take a particular velocity, agent i assumes that it will take one velocity from a large set and plans to avoid all of the velocities in that set. Agent j does the same with respect to i . These two sets are defined so that there is no pair of feasible velocities which can lead to collision within τ seconds. This solves the first problem. The second problem — oscillatory motion due to individual overreaction — is solved by defining the planes such that the amount of change to the relative velocity required to avoid collision is evenly distributed between the two agents. Each agent assumes the other agent will make an equal effort to avoid collision. The effort is evenly apportioned and no agent overreacts.

The ORCA half-plane can be constructed geometrically in the following manner. Assume agents i and j have velocities \mathbf{v}_i and \mathbf{v}_j , respectively, and that these velocities place them on a collision course (i.e. $\mathbf{v}_i - \mathbf{v}_j \in VO_{ji}$.) Let \mathbf{u} be the vector from $\mathbf{v}_i - \mathbf{v}_j$ to the closest point on the boundary of the velocity obstacle (Fig. 1(d).) More formally,

$$\mathbf{u} = (\operatorname{argmin}_{\mathbf{v} \in \partial VO_{ji}} \|\mathbf{v} - (\mathbf{v}_i - \mathbf{v}_j)\|) - (\mathbf{v}_i - \mathbf{v}_j) \quad (1)$$

is the minimum change in relative velocity between i and j necessary to guarantee no collision within τ seconds. To model the reciprocity, half of the minimum change is applied to each agent. So, we can define the ORCA velocity obstacle induced by agent j on agent i as:

$$ORCA_{ji} = \{\mathbf{v} | (\mathbf{v} - (\mathbf{v}_i + \frac{1}{2}\mathbf{u})) \cdot \hat{\mathbf{n}} < 0\}, \quad (2)$$

where $\hat{\mathbf{n}}$ is the normalized direction of \mathbf{u} . Proof of the guarantees can be seen in the original paper [1].

The ORCA algorithm has several desirable properties:

- It is very computationally efficient. For example, it has been used to simulate 35,000 agents in better than real time (2.6 seconds of simulation for 1 second of computation) [22].
- It has been shown to produce consistent and stable results for time steps as large as 0.2 s [23].
- Pedestrian simulations with this model have exhibited self-organizing behaviors (e.g. lane formation, arching, jamming, etc.) and reduced flow through a bottleneck [20].

- Microscopic interactions between individual agents have also been reproduced using these approaches [21].

However, there is an important phenomenon, which pedestrians exhibit, that ORCA agents do not: the fundamental diagram (see Fig. 2.) To be considered a viable pedestrian model, ORCA agents must reproduce the fundamental diagram. Our goal is to introduce this property while preserving ORCA’s desirable properties.

4 Adherence to the Fundamental Diagram

The fundamental diagram is the observed relationship between pedestrian density and speed [24]; as density increases, speed decreases. However, ORCA agents do not adhere to this pattern; the agents can move at arbitrarily high speeds even at extreme density. This is because the ORCA algorithm is only concerned with collisions. If the relative velocity between two agents is zero, no collision is possible, regardless of their relative positions. We illustrate this phenomenon by reproducing a real-world experiment with ORCA.

Seyfried et al. performed a simple experiment to examine the fundamental diagram [25]. They approximated a one-dimensional scenario with periodic boundary conditions by creating a narrow elliptical path, in which pedestrians were forced to walk single file. They performed a series of tests. In each test, they increased the number of subjects in the ring, implicitly increasing the density, and measured the speed of the subjects. We reproduced their experiment in a virtual context.¹ The results of the simulation are shown in Fig. 2. The black dots show the experimental data; as density increases, speed decreases. The white dots show the simulated results; agent speeds remain consistent, regardless of the density.² Human beings clearly exhibit a response to density which leads to changes in their walking speed which is not captured by the ORCA formulation.

To use ORCA as a pedestrian model, ORCA agents must, in addition to self-organizing behaviors, also show behavior consistent with the fundamental diagram. We propose a density-dependent behavior model, which models underlying physiological and psychological factors which lead to the fundamental diagram. In this section, we will discuss the evidence for what these factors may be, define a model based on those observations, and show the results of applying that model.

4.1 Model of Density-dependent Behavior

Biomechanists have performed extensive analysis on the human walking gait and have deduced the following [26]:

¹ We refer the reader to the original paper [25] for the full experimental parameters.

² The variation in simulation speed arises from normally distributed preferred speeds for the agents.




Fig. 2. Comparison of ORCA with experimental data in a simple one-dimensional experiment [25]. The experimental data (black dots) show decreasing speed as density increases. The simulated data (white dots) show consistent speed, regardless of density.

1. Humans choose a gait that minimizes the energy expenditure per unit distance traveled.
2. For any given speed, there is a “natural” stride length which minimizes the energy expenditure. Imposition of other stride lengths leads to a greater than optimal energy expenditure.
3. The “natural” stride length grows as walking speed increases.

In other words, to walk a particular speed, there is a natural stride length. But the converse is also true; for a particular stride length, there is a natural walking speed. We believe that this physiological property is a significant factor which leads to the fundamental diagram. As density increases, the space available to a pedestrian decreases. If it decreases sufficiently to limit stride length, speed must likewise decrease. To walk a certain speed with a “natural” gait, a minimum amount of space is required.

However, this physiological cause is not the sole contributing factor to the fundamental diagram. Researches have shown that sensitivity is also dependent on cultural factors [27]. In this experiment, the previous one-dimensional experiment was re-created with Indian subjects in place of German subjects. The Indian subjects proved to be less sensitive to density than their German counterparts. From this we deduce that there is a psychological/sociological component. This component serves to increase the amount of space required to walk at a particular speed beyond the strictly physical requirements.³ In addition to this cultural factor, we assert that a pedestrian’s ability to accurately estimate required and available space is imprecise and possibly even instinctual; the fact that pedestrians rarely collide suggests that this estimation is conservative.

We combine the biomechanical data with the social observations to create a model which accounts for the human factors which lead to the fundamental diagram. Specifically, we propose a linear model of a physiological constraint and a psychological constraint.

³ In ordered marching, as seen with soldiers, the fundamental diagram can be completely eliminated as walking can be coordinated to maintain speed at high density.

Physiological constraint: Biomechanists have researched the relationship between the dominant modes of human walking. Although the pedestrian dynamics community largely subscribes to the idea that the relationship between stride length (L) and walking speed (v) is linear [24], more recent research suggests the relationship is actually quadratic [26, 28]:

$$L(v) = \frac{H}{\alpha} \sqrt{v}, \quad (3)$$

where $\alpha = 1.57$ is a “stride factor”, as determined by Dean from experimental data [28] and H is the normalized height of the agent ($height / 1.72 \text{ m.}$)⁴

Psychological constraint: Psychologists and biomechanists have shown that humans have a strong sense of personal space — a “buffer” that extends beyond mere physical requirements [29, 27]. We are unaware of any mathematical models which relate changes in personal space to speed or cultural factors, but we recognize the fact that they do play a role and must be accounted for. In the absence of more specific data, we assume the extra buffer space (B) is a simple linear function of the physiological space requirement:

$$B(v) = \beta L(v). \quad (4)$$

When walking slowly and taking small steps, the buffer space is small. When walking quickly, and taking large steps, the buffer space grows. It is controlled by the stride buffer parameter, β . We will show that varying the buffer value is sufficient to reproduce differences observed across different cultures (Section 4.2.)

From these constraints, we can define our novel models of both: the space required (S) to move at a particular speed and its inverse, the natural speed for a given available space:

$$S(v) = B(v) + L(v) = (1 + \beta)L(v) = (1 + \beta) \frac{H}{\alpha} \sqrt{v}, \quad (5)$$

$$v(S) = \max \left(v^0, \left(\frac{S\alpha}{H(1 + \beta)} \right)^2 \right). \quad (6)$$

We use the equations to modify the *preferred speed* of the agent. The pedestrian models we’ve discussed in this paper (SF, CA and VO) all share a common paradigm. Each agent has a strategic intent which is adapted to local dynamic conditions by the pedestrian model. The strategic intent is typically the trajectory the agent would walk in the absence of other agents. Derived from a global navigation strategy, this might be the shortest path through a building to its goal traversed at some preferred speed. This strategy is communicated to the pedestrian model as a preferred velocity toward an immediate goal. The pedestrian model adapts the preferred velocity to account for nearby agents and

⁴ Dean originally presented the relationship between stride frequency (f) and speed (v) [28]. We have used the relationship $v = fL$ to reformulate Dean’s work as the relationship between stride length and speed.

obstacles such that the agent makes progress towards its goal without collisions. Most research has focused on this local adaptation (probabilities and floor fields in CA [5], repulsive forces in SF [12, 14, 13], and velocity obstacles in VO [2, 1].) However, we suggest that models may be improved if we consider the idea that the intent itself adapts to local conditions. At the very least, by allowing for this possibility, we have access to more degrees of freedom in simulating pedestrians.

By modifying the preferred velocity, we are hypothesizing that pedestrians do adapt their intentions. The data supports this hypothesis. Chattaraj et al. reported that the freespace walking speeds of the subjects in the one-dimensional experiment had a mean speed of 1.24 m/s with a standard deviation of 0.15 m/s [27]. However, in Fig. 3, the experimental subjects show an interesting behavior. At the lowest densities (0.5-1.0 people/m²), the speeds saturate to less than 1 m/s. At the sparsest, the agents should be able to walk faster than 1 m/s, but they choose not to. This is evidence that their intent has adapted to their circumstances.

4.2 Density-dependent Behavior in One Dimension

With this new model for preferred speed, we can repeat the one-dimensional experiment. We give each of the agents a default preferred speed according to the distributions reported above. We've introduced two new parameters: the physiological stride factor (α) and the psychological stride buffer (β). We use normally distributed values for these parameters. The mean for stride factor, $\mu_\alpha = 1.57$, is the value provided by Dean in his work. By varying the value with the standard deviation, $\sigma_\alpha = 0.15$, we simulate normally distributed heights. We selected the mean stride buffer value that fit the data well, which gave us a mean and standard deviation for β of $\mu_\beta = 0.9$ and $\sigma_\beta = 0.2$. All other simulation parameters were the same as the simulation that produced the data in Fig. 2. Fig. 3(a) shows the results using our density-dependent speed model. Unlike the results shown in Fig. 2, the simulated agents closely follow the fundamental diagram displayed by the German subjects.

To illustrate the full effect of the speed buffer parameter, we also reproduced the Indian experiment. The authors reported that the Indian subjects exhibited essentially the same freespace speeds as the German subjects. We make the reasonable assumption that the physiology is the same across cultures, so we should be able to reproduce the Indian subjects' behavior by changing *only* the psychological constraint. Reducing the stride buffer value should decrease the agents' sensitivity to density. Changing β to the distribution, $\mu_\beta = 0.45$, $\sigma_\beta = 0.125$, captured the Indian behavior, supporting the hypothesis (all other parameters were identical.) The results can be seen in Fig. 3(b).

5 Estimating Available Space

To apply (6) to two dimensional scenarios, we need to determine the scalar "space" available to the agent. We propose a method for computing the space




Fig. 3. Comparisons between measured and simulated data in the one-dimensional experiment described in [25, 27]. (a) shows the results with German subjects and (b) shows the same for Indian subjects.

available for an agent, based on the relative positions and velocities of its neighbors. Our method is based on the following principles:

- As personal space has been observed to be biased forwards in the pedestrian’s direction of travel [29], the space that matters most lies in the agent’s preferred direction of travel ($\hat{\mathbf{v}}^0 = \mathbf{v}^0 / \|\mathbf{v}^0\|$.)
- The presence of other agents reduces the amount of available space. The degree to which they limit it depends on their walking speed, their orientation and their position, relative to the preferred direction of travel.
- The total reduction of available space imposed by neighboring agents can be approximated by the single agent that reduces the available space the most.

The method relies on a metric called the “effective” distance. Effective distance (E_{ji}) represents the overall impact of agent j on agent i ’s perception of available space, based on the principles listed above. The perceived space available for agent i (S_i) is the smallest effective distance of all of agent i ’s neighbors. The effective distance of a neighbor, j , depends on its position relative to i , its orientation, and its speed. We define effective distance of agent j with respect to agent i (E_{ji}) as the combination of the Euclidian distance between agents i and j (d_{ji}), a directional penalty (Δ_{ji}), and an orientation penalty (O_{ji}), as follows:

$$S_i = \min \left(\sum_{j \in N_i} E_{ji} \right), \quad (7)$$

$$E_{ji} = d_{ji} + \Delta_{ji} - O_{ji}, \quad (8)$$

$$\delta_i = \frac{(1 + \beta)H}{2\alpha} \sqrt{v_{\max}}, \quad (9)$$

$$\Delta_{ji} = 0.15\delta(1 - (\hat{\mathbf{v}}_i^0 \cdot \hat{\mathbf{d}}_{ij})), \quad (10)$$

$$O_{ji} = \max \left(r_j, \frac{H\sqrt{v_j}(1 + \beta)|\hat{\mathbf{v}}_j \cdot \hat{\mathbf{d}}_{ji}|}{2\alpha} \right), \quad (11)$$

where N_i is the set of i 's neighboring agents, $\hat{\mathbf{d}}_{ji}$ is the unit-vector pointing from agent i to agent j , and δ_i is the maximum Euclidian distance at which another agent can cause agent i to slow down. Fig. 4 illustrates effective distance.

The *directional penalty*, Δ_{ji} , increases the effective distance of an agent based on agent j 's angular displacement from agent i 's preferred direction of travel. If agent j lies directly in the direction of travel, the directional penalty is zero and the effective distance remains the same as the Euclidian distance. As the direction to j moves away from the preferred direction, the effective distance increases. This term represents an approximation of the elliptical personal space observed in experiments [29]. Neighbors to the side of the agent must be closer than those in front to have the same constraining influence.

The orientation penalty computes how much of the distance between agent and neighbor is consumed by the neighbor's stride. The space the agent consumes is predominantly in the direction of its travel. If the agent is moving perpendicularly to the displacement vector ($\mathbf{d}_{ij} \cdot \mathbf{v}_j = 0$), then the closest point on the agent is its side and it consumes a distance equal to its radius. However, if it is walking in the direction of the displacement, it is occupying a portion of the space between i and j equal to half its stride length.




Fig. 4. An illustration of effective distance. (a) shows the position of agents i , a , b , c and d in Cartesian space. (b) shows the effective distance between agents a , b , c , and d with i based on their direction and orientation.

Fig. 4(a) shows the configuration of agent i and its four neighbors in Cartesian space. Agents a and b are closest based on simple Euclidian distance. However, Fig. 4(b) shows the impact of effective distance. The ellipses represent the space consumed by each agent's stride. Because the angle between preferred direction and direction towards agents a and b is large, agents a and b receive a large directional penalty. Agents c and d are physically the same distance from agent i and receive the same directional penalty. But because agent c 's velocity is nearly parallel to the displacement direction, its stride consumes more of the available space than d . Thus, the space available for agent i is E_{ci} .

5.1 Density-dependent Behavior in Two Dimensions

We evaluate the efficacy of the density-dependent behavior in two dimensions by comparing real world experiments with comparable simulated scenarios. Zhang

et al. performed experiments on the flow of groups of pedestrians in a corridor. They examined both uni-directional [30] and bi-directional flow [31]. We refer the reader to the papers for exact details on the spatial configuration of the experiment environment.

Simulation Parameters: It has been observed that humans can be bounded by an ellipse with major and minor axes of 0.24 and 0.15 m, respectively [24]. The area of this ellipse is 0.11 m^2 . We model the agents as circles with the same area as that ellipse: a circle with a radius of 0.19 m^2 . This represents a statistical model of the human more than a physical one.

As with the one-dimensional experiment, we use normally distributed values for three key parameters: preferred speed (v^0), stride factor (α), and stride buffer (β). Because one of our goals is to create a pedestrian model that doesn't require scenario-specific tuning, we applied the same values for the uni- and bi-directional flow experiments as we did for the one-dimensional experiment. All other simulation parameters (e.g. initial density, number of agents, etc.) are as specified in the experiment documentation.

5.2 Two-dimensional, Uni-directional Flow

In the uni-directional flow experiments, Zhang et al. controlled the flow of pedestrians into a corridor of uniform width [30]. They recorded the trajectories of the subjects over multiple iterations, each with an increased level of flow into the corridor (and a corresponding increase in density.) Once the maximum in-flow was achieved, they began constricting the exit width in subsequent iterations to further increase density. We reproduced this scenario. Please, refer to the original paper for the details [30]. The results can be seen in Fig. 5.




Fig. 5. Comparisons of the fundamental diagram for real-world and simulated data for the uni-directional flow experiment [30]. (a) The real-world pedestrian data. (b) The simulated agent data. Weidmann's model [24] is plotted for reference. (c) Comparison of the fundamental diagram exhibited by simulated agents in uni-directional and bi-directional flow. In the range of density between 1.0 and 2.0 m^{-2} , the bi-directional flow speed falls off more quickly. This same phenomenon was observed in real crowds [31].

There are several points we would emphasize with respect to Fig. 5. First, although the agents don't exactly reproduce the behaviors observed in the subjects of the experiment, they do produce a behavior consistent with Weidmann's model of the fundamental diagram. Second, the subjects themselves exhibit speeds far above what Weidmann's model would predict. It may be a property of the set of subjects (mostly students [30].) It may also be something unique to the experiment. This requires further study. Third, the agents achieve much higher density than do the experiment subjects. Nothing physical prevents the subjects from drawing closer together and achieving comparable density. We must assume that there is some psychological/cultural factor which keeps them sparsely distributed for which our model doesn't account.⁵

5.3 Two-dimensional, Bi-directional Flow

In a later paper, Zhang et al. performed similar experiments as those in the previous section, but this time included bi-directional flow [31]. One of the interesting phenomena observed in this experimental data is that, in the band of density between 1.0 and 2.0 m^{-2} the bi-directional flow is more sensitive to density and the speed decreases more quickly. We observed this same phenomenon in our simulated data, as shown in Fig. 5(c).⁶ At the same time, the previously noted density issues also arose in the simulation of bi-directional flow.

6 Analysis and Conclusion

We have presented a new model for pedestrian dynamics based on ideas from robotics. The velocity obstacle formulation provides an intuitive basis for simulating human pedestrians. Originally devised for robots, we have shown that it can be easily extended to conform to human behavior based on biomechanical models and empirical observations. We have further shown that the model's parameter space is well-behaved; it does not require per-scenario tuning. We performed all experiments with the same settings. Previous research has shown that this formulation leads to a wide range of self-organizing behaviors (e.g. lane formation, arching, vortices, etc.) [20] and it has also been shown to be a very stable approach, admitting time steps as large as 0.2 s [23].

There are still aspects of this formulation which can be improved. First, while there is very strong agreement between simulation and real-world data in the one-dimensional experiments, the agreement in two dimensions isn't as strong. While the agents exhibit behavior consistent with the fundamental diagram, it's not quite the same as the measured data. One possible source for this disagreement may be the the effective distance function (7). More analysis of the pedestrian data may suggest a better mapping from local agent configurations to "effective

⁵ In the experiment, the subjects began in a space with a density of 3 people/m^2 . If the initial density were higher, would it affect the observed density?

⁶ For visual clarity, we follow the example of the experiment authors in presenting those samples which lie within one standard deviation of the mean.

distance”, yielding better agent behaviors. More importantly, RVO agents exhibit a tendency to reach maximum density with little effort. While pedestrians are known to reach very high densities, it would seem that it requires significant systemic pressures to do so [22]. This requires greater understanding so that it can be incorporated into our model.

Acknowledgments

This research is supported in part by ARO Contract W911NF-10-1-0506, NSF awards 0904990, 1000579, 1117129, and 1142382, and Intel. The experimental data was made possible by DFG-Grant Nos. KL 1873/1-1 and SE 1789/1-1 and the “Research for Civil Security” program funded by German Federal Ministry of Education and Research.

References

1. Van den Berg, J., Guy, S.J., Lin, M., Manocha, D.: Reciprocal n-body collision avoidance. In: Inter. Symp. on Robotics Research. (2009)
2. Fiorini, P., Shiller, Z.: Motion planning in dynamic environments using velocity obstacles. *IJRR* **17**(7) (July 1998) 760–762
3. Blue, V.J., Adler, J.L.: Emergent fundamental pedestrian flows from cellular automata microsimulation. *Transportation Research Record: Journal of the Transportation Research Board* **1644** (1998) 29–36
4. Blue, V.J., Adler, J.L.: Cellular automata microsimulation of bidirectional pedestrian flows. *Transportation Research Record: Journal of the Transportation Research Board* **1678** (1999) 135–141
5. Burstedde, C., Klauck, K., Schadschneider, A., Zittartz, J.: Simulation of pedestrian dynamics using a two-dimensional cellular automaton. *Physica A: Statistical Mechanics and its Applications* **295** (2001) 507–525
6. Kirchner, A., Klupfel, H., Nishinari, K., Schadschneider, A., Schreckenberg, M.: Discretization effects and the influence of walking speed in cellular automata models for pedestrian dynamics. *J. Stat. Mech.* **2004** (October 2004)
7. Maniccam, S.: Traffic jamming on hexagonal lattice. *Physica A: Statistical Mechanics and its Applications* **321** (2003) 653–664
8. Yamamoto, K., Kokubo, S., Nishinari, K.: Simulation for pedestrian dynamics by real-coded cellular automata (rca). *Physica A: Statistical Mechanics and its Applications* **379** (2007) 654–660
9. Bandini, S., Federici, M., Manzoni, S., Vizzari, G.: Towards a methodology for situated cellular agent based crowd simulations. *Engineering societies in the agents world VI* (2006) 203–220
10. Sarmady, S., Haron, F., Talib, A.: A cellular automata model for circular movements of pedestrians during tawaf. *Simulation Modelling Practice and Theory* (2010)
11. Helbing, D., Molnar, P.: Social force model for pedestrian dynamics. *Phys. Rev. E* **51**(5) (May 1995) 4282–4286
12. Helbing, D., Farkas, I., Vicsek, T.: Simulating dynamical features of escape panic. *Nature* **407** (2000) 487–490

13. Yu, W.J., Chen, R., Dong, L.Y., Dai, S.Q.: Centrifugal force model for pedestrian dynamics. *Phys. Rev. E* **72** (Aug 2005) 026112
14. Chraibi, M., Seyfried, A., Schadschneider, A.: Generalized centrifugal-force model for pedestrian dynamics. *Phys. Rev. E* **82**(4) (2010) 046111
15. Moussaïd, M., Perozo, N., Garnier, S., Helbing, D., Theraulaz, G.: The walking behaviour of pedestrian social groups and its impact on crowd dynamics. *PLoS ONE* **5**(4) (2010)
16. Seyfried, A., Schadschneider, A., Kemloh, U., Chraibi, M.: Force-based models of pedestrian dynamics. *Networks and Heterogeneous Media* **6**(3) (2011) 425–442
17. Helbing, D., Molnár, P., Farkas, I., Bolay, K.: Self-organizing pedestrian movement. *Environment and Planning B: Planning and Design* **28**(3) (2001) 361–383
18. Van den Berg, J., Lin, M., Manocha, D.: Reciprocal velocity obstacles for real-time multi-agent navigation. *ICRA* (2008) 1928–1935
19. Van den Berg, J., Patil, S., Sewall, J., Manocha, D., Lin, M.: Interactive navigation of multiple agents in crowded environments. In: *I3D '08, New York, NY, USA, ACM* (2008) 139–147
20. Guy, S.J., Curtis, S., Lin, M.C., Manocha, D.: Least-effort trajectories lead to emergent crowd behaviors. *Phys. Rev. E* **85**(1) (2012) 016110
21. Guy, S.J., Lin, M.C., Manocha, D.: Modeling collision avoidance behavior for virtual humans. In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 2 - Volume 2.* (2010) 575–582
22. Curtis, S., Guy, S.J., Zafar, B., Manocha, D.: Virtual tawaf: A case study in simulating the behavior of dense, heterogeneous crowds. In: *1st IEEE Workshop on Modeling, Simulation and Visual Analysis of Large Crowds.* (2011) 128–135
23. Curtis, S., Snape, J., Manocha, D.: Way portals: Efficient multi-agent navigation with line-segment goals. *I3D '12* (2012) 15–22
24. Weidmann, U.: *Transporttechnik der Fussgaenger.* Schriftenreihe des IVT (1993)
25. Seyfried, A., Steffen, B., Klingsch, W., Boltjes, M.: The fundamental diagram of pedestrian movement revisited. *J. Stat. Mech.* (10) (October 2005)
26. Inman, V.T., Ralston, H.J., Todd, F., Lieberman, J.C.: *Human Walking.* Williams & Wilkins (1981)
27. Chattaraj, U., Seyfried, A., Chakroborty, P.: Comparison of pedestrian fundamental diagram across cultures. *Advances in Complex Systems* **12**(3) (2009) 393–405
28. Dean, G.A.: An analysis of the energy expenditure in level and grade walking. *Ergonomics* **8**(1) (1965) 31–47
29. Gérin-Lajoie, M., Richards, C.L., McFadyen, B.J.: The negotiation of stationary and moving obstructions during walking: Anticipatory locomotor adaptations and preservation of personal space. *Motor Control* **9** (2005) 242–269
30. Zhang, J., Klingsch, W., Schadschneider, A., Seyfried, A.: Transitions in pedestrian fundamental diagrams of straight corridors and t-junctions. *J. Stat. Mech.* **2011**(06) (2011) P06004
31. Zhang, J., Klingsch, W., Schadschneider, A., Seyfried, A.: Ordering in bidirectional pedestrian flows and its influence on the fundamental diagram. *J. Stat. Mech.* (02) (2012) P02002