A Predictive Collision Avoidance Model for Pedestrian Simulation

Ioannis Karamouzas, Peter Heil, Pascal van Beek, and Mark H. Overmars

Center for Advanced Gaming and Simulation, Utrecht University,
The Netherlands
{ioannis,pccheil,phbeek,markov}@cs.uu.nl

Abstract. We present a new local method for collision avoidance that is based on collision prediction. In our model, each pedestrian predicts possible future collisions with other pedestrians and then makes an efficient move to avoid them. Experiments show that the new approach leads to considerably shorter and less curved paths, ensuring smooth avoidance behaviour and visually compelling simulations. The method reproduces emergent behaviour like lane formation that have been observed in real crowds. The technique is easy to implement and is fast, allowing the simulation in real time of crowds of thousands of pedestrians.

Keywords: collision avoidance, interaction, pedestrian simulation.

1 Introduction

Virtual city environments are nowadays commonly used in computer games, simulations, training applications, and online communities. Although in many such applications the cities look very realistic, to become more lively and appealing, they need to be populated with a large number of simulated pedestrians. On one hand, these virtual pedestrians have to be visually attractive to the viewer. On the other hand, they should be able to move through the virtual world in a natural looking way.

Over the past years numerous models have been proposed to simulate individuals, groups and crowds of characters. Treuille *et al.* [1] have recently presented a dynamic potential-field approach that unifies global navigation and local collision avoidance into a single framework. However, their approach is limited to homogeneous groups of characters moving toward a common goal and cannot simulate individuals with distinct characteristics and goals. An alternative approach is to decouple global planning from local collision avoidance [2,3,4,5]. Typically, a graph-based method is used to direct the global motion of the character, whereas its local behaviour is governed by some sort of agent-based [6] or force-field approach [7].

At the local level, the way that virtual characters interact and avoid collisions with each other is essential for the realism of the generated motions. Helbing simulated the behaviour of pedestrians using (social) forces related to physics [7]. Since his original work, a number of models have been proposed to simulate

A. Egges, R. Geraerts, and M. Overmars (Eds.): MIG 2009, LNCS 5884, pp. 41–52, 2009.

crowds of pedestrians under normal and emergency situations (e.g. [8,9]). Nevertheless, in all these approaches, due to lack of anticipation and prediction, the characters interact when they get sufficiently close. Consequently, the resulting motions tend to look unnatural and contain undesirable oscillations. The problem becomes more obvious in large and cluttered environments, with pedestrians constantly changing their orientations, pushing each other and moving back and forth.

An alternative way for solving interactions between virtual humans is based on collision prediction. In these approaches, the agents' trajectories are linearly extrapolated and used to determine collisions in the near future. Based on this idea, Reynolds has introduced the *unaligned collision avoidance* behaviour [6], whereas Feurtey devised an elegant collision detection algorithm that predicts potential collisions within time and resolves them by adapting the speed and/or the trajectories of the agents. Inspired by Feurtey's work, Paris *et al.* [10] presented an anticipative model to steer virtual pedestrians without colliding with each other. Similarly, Shao and Terzopoulos [5] proposed a number of reactive behavioral routines to determine the avoidance maneuvers of the agents. More recently, van den Berg *et al.* [11] introduced the concept of *Reciprocal Velocity Obstacle*. Finally, Pettré *et al.* [12] proposed an egocentric model for local collision avoidance that is based on experimental studies and can be automatically calibrated from the experimental data.

Recently, example-based techniques have also been explored for simulating interacting virtual characters [13,14]. However, these approaches are too computationally expensive and cannot be used for real-time interactive applications.

Contributions. In this paper, we present a new local method for realistic character avoidance. It is based on the hypothesis that an individual adapts its route as early as possible, trying to minimise the amount of interactions with others and the energy required to solve these interactions. Building upon this hypothesis, in our model, a pedestrian computes with which other pedestrians is in collision course within a certain anticipation time. It calculates how such collisions will take place and then make an efficient move to avoid them. Consequently, the characters do not repel each other, but rather anticipate future situations avoiding all collisions long in advance and with minimal effort.

We show that the new approach leads to energy-efficient motions and considerably less curved paths for the pedestrians, resulting in a smooth, natural flow. The generated motions are oscillation-free. The method reproduces emergent behaviours, like lane formation, that have been observed in real crowds. Since it is a predictive rather than a reactive approach, it is somewhat similar in nature to the approaches of [5,10,11] mentioned above. However, it is based on a force-field approach and, hence, it is much easier in its formulation and implementation and considerably faster, allowing real-time simulations of crowds of thousands of pedestrians. Our approach bears also some resemblance with the avoidance behaviour proposed in [6]. In our model, though, the direction and the magnitude of the collision avoidance force is computed in a different way, ensuring smooth avoidance behaviour and visually pleasing simulations.

2 Pedestrian Interactions

In this section we present some key concepts regarding how pedestrians interact with each other in real life. In the next sections, we use these concepts to formulate a model of realistic collision avoidance.

Scanning and Externalization. Human interactions have been widely studied in the field of sociology. Of particular importance are the studies of Goffman [15] related to how people behave at a microscopic level. According to his observations, two processes govern the avoidance behaviour of pedestrians, the *externalization* and the *scanning*. In externalization, a pedestrian uses its body language to inform the others about its intentions, that is, to indicate its planned course. At the same time, it continuously scans the environment to gather the signals given out by other pedestrians. Eventually, a voluntary coordination takes place and a collision is resolved.

Personal Space. Another important concept in social interactions is the *personal space* that surrounds an individual. More formally, the personal space can be defined as the portable territory around an individual which others should not invade. It regulates the safe distance that an individual needs to maintain from others in order to feel comfortable. The size and the shape of the personal shape are constantly changing depending on the crowd density, as well as the travel speed of the pedestrian. According to Goffman [15], the personal space can be represented as an oval, narrow to the sides of the individual and long in front of him/her.

Principle of Least Effort. The *principle of least effort* originates from the field of psychology and states that given different possibilities of actions, people select the one that requires the least effort [16]. Consequently, we can assume that, upon interacting with each other, pedestrians try to avoid unnecessary detours and follow energy-efficient trajectories, that is paths that reduce the amount of movement and turning effort.

3 Pedestrian Simulation Model

In our problem setting, we are given a virtual environment in which n pedestrians $P_1, ...P_n$ have to navigate toward their specified goal positions $\mathbf{g}_i$ without colliding with the environment and with each other. For simplicity we assume that each pedestrian moves on a plane or a terrain and is modeled as a disc with radius r_i. At a fixed time t, the pedestrian P_i is at position $\mathbf{x}_i(t)$, defined by the center of the disc, has an orientation $\theta_i(t)$ and moves with velocity $\mathbf{v}_i(t)$. This velocity is limited by a maximum speed $u_i^{\max}$, that is $\|\mathbf{v}_i(t)\| \leq u_i^{\max}$. Hereafter, for notational convenience, we will not explicitly indicate the time dependence.

We propose a (social) force model to obtain realistic collision avoidance. The behaviour of each pedestrian P_i derives from the following assumptions:

1. At each step of the simulation the pedestrian P_i is trying to reach its desired goal position $\mathbf{g}_i$. Thus, a goal force $\mathbf{F}_g$ is applied that attracts the pedestrian to its goal.
2. The pedestrian P_i prefers to move toward its destination with a certain speed u_i^{pref}. It tends to reach this speed gradually within a certain time τ. Along with the first assumption, the goal force $\mathbf{F}_g$ can be defined as:

$$\mathbf{F}_g = \frac{1}{\tau}(u_i^{\text{pref}} \mathbf{n}_{gi} - \mathbf{v}), \tag{1}$$

   where $\mathbf{n}_{gi} = \frac{\mathbf{g}_i - \mathbf{x}_i}{\|\mathbf{g}_i - \mathbf{x}_i\|}$. Note that $\mathbf{F}_g$ is similar to the force described in Equation (2) of Helbing's social force model [7].
3. As P_i advances toward its goal, it also has to avoid collisions with the environment (e.g. building walls). Let $\mathcal{W}$ denote the set of walls that are present in the virtual environment. Then, a repulsive force $\mathbf{F}_w$ is exerted from each wall $w \in \mathcal{W}$ toward the pedestrian. This force can be defined as:

$$\mathbf{F}_w = \begin{cases} \mathbf{n}_w \dfrac{d_s + r_i - d_{iw}}{(d_{iw} - r)^{\kappa}}, & \text{if } d_w - r_i < d_s \\ \\ 0 & \text{otherwise,} \end{cases} \tag{2}$$

   where the constant κ indicates the steepness of the repulsive potential, $\mathbf{n}_w$ is the normal vector of the wall, d_{iw} defines the shortest distance between P_i and the wall and d_s denotes the safe distance that the pedestrian prefers to keep from the buildings.
4. The pedestrian P_i keeps a certain psychophysical distance ρ_i from other pedestrians in order to feel comfortable. This distance defines the so-called personal space of the pedestrian. For efficiency reasons, we model this space as a disc $B(\mathbf{x}_i, \rho_i)$ centered at the current position of the pedestrian and having radius $\rho_i > r_i$. In all of our simulations, this led to realistic behaviour and hence, the use of an elliptical personal space was not necessary.
5. The pedestrian P_i perceives the environment to detect imminent and future collisions with other pedestrians. A collision at some time $t_c \geq 0$ occurs when another pedestrian P_j invades into the personal space of P_i, that is:

$$\exists\, t_c \geq 0 \mid d_{ij} \leq \rho_i + r_j, \tag{3}$$

   where $d_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|$ denotes the distance between the pedestrians' centers.
6. Given a certain anticipation time t_α, the pedestrian P_i resolves potential collisions within this time in advance by adjusting its trajectory. As a result, it safely navigates toward its goal, evading other pedestrians with minimal effort. To simulate this behaviour, an evasive force $\mathbf{F}_e$ is applied on P_i. Note that the anticipation time t_α can vary between pedestrians.

4 Avoiding Collisions

The major novelty of our approach is the evasive force $\mathbf{F}_e$ that a pedestrian selects in order to avoid collisions and near collisions with other pedestrians. Given a pedestrian P_i, our collision avoidance algorithm consists of four steps.

4.1 Collision Prediction

In the first step of our algorithm we compute the set $CP_i^{t_\alpha}$ of pedestrians that are on collision course with the pedestrian P_i, given a certain anticipation time t_α. To achieve this, we first infer the desired velocity $\mathbf{v}_i^{\text{des}}$ of P_i. We compute $\mathbf{v}_i^{\text{des}}$ as the sum of its actual velocity and the velocity derived by virtually applying only the goal force and the wall repulsive forces:

$$\mathbf{v}_i^{\text{des}} = \mathbf{v}_i + \left(\sum \mathbf{F}_w + \mathbf{F}_g\right) \Delta t, \tag{4}$$

where Δt is the size of the time step of the simulation.

Then, we estimate the future position of P_i based on its current position $\mathbf{x}_i$ and desired velocity $\mathbf{v}_i^{\text{des}}$ as follows:

$$\mathbf{x}_i' = \mathbf{x}_i + t\mathbf{v}_i^{\text{des}} \tag{5}$$

Similarly, we predict the future motions of all the other pedestrians that P_i can see by linearly extrapolating their current trajectories. The viewing area of P_i is defined by its desired direction of motion and its field of view. Note that the pedestrian P_i can only estimate the actual velocities of the other pedestrians and not their desired ones, since it does not know their corresponding goal positions. Thus, from the perspective of P_i, the future position of a pedestrian P_j is given by:

$$\mathbf{x}_j' = \mathbf{x}_j + t\mathbf{v}_j \tag{6}$$

We can now determine whether the pedestrian P_i collides with another pedestrian P_j. According to (3) a collision occurs when the pedestrian P_j lies inside or intersects the personal space $B(\mathbf{x}_i, \rho_i)$ of P_i. By performing a Minkowski sum between the disc $B(\mathbf{x}_i, \rho_i)$ and the disc $B(\mathbf{x}_j, r_j)$ of P_j, the problem is reduced into a ray-disc intersection test that results in the following equation:

$$\|\mathbf{x}_j - (\mathbf{x}_i + \mathbf{v}t)\| = \rho_i + r_j, \tag{7}$$

where $\mathbf{v} = \mathbf{v}_i^{\text{des}} - \mathbf{v}_j$. Solving the above equation for t, we can estimate the possible collision times tc_{ij} between the personal space of the pedestrian P_i and the pedestrian P_j. If the equation has no solution or a single solution, then no collision takes place. If there are two solutions ($t1$ and $t2$), three cases can be distinguished:

- $t1, t2 \leq 0$: this is a past collision and can be ignored.
- $t1 < 0 < t2 \vee t2 < 0 < t1$: this is an imminent collision, i.e. $tc_{ij} = 0$, and hence, the pedestrian P_j is inserted into the set $CP_i^{t_\alpha}$.
- $t1, t2 \geq 0$: a collision will occur at time $tc_{ij} = min(t1, t2)$. If $tc_{ij} \leq t_\alpha$, the pedestrian P_j is inserted into the set $CP_i^{t_\alpha}$.

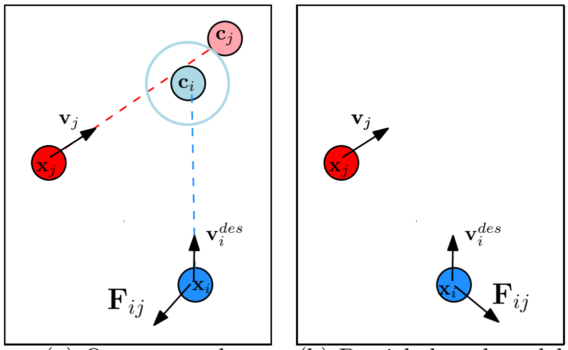

(a) Our approach (b) Particle-based model

Fig. 1. Example of the force used to avoid a future collision. Note that the direction of the force depends on the relative positions at the moment of the impact, rather than on the current positions of the pedestrians. The light blue disc represents the personal space of the pedestrian P_i.

4.2 Selecting Pedestrians

Having computed the set $CP_i^{t_\alpha}$, we sort it in order of increasing collision time and keep the first N pedestrians. Preliminary experiments (Section 6) have indicated that this number can be kept small (between 2 to 5 pedestrians), ensuring smooth avoidance behaviour. This not only reduces the running time of our algorithm, but also reflects natural human behaviour. In real-life, an individual tries to avoid a limited number of other pedestrians, usually those that are on collision course with him in the coming short time. Similarly, the virtual pedestrian P_i tries to evade the N pedestrians with which it will collide first.

4.3 Avoidance Maneuvers

We now show how the pedestrian P_i can avoid a potential collision with another pedestrian P_j that belongs to the set $CP_i^{t_\alpha}$. Let tc_{ij} be the time of collision between the two pedestrians. Let also $\mathbf{c}_i = \mathbf{x}_i + tc_{ij}\mathbf{v}_i^{\text{des}}$ and $\mathbf{c}_j = \mathbf{x}_j + tc_{ij}\mathbf{v}_j$ denote the locations of pedestrians P_i and P_j at time tc_{ij}; $\mathbf{c}_i$ and $\mathbf{c}_j$ derive from (5) and (6) respectively.

Based on these future locations, we select an evasive force $\mathbf{F}_{ij}$ for the pedestrian P_i, so that it can smoothly avoid the pedestrian P_j. The direction of the force is given by the unit vector $\mathbf{n}_{c_i c_j} = \frac{\mathbf{c}_i - \mathbf{c}_j}{\|\mathbf{c}_i - \mathbf{c}_j\|}$ pointing from $\mathbf{c}_j$ to $\mathbf{c}_i$. As an example, consider Fig. 1(a). The blue pedestrian will hit the red pedestrian at the back. Therefore, it makes a slight move toward the latter and eventually will pass behind it. In contrast, in a typical particle-based system the red pedestrian would have forced the blue one to move in the opposite direction, leading to a new collision in the near future (Fig. 1(b)).

The magnitude of the evasive force $\mathbf{F}_{ij}$ is approximated by a piecewise function f(D), where $D = \|\mathbf{c}_i - \mathbf{x}_i\| + (\|\mathbf{c}_i - \mathbf{c}_j\| - r_i - r_j)$ is the distance between the current

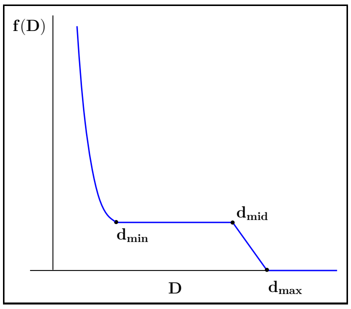

Fig. 2. The magnitude of the avoidance force as a function of the distance D

position $\mathbf{x}_i$ of pedestrian P_i and its future position $\mathbf{c}_i$ plus the distance between the two pedestrians at their time of collision. As can be inferred from Fig. 2, the function is defined for four intervals. The threshold $d_{\max}$ determines the start of the avoidance maneuver, whereas $d_{\min}$ defines the beginning of an impenetrable barrier between the pedestrians. The threshold d_{mid} regulates the start of the constant part of the function. This part is used to eliminate jerky behaviour. The three thresholds $d_{\min}$, d_{mid} and $d_{\max}$ can vary among the pedestrians to simulate different avoidance behaviours.

4.4 Computing the Evasive Force

In the last step of our algorithm, we compute the total evasive force $\mathbf{F}_e$ that is applied on the pedestrian P_i given its set $CP_i^{t_\alpha}$. Two approaches can be distinguished. The first is to take into account each one of the N pedestrians independently and compute a corresponding force $\mathbf{F}_{ij}$ as described above. Then, the evasive force can be defined as the weighted sum of those N forces:

$$\mathbf{F}_e = \sum_j^N w_{ij}\mathbf{F}_{ij}, \tag{8}$$

where the weighting factor w_{ij} gives a higher priority to the most imminent collisions.

The second approach, which differs from existing solutions, is to apply the forces sequentially (iterative approach). The idea here is as follows. Let P_j be a pedestrian in the set $CP_i^{t_\alpha}$ and let $\mathbf{F}_{ij}$ be the evasive force that is exerted on P_i. Let us also, for notational convenience, name this force $\mathbf{F}_\text{v}$. Then, we determine whether the pedestrian P_i still collides with each one of the remaining $N-2$ pedestrians in the set, after (virtually) applying the force $\mathbf{F}_\text{v}$. In other words, we first estimate the desired velocity of P_i by including $\mathbf{F}_\text{v}$ into the (4). Next, we iterate over the $N-2$ pedestrians in the set and compute the time of collision tc between P_i and each one of the selected $N-2$ pedestrians. If a collision still exists (i.e. $0 \leq tc \leq t_\alpha$), we compute the avoidance force as before and then add it to $\mathbf{F}_\text{v}$.

We repeat this process for each one of the N pedestrians in the set $CP_i^{t_\alpha}$. Then, the total evasive force $\mathbf{F}_e$ of pedestrian P_i is determined as the average of all the forces $\mathbf{F}_\text{v}$. Experiments have confirmed that by applying the forces sequentially a smoother and more realistic avoidance behaviour is achieved, as forces are only added if they are still required to avoid further collisions.

5 Implementation

This section provides implementation-specific details regarding our proposed pedestrian simulation model.

Efficient Collision Prediction. During each simulation step, we have to determine for each pedestrian P_i which other pedestrians are on collision course with him/her given a certain anticipation time. A naive implementation scheme would be to iterate over all other pedestrians checking whether a collision will take place. However, a more efficient implementation is to prune the search based on some cutoff distance, e.g. the maximum distance that the pedestrian P_i can travel given its anticipation time and maximum speed. Then, P_i only has to consider a limited number of pedestrians for potential collisions. Proximity computations to these pedestrians can be accelerated using a spatial hash data structure for answering nearest neighbor queries (see for example [5]).

Adding Variation. To increase the realism of the simulation some variation and irregularity is needed among the virtual pedestrians. Thus, a noise force $\mathbf{F}_n$ is also introduced in our model. This force, on one hand, allows us to take into account random variations in the individual behaviours, as well as incorporate mistakes that individuals make while avoiding each other. On the other hand, such a force is also needed to avoid artifacts that arise from symmetrical patterns, as well as to resolve extreme cases where two pedestrians have exactly opposite directions. More realism can also be achieved by varying several of the parameters of our model, such as the anticipation time of each pedestrian or the preferred speed.

Time Integration. The result from our proposed method is a system of positions, velocities and forces. To simulate this system, we first discretise our model in time by choosing a time step Δt for our simulation. Then, in each cycle of the simulation, we compute for each pedestrian P_i the sum of all the forces that are acted upon P_i and update its position and velocity using numerical integration. We also update the pedestrian's orientation θ_i by inferring it from the weighted average of its desired velocity and the current velocity.

6 Results

We have implemented our proposed method to test its applicability in real-time applications and validate the quality of the generated motions. All experiments

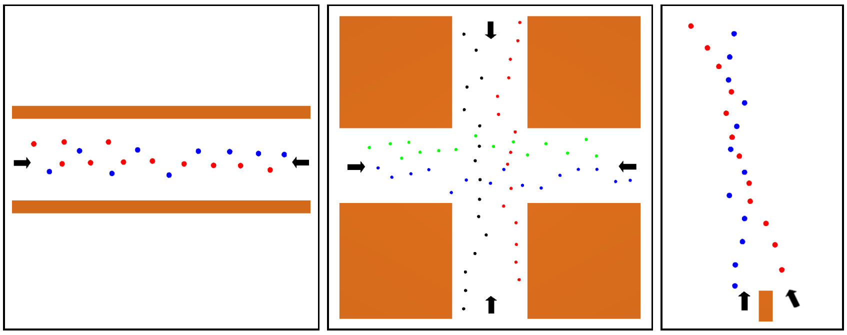

Fig. 3. Different simulation scenarios. Left: Hallway scenario. Middle: Four groups interacting at a crossing. Right: Crossing pedestrian flows.

were performed on a 2.5 GHz Intel Xeon CPU. Note that although this machine has four CPU cores, only one core was used for path planning.

We first calibrated the parameters of our model by performing a number of simple test-case scenarios as proposed in [17]. During these scenarios we recorded for each time step of the simulation, the position and the direction of each pedestrian. Then, a benchmark tool was implemented to analyse the quality of the generated motions and detect unrealistic behaviours. For that reason, a number of quantitative quality metrics have been devised. In particular, we used the integral of the square of the curvature to measure the smoothness of a pedestrian's path. We also computed the time it takes for a pedestrian to reach its goal, the length of the path it follows, as well as the average speed at which the pedestrian moves. Since we strive for energy-efficient motions, the total acceleration of the pedestrian and the total degrees it turned were also reported to indicate the amount of movement and turning effort spent by the pedestrian [17]. From the derived statistics, we determined optimal settings for our model. We experimentally confirmed that a smoother avoidance behaviour is achieved by applying the evasive forces iteratively, as explained in Section 4.4. The experiments showed that, using this approach, only a limited number of potential collisions N needs to be taken into account (2-5).

Having performed the initial calibration of our model, we ran a diverse set of simulation scenarios (Fig. 3). We refer the reader to `http://people.cs.uu.nl/ioannis/pam` for the results we obtained in simulating the aforementioned scenarios as well as for other simulations. In all of our experiments, we observed the phenomenon of dynamic lane formation that has been widely studied in pedestrian and crowd literature. In our model, though, the characters anticipate future collisions and adapt their motions in advance, evading others in a human-like manner. Even in crowded scenarios, our approach leads to a smooth and natural flow. Pedestrians scan the environment to find a free space to move into. Other individuals follow and a queue (lane) is formed that efficiently resolves the congestion.

We have also compared our approach to the Helbing's social force model [7] and the Reciprocal Velocity Obstacle (RVO) approach [11]. For a fair comparison,

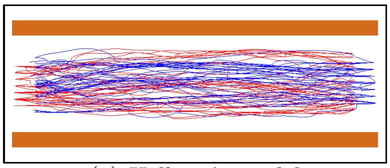
(a) Helbing's model

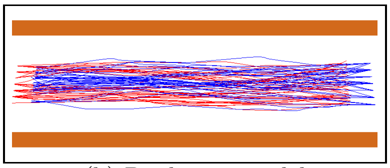
(b) Predictive model

Fig. 4. Example paths for the hallway scenario

we first optimised the two models following an approach similar to the one described above. In all test-cases, we noted that pedestrians using Helbing's model tend to adapt their motions rather late in order to avoid a potential collision. Thus, the resulting paths have a high curvature and the amount of unnecessary movements performed by the pedestrians is large. With the predictive approach, though, pedestrians plan early for collisions, avoiding unnecessary detours (see Fig. 4 for an example). In the RVO, each pedestrian adapts its velocity such that no collision will take place with other pedestrians and the environment. In low-density crowds, the method works very well. However, as soon as the environment becomes crowded, the RVO produces unrealistic motions. The pedestrians have to frequently speed up, slow down or change their orientations to safely reach their goals. On the contrary, our method ensures a smooth flow and more energy efficient motions.

We also quantitatively compared the models using our proposed evaluation metrics. Table 1 summarizes the corresponding statistics for the hallway scenario (note that similar results were obtained for all the other scenarios). A t-test was performed afterwards to determine how significant the results were. The analysis has shown that, indeed, our approach leads to shorter paths compared to the paths generated by Helbing's model, $p < 0.01$. The predictive approach ensures time-efficient motions allowing the pedestrians to move with significant higher speeds and follow much more smoother paths than pedestrians in Helbing's model (the differences in time, speed and smoothness are statistically significant, i.e. $p < 0.01$). In addition, using our approach, the number of interactions between the characters is reduced, which leads to more energy-efficient movements. Pedestrians have to spend less effort to avoid potential collisions and thus, the total acceleration is significantly lower compared to Helbing's model, $p < 0.01$. The same also applies to the turning effort of the pedestrians.

Statistical analysis has also confirmed our empirical observations regarding the RVO. Although no significant differences were found for the average travelling time and the speed of the pedestrians, our approach led to considerably less-curved paths than the RVO, $p < 0.01$. The analysis also indicated that the amount of the movement effort is lower using the predictive model compared to the amount spent by the pedestrians in the RVO, $p < 0.01$. Similarly, there is a significant difference in the amount of turning effort between pedestrians that employ our predictive avoidance method and pedestrians that use the RVO, $p < 0.01$.

Besides the quality of the generated motions, we are also interested in the performance of our proposed approach. To demonstrate its usability in real-time

Table 1. Statistics for the hallway scenario derived from the simulation of 100 agents

	Time		Path Length		Avg Speed		Smoothness		Total Accel		Degrees Turned	
	Mean	Stdev	Mean	Stdev	Mean	Stdev	Mean	Stdev	Mean	Stdev	Mean	Stdev
Helbing	30.33	3.71	37.12	2.34	1.23	0.09	1.08	1.08	55.73	28.37	323.22	225.79
Predictive	25.05	2.10	35.14	1.79	1.40	0.05	0.06	0.08	9.4	6.01	76.62	52.29
RVO	24.51	1.79	34.56	0.35	1.41	0.01	0.10	0.07	18.75	10.17	106.2	56.14

interactive applications, we simulated $1,000$ fully animated characters in a virtual park environment. Each character enters the park through a random gate and has to advance toward a randomly selected exit. To increase the realism of the simulation, we varied the preferred and maximum speeds among the characters, as well as the anticipation time of each character and the size of its personal space. The simulation was updated at 10 fps, whereas a parallel thread was used to render the characters and the scene at 30 fps. Our model led to smooth motions and thus, we did not have to decouple the simulation from the rendering. Since, on average, the CPU was busy for only 25% of each time step, we conclude that our model can be used for real-time simulation of thousands of characters. Finally, we also simulated the interactions of pedestrians at the crosswalks of an outdoor virtual city environment. Our approach resulted in a smooth avoidance behaviour and a visually pleasing simulation.

7 Conclusion and Future Work

In this paper, we presented a novel approach for simulating the collision avoidance behaviour of interacting virtual pedestrians. The intuition behind our approach is based on observed behaviour of real people. Each virtual character scans the environment to detect future collisions with other characters. It predicts how these collisions will take place and tries to resolve them in the current simulation step by making the most efficient move. In all of our experiments, the characters exhibit smooth behaviour and evade each other in a natural way. In addition, our proposed method is relatively easy to implement and yields to real-time performance.

For future work, we plan to automatically calibrate our model and validate our approach by exploiting existing video and motion capture data, as proposed in [12]. We would also like to combine our system with existing global path planning approaches. This will allow us to scale our method to very large and complicated environments. Furthermore, using these approaches, our framework can be easily extended to capture a wide variety of crowd characteristics. For example, we can include in our simulations coherent groups of characters, or characters that wander through a virtual environment without any specific goal.

Acknowledgments

This research has been supported by the GATE project, funded by the Netherlands Organization for Scientific Research (NWO) and the Netherlands ICT Research and Innovation Authority (ICT Regie).

References

1. Treuille, A., Cooper, S., Popović, Z.: Continuum crowds. ACM Transactions on Graphics 25(3), 1160–1168 (2006)
2. Lamarche, F., Donikian, S.: Crowd of virtual humans: a new approach for real time navigation in complex and structured environments. Computer Graphics Forum 23, 509–518 (2004)
3. Geraerts, R., Overmars, M.: The corridor map method: A general framework for real-time high-quality path planning. Computer Animation and Virtual Worlds 18, 107–119 (2007)
4. Sud, A., Gayle, R., Andersen, E., Guy, S., Lin, M., Manocha, D.: Real-time navigation of independent agents using adaptive roadmaps. In: ACM symposium on Virtual reality software and technology, pp. 99–106 (2007)
5. Shao, W., Terzopoulos, D.: Autonomous pedestrians. Graphical Models 69(5-6), 246–274 (2007)
6. Reynolds, C.W.: Steering behaviors for autonomous characters. In: The proceedings of Game Developers Conference, pp. 763–782 (1999)
7. Helbing, D., Molnár, P.: Social force model for pedestrian dynamics. Physical Review E 51, 4282–4286 (1995)
8. Helbing, D., Buzna, L., Johansson, A., Werner, T.: Self-organized pedestrian crowd dynamics: Experiments, simulations, and design solutions. Transportation Science 39(1) (2005)
9. Pelechano, N., Allbeck, J.M., Badler, N.I.: Controlling individual agents in high-density crowd simulation. In: SCA 2007: ACM SIGGRAPH/Eurographics Symposium on Computer Animation, pp. 99–108 (2007)
10. Paris, S., Pettré, J., Donikian, S.: Pedestrian reactive navigation for crowd simulation: a predictive approach. Computer Graphics Forum 26(3), 665–674 (2007)
11. van den Berg, J.P., Lin, M., Manocha, D.: Reciprocal velocity obstacles for real-time multi-agent navigation. In: ICRA, pp. 1928–1935. IEEE, Los Alamitos (2008)
12. Pettré, J., Ondrej, J., Olivier, A.H., Crétual, A., Donikian, S.: Experiment-based modeling, simulation and validation of interactions between virtual walkers. In: SCA 2009: ACM SIGGRAPH/Eurographics Symposium on Computer Animation, pp. 189–198 (2009)
13. Lerner, A., Chrysanthou, Y., Lischinski, D.: Crowds by example. Computer Graphics Forum 26, 655–664 (2007)
14. Lee, K., Choi, M., Hong, Q., Lee, J.: Group behavior from video: a data-driven approach to crowd simulation. In: SCA 2007: ACM SIGGRAPH/Eurographics Symposium on Computer Animation, pp. 109–118 (2007)
15. Goffman, E.: Relations in public: microstudies of the public order. Basic books, New York (1971)
16. Zipf, G.K.: Human Behavior and the Principle of Least Effort. Addison-Wesley, Reading (1949)
17. Singh, S., Kapadia, M., Faloutsos, P., Reinman, G.: Steerbench: a benchmark suite for evaluating steering behaviors. Computer Animation and Virtual Worlds (2009)