

Torso Crowds

Sybren A. Stüvel, Nadia Magnenat-Thalmann, Daniel Thalmann,
A. Frank van der Stappen, and Arjan Egges

Abstract—We present a novel dense crowd simulation method. In real crowds of high density, people manoeuvring the crowd need to twist their torso to pass between others. Our proposed method does not use the traditional disc-shaped agent, but instead employs capsule-shaped agents, which enables us to plan such torso orientations. Contrary to other crowd simulation systems, which often focus on the movement of the entire crowd, our method distinguishes between active agents that try to manoeuvre through the crowd, and passive agents that have no incentive to move. We introduce the concept of a *focus point* to influence crowd agent orientation. Recorded data from real human crowds are used for validation, which shows that our proposed model produces equivalent paths for 85 percent of the validation set. Furthermore, we present a character animation technique that uses the results from our crowd model to generate torso-twisting and side-stepping characters.

Index Terms—Crowd simulation, crowd animation, dense crowds, agent representation, holonomic motion

1 INTRODUCTION

THE shapes most often used to represent characters in crowd simulations are points and discs. In sparse crowd simulations, such a simple shape works well; the chosen representation does not have a large impact on the behaviour, as there is ample space around the agents. However, when the agents move very close to each other, motion follows shape. A common motion in dense crowds is the twisting of the torso, to squeeze through an opening between people. This is commonly seen at band performances, busy cocktail parties, or crammed lifts. Points and discs are ill suited for such situations, as the rotational symmetry prohibits planning of such twist. Instead, in this paper, we investigate an agent representation based on the torso. By employing a representation that is closer to the human shape, we expect to obtain more realistic human-like motions than disc-based crowd simulation methods.

Main contribution: In this article we introduce the Torso Crowd model for dense crowd manoeuvring, based on a novel capsule-shaped agent representation modelling the characters' torsos (see Fig. 1). Contrary to other crowd simulation systems, which often focus on the movement of the entire crowd, our method distinguishes between passive agents that have no incentive to move from their present location, and active agents that try to manoeuvre through the crowd towards a goal position. We introduce the concept of a *focus point* for crowd agents, which allows for more control and more realistic, social and complex behaviour. Furthermore,

we validate the active agent behaviour using ground truth data, obtained by motion capturing a real crowd.

We use the term *agent* to indicate an abstract crowd agent such as a point, disc, or capsule. The term *character* designates a humanoid virtual character, whereas the term *people* refers to real humans. We mostly consider the motions of the upper body, i.e., the torso. The lower body is only considered at the final visualization step, where humanoid body animation is generated. Torso *twist* is not defined as a rotation relative to the lower body, but as a rotation relative to the agent's trajectory.

Organization: The rest of the paper is organized as follows. Section 2 discusses related work. The overall setting of our crowd model is described in Section 3, followed by a description of the design of active (Section 4) and passive (Section 5) agents. We show our results, compare with a cylinder-based simulation method and with ground truth obtained from motion capture, and describe several simulated scenarios in Section 7. Section 8 describes the animation technique used to display the moving crowd agents as walking humanoid figures. Section 9 discusses future work, and concludes the article. The accompanying video can be found online at <http://stuel.eu/video/torso-crowds>.

2 RELATED WORK

For a general overview of crowd simulation techniques and topics, we refer to the books by Thalmann and Musse [1], and Pelechano et al. [2]. In the remainder of this section, we focus on work that is related to the simulation of dense crowds.

There are many approaches to simulating crowds, each leading to different behaviour. *Flow-based* methods are macroscopic, focusing on the crowd as a whole. They support very large crowds, as high-level coordination prevents obstructions. Common examples are fluid dynamics [3] or gas kinetics [4], which can be applied to particle-based crowds, and are particularly suitable for high-density simulations. However, such approaches model a global optimum, whereas humans generally behave less optimally and

- S.A. Stüvel, A.F. van der Stappen, and A. Egges are with the Virtual Human Technology lab, Utrecht University, Utrecht, The Netherlands. E-mail: {s.a.stuel, a.f.vanderstappen, j.egges}@uu.nl.
- N. Magnenat-Thalmann and D. Thalmann are with the Institute for Media Innovation, Nanyang Technological University, Singapore. E-mail: {nadiathalmann, danielthalmann}@ntu.edu.sg.

Manuscript received 24 Nov. 2015; revised 10 Mar. 2016; accepted 17 Mar. 2016. Date of publication 23 Mar. 2016; date of current version 26 May 2017. Recommended for acceptance by J. Lee.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below. Digital Object Identifier no. 10.1109/TVCG.2016.2545670




Fig. 1. Our crowd simulation model, showing different stages of the simulation. From left to right: the agent representation, calculation of the Voronoi diagram, planning a path towards a goal position, and finally the animation of virtual characters.

can even get stuck in very dense situations. *Cellular automaton* approaches such as the works by Chenney [5] and Alizadeh [6] discretize floor space into cells, where every character can occupy exactly one cell, and vice versa. Although such approaches are computationally inexpensive and thus support large crowds, they also result in even spacing between agents, which will appear unnatural when densities are high. Both flow-based and cellular automaton systems do not consider how people move, and thus generally do not result in believable human crowds. *Agent-based* methods employing social forces, such as the HiDAC model [7], or planning in velocity space, such as the RVO model [8] generally allow for high-density crowds while supporting individual behaviour of agents, and can be extended to support physical interaction with obstacles and the environment [9]. These approaches avoid agent collisions at all costs, even to the point where all agents stop moving. In contrast, in actual dense crowds people frequently bump into each other. This is reflected in our method, as our agents prioritize motion over collision avoidance. Our model switches agents to the RVO steering method when they are in low-density areas; such use of multiple models of agent representation for crowd navigation is also performed by Kapadia et al. [10].

Common representations used in path planning and crowd simulation are points [3], [4] and discs [8], [11]. Discs are the most commonly used representation due to their computational simplicity, and have proven to be suitable to simulate abstract (i.e., non-humanoid) crowds of any density. However, when using such a simple, rotationally symmetric representation, it becomes hard to animate more detailed human motion. This results in artefacts such as interpenetration of characters, unnatural distances between characters, and a lack of torso rotations. Furthermore, it has been shown that a disc does not accurately represent the actual volume occupied by the character in 3D space [12]. In Section 7 we show the importance of the agent shape in dense crowds, not just for realism of the motions, but also to support higher densities without getting stuck. Singh et al. use multiple discs [13] to represent an agent, and plan their motion using a footstep model. This approach allows for denser crowds than body-enclosing discs, and offers realistic walking animations. Their footstep model was made more robust by Berseht et al. [14], who also show the simulation of a densely packed crowd of agents. However, due to the lack of animated humanoid characters, it is hard to draw conclusions as to the realism of the result. The works

by Van Basten et al. [15] and Beacco et al. [16] can be used to animate such humanoid characters.

Finally, some crowd animation methods directly use animation data to drive the characters [17], [18], also producing realistically animated characters. However, due to the high-level planning of these methods, planning the motion of individuals, such as specific characters moving towards their respective goal positions, is much harder to do. The computer model SIMULEX [19] uses three circles to approximate the top view human shape: a larger one for the head, and two smaller ones for the shoulders. SIMULEX does not consider torso orientation when planning; it just uses the three circles to more accurately perform collision detection between pedestrians. Rotating polygonal shapes are supported by the Reciprocally-rotating Velocity Obstacles method by Giese et al. [20]. Even though the rotations are considered when planning motions, the planner does not aim at the simulation of humanoid behaviour. As a result, the agents translate at an angle in cases where a human would simply walk straight. To our knowledge, our method is the first to present the capsule as agent representation in crowd simulations.

Our method employs Voronoi diagrams for planning motions through the crowd. The edges of such a diagram represent the path of maximum clearance between agents; intuitively this corresponds well with the desire of people to minimize perceived effort when walking [21], [22]. Stüvel et al. [23] showed that in a dense crowd people indeed move along such paths. The Explicit Corridor Map method by Geraerts [24] uses city-scale generalized Voronoi diagrams for path planning. Sud et al. [25] also employ generalized Voronoi diagrams for path planning; contrary to our approach, they limit the use of this diagram to static obstacles only, and do not apply it for agent avoidance. Sud et al. [26] perform path planning based on 1st and 2nd order Voronoi diagrams, containing information about respectively the closest agent and the closest pair of agents. They employ a path scoring technique slightly resembling our proposed method. While their article promotes speed of computation, our approach focuses on a richer character representation, and validation against real crowd data.

3 SETTING AND PROBLEM FORMULATION

Our crowd simulation system considers the torso as the main moving element. The algorithm is based on the findings by Stüvel et al. [23], [27], who observed and recorded




Fig. 2. Two types of crowd agent representation. On the left a common crowd agent: a point with a radius. On the right our crowd agent: a line segment with a radius.

dense crowd behaviour. In their experiment, participants were given the task of manoeuvring through the crowd to predefined points. The movement of the crowd was recorded using a motion capture system, and these data serve as a ground truth for the behaviour of people actively manoeuvring through dense crowds. We refer to Stüvel et al. [23] for the details of the experiment. Our Torso Crowd model is designed to support the observed motions of the crowd-escaping participants, and believable simulation of essentially stationary people.

When observing dense crowds in general, and the previously mentioned recordings in specific, it is clear that torso rotations are critical when manoeuvring through a dense crowd. To support such rotations, our agents extend the common disc-based crowd agents, as shown in Fig. 2. The common agent is defined as a point with a radius r' ; our agent model extends this point to a line segment of length ℓ , with a (probably different) radius r . This extension eliminates the rotational symmetry, thereby making it possible to plan torso rotations.

People standing still in a crowd behave differently from people trying to reach a certain goal position. In order to model these differences in behaviour, two types of agents are used in our crowd simulation technique. *Active agents* move to reach their goal position, whereas *passive agents* mostly stay in place, only moving to make room for other agents. Section 4 describes the behaviour of the active agents, while Section 5 describes the passive agents. Different aspects of the *environment*, such as walls, doors, and other obstacles, are handled by our method in a unified way; this is described in Section 6. The goals of the agents are determined by a high-level planner, which is scenario-dependent and not described further in this article. When an active agent reaches its goal, depending on the intended scenario, the agent optionally switches to passive behaviour. Similarly, when the high-level planner provides a passive agent with a new goal position, that agent will switch to active behaviour.

The crowd consists of N agents A_i , $i \in [1, \dots, N]$. Each agent A_i in reference placement is defined as the Minkowski sum of a line segment of length ℓ_i centred around the origin, and a disc of radius r_i . The placement of an agent is represented by a pair (\mathbf{a}_i, θ_i) , where \mathbf{a}_i and θ_i are the agent's position and orientation. For ease of discussion, we denote the direction of the forward-facing normal of the torso of agent A_i in placement (\mathbf{a}_i, θ_i) by \mathbf{n}_i , the continuous set of points covered by its central axis by \mathbf{s}_i , and its linear velocity vector as $\dot{\mathbf{a}}_i$. These concepts will be detailed in the following sections.




Fig. 3. Top-down view of a real, motion captured crowd, with the generalized Voronoi diagram in white lines.

4 ACTIVE AGENTS

From observation of the previously mentioned ground truth data and dense crowds in general, we formulate the following assumptions as basis for our active agent model.

- People tend to choose a comfortable path, that is, maximize clearance, by avoiding areas of very high density. Occasionally, a less comfortable path may be chosen, when the discomfort is only temporary and the path leads to an area of larger clearance.
- People tend to minimize perceived energy use [21], and thus prefer short, straight paths.
- People generally move in the direction of their goal, but divert from the shortest path when it is obstructed or when an alternative path is significantly more comfortable.
- Averaging at 0.4 m/sec, the traversing speed through a dense crowd is relatively low [27]. This, combined with the dynamic nature of crowds and possibly a lack of overview of the situation, makes long-distance planning of exact paths to the goal impractical.

The generalized Voronoi diagram (GVD) is a partitioning of the plane. In our crowd model, a cell is defined for each agent, being the set of all points that is closest to that agent. The GVD is represented as a pair $\{V, E\}$ of vertices V and edges $E \subset V \times V$ that represent the boundaries of those cells, where the edges are arcs (possibly with zero curvature, i.e., line segments). Every point on an edge or a vertex is equidistant to its neighbouring crowd agents. These edges form the medial axis between the agents, and thus represent a more or less comfortable path that maximizes local clearance. Stüvel et al. [23] have observed that people in dense crowds indeed follow such paths (also see Fig. 3).

In real situations, observing surrounding people, planning a path, and manoeuvring through the crowd, are intertwined in a continuous process. However, people are not continually reconsidering all their options all the time, but rather make more or less discrete decisions. Our agents reflect this behaviour by replanning their actions at a moderate rate.

Similar to Sud et al. [28], all active agents use the same GVD for planning their motion. The passive agents use a




Fig. 4. The exact Generalized Voronoi Diagram of the capsules (black) and the approximated diagram of their central axes (white). The distance between corresponding edges depend on the difference in capsule radius.

slightly altered GVD, as described in Section 5. As a result, the GVD needs to be calculated at most twice per simulation update, regardless of the crowd size and frequency of planning. For simplicity of computation, we use the central axes s_i to calculate the GVD, rather than the agent shapes themselves. Due to the nature of the capsule, where the distance from the axis to the edge of the shape is constant, the approximated GVD is very similar to the exact GVD (see Fig. 4); the distances between the corresponding edges are in the order of magnitude of the differences between the agent radii. Such small differences, in our case in the order of centimetres [23], are unlikely to cause noticeable changes in the crowd's behaviour.

Our Torso Crowd method is suitable for simulating dense crowd manoeuvring, and based on experimental observations of such behaviour. However, in situations where the crowd is not dense, we have no proof of validity. As a consequence, our implementation switches to the RVO2 crowd simulation algorithm [8] when agents move out of the dense crowd, using a circular agent shape that is large enough to encompass the agent and its personal space. This can be seen in the accompanying video, when agents walk out of a lift and into an empty hallway. We reuse the definition of *dense* situations by Stüvel et al. [23], namely those situations where there is at least three humans per square metre, as measured by the area of their Voronoi cell. Since we can measure this density on a per-agent basis, this decision is also made for each agent individually.

In the next sections, we discuss the planning and execution of the agent's movement. Firstly, similar to real people, a desired position is planned, taking into account potential torso twists needed to reach that position. Since the available clearance at the planned position poses a bound on the torso orientation, this orientation is planned in a second step.

4.1 Limited-Horizon Path Planning

To plan the movement of an active agent, the following steps are taken:

- 1) Find paths by exploring the vicinity in the GVD of the Voronoi cell containing the agent.
- 2) Calculate a score for each path, and determine the best-scoring path.
- 3) Calculate the desired agent orientation at the start of the path, accounting for available clearance.




Fig. 5. Candidate paths in white lines, with the best-scoring path as a thick, red line. The circled edge was shorter than H_ϵ . The paths are extended to the agent position. The goal position is bottom left outside the frame. Note that the paths along curved Voronoi edges are just drawn as straight line segments for simplicity.

The GVD provides proximity information in a natural way; the cell of the active agent represents its proximity, and the outgoing edges of that cell's vertices form paths between agents in its direct vicinity. The search is initialized by taking these outgoing edges, i.e., the edges that only have a single vertex incident to the active agent's Voronoi cell. This set is then extended, parameterized by given values for Euclidean distances H_D and H_ϵ , and edge count limit H_C , as follows: the outgoing edges are followed depth-first until either distance H_D , or edge count limit H_C is reached. For the latter limit, edges shorter than H_ϵ are ignored (as circled in Fig. 5). Such short edges occur, for example, when four agents are almost equidistant, and the clearance between the agents would likely be perceived as a single space. Hence, such edges are unlikely to correspond to human perception. Even though this approach could theoretically lead to a path consisting of an arbitrarily large number of edges, such a situation does not occur in dense crowds when using crowd agents of more or less realistic human-like sizes. The resulting path P consists of a sequence of GVD edges $e \in P$; following the path should bring the agent closer to its goal.

After a set of candidate paths is found, each path is given a score. The agent will attempt to use the path with the highest score. The composite score function $S(i, P)$ takes agent A_i and path P . It enforces the behaviour of real people in dense crowds, based on the observations by Stüvel et al. [23]. As all score functions should be balanced to make a final decision as to the best possible path, they are combined into a weighted sum:

$$S(i, P) = w_g S_g(i, P) + w_c S_c(P) + w_l S_l(P) + w_m S_m(i, P),$$

where $S_g(i, P)$, $S_c(P)$, $S_l(P)$ and $S_m(i, P)$ are score functions, and w_g , w_c , w_l and w_m are weights given to these sub-scores. Values for these weights are determined in Section 7.1. In the following descriptions of each score function, \mathbf{p}_0 and \mathbf{p}_f respectively indicate the initial and final

vertex positions of path P . Note that the path's final point \mathbf{p}_f does not necessarily correspond to the agent's goal position \mathbf{g}_i , due to the limit on the path length described earlier.

Score function $S_g(i, P)$ drives the agent towards its goal. It measures how well the path leads to the goal position \mathbf{g}_i , expressing the distance, from the end of the path to the goal, as a ratio of the total Euclidean distance to the goal. This normalization ensures that the resulting score is independent of the absolute distance to the goal:

$$S_g(i, P) = 1 - \frac{|\mathbf{g}_i - \mathbf{p}_f|}{|\mathbf{g}_i - \mathbf{a}_i|}.$$

Score function $S_c(P)$ measures the clearance radius along the path, ensuring that the agent prefers comfortable routes with large clearances. It consists of two components. The first component stems from the moderate rate replanning principle. It assumes that a person plans a motion towards a more spacious area; when this area is reached, a new decision can be made. The second component prefers motion along paths with as much clearance as possible. The GVD structure enables efficient calculation of clearance radius $C(\mathbf{x})$ at any point $\mathbf{x} \in \mathbb{R}^2$

$$S_c(P) = w_c^F C(\mathbf{p}_f) + w_c^A \frac{1}{|P|} \int_{\mathbf{x} \in P} C(\mathbf{x}) d\mathbf{x},$$

where $|P|$ indicates the total arc length of P , $\mathbf{x} \in P$ are the collection of all points along path P , and w_c^F and w_c^A are weights for respectively the final and the average clearance of the path. Due to implementation details of our GVD library, we only had access to the minimal clearance of edges, and approximated the integral using discretized summation. This score function also serves as a term to minimize the relative rotation of the torso with respect to the motion trajectory, due to the way the clearance information is used to plan torso orientations (see Section 4.2)

The conservation of energy can be broken down into two components: the minimization of the distance travelled, and the effort required to travel that distance. Score function $S_l(P)$ models the first component, and measures path length. This function combines with $S_g(i, P)$ into the preference of short paths leading to the goal

$$S_l(P) = - \sum_{e \in P} |e|.$$

Score function $S_m(i, P)$ represents the second component of energy conservation, by penalizing changes in momentum, i.e., sharp turns. Since we can safely assume that the mass of the agent is constant, any change in momentum is explained by a change (in the direction of) the velocity vector, which in turn can be modelled by the cosine similarity of the current velocity and the direction towards the path:

$$S_m(i, P) = \frac{\dot{\mathbf{a}}_i \cdot \mathbf{e}_0}{|\dot{\mathbf{a}}_i| |\mathbf{e}_0|},$$

where $\mathbf{e}_0 = \mathbf{p}_0 - \mathbf{a}_i$, the vector connecting the agent to the starting point of the path.

A more elaborate alternative for $S_m(i, P)$ could calculate the weighted integral of the curvature along \mathbf{e}_0 and P , with the weight inversely proportional to the distance from the

agent. This would take the curvature of the entire path into account, emphasizing more immediate momentum changes. However, our proposed approach is simpler, and seems to be sufficient in practice. Furthermore, due to the agent replanning while it is en route to its goal, effectively the curvature of the entire path is taken into account.

4.2 Torso Rotation Planning

Once the best path P has been chosen, which determines the next torso position, the torso orientation T_o is determined. For this we use the torso normal \mathbf{n}_i of agent A_i . Torso orientation T_o consists of two components: heading T_h and torso twist T_t . The first component, T_h , represents a common nonholonomic walking motion along the start of the path. Its computation is trivial and not described here. The second component, torso twist T_t , adjusts for the minimal clearance along the start of the path. The clearance at later parts of the path is of less importance for the current torso twist planning, due to the moderate-rate replanning principle. The first edge of the path lies between two neighbouring agents, and ends at a point of local maximum clearance behind those two agents. It is this part of the path that is used for the planning of the torso twist. The clearance at a point indicates the distance from that point to the nearest agent capsule. To maximize the time for the agent to smoothly change its torso orientation towards the desired twist, we calculate the minimal clearance c along the first edge of path P . The torso twist T_t can then be expressed in radians as

$$T_t = \begin{cases} 0 & : c - r_i \geq w_i \\ \cos^{-1} \left(\frac{c - r_i}{w_i} \right) & : 0 < c - r_i < w_i \\ \frac{\pi}{2} & : c - r_i \leq 0, \end{cases}$$

where $w_i = \ell_i/2 + r_i$ is the half-width of the capsule. This results in two possible orientations, both of which will fit the available clearance equally well: $T_h + T_t$ and $T_h - T_t$. The choice for the final orientation is based on the findings by Stüvel et al. [27]. They observe that, while manoeuvring through a dense crowd, people tend to aim their torso normal towards their goal position. The absolute angle between the torso normal and the vector to their goal is limited to 90° for 90 percent of the time, and never more than 120° . Consequently, we choose $T_o = T_h \pm T_t$ such that this angle is minimized. Once the desired position \mathbf{p}_0 and orientation T_o have been calculated, each agent employs proportional-derivative controllers to steer towards the planned configuration.

So far we have discussed the general approach for an active agent. Based on observations from real crowds, we deviate from this approach when a character starts to move towards a goal. Stüvel et al. [23] observed that, before they start manoeuvring, people orient their torso towards their goal. Similar behaviour is incorporated into our crowd model. When an agent becomes active and starts planning its movements, it performs the same planning steps as described in the previous sections. However, it discards the planned position \mathbf{p}_0 , and rotates on the spot towards the planned orientation T_o . Subsequent planning steps are performed as described earlier.

5 PASSIVE AGENTS

In this section, we discuss the behaviour of *passive* crowd agents, which, in contrast to active agents, do not have an explicit target to navigate to. We consider two sometimes contradictory motivations for their placement: finding local comfort, and rotation towards a focus point. Our passive agents locally optimize their placement, making themselves as comfortable as possible, i.e., maximize the clearance around them, given the constraints of their immediate surroundings. The translation \mathbf{t}_S to reach a more comfortable placement is described in Section 5.1. When the geometry of the environment and the configuration of the crowd allow for it, a trade-off is made between rotating to a comfortable orientation and a rotation towards a focus point. This can be the centre of a chatting group of people, the charismatic front man of a performing band, or simply the floor number display of the lift. To our knowledge, we are the first to use such a focus point in a crowd simulation system. The rotation ϕ_S from the current to the desired orientation is described in Section 5.2. Passive members of a crowd temporarily accept a less comfortable position in order to make way for someone else to pass; this avoidance by translating (\mathbf{t}_A) and rotating (ϕ_A) is described in Section 5.3. In Section 5.4 we show how these desires are combined into the agent's motion.

5.1 Space Finding

Passive agents try to coarsely maintain their position. For example, even when a lift is crowded, the door is open, and outside the lift is a plethora of space, agents waiting in the lift will remain in that lift. Manoeuvring to a different area, such as stepping out of the lift, is considered active behaviour, and is described in the previous section. We use walls and doors (see Section 6) to delineate areas in the environment. To restrict the space finding algorithm to the agents' current area, our passive agents consider all doors as closed, regardless of their actual state. However, the agents do search for a better place to stand in their direct vicinity; this is what we call *space finding* behaviour. This results in a translation vector \mathbf{t}_S from their current position to a more spacious position. Effectively it is a combination of comfort optimization and avoidance of passive agents.

Whether the space finding algorithm is engaged depends on the space around the agents. We assume that our passive agents like to stand in a spot where there is enough space surrounding them. When that is the case, i.e., the distance to the nearest neighbouring agent or obstacle is larger than a certain threshold, they remain stationary, even though there may be even more space available to them; the agent is marked as *happy* with its current placement, and will not engage the space finding algorithm (so $\mathbf{t}_S = \mathbf{0}$). This threshold can be configured individually for each agent, and can be a function of culture, scenario, or the geometry of the surroundings.

In tighter situations, our passive agents move to maximize their comfort. To obtain nearby candidate positions of maximal comfort, agents consider points of maximum clearance between their surrounding neighbours. By definition, such points correspond with vertices of a Generalized Voronoi Diagram (GVD, see Section 4) of those neighbours. Such a *local Generalized Voronoi Diagram* \mathcal{L}_i of agent A_i is the GVD




Fig. 6. Example of a local Generalized Voronoi Diagram (GVD), with points of maximal local clearance, in orange. The GVD of the crowd is shown in white. The features that define the local GVD are shown in magenta. The dashed circle shows the clearance of the agent.

defined by \mathcal{N}_i , where \mathcal{N}_i is the set of neighbouring agents and obstacles of agent A_i . \mathcal{N}_i can be efficiently extracted from the GVD of the entire crowd, by iterating over the edges of the cell containing A_i , and taking the agents or obstacles on the opposite side of the edges. Note that agent A_i itself is *not* included in \mathcal{L}_i (see Fig. 6). The vertices of \mathcal{L}_i correspond to local clearance maxima, and thus potentially comfortable positions for the agent to move to.

People try not to spend too much energy [21], and will accept a marginally more cramped situation when walking to a better spot would take a significant effort. We use the following energy minimization function to balance the gain (more available space) with the expended energy (the distance to travel to that space). All vertices $\mathbf{v}_j \in \mathcal{L}_i$ are considered potential better positions, and are given an energy cost

$$E(\mathbf{a}_i, \mathbf{v}_j) = \frac{|\mathbf{v}_j - \mathbf{a}_i|}{C(\mathbf{v}_j) - C(\mathbf{a}_i)},$$

$$\mathbf{v}_d = \operatorname{argmin}_{\mathbf{v}_j \in \mathcal{L}_i} E(\mathbf{a}_i, \mathbf{v}_j),$$

$$\mathbf{t}_S = \mathbf{v}_d - \mathbf{a}_i,$$

where $C(\mathbf{x})$ indicates the clearance around \mathbf{x} ; \mathbf{v}_d denotes the vertex with the lowest energy cost, and determines the agent's space finding translation vector \mathbf{t}_S . This scoring is efficient; we have found that, in practice, 89 percent of the time \mathcal{L}_i contains no more than three vertices, with an average of 2.7 vertices.

5.2 Orientation Finding

When agents are squeezed into a small area, they rotate themselves to fit the available space. However, if the constraints allow for it, the agents focus on a given point (a performing band on a stage, floor number display of a lift, etc.). This results in a rotation ϕ_S from the current orientation of the agent towards a desired orientation. The *focus point* is environment- and scenario-dependent, and can of course change over time and be different for each person or agent. It is denoted as \mathbf{f}_i for agent A_i . The accompanying video shows the effect of this focus point. A group of agents have a focus point in the centre, and the video demonstrates the effect of increased density on this group: the group stays together, even though the focus point has no direct influence on the position of the agents (see Fig. 10).

In the remainder of this section, p is the passive agent's index number, so \mathbf{a}_p indicates its position. The angle between the agent's torso normal \mathbf{n}_p (see Section 3) and the vector to its focus point \mathbf{f}_p is defined as

$$\alpha_f = \angle(\mathbf{f}_p - \mathbf{a}_p, \mathbf{n}_p),$$

where $\angle(\mathbf{x}, \mathbf{y})$ indicates the signed angle between two vectors on the interval $(-\pi, \pi]$.

When marked as *happy* with their current placement (which depends on the available clearance, as described in Section 5.1), our crowd agents rotate such that their torso normal points towards their focus point. In this case, we take $\phi_S = \alpha_f$.

The shape of the available space is the dominant factor in someone's orientation when that space is tight; one rotates to fit the little space available. The narrower the space, the less important any focus point becomes. To include this behaviour in our model, we inspect the shape of the agent's Voronoi cell. Since this cell contains all points that are closer to the agent than to any other agent, it is a good model for their available space. The *width* of the cell is defined as the minimal distance between two parallel lines that enclose the cell. The direction of these lines are a common measure for the oblong direction of the cell. However, this direction is not stable under small variations in agent configurations, so we use a more elaborate approach. To obtain a vector that indicates the overall orientation of the space, a Principal Component Analysis (PCA) [29] is applied. Such an analysis is applied to a point cloud, to determine its dominant direction. Since it cannot be applied to continuous shapes, intuitively we could sample the interior of the Voronoi cell to obtain such a point cloud. However, to increase computational performance, we limit this approach to the sampled cell edges; considering the results this is sufficient. The result of the PCA consists the eigenvectors and eigenvalues of a covariance matrix; when ordered from large to small by their absolute eigenvalues c_1 and c_2 , the eigenvectors indicate the first and second principal components \mathbf{C}_1 and \mathbf{C}_2 . In the remainder of this section, we assume that the absolute eigenvalues are ordered by magnitude, i.e., c_1 belongs to eigenvector \mathbf{C}_1 .

When the Voronoi cell of a passive agent has no clear orientation, the eigenvectors hold little information, and the eigenvalues will be more or less equal. In this case, the agent rotates towards the focus point. When the shape of the cell is elongated, and thus relevant for the orientation of the crowd agent, the first principal component aligns with the cell's shape. This relevance is indicated by a large difference between the first and second eigenvalue of the covariance matrix, i.e., $c_1 - c_2 \geq \epsilon_2$ (ϵ_1 will be introduced later as a lower bound). In this case, there are two possible orientations for the agent, in which the agent's central axis \mathbf{s}_p aligns with either \mathbf{C}_1 or $-\mathbf{C}_1$; the agent chooses the orientation that minimizes α_f . If there is no focus point, α_f is not defined, and the agent chooses the orientation that requires the smallest rotation from its current orientation

$$\alpha_c = \angle(\pm \mathbf{C}_1, \mathbf{s}_p).$$

To ensure smooth transition between α_c and α_f , we blend between them depending on the eigenvalue difference:

$$\phi_S = \begin{cases} \alpha_c & \text{if } \epsilon_2 \leq c_1 - c_2 \\ I(\alpha_c, \alpha_f, t) & \text{if } \epsilon_1 \leq c_1 - c_2 < \epsilon_2 \\ \alpha_f & \text{if } c_1 - c_2 < \epsilon_1, \end{cases}$$

where $\epsilon_1 < \epsilon_2$, $I(\alpha_c, \alpha_f, t)$ indicates angular linear interpolation along the shortest arc for $t = (c_1 - \epsilon_1)/(\epsilon_1 - \epsilon_2)$. In our implementation, we use $\epsilon_1 = 0.015$ and $\epsilon_2 = 0.045$.

5.3 Avoidance of Active Agents

The behaviour of passive and active agents is quite different. Passive agents move slower, and try to divide the available space between them. Active agents move faster (when allowed by the constrained environment), and, more importantly, try to reach a specific goal. These differences are also reflected in the way that passive agents perform agent avoidance. This section describes how they avoid *active* agents.¹

Since far away agents have negligible chance of colliding with the passive agent, only those nearby are avoided. Of the active agents that are within an avoidance distance d_i of the passive agent, measuring distance between the agents' capsules, the nearest K are considered for avoidance. In our implementation we used $d_i = 0.4r_i$ and $K = 4$. The avoidance distance d_i can be varied to model observant (larger) or unaware (smaller) behaviour, and is not necessarily related to the agent's radius. In the following description of the avoidance behaviour, we denote the index of the active agent that is to be avoided as $i \in \{i_1, \dots, i_K\}$, and the index of the passive agent as p . Agents that move away from the avoiding agent, i.e., where $(\mathbf{a}_i - \mathbf{a}_p) \cdot \dot{\mathbf{a}}_i > 0$, are safely ignored, as their motion is sufficient to avoid any collisions.

The avoidance behaviour consists of two components, a rotation ϕ_A and a translation \mathbf{t}_A . The passive agent rotates to minimize its width in the active agent's direction of movement, and it translates to move out of the way. The active agent's position \mathbf{a}_i and velocity vector $\dot{\mathbf{a}}_i$ are used to determine a first-order approximation of its future trajectory.

Passive agent A_p rotates to reduce its width perpendicular to $\dot{\mathbf{a}}_i$, allowing A_i as much space as possible to pass. ϕ_A is chosen such that the central axis \mathbf{s}_p aligns with either $\dot{\mathbf{a}}_i$ or $-\dot{\mathbf{a}}_i$, depending on which produces the smallest rotation:

$$\phi_i = \angle(\pm \dot{\mathbf{a}}_i, \mathbf{s}_p).$$

The final rotation ϕ_A is the sum of the individual rotations ϕ_i . This summation is very simple; we are interested in a more refined approach, such as computing the rotation to avoid the one agent that is most likely to collide, based on its position and velocity. The avoidance of other agents could then be performed once that agent has been avoided. The investigation of more elaborate methods is left as future work.

To step out of the way of agent A_i , the passive agent translates perpendicular to the velocity vector $\dot{\mathbf{a}}_i$ (see Fig. 7). For the active agent, we determine the line through \mathbf{a}_i and oriented along $\dot{\mathbf{a}}_i$. For the passive agent, we determine the line orthogonal to $\dot{\mathbf{a}}_i$ and intersecting \mathbf{a}_p . The intersection point \mathbf{x}_i of those lines determines translation vector \mathbf{t}_i :

$$\mathbf{t}_i = \frac{1}{\delta_i} \frac{\mathbf{a}_p - \mathbf{x}_i}{|\mathbf{a}_p - \mathbf{x}_i|},$$

1. Avoidance of passive agents is handled by the space-finding algorithm, which is described in Section 5.1.




Fig. 7. Active agent avoidance; the passive agent A_p (green) will move to avoid the active agent A_i (cyan). Arrow \mathbf{t}_i indicates the resulting avoidance vector.

with dampening factor $\delta_i > 0$. The dampening factor can be agent-specific, to allow for different personality traits. A high dampening factor will make the agent slower to respond than a low dampening factor. In the accompanying video we used $\delta_i = 200$ for all agents. The final agent avoidance translation vector \mathbf{t}_A is the sum of individual translations \mathbf{t}_i .

5.4 Turning Desire into Action

The previous sections described methods to obtain a vector towards more space \mathbf{t}_S , a rotation ϕ_S towards a focus point or to align with the available space, and translation \mathbf{t}_A and rotation ϕ_A to avoid active agents. This section describes how our method selects which translation and rotation to use to produce the agent's motion.

The space finding translation vector \mathbf{t}_S is only applied when certain conditions are met. Firstly, based on the principle of energy minimization, we assume that people accept a marginally worse situation when manoeuvring into a better spot would use significantly more effort than standing still. In our algorithm, the clearance at the found point must be significantly better than the agent's current situation; we use a threshold value of 125 percent of the agent's current clearance. Not only does this produce more natural results (an irregular distribution of free space among the crowd), it also prevents oscillation between points of similar clearance. Secondly, when making space for someone to pass (see Section 5.3), people generally accept a worse situation, as it will only be temporarily. However, people try to move towards an open space if one is available and can be reached while still allowing someone to pass, since this will make it both easier for the passing person and more comfortable for the avoiding person. To model this, space finding vector \mathbf{t}_S is only applied when agent avoidance and space finding result in a translation in roughly the same direction; in other words, when the dot product $\mathbf{t}_S \cdot \mathbf{t}_A > 0$. When these are more or less opposite, only the agent avoidance is performed. The same approach is taken for ϕ_S and ϕ_A ; if both rotate in the same direction, they are combined, otherwise only ϕ_A is applied

$$\mathbf{p}_0 = \mathbf{a}_p + \mathbf{t}_A + \begin{cases} \mathbf{t}_S & \text{if } \mathbf{t}_S \cdot \mathbf{t}_A > 0 \\ \mathbf{0} & \text{otherwise} \end{cases}$$

$$T_o = \theta_p + \phi_A + \begin{cases} \phi_S & \text{if } \phi_A \phi_S > 0 \\ 0 & \text{otherwise,} \end{cases}$$

where θ_p is the passive agent's current orientation, and \mathbf{p}_0 and T_o are respectively the planned position and torso orientation as described in Section 4. The movement of the agent is controlled in the same way as described in that section.

6 WALLS, DOORS, AND OTHER OBSTACLES

In order to model realistic scenarios, our method supports walls, doors and polygonal obstacles. To integrate these into the crowd behaviour, they are all modelled as line segments and included as additional sites in the generalized Voronoi diagram (GVD). As a result, the GVD contains line segment sites for agents, walls, doors and obstacles. All these are interpreted by the crowd agents as impenetrable obstacles.

Doors are modelled as special wall segments that can be enabled when the door is closed, and disabled when the door is opened. As described in Section 5.1, doors are interpreted differently by active and passive crowd members. When a door is open, its line segment simply is not inserted into the active agents' GVD at the next simulation update. The GVD for the passive agents always inserts door line segments, to ensure that the space finding algorithm does not cross area boundaries.

In real life, people anticipate the movement of others. Anticipation in crowd simulation has been studied before [30], [31], [32]; in these works, crowd agents predict other agents' movements, and adapt their own motion to avoid collisions. Our crowd model takes the opposite approach; our active agents place information in the environment to notify passive agents of their intentions, similar to the approach by Yeh et al. [33]. As a real-life example of the intended behaviour, consider a person entering a lift; people appear to mentally model the space required for that person, and make space accordingly. Since the final orientation of the person is not known a-priori, a point would be sufficient to model this. In our simulation, active agents insert point obstacles in the passive GVD, at their goal position \mathbf{g}_i . As a result, the passive agents make space around this position, sooner than the avoidance behaviour would. This is only applicable in situations where the active agent's behaviour is predictable, such as when entering or exiting a lift or bus, which is why it is an optional feature of our crowd simulation method.

7 RESULTS

In this section, we validate our Torso Crowd model against a real crowd, in order to find values for parameters that result in human-like behaviour. Furthermore, we investigate our model by looking at several scenarios. We also compare our model with a disc-based crowd simulation: Reciprocal Velocity Obstacles [8].

7.1 Validation and Parameter Optimization Using a Real Crowd

To validate our crowd model behaviour, we used motion capture data of a real crowd [23]. This data set contains the torso width and thickness of each participant, and a recording of their locations and torso orientations during each of 47 trials. These recorded motions represent human behaviour in a real situation, and thus form suitable ground truth for our parameter optimization and model verification. We do

TABLE 1
The Path Planner Parameters Obtained from Our Comparison with Our Ground Truth Data


category	parameter	value	parameter	value
Planner horizon	H_C	3	H_D	1.50
	H_ϵ	0.05 m		
Score function weights	w_c	2.30	w_g	1.41
	w_l	0.21	w_m	1.00
Clearance weights	w_c^F	0.1	w_c^A	0.9

All values were obtained by manual optimization.

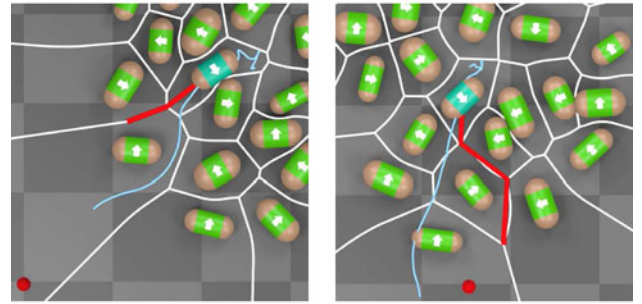
note that the recordings were performed in a controlled environment, and thus may not be a faithful representation of day to day scenarios. We leave evaluation using real crowds, for example using video analysis, to future work. The data set is used to validate the behaviour of the active agents, and the passive agents in the interior of the crowd. In the experiment, the active participants had the concrete, reasonably realistic task of manoeuvring through a crowd to a given point. The rest of the crowd had to stand still in a dense configuration, which was necessarily synthetic for the participants at the edge of the crowd due to the set-up of the experiment. To compare the behaviour of the active participants with our crowd simulation system, we look at topological equivalence, rather than Euclidean distance between paths, as the exact positions of the paths are highly dependent on the behaviour of the passive crowd members. The parameters for the passive agents are simpler and more intuitive than those for the active agents, and were chosen based on visual inspection of the simulation results of the scenarios described in Section 7.2.

Our crowd model uses a number of parameters that determine the behaviour of the active agents, as described in Section 3. These parameters, with their optimized values, are shown in Table 1. To optimize these parameters, we used the following approach:


- 1) *Conversion*: The motion capture data is converted to our abstract agent representation, enabling us to input captured situations into our crowd simulation method.
- 2) *Test sets*: We choose N random frames from the recorded motion capture data. We ensure that each of the chosen frames represents a different situation. The set of frames is separated into two distinct, equally sized, randomly chosen subsets \mathcal{T} for parameter tweaking and \mathcal{V} for verification.
- 3) *Parameter tweaking*: For each frame in \mathcal{T} , the choices of the path planning algorithm are compared with the choices of the participant. We adjust parameters and repeat the comparison, until either all choices made by the path planner are equal to the choices made by the participants, or no more improvements can be made. When the planned path passes between the same agents as the participant's motion, they are considered equal.
- 4) *Verification*: For each frame in \mathcal{V} , the same type of comparison is performed, as a verification of the parameters. We also measure the difference in planned and recorded torso twist.



(a) We consider the blue and red paths as topologically equal, since they choose the same route between the same agents.



(b) We consider those two paths as topologically unequal, since they choose a different route towards the goal, but equally plausible routes towards the goal position.



(c) In the left image, the planner took an unnatural decision, due to noise in the recorded velocity vector. However, 1/60 second later (right image) the planner made the same choice as the participant.

Fig. 8. Comparison between the planned paths (thick red line) and the recorded motion capture data (thin blue line).

We used $N = 80$ to tweak and verify our parameters. Little adjustment was needed during the tweaking phase, resulting in the parameters displayed in Table 1. To prevent over-fitting to our motion capture data, we also validated against the behaviour observed in the simulations seen in the accompanying video. During the verification phase, the path planner chose a path that was topologically equivalent to the participants in 85 percent of the cases. Fig. 8a shows examples of such correctly planned paths. In four of the six cases where the planner diverted from the recorded data, the planned path was equally plausible (see Fig. 8b). In the recordings of the other two cases, at the exact frame used for validation, the participant shifted weight from one foot to the other while otherwise stationary, which resulted in a large change in the instantaneous momentum vector and thus in a different path being chosen (see Fig. 8c); within 1/30 second after the test frame, the planner chose the same path as the participant in both cases. Of course this is not an issue when using simulated data, as our system does not model this weight shifting.




Fig. 9. Stills of the “small lift” scenario. Three agents enter the lift, while the others make space.

The verification of our model also includes a comparison between the planned and recorded torso orientations for the 34 test cases where the predicted path was topologically equivalent to the path of the recorded participant. To remove the influence of local path variations, we compare the torso *twists*, since these are relative to the agent’s and participant’s own paths. The twist is defined as the angle between the torso normal and the torso’s instantaneous velocity vector (as described in Section 4.2). For each verification frame, our Torso Crowd method is used to plan the agent’s next short-term target position \mathbf{p}_0 and torso twist T_t . The recording is then forwarded to the time where the participant reaches \mathbf{p}_0 , after which his/her torso twist T'_t is determined. The error is then expressed as the signed difference between the twists: $E = T_t - T'_t$ where the sign of error E indicates whether our planner over-estimates (positive) or under-estimates (negative) the required twist. In 10 cases we under-estimated the required twist. We classify one of those cases as outlier; it showed a -42° difference due to the participant moving at that angle even though it was not needed given the available space. In the other under-estimated cases the average error was small at -12° , and the error was never more than -16° . In 25 of the 34 cases, we over-estimated the required twist. This is easily explained by the fact that the plan is based on the GVD, which represents the current situation. In the recorded data, it is clear to see that the passive participants make space for the active participant, resulting in more available space, hence less torso twist is required. The average error when over-estimating was 22° . The largest error in the predicted twist was 76° . However, in this case the planned global torso orientation was reached within 0.8 seconds after reaching the planned position. Note that we compare the torsos at the moment in time where the distance between the recorded participant and the planned position is minimal. In 16 of the 34 test cases, either the planned torso twist T_t or the global torso orientation T_o is approached (within a 2° error margin) within 0.5 seconds from that moment in time. This indicates that the recorded participant rotates at a slightly different rate, but still assumes the planned configuration shortly before or after. The average of the absolute error is quite small at 19° , and the median of 16° indicates that more than half of the predictions have a smaller-than-average error. We can conclude that our method for simulating active agents corresponds well with the ground truth data.

The avoidance behaviour of the passive participants was also investigated, to confirm that they show the alignment behaviour we model in Section 5.3. Since our aim is the simulation of dense crowds, we discarded the participants at the edge of the crowd, and limited this analysis to those that are in a dense situation as per the metric described by Stüvel et al. [23]. Their continuous motion was segmented into

avoidance actions, which are defined as a period in which the participant shows a translation and/or rotation in order to make way for the active participant. In our data set, all avoidance actions consisted of at least a translation (average 0.09 m, $\sigma = 0.07$ m), which allowed us to find the peak in translation speed, and use the local minima around this peak to define the start and end timekeys of each avoidance action. At both timekeys, we measured the angle between the passive participant’s torso segment \mathbf{s}_p and the active participant’s velocity vector $\hat{\mathbf{a}}_i$. By analysing 94 avoidance actions, we found that at the start of the avoidance action, the average angle was 42° ($\sigma = 25^\circ$), and at the end timekey it was 30° ($\sigma = 22^\circ$). A paired-samples T-test on the angles shows that this is a strong significant difference ($p < 0.0001$), indicating that there is indeed a trend to align with the active agent’s velocity vector. The specific values of the observed averages are of relative importance, as we did not account for any anticipation or other temporal effects. Doing so may produce stronger results, which is left for future research. The simulated avoidance behaviour is parameterized, and can be adjusted to mimic these findings.

7.2 Examples and Comparison with Disc-Based Simulation

We have modelled several scenarios to test our crowd simulation method. As we focus on situations where a large part of the crowd stands still, typical tests where the entire crowd moves do not suffice. Furthermore, in dense crowds people often bump into each other, so a benchmarking method that penalizes collisions, such as SteerBench [34], will produce unrealistic scores. Instead, we have chosen to use a lift and a hallway to model crowded spaces. All scenarios are simulated at real-time on a single CPU core of a modern PC (Intel Core i7 at 4 GHz). With 44 agents, the *Hallway* scenario ran at 23 frames per second; note that the crowd simulation algorithm was implemented in Python, and the exact Voronoi diagram was computed on the same CPU core using the VRONI library [35]; a GPU-based implementation of the GVD algorithm [36] and an optimized implementation in C of the crowd simulation algorithm will result in a better performance; we leave this for future work. The scenarios are shown in the accompanying video. For each scenario, we first show the simulated agents, and then animated characters that follow the motions of those agents. The set of parameters obtained in Section 7.1 was used in all scenarios, with the exception of the clearance parameter w_c of the active agents. We use $w_c = -0.5$ when entering the densely crowded lift to encourage the agents to move from the low-density hallway into the high-density lift.

Small lift: In this scenario, the lift visits various floors, and on each floor agents get in or out of the lift (see Fig. 9). This




Fig. 10. Stills of the “large lift” scenario. The three green agents have a common focus point (the red dot).


scenario shows the typical division of the available space seen in lifts: one person by itself stands more or less in the middle of the lift, while the space gets divided when more people enter. While waiting for their floor, the agents turn towards a common focus point: the floor indicator panel above the door. When agents leave the lift, the remaining space is used by the remaining agents. Note that, mimicking real life, the space is *not* optimally divided amongst the agents. Instead, agents around a gap, where an agent stood before leaving the lift, benefit most from the newly available space. The effect of the insertion of the immediate goal of active agents, as described in Section 6, can clearly be seen when passive agents make space as an active agent enters the lift.

Large lift: This scenario demonstrates what happens when a group of agents share a focus point, and the density of the crowd increases. The three green agents (see Fig. 10) share a focus point that is positioned at the centroid of their positions. The other agents in the simulation do not have a focus point. The behaviour of the agents entering the lift is not necessarily natural, since half of them have been scripted to move to the back of the lift. This behaviour is more disruptive to the agents already present, and thus forms a more interesting scenario. Even though the focus point has no direct influence on the agents positions, the three agents stay together.


Hallway: In this scenario we show a character manoeuvring through a larger crowd in a hallway. We use this scenario to compare the behaviour of our Torso Crowd model with a widely accepted crowd simulation model: Reciprocal Velocity Obstacles [8] (RVO). This comparison does not aim specifically at RVO; we just use RVO as a good example of a disc-based crowd simulation model. In this comparison, the Torso Crowd agents share the same focus point, out of view on the right-hand side. One agent tries to manoeuvre towards its goal position, while the remainder of the crowd is stationary; those agents have a zero preferred velocity.

Since RVO models agents as discs, we need to convert our capsule representation. We keep in mind that the agents actually represent humanoid shapes; making the RVO agents narrower will result in many undetected intersections. Therefore, the radius is chosen such that the disc encloses the torso capsule, as shown in Fig. 11b. The blue line shows how far the agent was able to move: in such a dense, stationary crowd, the disc-shaped agents are too big to manoeuvre, while this density is not a problem for Torso Crowds (see Fig. 11a). One of the underlying issues is that RVO agents only make space to avoid collisions. When the active agent slows down to avoid a collision, the surrounding agents only move with half the speed necessary to avoid the collision. This forces the


active agent to slow down even more, finally forcing it to stand still. Its velocity vector then becomes zero and holds no information, and the agents in its surroundings will no longer move.



(a) Torso Crowd finds a path to the goal.



(b) RVO with the same width as the capsules does not find a path to the goal.



(c) RVO with the same surface area as the capsules to allow manoeuvring.

Fig. 11. Motion paths of a Torso Crowd agent, and RVO approaches. The RVO agents are displayed with a capsule shape overlay, to visualize intersections between agent-driven humanoid characters.

TABLE 2
Agent Diameters Used in the Comparative Scenario, in Metres

Simulation	shape	min	max	mean
Torso Crowds	capsule	0.382	0.504	0.443
Regular RVO	disc	0.382	0.504	0.443
Same-area RVO	disc	0.280	0.399	0.345

For capsule agents the diameter is defined as $2r_i + \ell_i$, whereas for disc agents this is $2r_i$.


To give the RVO agents more space, we reduce the agent radii, such that the surface area of the agent's ground projection is equal to that of the capsule. This makes the RVO agents narrower but still thicker than the Torso Crowd agents, and results in an equal ground coverage percentage for RVO and Torso Crowds. The RVO agent can then successfully navigate the crowd, at the expense of intersections between the characters. Fig. 11c shows this situation, with red capsules to visualize the character torsos. Statistics on our choice of agent sizes are shown in Table 2; the average width of 0.44 metres matches the average torso width (measured shoulder to shoulder) reported by Stüvel et al. [23]. We can further increase the crowd density; even the smaller RVO agents move slowly, and eventually do not find a path to the goal (Fig. 12b). Our Torso Crowd model still handles this situation, and allows the agent to manoeuvre to its goal position (Fig. 12a).

We can observe more differences. The Torso Crowd agent takes a longer path through the crowd, as it has been configured to avoid areas of low clearance (i.e., agents that stand close together). The RVO agent tries to maintain the shortest path by preferring velocities directly towards the goal position. Another difference is that RVO agents are limited to non-holonomic behaviour; an agent cannot take a step backward or to the side to make room for a passing agent, resulting in unrealistic instantaneous rotations when a human character is animated in its place. Where the passive Torso Crowd agents fill up the space in the wake of the blue agent to make themselves more comfortable, the green RVO agents remain stationary. These results show that the disc shape is not suitable for the simulation of dense crowds. We can also conclude that our Torso Crowd model shows a wider range of motions.


8 CHARACTER ANIMATION

In order to display a humanoid crowd, the motions of the crowd agents need to be mapped to humanoid characters. This poses an under-specified problem. Since only the torso motion is simulated, the lower body orientation needs to be reconstructed before further body animation is possible. In our method, we first describe our lower body estimation method, and then the proposed skeletal animation method. This section briefly describes our approach; for further details we refer to the work by Stüvel [37].

The lower body orientation is estimated based on two observations. Firstly, when manoeuvring, the lower body is oriented more or less in the same direction as the torso, and slightly turned towards the direction of motion. Secondly, the lower body cannot instantly change its orientation. By applying smoothing and computing the angle between the smoothed torso orientation and the trajectory of the motion we determine the lower body orientation.



(a) Torso Crowd finds a path to the goal.



(b) RVO with the smaller agents, with the same surface area as the capsules, does not find a path to the goal.

Fig. 12. Motion paths through an even denser crowd. RVO does not find a path to the goal, while Torso Crowds does. The RVO agents are displayed with a capsule shape overlay, to visualize intersections between agent-driven humanoid characters.

Once the lower body orientation is determined, we can animate the skeletal structure that determines the character's pose. A commonly used technique for crowd animation is the use of a single walk cycle to animate characters at various speeds, where the animation playback rate depends on each character's walking speed. Such an approach is simple to implement, but does not support holonomic motion (such as side-stepping). Furthermore, it results in a direct dependency between walking speed and cadence (steps per minute). However, when people change their walking speed, both the cadence and stride length change [38]. This change in stride length cannot be captured in a single walk cycle, producing unnatural results. To address these issues, the basis for our animation technique is two sets of ten gender-specific walk cycles, consisting of an idle animation (0.00 m/sec), eight slow (0.45 m/sec) walk animations in different directions, and a faster (1.00 m/sec) straight forward walk. The eight slow animations consist of straight forward and backward walking, left and right sidestepping, and diagonal steps in four directions. The speed of 0.45 m/sec was chosen for those animations as it was found to be the average speed when manoeuvring through a dense crowd [27]. To produce a character that walks at the correct speed, the joint angles of the animations are blended using weights that depend on the speed of the crowd agent and the lower body orientation. Constraints are placed on the spine bones to incrementally rotate the torso to produce the required torso twists. An example is shown in Fig. 13.



Fig. 13. A crowd of animated, human characters in a lift. The man in the blue clothing is the active character, whose torso twist is clearly visible. The path is shown in blue.

9 CONCLUSION

In this article we have introduced a novel crowd simulation method, based on the manoeuvring of a number of people in otherwise stationary dense crowds. By extending the common disc-based agent representation to capsules, we are able to plan upper body twisting based on available clearance. Such torso twisting is critical for believable dense crowd manoeuvring.

Our method has been validated against data obtained from real crowd behaviour. The active agent behaviour matches paths chosen by humans in 85 percent of the cases, and produces different but equally plausible paths in 10 percent of the cases. The method's parameter values were manually optimized; it would be interesting to investigate automatic parameter tuning such as proposed by Wolinski et al. [39] and Berseth et al. [40]. Even though we used a simplified Voronoi diagram, the resulting behaviour is a close match to the ground truth (as shown in Section 7.1). The majority of our validation focused on the behaviour of active agents; further comparison, with different ground truth data, could improve realism of the passive crowd members as well, and could strengthen our design decisions, such as the space-finding behaviour, the assumption that passive agents do not move to different rooms, and show what role focus points play in real crowds.

Regardless of the method to obtain the parameters, it is likely that their scope is limited to high-density situations. Since the planning of torso twist is no longer a necessity in lower-density crowds, our crowd simulation system switches between our proposed method and a different agent-based method aimed at regular locomotion, depending on the density of the crowd. Alternatively, our system could be extended to handle lower densities, by employing density-dependent parameter values; this should be relatively straight-forward, since our density metric is agent-oriented, and the parameters are already adjustable for each agent. It would also be interesting to add velocity-based path planning to our method; for example, the change of clearance over time could be used to

prefer small-but-growing openings in the crowd over larger-but-shrinking ones.

All parameters in our model can be personalized per agent, for example by sampling them from a given random distribution or by adding noise. More human-like behaviour could be introduced, such as hesitation when one is unsure about the path to take; hesitation could be simulated in our model when different paths have similar scores. Such diversification of the crowd simulation is left for future work.

The passive agents use a generalized Voronoi diagram to find comfortable places to stand. By definition, such a diagram is symmetric, in that there is no distinction between agents and walls, or the front or rear of agents. This symmetry results in artefacts, such as agents standing too far away from walls. A possible solution may be found in a multiplicatively or additively weighted generalized Voronoi diagram [41], which might also be useful to model the asymmetrical nature of people's personal space [42]. However, since there are no suitable, robust implementations available, we are unable to implement such an approach at this time, and leave this to future work.

In our scenarios, each active agent was appointed a fixed, scenario-specific goal position. When that goal is reached, the agent switches to passive behaviour. Due to the dynamic nature of the crowd, the scripted goal position may not be the most comfortable (see Section 5.1), and the agent will move to a desirable point after reaching the goal. When approaching the goal, the active agent could use a local GVD to find a comfortable position in the goal area, and actively move there before switching to passive behaviour.

We have used an animation system that shows walking characters in a crowd using the motions obtained from the simulation. Our system uses a kinematic approach, hence it does not respond to inter-character collisions. Due to the density of the crowd, however, such collisions are likely to occur. We are currently investigating a method employing physics-based characters that follows our torso planning method [43]. Such a system would be able to respond to collisions in a physically correct way, and be used to plan lower-body motion. Another interesting way to extend our model is based on the observation that in dense crowds people often use their arms for navigation. Not only are they used to physically make space, but also for notification as to the intent to pass between people, and as a tactile addition to visual information about one's neighbours in the crowd. A different approach to improving the result of the animation system would be the integration of a footstep-based method, such as described by Singh et al. [13] and subsequently improved by Berseth et al. [14]; we expect that planning both footstep positions and torso orientations may lead to more natural results. It would also be interesting to extend motion editing methods [44], [45], [46] such that torso orientations are taken into account.

Further research could extend the Torso Crowd model to allow for a crowd of mostly active agents. It would be interesting to add a velocity component similar to RVO to the planner. Furthermore, the Torso Crowd representation could be employed to reduce the energy needed to manoeuvre a crowd for other crowd simulations. For example, our

passive agents anticipate the motions of the active agents, and move aside and twist their torso to make space. Such behaviour can also be observed in less dense crowds, in cases where making twisting the torso is not a geometric necessity for someone to pass, but does provide them with a more energy-efficient path. This happens, for example, when making space for someone running towards a train. This shows that torso planning is not limited to dense crowds.

ACKNOWLEDGMENTS

This research is supported by the Dutch nationally funded project COMMIT/. Part of this work has been performed during the stay of Sybren A. Stüvel at Nanyang Technological University, Singapore. We thank the group of Roland Geraerts for the use of their Explicit Corridor Map framework, which we used to calculate the Voronoi diagrams. We also thank the Blender community for their help.

REFERENCES

- [1] D. Thalmann and S. Musse, *Crowd Simulation*, 2nd ed. New York, NY, USA: Springer-Verlag, 2013.
- [2] N. Pelechano, J. Allbeck, and N. Badler, *Virtual Crowds: Methods, Simulation, and Control*. San Rafael, CA, USA: Morgan & Claypool Publishers, 2008.
- [3] L. Henderson, "The statistics of crowd fluids," *Nature*, vol. 229, no. 5284, pp. 381–383, 1971.
- [4] W. Kerr and D. Spears, "Robotic simulation of gases for a surveillance task," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Aug. 2005, pp. 2905–2910.
- [5] S. Chenney, "Flow tiles," in *Proc. ACM SIGGRAPH/Eurographics Symp. Comput. Animation*, 2004, pp. 233–242.
- [6] R. Alizadeh, "A dynamic cellular automaton model for evacuation process with obstacles," *Safety Sci.*, vol. 49, no. 2, 2011.
- [7] N. Pelechano, J. Allbeck, and N. Badler, "Controlling individual agents in high-density crowd simulation," in *Proc. ACM SIGGRAPH/Eurographics Symp. Comput. Animation*, 2007, pp. 99–108.
- [8] J. van den Berg, S. Guy, M. Lin, and D. Manocha, *Reciprocal n-Body Collision Avoidance*. Berlin, Germany: Springer, 2011, vol. 70, pp. 3–19.
- [9] S. Kim, S. Guy, and D. Manocha, "Velocity-based modeling of physical interactions in multi-agent simulations," in *Proc. 12th ACM SIGGRAPH/Eurographics Symp. Comput. Animation*, 2013, pp. 125–133.
- [10] M. Kapadia, A. Porres, and F. García, "Multi-domain real-time planning in dynamic environments," in *Proc. ACM SIGGRAPH/EUROGRAPHICS Symp. Comput. Animation*, 2013, pp. 115–124.
- [11] S. Lemercier, A. Jelic, R. Kulpa, J. Hua, J. Fehrenbach, P. Degond, C. Appert-Rolland, S. Donikian, and J. Pettré, "Realistic following behaviors for crowd simulation," *Comput. Graph. Forum*, vol. 31, no. 2pt2, pp. 489–498, 2012.
- [12] S. Stüvel, N. Magnenat-Thalmann, D. Thalmann, A. Egges, and F. van der Stappen, "Hierarchical structures for collision checking between virtual characters," *Comput. Animation Virtual Worlds*, vol. 25, no. 3–4, pp. 333–342, 2014.
- [13] S. Singh, M. Kapadia, G. Reinman, and P. Faloutsos, "Footstep navigation for dynamic crowds," *Comput. Animation Virtual Worlds*, vol. 22, no. 2–3, pp. 151–158, 2011.
- [14] G. Berseth, M. Kapadia, and P. Faloutsos, "Robust space-time footsteps for agent-based steering," in *Proc. Comput. Animation Social Agents*, 2015, pp. 115–124.
- [15] B. Van Basten, S. Stüvel, and J. Egges, "A hybrid interpolation scheme for footprint-driven walking synthesis," in *Proc. Graph. Interface*, 2011, pp. 9–16.
- [16] A. Beacco, N. Pelechano, M. Kapadia, and N. Badler, "Footstep parameterized motion blending using barycentric coordinates," *Comput. Graph.*, vol. 47, pp. 105–112, 2015.
- [17] K. Hyun, M. Kim, Y. Hwang, and J. Lee, "Tiling motion patches," *IEEE Trans. Vis. Comput. Graph.*, vol. 19, no. 11, pp. 1923–1934, Nov. 2013.
- [18] K. Jordao, J. Pettré, M. Christie, and M.-P. Cani, "Crowd sculpting: A space-time sculpting method for populating virtual environments," *Comput. Graph. Forum*, vol. 33, no. 2, pp. 351–360, 2014.
- [19] P. Thompson and E. Marchant, "A computer model for the evacuation of large building populations," *Fire Safety J.*, vol. 24, no. 2, pp. 131–148, 1995.
- [20] A. Giese, D. Latypov, and N. Amato, "Reciprocally-rotating velocity obstacles," in *Proc. IEEE Int. Conf. Robot. Autom.*, May 2014, pp. 3234–3241.
- [21] G. Zipf, *Human Behavior and the Principle of Least Effort*. Reading, MA, USA: Addison-Wesley, 1949.
- [22] S. Guy, J. Chhugani, S. Curtis, P. Dubey, M. Lin, and D. Manocha, "PLEdestrans: A least-effort approach to crowd simulation," in *Proc. Symp. Comput. Animation*, 2010, pp. 119–128.
- [23] S. Stüvel, M. de Goeij, A. van der Stappen, and A. Egges, "An analysis of manoeuvring in dense crowds," in *Proc. 8th ACM SIGGRAPH Conf. Motion Games*, 2015, pp. 85–90.
- [24] R. Geraerts, "Planning short paths with clearance using explicit corridors," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2010, pp. 1997–2004.
- [25] A. Sud, R. Gayle, E. Andersen, S. Guy, M. Lin, and D. Manocha, "Real-time navigation of independent agents using adaptive roadmaps," in *Proc. ACM Symp. Virtual Reality Softw. Technol.*, 2007, pp. 99–106.
- [26] A. Sud, E. Andersen, S. Curtis, M. Lin, and D. Manocha, "Real-time path planning in dynamic virtual environments using multi-agent navigation graphs," *IEEE Trans. Visual. Comput. Graphics*, vol. 14, no. 3, pp. 526–538, May 2008.
- [27] S. Stüvel, M. de Goeij, F. van der Stappen, and A. Egges, "An analysis of manoeuvring in dense crowds," Dept. of Information and Computing Sciences, Utrecht Univ., Tech. Rep. UU-CS-2015-006, 2015.
- [28] A. Sud, E. Andersen, S. Curtis, M. Lin, and D. Manocha, "Real-time path planning for virtual agents in dynamic environments," in *Proc. ACM SIGGRAPH Classes*, 2008, Art. no. 55.
- [29] L. Hyvärinen, "Principal component analysis," in *Mathematical Modeling for Industrial Processes*, Berlin, Germany: Springer, 1970.
- [30] S. Paris, J. Pettré, and S. Donikian, "Pedestrian reactive navigation for crowd simulation: A predictive approach," *Comput. Graph. Forum*, vol. 26, pp. 665–674, 2007.
- [31] I. Karamouzas, P. Heil, P. van Beek, and M. Overmars, "A predictive collision avoidance model for pedestrian simulation," in *Motion in Games*, Berlin, Germany: Springer, 2009, vol. 5884, pp. 41–52.
- [32] Y. Suma, D. Yanagisawa, and K. Nishinari, "Anticipation effect in pedestrian dynamics: modeling and experiments," *Physica A, Statist. Mech. Appl.*, vol. 391, no. 1–2, pp. 248–263, 2012.
- [33] H. Yeh, S. Curtis, S. Patil, J. van den Berg, D. Manocha, and M. Lin, "Composite agents," in *Proc. ACM SIGGRAPH/EUROGRAPHICS Symp. Comput. Animation*, 2008, pp. 39–47.
- [34] S. Singh, M. Kapadia, P. Faloutsos, and G. Reinman, "Steerbench: A benchmark suite for evaluating steering behaviors," *Comput. Animation Virtual Worlds*, vol. 20, no. 5–6, pp. 533–548, 2009.
- [35] M. Held, "VRONI: An engineering approach to the reliable and efficient computation of Voronoi diagrams of points and line segments," *Comput. Geom.*, vol. 18, no. 2, pp. 95–123, 2001.
- [36] A. Sud, N. Govindaraju, R. Gayle, and D. Manocha, "Interactive 3D distance field computation using linear factorization," in *Proc. Symp. Interactive 3D Graph. Games*, 2006, pp. 117–124.
- [37] S. Stüvel, "Dense Crowds of Virtual Humans," Ph.D. dissertation, Universiteit Utrecht, Utrecht, The Netherlands, 2016.
- [38] C. Kirtley, M. Whittle, and R. Jefferson, "Influence of walking speed on gait parameters," *J. Biomed. Eng.*, vol. 7, no. 4, pp. 282–288, 1985.
- [39] D. Wolinski, S. J. Guy, A.-H. Olivier, M. Lin, D. Manocha, and J. Pettré, "Parameter estimation and comparative evaluation of crowd simulations," *Comput. Graph. Forum*, vol. 33, no. 2, 2014.
- [40] G. Berseth, M. Kapadia, B. Haworth, and P. Faloutsos, "SteerFit: Automated parameter fitting for steering algorithms," presented at the ACM SIGGRAPH/Eurographics Symp. Computer Animation, New York, NY, USA, 2014.
- [41] A. Okabe, B. Boots, K. Sugihara, and S. Chiu, *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. Hoboken, NJ, USA: Wiley, 2009, vol. 501.
- [42] L. Hayduk, "The shape of personal space: An experimental investigation," *Canadian J. Behavioural Sci./Revue Canadienne Des Sci. Du Comportement*, vol. 13, no. 1, p. 87, 1981.

- [43] Z. Kavafoglu, E. Kavafoglu, and A. Egges, "Robust balance shift control with posture optimization," in *Proc. 8th ACM SIGGRAPH Conf. Motion Games*, 2015, pp. 183–192.
- [44] T. Kwon, K. Lee, J. Lee, and S. Takahashi, "Group motion editing," *ACM Trans. Graph.*, vol. 27, no. 3, pp. 80:1–80:8, 2008.
- [45] M. Kim, K. Hyun, J. Kim, and J. Lee, "Synchronized multi-character motion editing," *ACM Trans. Graph.*, vol. 28, no. 3, pp. 79:1–79:9, 2009.
- [46] J. Kim, Y. Seol, T. Kwon, and J. Lee, "Interactive manipulation of large-scale crowd animation," *ACM Trans. Graph.*, vol. 33, no. 4, pp. 83:1–83:10, 2014.




Sybren A. Stüvel received the master's degree in computer science and the PhD degree from the Virtual Human Technology Lab in the Department of Information and Computing Sciences, Utrecht University, The Netherlands. His main research topic is the animation of dense crowds of virtual characters. He developed a method for the generation of human locomotion based on footstep positions. He is co-author of the IEEE Handbook of Digital Games, and is part of the COMMIT/ project.



Nadia Magnenat-Thalmann received the bachelor's and master's degrees in various disciplines (psychology, biology and biochemistry) and the PhD degree in quantum physics from the University of Geneva in 1977. She has pioneered research into virtual humans over the last 30 years. She was professor at the University of Montreal in Canada, and published more than 500 papers and books on research topics such as 3D clothes, hair, body gestures, emotions modelling, and on medical simulation. Besides

directing her research group MIRALab in Geneva, she is currently a visiting professor and the director of the Institute for Media Innovation, Nanyang Technological University, Singapore.



Daniel Thalmann received the PhD degree in computer science from the University of Geneva in 1977 and an Honorary Doctorate from University Paul-Sabatier in Toulouse, France. He is with the Institute for Media Innovation, Nanyang Technological University, Singapore. He is a pioneer in research on virtual humans, and currently working on real-time virtual humans in virtual reality, crowd simulation, and 3D Interaction. He has published more than 500 papers in graphics, animation, and virtual reality, and is coeditor of 30

books. He also received the Eurographics Distinguished Career Award in 2010 and the 2012 Canadian Human Computer Communications Society Achievement Award.



A. Frank van der Stappen received the MSc degree from the Eindhoven University of Technology, Eindhoven, The Netherlands, in 1988 and the PhD degree from Utrecht University, Utrecht, The Netherlands, in 1994. He is currently an associate professor in the Department of Information and Computing Sciences, Utrecht University. His research interests include manipulation, motion planning, grasping, simulation, animation, and geometric algorithms.



Arjan Egges received the PhD degree from MIRALab - University of Geneva, Switzerland, on the topic of emotion and personality models, in combination with automatically generated face and body motions using motion capture data. He is an associate professor at the Virtual Human Technology Lab in the Department of Information and Computing Sciences, Utrecht University, The Netherlands. His current research focuses on crowd animation and motion perception as a part of the COMMIT/ and TARDIS projects. He is also

an associate editor of the Computer Animation and Virtual Worlds journal published by Wiley and he is one of the founders of the annual Motion in Games conference.

▷ **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.**