

An Analysis of Manoeuvring in Dense Crowds

S.A. Stüvel, M.F. de Goeij, A.F. van der Stappen, A. Egges
VHTLab, Universiteit Utrecht *

Pre-print edition, will be published in the Motion in Games 2015 conference proceedings.



Figure 1: Stills of one of the trials, showing a participant escaping the crowd.

Abstract

In high-density crowds, one can observe torso twists; people rotate their upper body to decrease their width perpendicular to the motion path, in order to squeeze through narrow spaces between other crowd members. In this paper we investigate such behaviour, by recording and analysing dense crowds. Apart from the common approach, where only the position of each person in the crowd is recorded, we also record and analyse the torso orientations. To the best of our knowledge, this has not been done before in the context of dense crowds. We show that the paths chosen by the participants can be predicted by Generalized Voronoi Diagrams based on line segment representations of the participants' torsos, and attest that the medial axis of a capsule-shaped representation of the torso is a good choice for such line segments.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation I.6.4 [Simulation and Modeling]: Model Validation and Analysis— [I.6.5]: Simulation and Modeling—Model Development

Keywords: crowd behaviour, dense crowd, human motion, Voronoi, experiment, crowd simulation

1 Introduction

Crowd simulation is a widely researched topic. Many aspects of the macroscopic human behaviour have been successfully captured in cylinder-based mathematical models; see [Thalmann and Musse 2013] for a recent survey. However, such models cannot capture the torso-twisting behaviour observed in real, high-density crowds. In our daily lives, we often face crowded situations, such as a busy lift, bus or pop concert. In such cases, the members of the crowd are often stationary, only moving when someone wants to pass. While walking through a crowd, a person balances several desires, such as reaching some goal position, avoiding collisions with others, and conserving energy. Crowd models generally try to mimic this behaviour by planning short paths that avoid collisions. However, when the crowd density increases, choosing a collision-free path

becomes more difficult. In such high-density crowds, one can observe torso twists, regular steps and side-stepping (i.e. holonomic behaviour); people rotate their upper body to decrease their width perpendicular to the motion path, in order to squeeze through the narrow spaces between other crowd members. Rather than *walking*, such motions are better described as *manoeuvring* through the crowd.

In this paper, we describe an experiment aimed at understanding such dense crowd manoeuvring behaviour. The underlying motivation is to obtain a ground truth that supports an implementation of such behaviour in a crowd simulation algorithm. In our experiment, the participants form a crowd of such density that it is sparse enough to manoeuvre through, but dense enough to require torso rotations in order to do so (see Figure 1). In our analysis we explicitly ignore the lower body. Not only is it hard to see in a dense crowd, it has also been shown that the shoulder movements through space are a good indication of the movement of the entire body [Arechavala et al. 2008a] Our accompanying video at <http://stuvel.eu/video/dense-crowd-manoeuvring> shows one of those trials, including the capsule-based representation and resulting Voronoi diagram.

Main contribution We investigate properties of dense crowd manoeuvring. We show that in dense crowds people follow Generalized Voronoi Diagrams based on a line-segment representation of the agents. Two possible methods for selecting those line segments are investigated, and we show that the medial axis of a capsule representation of the torso is a good choice for such line segments.

Organization The rest of the paper is organized as follows. Section 2 discusses related work. The details and execution of the experiment are described in Section 3, with analysis of the results in Section 4. The implications and possible future work are discussed in Section 5, which concludes the article.

2 Related work

The study of crowd behaviour spans a large research area, ranging from computer vision techniques to assess human behaviour [Zhao and Li 2014] to the simulation of evacuation scenarios [Zheng et al. 2009], and application in (serious) games. In this section we focus on the works relevant to the *study of dense crowds*. For a general

*email: {s.a.stuvel,a.f.vanderstappen,j.egges}@uu.nl

overview of crowd simulation techniques and topics, we refer to the book by Thalmann and Musse [2013].

Two possible ways of simultaneously capturing the motions of multiple individuals in a crowd are video recordings and motion capture systems. *Video recordings* are most often used; as these require no markers to be attached to the participants, the data are easier to obtain. A common approach to process a video stream is either pure visual comparison [Narain et al. 2009], or the use of feature point tracking in order to determine movement of people in a crowd [Saxena et al. 2008; Charalambous et al. 2014]. Such feature point tracking is also used by Lee et al. [2007] to *train* a crowd simulation system, such that it exhibits behaviours imitating real human crowds. All these works use the video data to determine the positions of the recorded people; under the assumption of non-holonomic motion, the orientation is determined from the velocity vector. Such techniques are not suitable for studying the torso twisting behaviour in crowds, as firstly the position alone is not enough to study the torso twist, and secondly the nonholonomy assumption is no longer valid for dense crowds [Mombaur et al. 2010]. Furthermore, Jacques et al. [2010] state that, due to severe occlusions, it becomes difficult to accurately track individuals in crowded situations. They found that there is actually a strong negative relation between the density of the crowd and the accuracy of tracking methods. For these reasons, we have chosen to employ a *motion capture system* for our experiment.

Motion capture systems have been used by Wolinski et al. [2014] to optimize parameters for various crowd simulation systems to increase similarity between the simulated and recorded crowd behaviour. Jelic et al. [2012] and later Lemerrier et al. [2012] used motion capture to study following behaviour in crowds, by recording people following each other in a circular fashion. In all those works, the participants are represented as moving discs on the ground plane. Hence, the resulting data are unsuitable for the study of torso twisting behaviour. Truong et al. [2010] used a motion capture system extend the work by Arechavaleta et al. [2008b] to include holonomic motion. They distinguish between those two behaviours by observing the orientation of the shoulders with respect to the velocity vector of the body. Truong et al. only observed a single participant at the time, allowing for an unobstructed view from the motion capture cameras. We use a similar approach, except that in our experiment we actually study a crowd, so any marker below the shoulder has a high chance of being occluded by the other participants. Therefore, we only use markers on the shoulders.

In our experiment, we compare the motions of people in dense crowds to the motions predicted by a Generalized Voronoi Diagram. It has been shown that such diagrams can be used to steer medium-density crowds [Sud et al. 2008]. However, it is unknown whether people in dense crowds also follow a Voronoi Diagram. To our knowledge, we are the first to investigate this.

3 Experiment

This section describes the experiment design, execution, and results. In short, the goal of the experiment is to obtain information on how people manoeuvre through dense crowds. We observe and analyse torso motions; even though we do not explicitly consider hip rotations, their effect is encoded in the motion of the torso.

3.1 Experiment design

The experiment consists of a repetition of trials. In each trial, the person in the centre of the crowd, the designated *walker*, manoeuvres towards a predefined goal position, while the other participants remain more or less stationary. The experiment is aimed at pro-

viding ground truth data, which can be analysed to further understand crowd behaviour, as well as provide empirical data to improve crowd simulation methods. We are particularly interested in a geometric model that predicts the chosen path through the crowd.

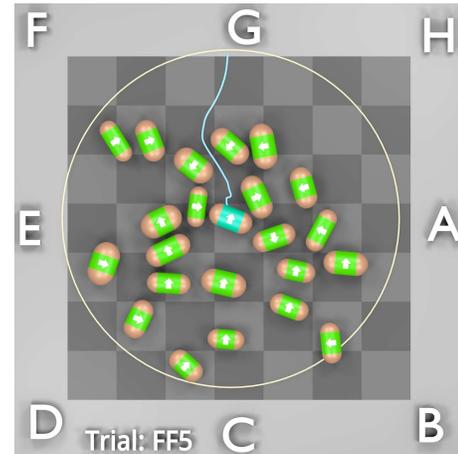


Figure 2: A situation of a task, showing the walker in the starting configuration and letter G as the goal.

Participants are adorned with three motion capture markers: one on the left shoulder, and two on the right. Each participant started at a predetermined position, simplifying the identification and labelling during the post-processing of the recorded data. A circle was drawn in the middle of the motion capture studio. All participants except the walker were prohibited from leaving this circle during each trial, to allow for a consistent overall crowd density and a symmetric distribution of the participants around the centre. The location and size of the circle depended on the limitations of the motion capture system as well as the intended crowd density. Around the circle, the letters A-H were hung 45 degrees apart, as shown in Figure 2. The motion capture system employs fourteen cameras recording at 100 Hz.

The experiment trials were grouped into sets. For each set, a single participant was chosen for the role of walker. For each walker, we recorded a set of seven trials. The walker received a randomized set of task cards: a random subset of six of the eight letters, and a question mark. He/she was instructed to keep these cards hidden from the other participants. As two letters were excluded for each walker, the other participants would not be able to predict the tasks.

The non-walker participants were asked to treat the situation as a densely packed bar, and to crowd together to such a degree that it would be non-trivial but still possible to manoeuvre through. As for dealing with the walker, we asked participants to let him/her through as they would have in similar situations in real life and not anticipate their movement too much.

3.2 Execution

Each participant was measured for their shoulder width, chest thickness, and distance between the left and right shoulder markers. The latter measurement allows us to model the participants as a line segment translating and rotating on the ground plane, which we used for most of our analyses (see Section 3.3). The other dimensions are used to define an alternative capsule-shaped representation, which will be discussed in Section 4.2.

Each *walker* was hand-picked from the participants on the basis of height and gender. We chose participants of average height (com-

pared to the participants, not the country or world population), on the assumption that being too short or too tall will influence the behaviour – we leave the influence of height on the behaviour in crowd to future work. We alternated the gender of the walker to eliminate gender bias.

Each trial consisted of the following steps:

1. All participants move to their predefined starting positions. Recording starts.
2. The *walker* moves to the centre of the circle and rotates to face letter G, while the other participants walk around in a more or less random fashion. This ensures the crowd is different in each trial, and gives the walker ample time to covertly inspect his/her set of cards to determine the task to perform.
3. When the crowd is sufficiently randomized and evenly distributed around the walker, a verbal signal is given. The participants move to fill possible gaps in their vicinity, and then stop walking.
4. The walker manoeuvres through the crowd. Depending on the task, he/she tries to reach the goal letter, or, when the task card indicates a question mark, tries to exit the crowd in the easiest direction of his/her own choice.
5. When the task is complete, recording stops, and the next trial begins.

A total of 23 people (16 male, 7 female) participated in the experiment, with an average age of 24 years ($\sigma = 8.4$). Their average chest width was 0.44 metre ($\sigma = 0.03$), and chest thickness 0.23 metre ($\sigma = 0.03$). 7 participants took the role of walker, and a total of 47 usable trials were recorded; two recordings were rejected due to a technical issue and a participant not adhering to the task.

3.3 Representation

After the experiment, all motion capture data was mapped to an abstract agent representation for analysis purposes. Each agent consists of a line segment, defined by the ground projections of the left shoulder marker and the centroid of the two right shoulder markers. The centroid of the segment defines the position of the abstract agent. The *torso normal* of the agent is defined as the normal of the line segment facing the front of the torso. Hence, it is always perpendicular to the line segment, regardless of the direction of motion. This representation has been shown to be suitable for human motion analysis [Arechavaleta et al. 2008a].

4 Results

In this section we describe our experimental results, based on an analysis of the abstract representation of the motion capture data. At the start of each trial, the walker turns towards the target position, investigates the situation, and shifts balance in order to start manoeuvring. This initial turning towards the target position occurred naturally, and was not instructed. What follows is a period of *dense manoeuvring*, until he/she exits the crowd; the recordings from this period are the subject of our analysis.

In this paper we discuss our most prominent results. For more details and additional analyses we refer to Stüvel et al. [2015].

4.1 Modelling direction of movement

It is a known fact that people minimize their perceived energy use when walking [Zipf 1949; Guy et al. 2010]. We predict that this holds true for dense crowds as well, and hypothesize that in such

a crowd people no longer attempt avoiding single individuals, but choose an opening between pairs of individuals to pass through, and thus:

1. move towards the *midpoint* of the opening, as this provides a path with least chance of collision;
2. move towards the *area of largest clearance* behind that midpoint when the planning agent is closer to the other individuals than the individuals are to each other.

A *midpoint* is defined as the middle of the opening between two agents. In our model, we use the middle of the line segment that represents the shortest distance between the agents' line segments. By definition, this point lies on an edge of a Generalized Voronoi Diagram (GVD) [Lee and Drysdale III 1981], or on an extension of such an edge when a third agent is close by. In such a case, the midpoint would lie on a GVD edge if the third agent were to be removed; see the right-hand image in Figure 3 for an example. The GVD is defined as a pair (V, E) of vertices V and arc-shaped edges E that partition the ground plane. Each cell of this partition is defined by an agent, and consists of all points that are closer to that agent than to any other. The centre of the *area of largest clearance* behind the aforementioned opening corresponds to a vertex of the GVD.

In our analyses we simplify the GVD by using the agents' line segments (see Section 3.3), rather than using their exact shape. In this section, we will use the line segments defined by the shoulder markers. In Section 4.2 we will investigate an alternative line segment that produces even better results.

To verify our expectations, we have analysed the difference between the actual direction of movement and the direction predicted by the expectations described above. In our analysis, we use three different methods to test our expectations:

MIDPOINT Only *midpoints* are considered, regardless of the distance to the planning agent. This corresponds to only testing the first hypothesis.

VERTEX Only *vertices* of the Voronoi cell defined by the planning agent are considered. This corresponds to only testing the second hypothesis, regardless of the distances.

LIMITED MIDPOINT *Midpoints* are considered, but limited to the closest Voronoi *vertex* on their corresponding edges. This corresponds to both hypotheses.

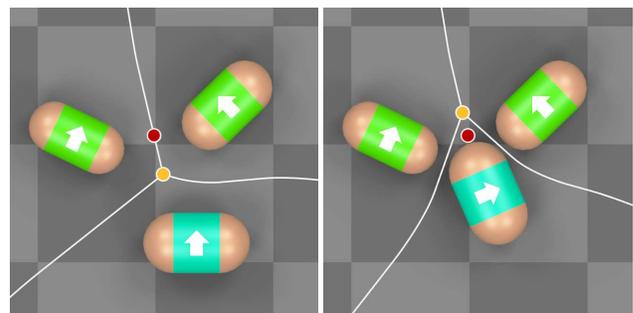


Figure 3: Two situations that show the difference between the midpoint (dark red) of the two green agents, and the vertex (light orange) that indicates the start of the path between them. The planning agent is displayed in cyan.

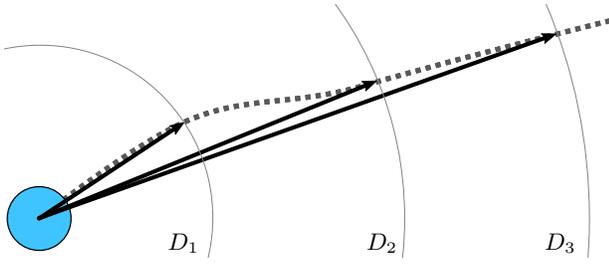


Figure 4: Direction vectors determined by three different plan-ahead distances. The cyan dot indicates the position of the participant, and the dotted line shows their recorded path.

Figure 3 shows the difference between the *midpoint* of the two green agents in dark red, and the *vertex* that indicates the start of the path between them in light orange. The LIMITED MIDPOINT approach would use the dot that is furthest away from the cyan planning agent.

We aim to compare the directions predicted by the methods described above with the short-term direction of movement of the participant, in the order of magnitude of up to one second to capture short-term human decision making; we are not interested in the long-term overall direction, as we know it will approach the vector from the starting point to the goal position. This direction cannot directly be determined by the instantaneous velocity vector, since this vector varies too rapidly. Such rapid changes are especially noticeable at low velocities, where merely shifting weight can rotate the velocity vector by 180° . To filter out these variations, we consider the recorded path of the agent, and obtain the vector from the current position to the path at a given Euclidean distance of $D \in [0.05, 1.5]$ metres (see Figure 4). The participant’s position at this distance D is uniquely defined due to the nature of our recordings, as the participants move more or less monotonically towards their target. By choosing this distance too low, the resulting direction vector will mimic the instantaneous velocity vector and thus vary rapidly and be an ill match for the overall direction of the participant. On the other hand, choosing this distance too large will result in a direction vector that matches the long-term behaviour of the participant, instead of the short-term behaviour we are interested in. The maximum velocity we have measured during the experiment is 1.4 m/s; we are interested in short-term decisions in the order of one second, resulting in an upper bound of 1.4 metre. We have chosen our interval slightly larger than strictly necessary to ensure that the optimal distance falls within the parameters of the experiment. The search space $D \in [0.05, 1.5]$ is sampled uniformly into 10 bins; we have observed that dividing into more fine-grained bins does not produce significantly different results. Distance D is named the *plan-ahead distance*, as it could indicate the distance people consider in order to plan their direction. An example is shown in Figure 4. To our best knowledge, the relation between the agent’s instantaneous speed and the preferred plan-ahead distance has not been studied in the context of dense crowds.

As mentioned in the introduction of this section, our analysis will be limited to the timespan of *dense manoeuvring*, which ends when the participant exits the dense crowd. We continually observe the participant’s Voronoi cell; when the cell is incident to a Voronoi cell defined by the bounding box, we no longer consider the participant to be in the dense crowd. This intuitive definition works well in practice. An additional criterion makes this definition slightly stricter; by limiting the surface area of the Voronoi cell to 0.32 m^2 we truncate a small number of trials and ensure that all analysed data describe dense manoeuvring. An important aspect of our definition is that it does *not* consider the aggregate crowd density, but

categorizes the situation of individual agents.

Now that we can compute a direction of movement and define the timespan of dense manoeuvring, we can test our hypothesis. Since there is a correlation between crowd density and walking speed [Daamen 2004], we have included the speed of the agent as a variable in our analysis. Figure 6a shows the graphs for the LIMITED MIDPOINT testing method; for the graphs of the MIDPOINT and VERTEX methods we refer to Stüvel et al. [2015]. For each frame F in our recordings, we determine the instantaneous speed of the walker $s(F)$. We discretize the measured speeds into twelve bins, allowing us to average similar results from different frames. This number of bins was chosen by balancing a high enough bin count to obtain detailed data, and low enough count to ensure a sufficiently large number of data points in each bin. Speed $s(F)$ determines the bin column in which the data for that frame is plotted. We then vary plan-ahead distance D for each frame F incrementally to determine direction of movement vectors \vec{d} , an example of which is shown in Figure 4. This plan-ahead distance determines the row in the graph. We then calculate error $e(\vec{d}, P)$ with respect to a prediction P , as described in the next paragraph. The error is stored at the appropriate bin; the graphs in Figure 6 show the average error (left) and standard deviation (middle) of the binned errors, and the number of data points in each bin (right).

The error $e(\vec{d}, P)$ depends on prediction P , which is the set of points considered by the different prediction methods MIDPOINT, VERTEX and LIMITED MIDPOINT. The error is defined as the minimal angle between the direction of movement \vec{d} and the prediction points in P :

$$e(\vec{d}, P) = \min_{\vec{p} \in P} \left[\arccos \left(\frac{\vec{d}}{\|\vec{d}\|} \cdot \frac{\vec{p} - \vec{x}}{\|\vec{p} - \vec{x}\|} \right) \right]$$

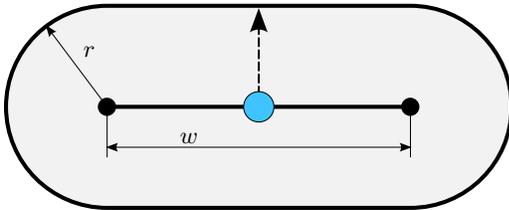
For each binned speed we find the plan-ahead distance that results in the smallest predicted error e . These bins are marked with a white dot in Figure 6, and indicate the most likely plan-ahead distance L for each speed v . As we try to predict future motions based on the static situation in the current frame F , and the participants can obviously change their trajectory dynamically at any moment in time, it is impossible to obtain a zero error. However, the low average error in these bins (≈ 7.2 degrees) shows that there is only a small difference between the direction of movement of the participants, and the directions predicted by the Generalized Voronoi Diagram of the crowd.

The white lines in Figure 6 indicate the best linear regressions through the bins of minimal error, weighted by the sample density. They show a more or less constant plan-ahead distance of approximately 0.40 metres, indicating that people in dense crowds generally plan for a constant distance regardless of their speed. For results on the MIDPOINT and VERTEX methods we refer to Stüvel et al. [2015].

4.2 Capsule-shaped agent representations

In the previous section, we showed that people manoeuvring in dense crowds tend to follow a Generalized Voronoi Diagram (GVD). We used a simple representation for each participant: a line segment spanned by the motion capture markers on the shoulders. Even though this is an accepted way to capture human motion [Arechavaleta et al. 2008a], it is based on more or less arbitrary points on the shoulders, and the distance from the line segment to the edge of the participant’s torso is not constant in all directions. Hence, the GVD of the participants will differ from the GVD defined by those line segments. This difference can be reduced by

slightly changing the length of the used line segment such that the distance from the edge of the torso to the line segment is more or less constant. The result is a capsule-shaped representation of the torso, as shown in Figure 5a. When using this representation, the inter-agent distances only differ by a constant amount from the distances between the line segments, in all directions.



(a) Schematic view of the capsule representation. The dashed arrow indicates the torso normal of the agent.



(b) Top view of a motion captured situation. It can clearly be seen how the capsule model fits the torso of the characters. We can only show this for our character models, as we did not have an orthographic overhead camera during the experiment.

Figure 5: Our abstract crowd agent representation, consisting of an oriented line segment with a radius.

Each capsule consists of a line segment of length w with a buffer of radius r , as shown in Figure 5a. For each participant, the measured chest thickness T determines radius $r = T/2$, and the measured chest width W determines the line segment length $w = W - 2r$. Figure 5b shows an overhead view of a motion captured situation, and demonstrates how well the capsule shapes fit the torsos. Of course, the legs extend from the capsule, but this is a common property of cylinder-based crowd simulations too, hence not an issue specific to the chosen representation. This approach resulted in a shortening of the line segments by an average of 10 cm ($\sigma = 3$ cm).

We have repeated the same analysis as described in Section 4.1, replacing the line segments based on motion capture markers by line segments based on the capsule representation. Figure 6b shows the graphs for the LIMITED MIDPOINT testing method; for the graphs of the MIDPOINT and VERTEX methods we refer to Stüvel et al. [2015]. As expected, the capsule-based representation results in a smaller error, hence a better model for the participants' motions. The best results are obtained with the LIMITED MIDPOINT method, resulting in an error of only 6.65 degrees ($\sigma = 5.33$); this translates to an 8% smaller error and 30% smaller standard deviation. This is corroborated by visually comparing the graphs shown in Figure 6; we can observe that this representation and method also result in the smoothest distribution of the error, and smallest standard deviation with the smoothest distribution, hence produces the most reliable results.

5 Discussion & Conclusion

Our most prominent result is the correlation between dense crowd manoeuvring patterns with Generalized Voronoi Diagrams. We generated this Voronoi diagram by modelling the participants as line segments, and analysed the motions of the agents. Our results show that with an average error of less than 7° ($\sigma = 5^\circ$) our capsule-based LIMITED MIDPOINT method successfully matches the paths of our participants. In comparison, representing agents using shoulder marker line segments results in not only a higher average error, but also higher fluctuations in the standard deviation. We expect that this Voronoi-based model allows for a crowd simulation algorithm that mimics people's behaviour in dense crowds. A capsule-based line segment representation will result in more natural, human behaviour in such a simulation.

An average error of less than 7° seems small. However, given that we introduced a new method of crowd analysis, the question whether this error is good or bad remains to be answered. A perception study may give more information about usable bounds on this error, where the resulting crowd behaviour is considered humanoid.

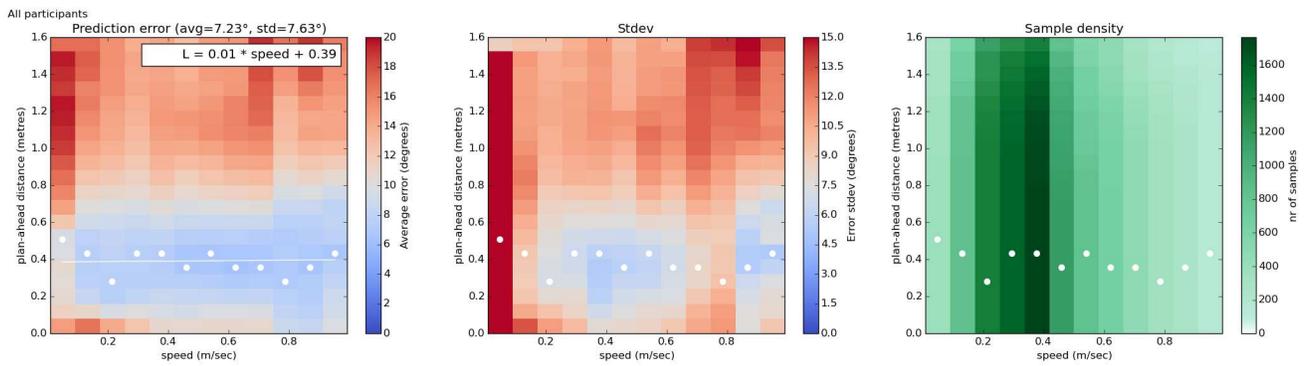
As we focus on torso orientations, the orientation of the head was not recorded during our experiment. A future experiment, employing a cap with motion capture markers, could record head movement. Such recordings could produce more natural results when applying full-body animation, as well as provide interesting data on the impact of vision on manoeuvring behaviour.

Acknowledgements

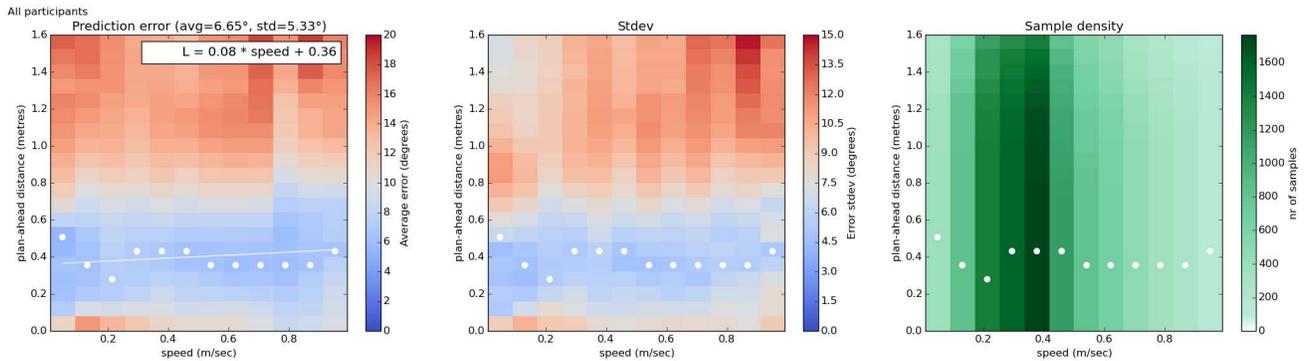
We want to thank everybody who took the time to participate in our experiment. This research is supported by the Dutch nationally funded project COMMIT.

References

- ARECHAVALETA, G., LAUMOND, J.-P., HICHEUR, H., AND BERTHOZ, A. 2008. On the nonholonomic nature of human locomotion. *Autonomous Robots* 25, 1-2, 25–35.
- ARECHAVALETA, G., LAUMOND, J.-P., HICHEUR, H., AND BERTHOZ, A. 2008. An optimality principle governing human walking. *Robotics, IEEE Transactions on* 24, 1 (Feb), 5–14.
- CHARALAMBOUS, P., KARAMOUZAS, I., GUY, S., AND CHRYSANTHOU, Y. 2014. A data-driven framework for visual crowd analysis. *Computer Graphics Forum* 33, 7 (Oct.), 41–50.
- DAAMEN, W. 2004. *Modelling passenger flows in public transport facilities*. Delft University Press, 9 Sept.
- GUY, S., CHHUGANI, J., CURTIS, S., DUBEY, P., LIN, M., AND MANOCHA, D. 2010. Pedestrians: a least-effort approach to crowd simulation. In *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, SCA '10, 119–128.
- JACQUES JR., J., RAUPP MUSSE, S., AND JUNG, C. 2010. Crowd analysis using computer vision techniques. *Signal Processing Magazine, IEEE* 27, 5 (Sept), 66–77.
- JELIC, A., APPERT-ROLLAND, C., LEMERCIER, S., AND PETTRÉ, J. 2012. Properties of pedestrians walking in line: Fundamental diagrams. *Physical Review E* 85 (Mar), 036111.



(a) Results for the motion capture marker based approach described in Section 4.1



(b) Results for the capsule-based approach described in Section 4.2

Figure 6: Results of our analyses using the LIMITED MIDPOINT method. The agent's instantaneous speed is shown on the x-axis; the chosen plan-ahead distance D is shown on the y-axis. The colour indicates the average error over all recorded trials (left), standard deviation (middle) and number of data points (right). The white dots indicate the bins with the smallest error.

- LEE, D. T., AND DRYSDALE III, R. L. 1981. Generalization of voronoi diagrams in the plane. *SIAM Journal on Computing* 10, 1, 73–87.
- LEE, K., CHOI, M., HONG, Q., AND LEE, J. 2007. Group behavior from video: A data-driven approach to crowd simulation. In *Proceedings of the 2007 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, SCA '07, 109–118.
- LEMERCIER, S., JELIC, A., KULPA, R., HUA, J., FEHRENBACH, J., DEGOND, P., APPERT-ROLLAND, C., DONIKIAN, S., AND PETTRÉ, J. 2012. Realistic following behaviors for crowd simulation. *Computer Graphics Forum* 31, 2pt2, 489–498.
- MOMBAUR, K., TRUONG, A., AND LAUMOND, J.-P. 2010. From human to humanoid locomotion—an inverse optimal control approach. *Autonomous Robots* 28, 3, 369–383.
- NARAIN, R., GOLAS, A., CURTIS, S., AND LIN, M. 2009. Aggregate dynamics for dense crowd simulation. *ACM Trans. Graph.* 28, 5 (Dec.), 122:1–122:8.
- SAXENA, S., BRÉMOND, F., THONNAT, M., AND MA, R. 2008. Crowd behavior recognition for video surveillance. In *Advanced Concepts for Intelligent Vision Systems*, J. Blanc-Talon, S. Bourennane, W. Philips, D. Popescu, and P. Scheunders, Eds., vol. 5259 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 970–981.
- STÜVEL, S., DE GOEIJ, M., VAN DER STAPPEN, F., AND EGGES, A. 2015. An analysis of manoeuvring in dense crowds. Tech. Rep. UU-CS-2015-006, Department of Information and Computing Sciences, Utrecht University.
- SUD, A., ANDERSEN, E., CURTIS, S., LIN, M., AND MANOCHA, D. 2008. Real-time path planning for virtual agents in dynamic environments. In *ACM SIGGRAPH 2008 Classes*, ACM, New York, NY, USA, SIGGRAPH '08, 55:1–55:9.
- THALMANN, D., AND MUSSE, S. 2013. *Crowd Simulation*, 2 ed. Springer-Verlag London.
- TRUONG, T. 2010. Unifying nonholonomic and holonomic behaviors in human locomotion.
- WOLINSKI, D., J. GUY, S., OLIVIER, A.-H., LIN, M., MANOCHA, D., AND PETTRÉ, J. 2014. Parameter estimation and comparative evaluation of crowd simulations. *Computer Graphics Forum* 33, 2, 303–312.
- ZHAO, F., AND LI, J. 2014. Pedestrian motion tracking and crowd abnormal behavior detection based on intelligent video surveillance. *Journal of Networks* 9, 10.
- ZHENG, X., ZHONG, T., AND LIU, M. 2009. Modeling crowd evacuation of a building based on seven methodological approaches. *Building and Environment* 44, 3, 437–445.
- ZIPF, G. 1949. *Human behavior and the principle of least effort*. Addison-Wesley Press.