


# Imperceptible Relaxation of Collision Avoidance Constraints in Virtual Crowds

Richard Kulpa\*  
M2S - Univ. Rennes 2  
Gol em S.A.

Anne-H el ene Olivier†  
INRIA Rennes

Jan Ondr ej‡  
INRIA Rennes

Julien Pettr e§  
INRIA Rennes



**Figure 1:** Hundreds of characters are causing collisions in these crowd scenes. They are however viewed under some conditions that prevent from easily detecting them (a). We performed perceptual studies (b) to define a LOD selection function that chooses the characters (in red) for which collisions detection can be avoided (c). Important computation-time savings are obtained from applying the LOD selection function to the design of crowd simulators based on level-of-details methods (d).

## Abstract

The performance of an interactive virtual crowd system for entertainment purposes can be greatly improved by setting a level-of-details (LOD) strategy: in distant areas, collision avoidance can even be stealthily disabled to drastically speed-up simulation and to handle huge crowds. The greatest difficulty is then to select LODs to progressively simplify simulation in an imperceptible but efficient manner. The main objective of this work is to experimentally evaluate spectators' ability to detect the presence of collisions in simulations. Factors related to the conditions of observation and simulation are studied, such as the camera angles, distance to camera, level of interpenetration or crowd density. Our main contribution is to provide a LOD selection function resulting from two perceptual studies allowing crowd system designers to scale a simulation by relaxing the collision avoidance constraint in a least perceptible manner. The relaxation of this constraint is an important source for computational resources savings. Our results reveal several misconceptions in previously used LOD selection functions and suggest yet unexplored variables to be considered. We demonstrate our function efficiency over several evaluation scenarios.

**CR Categories:** I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Radiosity;

**Keywords:** Crowd simulation, collision avoidance, performance, believability, perception, experimentation

**Links:** [DL](#) [PDF](#)

\*e-mail: richard.kulpa@univ-rennes2.fr

†e-mail: anne-helene.olivier@inria.fr

‡e-mail: jan.ondrej@gmail.com

§e-mail: julien.pettre@inria.fr

### ACM Reference Format

Kulpa, R., Olivier, A., Ondr ej, J., Pettr e, J. 2011. Imperceptible Relaxation of Collision Avoidance Constraints in Virtual Crowds. *ACM Trans. Graph.* 30, 6, Article 138 (December 2011), 10 pages. DOI = 10.1145/2024156.2024172 <http://doi.acm.org/10.1145/2024156.2024172>.

### Copyright Notice

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

  2011 ACM 0730-0301/2011/12-ART138 \$10.00 DOI 10.1145/2024156.2024172 <http://doi.acm.org/10.1145/2024156.2024172>

## 1 Introduction

Interactive virtual crowds require high-performance simulation, animation and rendering techniques to handle numerous characters in real-time. Moreover these characters must be believable in their actions and behaviors. Believability is defined as a trade-off between realism and performance which aims at lowering the complexity of simulation models in favor of performance. The main challenges are to remove the least perceptible details first, to preserve the global visual aspect of results at best and meanwhile, to significantly improve computation times. Level-of-details (LODs) strategies were proposed to combine several realism-performance trade-offs together in a single application. This enables designers to progressively and locally adapt these levels with respect to the visual importance of objects in the displayed scene. A LOD selection function automatically determines the most efficient level to be used for each of the displayed object with respect to several criteria, such as visibility, viewpoint, saliency, etc. Perceptual studies are pertinent answers to the problem of designing LOD selection functions. They were used in previous work on virtual crowds to address some rendering and animation aspects. This paper addresses the problem of scaling crowd simulation models. More specifically, we question the need to solve every single collision within a large moving virtual crowd.

Collision avoidance is a prevailing constraint in the formulation of microscopic crowd simulation models. The absence of interpenetration between bodies ensures that a nominal level of realism is reached when, for example, simulating pedestrian traffic for architectural design. However, in the context of believable crowds, absence of collision is not necessarily the best criterion to guarantee satisfying results and is even sometimes responsible for some strange, but collision-free, maneuvers which are particularly detectable by spectators. In addition, avoidance needs to iteratively check and solve collisions by time-consuming algorithms. Relaxing - in a visually imperceptible manner - the collision avoidance constraint wherever possible simply represents an opportunity for saving important computational resources. Although this artifice has been used in previous work, it was never given an attempt to search for the bounds of collision perception and validate the proposed LOD selection function.

Motivated by these assessments, we propose two successive perceptual studies to inspect the effect of various factors on the visual perception of collisions. The first study focuses on pairwise collision situations out of the context of crowds. Factors related to

observation conditions are studied, such as the distance to camera, the camera tilt angle or the pan angle relatively to the characters trajectory. The second study evaluates the effect of visual complexity implicitly induced by crowd scenes on collision perception.

The main paper contribution is a LOD selection function that determines where collisions between characters do require to be solved or not over a crowd scene. We show that previously proposed functions were suboptimal. As an example, distance to camera, which was mainly used in previous works, is not a more important factor than camera tilt angle or perspective. Secondary contribution is the integration of this LOD into existing crowd simulators, which allows us to demonstrate their efficiency on several examples.

## 2 Related Work

Crowd simulation is an active topic particularly promoted by the needs of both the architecture and the entertainment fields. Crowd simulation generally implies to compute the global motion of numerous characters gathered in a same area at a certain level of density, each being propelled by individual or common goals. A main challenge is then to model interactions between characters. The required level of realism of interactions varies with respect to the aimed application field. Several types of approaches emerged: cell-automaton [Burstedde et al. 2001] or physics-based [Helbing and Molnar 1995] models were used in the domain of architecture. Our work focuses on interactive applications for entertainment which requires performance and global motion believability. Continuum-based [Treuille et al. 2006], agent-based [Reynolds 1987] or geometry-based [Paris et al. 2007; van den Berg et al. 2008] solutions were preferred in this field.

A great amount of work has focused on improving simulation performance to enable real-time virtual crowds. Naively, each character in a crowd potentially interacts with all the remaining ones ( $1-to-n$  interactions) which tends to a quadratic algorithmic complexity. This statement applies to geometry or physics-based approaches. This complexity is a bottleneck preventing from widely increasing crowd size. However, various strategies were proposed to reach impressive real-time results. One solution is to consider interactions between neighbor characters only, at the cost of numerous nearest neighbor searches; this problem is however a classical optimization one with efficient solutions [Samet 2005]. 1st and 2nd order Voronoi diagrams were combined to efficiently simulate a crowd of independent agents [Sud et al. 2008]. Another strategy is to use an intermediate layer to model interactions: this is the key-idea brought by Treuille and colleagues [Treuille et al. 2006]. A grid gathers both static (goals) and dynamic (density) simulation data, which are modeled as discrete potential fields: characters then move according to the gradient.  $1-to-n$  are then reduced to  $1-to-layer$  interactions, which results in a drastic reduction of complexity. Limitations however remain, characters have to share common goals and artifacts are observed when density becomes high. Limitations were partly solved by introducing hybrid approaches [Narain et al. 2009].

Nevertheless, higher performance is generally obtained at the cost of limitations or artifacts. Various types of motion defects are recurrently observed [Kapadia et al. 2009]. To scale a simulation, one key idea is to mix accurate and efficient models together in a same simulation system. Quality of results is then locally and progressively degraded in favor of performance. Such a strategy is set by LOD techniques. These techniques were first used in the context of visualization problems: they were for example applied to the real-time rendering of crowds [Tecchia et al. 2002; Dobbyn et al. 2005]. They were more recently applied to crowd simulation itself in [Niederberger and Gross 2005; Pettré et al. 2006; Paris et al.

2009; Kistler et al. 2010]. All of these papers suggest to adapt the collision avoidance behavior with respect to the selected LOD, and some even stop solving collisions where the lowest level of quality is applied: characters are steered in a basic manner and randomly spread to maintain some visual illusion. However, the proposed LOD selection functions were manually designed. They determine the quality level required for simulating each character with respect to: its distance to camera, its centrality in screen, or the local crowd density. No clear evaluation was yet proposed to validate LOD selection functions concerning collision avoidance.

How to evaluate the employed believability trade-offs and validate their selection in LOD-based methods? Perceptual studies are an adequate answer to this question. They were previously used in the frame of interactive crowds to evaluate the required level of geometrical human representations [McDonnell et al. 2005], variety of visual aspect [McDonnell et al. 2008], representation of emotional content [McHugh et al. 2010], variety of motion [McDonnell et al. 2009], etc. Such results help the crowd systems designers to focus the available computational resources only where needed in a relatively optimal manner.

Whereas the previously mentioned studies focus on the individual properties of characters (visual aspect and motion), our paper focuses on interactions between characters. The simulation of interactions represents a great proportion of computational resources needed by crowd systems, and can hardly exploit simple tactics to improve performance (e.g., pre-computations [Yersin et al. 2009]). This paper explores the perception of collisions in the aim of solving them only where required.

## 3 Overview

### 3.1 Experiments

Are spectators able to detect collisions in a moving crowd? The factors that may affect the perception of collisions are numerous. We focused on the most relevant ones in the task of designing LOD selection functions: they have to be local and evaluated at low computational cost. We chose factors related to the conditions of observation, such as the distance to camera and the camera angles, or to the local conditions of simulation, such as crowd density. The importance of remaining factors is not excluded, but their role is not studied in this paper. We kept them constant: hardware to display stimuli, animation and rendering techniques, visual aspect of the scene, crowd motion, etc. For that reason, perceptual experiments were based on uniform gray-scale characters to avoid effects of texture and color.

The paper proposes two successive perceptual studies followed by an evaluation experiment. The first study focuses on the perception of collisions between two characters under varying conditions of observation, whereas the role of the visual complexity induced by a crowd motion is addressed during the second study. As numerous true stimuli (collisions) are displayed to participants during the second experiment, we performed a complementary study to distinguish whether false answers were due to a wrong evaluation of the displayed situation or to the saturation of participants' perception-action loop. While the first experiment reveals that participants accurately detect collisions beyond expectations, the second experiment provides information about the influence of crowd density. A LOD selection function was defined from these results to make the perception of collisions difficult in crowds. This function was then validated in evaluation part by using more realistic crowd simulations. To go further, we also showed that in our conditions there is no influence of textures since the collision detection was similar between gray-scale characters and textured ones. Nevertheless, a

dedicated study is required to exhaustively confirm that statement.

## 3.2 Experimental setup



Figure 2: Experimental setup

All the proposed perceptual studies asked participants to detect collisions among characters. They were conducted based on video stimuli. Video resolution was  $1920 \times 1200$  pixels, screen was 24 inches with the same resolution. Participants were asked to click on any collision they detect by rapidly clicking where it occurred using a classical mouse device (mouse setup parameters were kept constant). They were seated in front of a desk and free to move (see Figure 2). Video showed moving crowds with controlled characteristics: we particularly controlled the presence of collisions, their number, as well as their distribution in space and time.

We elaborated a simplistic crowd model to accurately control stimuli content. We generated bidirectional crowd flows as illustrated in Figure 1, Figure 14 and in the companion video: characters follow in both directions linear trajectories aligned with a unique axis. We voluntarily chose this particular case of parallel trajectories because it is the easiest situation for participants to detect collisions. As a result, the LOD selection function proposed in this paper is based on the most challenging conditions. Characters walk at a unique and constant speed ( $1.4m.s^{-1}$ ). 3 DOF completely define the motion of each character: the first is a spatial coordinate that determines where is located the trajectory (e.g., the Y coordinate value if characters walk along the X world axis), the second is a temporal coordinate that determines when the character appears at the bounds of the scene, and finally the third is the walk direction. Once the desired characteristics of a stimulus are defined, the corresponding crowd motion is computed according to an iterative 3-step process: 1) choose the 3 DOF values at random 2) check if the resulting trajectory satisfies the desired characteristics 3) accept or reject the trajectory accordingly and go to step 1. The process is stopped when all the desired characteristics are satisfied (e.g., the desired crowd density is reached). Crowd rendering is OGRE-based. Characters geometry and appearance were identical as illustrated in Figure 4. The walking motion is obtained from a cyclic motion captured locomotion played in loop.

## 4 Pairwise Collision Perception

### 4.1 Objective and Method

Collisions between characters can be detected by spectators under bounded conditions of observation. The objective of the Pairwise Collision Perception experiment is to quantitatively evaluate these bounds and to determine the relative effect of the viewpoint parameters on detection. Pairwise Collision Perception is a forced choice




Figure 3: Illustration of the factors studied in the Pairwise Collision Perception experiment.




Figure 4: Snapshots extracted from video stimuli. They show the parameters used for Pairwise Collision Experiment.

experiment. Stimuli were a set of 3 seconds long videos showing two characters walking in opposite directions along straight and parallel trajectories. In the middle of the sequence, characters reach a minimal distance. The **in-between distance**  $d$  was set according to 5 values in order to generate colliding and collision-free situations.  $d$  is in meters (characters height is  $1.70m$ ). Distance  $d_1 = 0.16m$  provokes a collision between characters' heads,  $d_2 = 0.35$  between their shoulders and  $d_3 = 0.54m$  between their arms and hands only.  $d_4 = 0.81$  and  $d_5 = 1.08m$  allow characters to have a collision-free motion. Resulting situations are illustrated in Figure 4. Remaining factors concern the viewpoint parameters. **Distance to camera** was set using 5 different values. In order to remain invariant to other camera parameters, we express distance to camera as the resulting characters vertical height  $h$ , in pixels, once rendered to the screen using an horizontal camera axis:  $h_1 = 120$ ,  $h_2 = 60$ ,  $h_3 = 30$ ,  $h_4 = 15$  and  $h_5 = 7$  pixels were studied. **Camera tilt angle** was set using the following values:  $a_1 = 0^\circ$  (camera axis is horizontal),  $a_2 = 20^\circ$ ,  $a_3 = 40^\circ$ ,  $a_4 = 60^\circ$  or  $a_5 = 80^\circ$ . Finally, situations were shown under 2 **perspectives**, front or side view, by changing the camera pan angle relatively to the characters' walking trajectories. These conditions resulted into 250 possible combinations and stimuli.

18 participants took part in this experiment (12M, 6F) aged from 23 to 41 years old ( $30.5 \pm 6.2$ , *mean*  $\pm$  *SD*). They saw the 250 video stimuli with 3 repetitions (in total 750) in a random order and

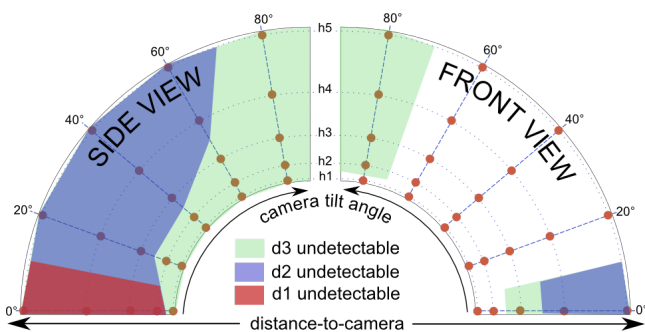
were asked to indicate if there was a collision (see Section 3.2 for technical details).

## 4.2 Results

**Accuracy thresholds** Participants' answers were considered: true positive  $TP$  when a collision occurred ( $d = d_1, d_2$  or  $d_3$ ) and participant clicked, false positive  $FP$  when collision did not occur and participant clicked, true negative  $TN$  when collision did not occur ( $d = d_4$  or  $d_5$ ) and participant did not click, and finally false negative  $FN$  when collision occurred and participant did not click. Accuracy is then defined as follows:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$

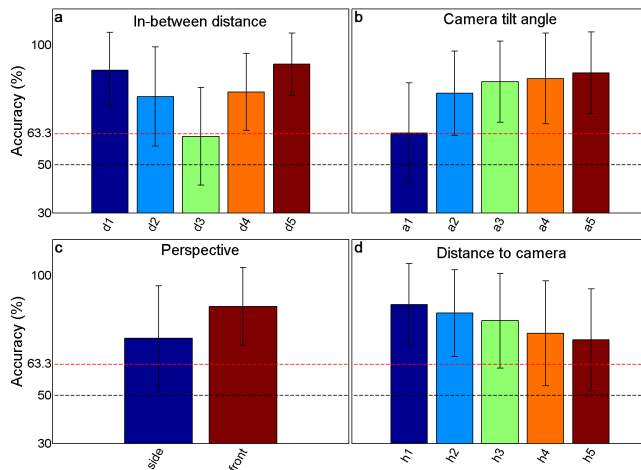
For one given condition, the size of the statistical sample of the measured accuracy is 54 trials (18 participants, 3 repetitions). We thus deduce the minimum accuracy value to conclude that a situation is correctly perceived. This value is 63.3% at 95% confidence (respectively 67.5% at 99%). Such a threshold allows us to determine the limits of accurate collision perception with respect to the three viewpoint factors as illustrated in Figure 5. Colored areas correspond to factors values for which situation is not correctly perceived with respect to the level of interpenetration between characters.



**Figure 5:** Accuracy thresholds with respect to viewpoint parameters. Each colored area contains viewpoints that prevent accurate detection of collision between two characters with respect to interpenetration distance.

**Relative effect of factors** We computed the participants accuracy over the three repetitions for each of the set of conditions (Figure 6) and performed a 4-way analysis of variance (ANOVA) with repeated measures. We performed post-hoc Tukey tests to evaluate the significance of difference between experimental results under changing conditions. Results showed that there was a large effect of the in-between distance factor  $d$  ( $F_{4,64} = 10.9, p < 0.001, \eta^2 = 0.27$ ).  $d_1$  (greatest interpenetration) and  $d_5$  (greatest separation) values significantly increase accuracy. Accuracy was 89.4% and 91.9% under  $d_1$  and  $d_5$  conditions respectively, whereas it was 61.9% in the case of a collision between characters' arms only ( $d = d_3$ ). We remind that randomly answer to stimuli would tend to have sample composed of 50% of true positive answers, and that accuracy threshold is 63.3% according to the sample size. The camera tilt angle has a medium effect on participants' accuracy ( $F_{4,64} = 156.7, p < 0.001, \eta^2 = 0.20$ ): an horizontal camera axis lowers the average accuracy to 63.3% whereas it reaches 88.3% when the tilt angle is set to 80°. Choosing a side or a front view has a medium effect as well ( $F_{1,16} = 53.4, p < 0.001, \eta^2 = 0.11$ ) as accuracy is 73.8% and 87.0% respectively. Distance to camera has a low effect on accuracy ( $F_{4,64} = 60.6, p < 0.001, \eta^2 = 0.07$ ): it

progressively decreases from 87.8 to 73.1% whereas the difference of distance to camera between  $h_1$  and  $h_5$  is high. Finally, analysis revealed that all the considered independent variables have significant interactions with a small effect size.



**Figure 6:** Main effect of the evaluated factors on accuracy (accuracy threshold is 63.3%): (a) in-between distance  $d$ , (b) camera tilt angle  $a$ , (c) front or side view and (d) distance to camera  $h$ .

## 4.3 Discussion

The Pairwise Collision Perception experiment shows that the effect of the distance to camera is less important than expected: participants were accurate in detecting collision even when characters are only 7 pixels high: the average accuracy remains above the 63.3% threshold with an average value of 73.1% upon the  $h_5$  condition. We remind that our goal is to provide a metrics to design crowd simulation LOD selection functions. In this situation, we conclude that distance to camera is not the most prevailing selection parameter according to our results, whereas it is classically used in the systems we are aware of, following the example of graphics rendering LOD selection functions [Niederberger and Gross 2005; Paris et al. 2009].

Figure 6 shows that the dispersion of accuracy is relatively high, the average coefficient of variation is 0.23. Dispersion can be explained by interaction between experimental conditions. Our analysis revealed for example a significant interaction between viewpoint (side or front) and in-between distance  $d$  ( $F_{4,64} = 15.5, p < 0.001, \eta^2 = 0.07$ ). For example, average accuracy is 47.7% under  $d = d_3$  with side view, whereas it reaches 94.7% under  $d = d_1$  with front view. On one hand, variance of results is mainly explained by the effect of independent variables. On the other hand, interactions cannot be neglected: this may make more difficult the calibration of LOD selection function parameters.

There are however promising ways to rapidly stop solving collisions to improve performance. Results show that using eye-level viewpoints with horizontal camera axis (e.g., first-person perspective) is an efficient way to prevent spectators from detecting collisions between characters. Bird's-eye perspectives are the most delicate to handle. Nevertheless, these two types of perspective correspond to different types of applications. Some allow spectators to embody a virtual walker and will implicitly provide eye-level viewpoints: few collisions will have to be solved to obtain believable results. Other applications may require top views to observe the crowd as a whole. In this latter type of viewpoint, a great distance to the ground is required to observe a large proportion of the crowd,



which lowers collision detectability. Finally, the use of side and front perspectives is a promising notion to be introduced as a LOD selection criterion. It plays an important role in collision detection but was not yet considered in previous work as a determinant variable.

Switching from one LOD to another may provoke popping effects that are particularly detectable by spectators. Our results show a strong relation between the level of interpenetration between character bodies and the level of detection of collision by spectators. This effect can be positively exploited to introduce a progressive enabling and disabling of collision solvers, to allow smooth transitions between LODs and prevent from perceptible side effects. Note that most of microscopic simulation models, as the ones cited in Section 2, take characters size into account (generally, as the radius of a bounding circle). This enables smoothing transitions by progressively lowering this parameter value to finally disable collision solving.

## 5 Collision Perception in Dense Situations

Previous studies showed that relaxing the collision avoidance constraint in a crowd simulation cannot be easily achieved in an imperceptible manner by only playing on viewpoint parameters. Nevertheless, the implicit visual complexity of virtual crowds is an interesting path to prevent spectators from detecting collisions. We assume in the Dense Situations experiment that this complexity mainly results from crowd density: we thus focus on relationships between collision detection frequency and crowd density. Concerning viewpoints, we consider the same as in previous study, but focus on those for which participants were still able to accurately detect collisions. This study needs stimuli that contain numerous collisions, up to tens of thousands. We cannot reasonably expect that participants will click on all of them. A preliminary study described in the next section has been done to size the expectable maximum click frequency, and to provide a reference scale to interpret the results of the Dense Situations Experiment.


### 5.1 Preliminary study

#### 5.1.1 Objective and Method

How many obvious collisions participants are able to click on? Our experiment answers this question and searches for the saturation frequency of this perception-action loop. We use stimuli showing a moving crowd (Figure 7a). We compute the characters trajectories so that collisions occur between them. Collisions are uniformly distributed in the screen space, but occur with increasing frequency in time. The time-interval between two collisions linearly decreases from 3.3 seconds to 0.1 second. 33 collisions are shown. Each collision is highlighted by drawing a transparent red circle over the two concerned characters. 26 participants (13M, 13F) aged from 23 to 50 years old ( $29, 7 \pm 7.7$ ) took part in this experiment. They were asked to click as fast and precisely as possible on the detected collision. We prepared 3 different stimuli with these characteristics.

#### 5.1.2 Results


Graph in Figure 7b shows the total number of collisions for which no click by participants was recorded (over a total of  $3 \times 17 = 51$  trials for each of the 33 collisions). All participants were able to click on the displayed collisions until collision #25, for which the time interval between two collisions was 0.9 seconds. After collision #29 (time interval 0.5 seconds), the number of non-clicked collisions significantly increases. Spatial accuracy of participants is quite constant until collision #29: most of participants clicked



**Figure 7:** (a) Partial snapshot of a video stimuli; (b) Number of non-clicked collisions with respect to the collision number.

within the displayed circle and at 25 pixels distance from the circle center on average (circle radius is 41 pixels). The mean delay between the appearance of the highlight circle and the time the mouse click was 1 second (ranging from 0.5 to 1.5 seconds). This delay noticeably starts increasing after collision #30 (time interval 0.4 seconds). In conclusion, we can expect participants to be able to indicate detected collisions with an accuracy below 50 pixels and a delay below 1.5 seconds.


### 5.2 Dense Situations Experiment



**Figure 8:** Illustration of the factors studied in the Dense Situations experiment.

#### 5.2.1 Objective and Method

Does the visual complexity induced by crowd density prevent spectators from detecting collisions? The Collision Perception in Dense Situations experiment answers this question using a perceptual study. As shown in Figure 8, we studied the joint effect of crowd density with previous factors: distance to camera, camera tilt and perspective. Stimuli are 32 seconds long videos displaying crowds of characters as shown in Figure 9. Trajectories were all parallel, followed by characters at same and constant speed but in opposite directions. Crowd was simulated in a standardized  $40 \times 40$  meters area, camera is pointing at its center. **Distance to camera**  $h$  was defined by the resulting height in pixels of a character rendered on the display and located in the middle of the area.  $h_1 = 30$  and  $h_2 = 15$  pixels values were used. **Camera tilt angle** values were set to  $a_1 = 20^\circ$ ,  $a_2 = 40^\circ$ ,  $a_3 = 60^\circ$  and  $a_4 = 80^\circ$ .



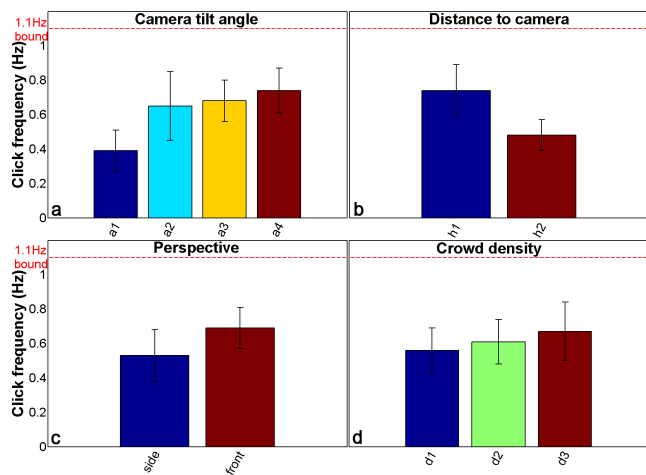
**Figure 9:** Snapshots extracted from video stimuli. They show the parameters used for the Dense Situations Experiment.

Scene was shown under the front and side **perspectives** as defined before. **Crowd density**  $d$  was set to  $d_1 = 0.25$ ,  $d_2 = 0.5$  and  $d_3 = 1 \text{ people.m}^{-2}$ . Examples of resulting stimuli are illustrated in Figure 9.

19 participants took part in this experiment (13M, 6F) aged from 22 to 44 years old ( $28.5 \pm 5.9$ ). They saw once the 48 videos resulting from the factors combination in a random order on a 24 inches wide screen. Seated in front of the screen, they were asked to click on all collisions they could detect.


It is not realistic to expect that participants will click on all the collisions displayed in stimuli as they contain from 2766 up to 45942 collisions in 32 seconds of time. As a result, we based our analysis on collision detection frequency assuming that the more collisions are perceptible, the more frequently participant will click on collisions. Preliminary experiment showed that we can expect a maximum detection frequency of 1.1Hz, this bound is our reference scale to discuss the results of this experiment.

## 5.2.2 Results



**Figure 10:** Main effect of the evaluated parameters on frequency of clicks: (a) camera tilt angle  $a$ , (b) distance to camera  $h$ , (c) front or side views and (d) crowd density  $d$ .

**Respective influence of the studied factors** We performed a 4-way analysis of variance (ANOVA) with repeated measures. The response variable is each participant's clicking frequency averaged over each stimuli. Results show a large effect of the camera angle again ( $F_{3,54} = 5.47$ ,  $p < 0.001$ ,  $\eta^2 = 0.30$ ): whereas the mean detection frequency is 0.4Hz when camera angle is set to  $20^\circ$ , detection rate reaches 0.75Hz at  $80^\circ$ . Distance to camera has also a large effect and explains 28% of the variance of results ( $F_{1,18} = 99,68$ ,  $p < 0.001$ ,  $\eta^2 = 0.28$ ): mean detection frequency decreases from an average of 0.75Hz for the close distance  $h_1$  to 0.5Hz for the far one  $h_2$ . As concerning previous experimental results, the perspective has quite an important role on collision perception ( $F_{1,18} = 19,87$ ,  $p < 0,001$ ,  $\eta^2 = 0.1$ ): detecting collisions in side view is more difficult than in front view. More surprisingly, density has a low effect size ( $F_{2,36} = 5.63$ ,  $p < 0.01$ ,  $\eta^2 = 0.03$ ) as illustrated in Figure 10. We finally observed a first order interaction between distance to camera and density  $h * d$  ( $F_{2,36} = 15,26$ ,  $p < 0,001$ ,  $\eta^2 = 0.08$ ): a short distance to camera combined with a high crowd density eases collision detection. Note that no interaction was found between camera tilt angle and distance to camera  $a * d$  ( $F_{3,54} = 2.56$ ,  $p = 0.063$ ).



**Figure 11:** Spatial distribution of clicks in the simulation floor space. (a) click density for front view group: participants detected collisions in the center of the area; (b) click density for side view group: participants detected collisions in the foreground area.

**Spatial distribution of detected collisions** We analyzed the spatial distribution of participants clicks by projecting the recorded click coordinates from the screen space to the simulation space. Analysis showed that the density of clicks is higher in the foreground and central areas. According to perspective (Figure 11), we show that in the case of front view, detected collisions are concentrated in the center of the simulation area (equivalent to the center of the screen) whereas in the case of side views, clicks are mainly concentrated in the foreground. Again this reveals the difficulty of handling front views: clicks cover a wider area than in the case of side views. Participants mainly focused on the center, but some also focused on the lateral borders of the area: looking carefully at the density plot, one can see an increase of the density in the corresponding zones. Side views deeply change this statement: participants intensively focused on the foreground. Detailed analysis revealed that clicks were also obtained in the center area when diving camera tilt angles are used ( $a = a_3$  and  $a_4$ ).

## 5.2.3 Discussion

The Dense Situations Experiment is based on the assumption that the more collisions are perceptible, the more frequently participant will click on collisions. We first notice that the average measured click frequency is always below the 1.1Hz bound deduced from the Preliminary Experiment. This observation corroborates our hypothesis because it seems that participants responses are not limited by mechanical aspects. Note that this "Click on as many collisions as possible" experiment procedure is a situation in which all the attention of the participants is focused on the task of detecting collisions.

This situation is thus the worst case when trying to hide collisions. We could obtain even better results by making hypotheses about the spectators' level of attention.

The low effect of crowd density on collision perception should be carefully considered. On one hand, increasing density makes the visual aspect of crowds more complex and may increase the difficulty of perceiving collisions. On the other hand, this also increases the total number of collisions as well as their spatial density and frequency. This may tend to ease collisions perception. For example, the total number of collisions in each of the 32 seconds video stimuli is approximately 2800 when  $d = d_1$ , 11200 when  $d = d_2$  and 44000 when  $d = d_3$ . These two effects seem to combine and compensate each other in our study.

Nevertheless, in spite of the visual complexity induced by a high level of crowd density, collisions remain quite easily detectable by spectators when scene is observed from a bird's-eye viewpoint and when view axis is aligned with the flow. However, a more detailed analysis of the spatial distribution of clicks shows that participants tend to follow one given character in the scene and to check when this specific character is entering into collision with other characters. In spite of their physical similarity, characters do not have the same role in the stimuli: their depth position makes them more or less important in the scene and participants tend to follow the foreground characters. However, this notion of foreground and background characters disappears in the case of front views. Then, the relative importance of one given character seems to be due to its centrality in the screen, or to the fact that he walks at the border of the simulation area. Nevertheless, most of the detected collisions were situated in the center of the scene.

## 6 Guidelines and LOD selection function

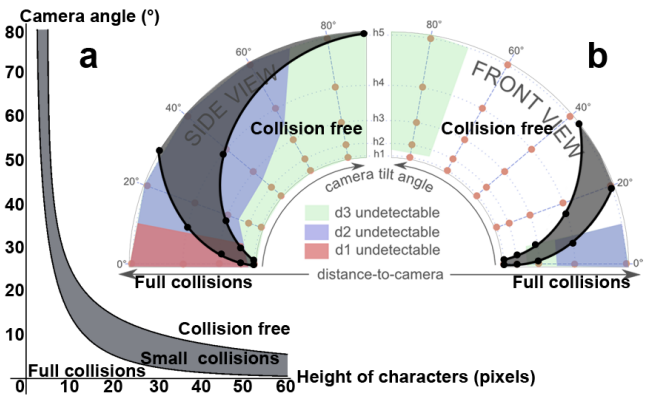
Both the Dense and the Pairwise experiments show how important is the effect of camera tilt angle and distance to camera<sup>1</sup> on collision detection. These parameters are independent (no interaction in our analysis). Thus, we propose a LOD selection function which defines the distance above which collision avoidance can be relaxed given the camera tilt angle under which the scene is perceived. Furthermore, the Dense Situations experiment reveals that crowd density has a low influence on collision detection. Thus, our selection function is parameterized on the quantitative results of the Pairwise experiment (Figure 5).

More precisely, our LOD selection function defines two bounds (Figure 12a): an upper bound under which collision avoidance can be progressively relaxed, and a lower bound under which collisions are imperceptible. As a result, collisions are fully avoided from the camera up to the upper bound, collisions are then progressively allowed between the upper and lower bounds, and no more checked beyond the lower bound. Ideally, with respect to the Pairwise experimental conditions (Figure 5), the upper bound of the LOD selection function should fit the d2 detectability threshold and the lower bound the d1 detectability threshold. We fit these thresholds by using a simple exponential function:

$$f(h_i) = a \cdot e^{-b \cdot (h_i - c)^d} \quad (1)$$

where  $a$ ,  $b$ ,  $c$  and  $d$  are parameters which depend on the considered bound (upper or lower) and on  $p_i$ . Most important parameters are  $c$  and  $d$ .  $c$  represents the character height under which collisions are

<sup>1</sup>As mentioned earlier, height can be transformed into distance to camera once virtual camera parameters and display resolution are known. One given  $h$  value corresponds to one distance to camera (computed for horizontal camera).



**Figure 12:** LOD selection function is based on two bounds: an upper bound under which collision avoidance can be progressively relaxed, and a lower bound under which collisions are imperceptible (a); These bounds are tuned according to the experimental data (b).

no more solved whatever the camera angles.  $d$  defines the lowest acceptable camera tilt angle depending on  $h_i$ : increasing  $d$  means activating collision avoidance for lower camera tilt angles.  $a$  and  $b$  are then jointly tuned to refine the shape of our LOD selection function:  $a$  globally scales the LOD selection function (e.g., radians to degrees) while  $b$  mainly modifies its curvature.

Given the important effect of the perspective condition (side or front view), we define two different sets of parameters values ( $a$ ,  $b$ ,  $c$  and  $d$ ) to fit these two situations. Intermediary perspective angles can be handled by interpolating between these two sets accordingly. For example, if  $p_i = 30^\circ$ , the bounds should first be computed as  $a_{30}^{h_i} = a_0^{h_i} + \frac{30}{90}(a_{90}^{h_i} - a_0^{h_i})$ , etc. Table 1 presents these 2 sets of parameters values for computing lower and upper bounds according to front and side perspectives.

	SIDE		FRONT	
Upper bound	$a_{90}^{h_i} = 200$ $c_{90}^{h_i} = 6$	$b_{90}^{h_i} = 0.8$ $d_{90}^{h_i} = 0.34$	$a_0^{h_i} = 110$ $c_0^{h_i} = 5$	$b_0^{h_i} = 0.9$ $d_0^{h_i} = 0.3$
Lower bound	$a_{90}^{l_0} = 80$ $c_{90}^{l_0} = 5$	$b_{90}^{l_0} = 0.65$ $d_{90}^{l_0} = 0.45$	$a_0^{l_0} = 100$ $c_0^{l_0} = 2.5$	$b_0^{l_0} = 0.7$ $d_0^{l_0} = 0.5$

**Table 1:** LOD selection function parameters value with respect to perspective (front or side view).

Figure 12b shows the superimposition of the LOD selection function on the detection thresholds. For side view, the parameters were tuned to ensure that our LOD selection function is conservative (below the d2 threshold). The computation cost savings obtained in the validation experiments could then be further improved by playing on new simulation-related factors not studied in this paper. For front view, we kept the mathematical definition of the function but changed parameters values. The LOD selection function is less conservative this time since it is above the d2 threshold. However, according to the distribution of clicks shown in Figure 11, the presence of foreground characters should prevent spectators to detect collision as easily as in the conditions of Pairwise experiment.

## 7 Evaluation

We evaluated the proposed LOD selection function in a final experiment. As in previous experiments, we showed to participants stimuli of a bidirectional flow of pedestrians and asked them to click




when they perceive a collision. We however explored more practical situations in comparison with previous setups. The bidirectional flow was simulated using an existing crowd simulation model that was adapted to enable or disable collision solving with smooth transitions. We also animated virtual humans looking more natural on top of simulated trajectories.

## 7.1 Crowd Simulator Adaptation

The LOD selection function proposed in the previous section determines whether collisions should be solved between virtual humans with respect to their position and the user's viewpoint. We successfully adapted two existing crowd simulation models, RVO2 [van den Berg et al. 2008] and Tangent Model [Pétré et al. 2009], to collision solving enabling/disabling. To this end, we played on two elements as illustrated in Figure 13. First is the set of agents  $S_a$  actually simulated by the considered model. By including (resp. excluding) a given agent from  $S_a$ , collision checking and solving is enabled (resp. disabled). Second, to avoid the negative effect of simulation discontinuities, we perform smooth transitions between areas where avoidance is enabled and disabled. When an agent walks from the upper bound to the lower one as defined by the proposed LOD selection function, its personal space is progressively increased from 0 to its nominal value. The nominal value is the personal space the virtual human would normally occupy in a classical simulation. In RVO2 and Tangent Model, this personal space is adjustable with a radius  $r_i$ : nominal radius is denoted  $r_{nom}$ . As a result, a crowd simulation is adapted by following the steps described below:

- for each agent  $i$ , compute  $h_i$ ,  $a_i$  and  $p_i$  (note that, due to the effect of perspective, these angles may change from one to another virtual human according to its position in the screen), compute the corresponding function parameters  $a_{p_i}^{hi/lo}$ ,  $b_{p_i}^{hi/lo}$ ,  $c_{p_i}^{hi/lo}$  and  $d_{p_i}^{hi/lo}$ , and finally the upper and lower bounds  $f_{p_i}^{hi}(h_i)$  and  $f_{p_i}^{lo}(h_i)$ ;
- if  $a_i > f_{p_i}^{hi}$ , set  $i$  into  $S_a$  and set  $r_i = r_{nom}$ ,
- else if  $f_{p_i}^{lo} < a_i < f_{p_i}^{hi}$ , include  $i$  to  $S_a$  and set  $r_i = (a_i - f_{p_i}^{lo}) / (f_{p_i}^{hi} - f_{p_i}^{lo})$ ,
- else if  $a_i < f_{p_i}^{lo}$ , exclude  $i$  from  $S_a$ .
- run one time step of the model for all agents in  $S_a$ .
- steer all agents out of  $S_a$  toward their goal (linear move, no collision check).



**Figure 13:** Adaptation of the crowd simulator model to handle collision LODs.

## 7.2 Evaluation Scenarios

We evaluated the proposed LOD selection function by showing participants 5 different stimuli of one minute long each, some are illustrated in Figure 14. All stimuli were prepared by simulating a bidirectional flow of 5,000 virtual humans with an average density of  $0.5 \text{ people.m}^{-2}$ . The animation of characters along these trajectories was made by using Golæm SDK software. Description is detailed below. Companion video illustrates each of them.

**Scenario #1** is made of 2 stimuli. They show the simulated crowd from the side. LOD selection function is applied and collisions are solved for only part of the crowd. On average, collisions are solved for 57% of virtual humans (22% with full collision avoidance, 35% in transition area, 43% without avoidance). The computation cost is then reduced by 42%. The difference between the 2 stimuli is the physical representation of virtual humans: first stimulus is prepared identically to previous experiments, i.e. with uniform gray-scale characters (Figure 14a), whereas the second stimulus uses 38 virtual humans with more realistic aspect (shape and texture) as shown in Figure 14b. This more realistic representation using textures is also used in scenarios #2 and #3.

**Scenario #1-bis** is a control situation. One stimulus was prepared under the same conditions as Scenario #1 using textured characters, except that the LOD selection function was not applied: the avoidance model is used everywhere in the simulation and, as a result, stimulus is collision free.

**Scenario #2** shows the crowd from front view. Previous experiment showed that this situation is more challenging since collisions are more perceptible. Our LOD selection function still allow to avoid checking collisions for 44% of virtual humans in average (39% with full collision avoidance and 17% in transition). The computation cost is then reduced by 41%.


**Scenario #3** shows the crowd with a  $45^\circ$  angle relatively to the direction of the crowd motion. We here check that, even out of our experimental conditions (only side and front views were tested), our LOD selection function with interpolated parameters still works. The camera tilt angle is relatively flat and collisions are not solved close to the camera: this allows to obtain 68% of virtual humans without collision avoidance (16% with full collision avoidance and 16% in transition). The computation cost is then reduced by 70%.

## 7.3 Results and Discussion

20 participants (10M, 10F), aged from 22 to 59 years old ( $28 \pm 8.9$ ) saw the stimuli and were asked to click on all the collisions they could detect. Wilcoxon Signed-Rank tests were performed to compare our results. We compare the number of clicks in situations with and without collisions to evaluate the efficiency of our LOD selection function to deceive participants. In Scenario #1, no significant difference was observed between the two stimuli ( $W = 21$ ,  $p = 0.25$ ), i.e., the visual aspect of characters has no significant effect here. Among the 20 participants, 8 participants clicked some collisions when the stimulus with gray-scale characters was shown: on the thousands of displayed collisions, only 28 clicks were recorded ( $1.4 \pm 2.5$  clicks per participant). With textured characters, only 7 participants clicked 15 times in total,  $0.75 \pm 1.4$  per participant, when the stimulus was shown. The conclusions of our previous experiments can reasonably be extrapolated to natural-looking crowds with textured characters.

More interestingly, 6 participants clicked when the Scenario #1-bis





**Figure 14:** All scenarios consider a bidirectional crowd flow rendered and perceived under changing conditions. (a, b) Scenario #1 explores the influence of rendering crowd in a more natural way using virtual humans with various representations; (c) Scenario #2 checks LOD selection function under the challenging condition of a frontal view; (d) Scenario #3 provides a situation where our LOD system is particularly efficient.

was shown (without collision), with a total number of 12 clicks ( $0.6 \pm 0.9$ ). No significant difference was revealed between Scenario #1 and #1-bis ( $W = -4, p = 0.84$ ). That means that for a same situation (viewpoint, camera angle, distance to camera, density), participants clicked as many times when there is no collision or when the LOD selection function is activated. Results are similar for other situations. No significant difference was neither found between Scenarios #1-bis and #2 ( $W = -31, p = 0.13$ ) nor between #1-bis and #3 ( $W = -2, p = 0.84$ ). On average, each participant clicked  $1.25 \pm 1.7$  (resp.  $0.7 \pm 1.5$ ) times in Scenario #2 (resp. #3). Moreover, no significant difference was found between Scenarios #1 and #2 ( $W = -27, p = 0.24$ ) or between #1 and #3 ( $W = 13, p = 0.49$ ). Therefore, we can argue that participants globally clicked as many times, whatever the shown stimulus. As a general conclusion, our LOD selection function is efficient whatever the situation.

These results are promising even more since participants were asked to carefully look for collisions. In total, very few collisions could be detected by participants, whereas thousands were visible. Our function to select when and where to solve collisions showed its efficiency to remove these constraints in an imperceptible way. Using average crowd density and classical field-of-view camera setup, we were able to drastically limit the number of virtual humans that require to satisfy collision avoidance constraints to few thousands, which meets the real-time performances of available crowd simulator. We successfully adapted our technique to two existing models.

## 8 General Discussion

Previous section shows promising results. While we push the proposed guidelines to their limits, we confirm the efficiency of the factors revealed to be the most relevant according to our experimental results (such as camera tilt angle). We focused our studies on some factors that can be rapidly evaluated at a local scale in the aim of applying our results to the design of LOD selection functions. However, experimental limits prevented us to explore many other factors. Experiments also required standardized situations. The possible extension of protocols draws up future works directions. Some are discussed in the sections below.

**Trajectory Parallelism** Our work focused on bidirectional linear crowd flows. This situation is meaningful (e.g., crowds in corridors or streets) but is also the worst case study as collision detection is

eased. Anticipation plays a great role in human perception. When two characters move closer with parallel trajectories, the distance between their parallel paths determines the existence of a collision. A spatial estimation of this distance is enough to anticipate the situation. On the contrary, in the case of secant trajectories, characters have to reach the followed paths intersection approximately at the same time. A spatio-temporal estimation is then required. While experiments revealed the importance of the effect of the camera-trajectory angle, more complex situations of secant trajectories perceived under various angles need to be addressed. A greater difficulty to perceive collisions can be expected, which makes this factor an important one to consider.

**Characters uniformity** We used gray-scale characters in the Pairwise and Dense Situation experiments to avoid effects of texture and color on collision perception. To go further, our evaluation scenarios were prepared using textured and varied characters. In the conditions proposed by our scenarios, results showed that there is no significant effect on collision detection compared to standardized gray-scale characters. However, the impact of adding variety to crowds (colors, textures, motion or shape) on perception of collision should be more carefully discussed. For example, introducing appearance variety may have two possible and contradictory effects. On one hand, visual complexity of the crowd scene is globally increased which may lower collision detectability, but on the other hand, some characters may become prominent. The prominence of some characters can however be positively exploited to estimate spectators visual attention and adapt LODs distribution accordingly. A combination with visual attention models [Hillaire et al. 2010] should then be considered.

**Local steering methods and Macroscopic crowd simulation models** Our LOD selection function ensures a smooth transition between areas with or without collision avoidance. Despite this essential transition, each of those areas has distinct individual steering methods. The animation of characters in both areas should then be coherent to have a believable crowd. First, at the local scale, collision avoidance results in individual manoeuvres which introduce some motion jerkiness. Should we preserve such motion attributes? Answering also requires an evaluation. Nevertheless, reproducing motion jerkiness when collision avoidance is disabled asks for extending steering methods accordingly. Second, the absence of collision avoidance may result in linear motion with unbelievable spatial distribution of characters. Macroscopic crowd simulation models answer this problem as some provide crowd probability density functions as an output [Maury et al. 2010].

**Crowd Rendering and Display** Crowd can be rendered by using 3D characters, such as in our experiments, or impostors [Dobbyn et al. 2005]. The notion of collision, interpenetration and depth order can be different when using impostors. Indeed, impostors cannot interpenetrate each other. A collision between two impostors perceived from side is then difficult to detect. However, in front view, it results in a sudden change of their depth order. Such an instantaneous change generates high-frequency visual artifacts that emphasize collisions. Specific perceptual experiments based on impostors should then be considered. Nevertheless, the development of 3D screens (that change the perception of distances and in particular depths) encourages following perceptual experiments on 3D characters.

## 9 Conclusion

In this paper, we conducted two perceptual studies in order to determine the conditions for collisions to be perceived by spectators in

the context of interactive virtual crowds. First experiment considered situations between two characters out of the context of a crowd and inspected the effect of the observation conditions on the perception of collisions. Results enabled us to bound the conditions of a second experiment where the role of the visual complexity induced by crowds was jointly examined. Non-trivial results from these two experiments revealed that crowd density influence is less important than expected and a LOD selection function cannot be based only on distance to camera. These results should encourage to revisit the existing LOD selection functions for real-time crowds. Meanwhile, the importance of other factors not yet considered emerged, providing new paths to improve these functions: in particular, the camera tilt angle and the angle under which a crowd flow is perceived. We drew up the corresponding guidelines and the resulting LOD selection function and provided an evaluation to ensure the strength of the proposed directions. This evaluation required to adapt existing crowd simulation models: we demonstrated that this adaptation can be done easily using distance between characters. We could generate video sequences showing thousands of visible collisions that could hardly ever be detected by several people carefully looking for them. We aimed our work to improve the performance of interactive virtual crowd systems since the computation cost of our evaluation scenarios for example was reduced up to 70%. The importance of such an application is continuously reinforced as the size of common virtual environments (e.g., Google Earth, City Engine, etc.) is always growing and will not remain empty of population for much longer.

## Acknowledgements

We wish to thank the reviewers for their useful comments. This research has been supported by the European project TANGO.

## References

- BURSTEDDE, C., KLAUCK, K., SCHADSCHNEIDER, A., AND ZITTARTZ, J. 2001. Simulation of pedestrian dynamics using a two-dimensional cellular automaton. *Physica A: Statistical Mechanics and its Applications* 295, 3-4, 507 – 525.
- DOBBYN, S., HAMILL, J., O'CONNOR, K., AND O'SULLIVAN, C. 2005. Geopostors: a real-time geometry / impostor crowd rendering system. *Proc. Symposium on Interactive 3D Graphics and Games*.
- HELBING, D., AND MOLNAR, P. 1995. Social force model for pedestrian dynamics. *Physical Review E* 51, 4282.
- HILLAIRE, S., BRETON, G., OUARTI, N., COZOT, R., AND LÉCUYER, A. 2010. Using a Visual Attention Model to Improve Gaze Tracking Systems in Interactive 3D Applications. *Computer Graphics Forum* 19 (6), 1830–1841.
- KAPADIA, M., SINGH, S., ALLEN, B., REINMAN, G., AND FALOUTSOS, P. 2009. Steerbug: an interactive framework for specifying and detecting steering behaviors. *Proc. ACM SIGGRAPH/Eurographics Symposium on Computer Animation*.
- KISTLER, F., WISSNER, M., AND ANDRÉ, E. 2010. Level of detail based behavior control for virtual characters. In *Intelligent Virtual Agents*, vol. 6356 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 118–124.
- MAURY, B., ROUDNEFF CHUPIN, A., AND SANTAMBROGIO, F. 2010. A macroscopic crowd motion model of gradient flow type. *Mathematical Models and Methods in Applied Sciences* 20, 10, 1787–1821.
- MCDONNELL, R., DOBBYN, S., AND O'SULLIVAN, C. 2005. LOD Human Representations: A Comparative Study. In *Int. Workshop on Crowd Simulation (V-CROWDS'05)*, 101 – 115.
- MCDONNELL, R., LARKIN, M., DOBBYN, S., COLLINS, S., AND O'SULLIVAN, C. 2008. Clone attack! perception of crowd variety. *ACM Transactions on Graphics* 25 (3).
- MCDONNELL, R., LARKIN, M., HERNÁNDEZ, B., RUDOMIN, I., AND O'SULLIVAN, C. 2009. Eye-catching crowds: saliency based selective variation. *ACM Trans. Graph.* 28, 3, 1–10.
- MCHUGH, J., MCDONNELL, R., NEWELL, F. N., AND O'SULLIVAN, C. 2010. Perceiving emotion in crowds: the role of dynamic body postures on the perception of emotion in crowded scenes. *Experimental Brain Research* 204 (3), 361–372.
- NARAIN, R., GOLAS, A., CURTIS, S., AND LIN, M. 2009. Aggregate dynamics for dense crowd simulation. In *SIGGRAPH Asia '09: ACM SIGGRAPH Asia 2009 papers*.
- NIEDERBERGER, C., AND GROSS, M. 2005. Level-of-detail for cognitive real-time characters. *Visual Computer* 21, 188–202.
- PARIS, S., PETTRÉ, J., AND DONIKIAN, S. 2007. Pedestrian reactive navigation for crowd simulation: a predictive approach. *Eurographics'07: Computer Graphics Forum* 26, (3), 665–674.
- PARIS, S., GERDELAN, A., AND O'SULLIVAN, C. 2009. Calod: Collision avoidance level of detail for scalable, controllable crowds. *Proc. International Workshop on Motion in Games*.
- PETTRÉ, J., CIECHOMSKI, P. D. H., MAÏM, J., YERSIN, B., LAUMOND, J.-P., AND THALMANN, D. 2006. Real-time navigating crowds: scalable simulation and rendering: Casa 2006 research articles. *Comput. Animat. Virtual Worlds* 17, 445–455.
- PETTRÉ, J., ONDŘEJ, J., OLIVIER, A.-H., CRÉTUAL, A., AND DONIKIAN, S. 2009. Experiment-based modeling, simulation and validation of interactions between virtual walkers. *Proc. ACM SIGGRAPH/Eurographics Symp. on Computer Animation*.
- REYNOLDS, C. W. 1987. Flocks, herds and schools: A distributed behavioral model. In *SIGGRAPH '87: Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, 25–34.
- SAMET, H. 2005. *Foundations of Multidimensional and Metric Data Structures*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- SUD, A., ANDERSEN, E., CURTIS, S., LIN, M., AND MANOCHA, D. 2008. Real-time path planning for virtual agents in dynamic environments. *ACM SIGGRAPH 2008 classes*.
- TECCHIA, F., LOSCOS, C., AND CHRYSANTHOU, Y. 2002. Image-based crowd rendering. *IEEE Comput. Graph. Appl.* 22 (March), 36–43.
- TREUILLE, A., COOPER, S., AND POPOVIĆ, Z. 2006. Continuum crowds. *ACM Trans. on Graphics (SIGGRAPH 2006)* 25 (3).
- VAN DEN BERG, J., PATIL, S., SEWALL, J., MANOCHA, D., AND LIN, M. 2008. Interactive navigation of individual agents in crowded environments. *Symposium on Interactive 3D Graphics and Games (I3D 2008)*.
- YERSIN, B., MAÏM, J., PETTRÉ, J., AND THALMANN, D. 2009. Crowd patches: populating large-scale virtual environments for real-time applications. *Proc. Symposium on Interactive 3D Graphics and games*, 207–214.