

A Synthetic-Vision Based Steering Approach for Crowd Simulation

Jan Ondřej*
INRIA

Julien Pettré*
INRIA

Anne-Hélène Olivier*
INRIA

Stéphane Donikian*
INRIA
Golaem S.A.




Figure 1: Animations resulting from our simulations. Emergent self-organized patterns appear in real crowds of walkers. Our simulations display similar effects by proposing an optic flow-based approach for steering walkers inspired by cognitive science work on human locomotion. Compared to previous approaches, our model improves such an emergence as well as the global efficiency of walkers traffic. We thus enhance the overall believability of animations by avoiding improbable locking situations.

Abstract

In the everyday exercise of controlling their locomotion, humans rely on their optic flow of the perceived environment to achieve collision-free navigation. In crowds, in spite of the complexity of the environment made of numerous obstacles, humans demonstrate remarkable capacities in avoiding collisions. Cognitive science work on human locomotion states that relatively succinct information is extracted from the optic flow to achieve safe locomotion. In this paper, we explore a novel vision-based approach of collision avoidance between walkers that fits the requirements of interactive crowd simulation. By simulating humans based on cognitive science results, we detect future collisions as well as the level of danger from visual stimuli. The motor-response is twofold: a reorientation strategy prevents future collision, whereas a deceleration strategy prevents imminent collisions. Several examples of our simulation results show that the emergence of self-organized patterns of walkers is reinforced using our approach. The emergent phenomena are visually appealing. More importantly, they improve the overall efficiency of the walkers' traffic and avoid improbable locking situations.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation I.6.5 [Simulation and Modeling]: Types of Simulation—Animation

Keywords: crowd simulation, steering method, collision avoidance, synthetic vision

1 Introduction

Crowd simulation research has significantly grown in importance these past two decades. Its field of application is wide and ranges from security and architecture domains to film and interactive entertainment industries. The visually impressive self-organized patterns that emerge at a large-scale from the combination of all local actions and interactions in crowds is probably a major reason for the level of attention paid in Computer Animation on this topic. Reynolds' seminal work on flocks of *boids* showed fascinating global motions obtained from simple local interaction rules [Reynolds 1987]. However, the proposed rules explicitly stick boids together to obtain emerging flocks. Moreover, boids motion rules are not directly transposable to human walkers.

Human crowds are a place of numerous and various interactions. In this paper, we focus on crowds of *individually walking humans* where interactions are limited to *collision avoidance*. Our motivation is to design a local collision avoidance method as similar to the real human behavior as possible, while displaying emerging self-organized patterns as witnessed in real crowds. This objective is representative of our bottom-up approach: specific large-scale formations are expected from realistic local interactions between walkers. Simulating emerging formations is crucial in order to obtain believable crowd animations. Obtaining them from individually steered walkers avoiding each-other, and thus, simulating self-organization, is particularly challenging.

ACM Reference Format

Ondřej, J., Pettré, J., Olivier, A., Donikian, S. 2010. A Synthetic-Vision Based Steering Approach for Crowd Simulation. *ACM Trans. Graph.* 29, 4, Article 123 (July 2010), 9 pages. DOI = 10.1145/1778765.1778860 <http://doi.acm.org/10.1145/1778765.1778860>.

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.
© 2010 ACM 0730-0301/2010/07-ART 123 \$10.00 DOI 10.1145/1778765.1778860
<http://doi.acm.org/10.1145/1778765.1778860>

*e-mail: {jan.ondrej,julien.pettré,anne-helene.olivier,donikian}@irisa.fr

Collision avoidance has recently received much attention. Several types of approaches were proposed (cf. Section 2 for an overview). Most recent *agent-based* techniques are based on geometrical models. Their common point is to explicitly compute admissible velocities that allow avoiding future collisions: efforts are focused on reaching the highest performance in order to handle large crowds. Then, the challenge is to steer walkers with believable trajectories, while remaining within an admissible velocity domain. However, geometrical models are disconnected from reality since humans unconsciously react to perceived obstacles to avoid collisions. This raises a fundamental question: can simpler perception/action control loops - probably closer to reality - steer virtual walkers and allow them to avoid collisions, even in complex situations? *Rule-based* techniques explored such a question. However, artifacts occur in most complex situations because of the difficulty in combining rules. *Particle-systems* and *continuum-based* methods ease the combination of interactions and are able to handle even larger crowds. However, they have drawbacks as well. The former sometimes fails to simulate emerging patterns of walkers while the latter may lead to unrealistic local motions, such as unfeasible accelerations or velocities.

In contrast with previous approaches, we steer walkers according to the visual perception they have of their own environment. Thus we formulate our collision avoidance solution as a visual-stimuli/motor-response control law. Our model is inspired by the work of Cutting and colleagues [1995] on human locomotion in the field of cognitive science. They stated that humans extract two major elements from their optic flow to achieve collision-free navigation. First is the derivative of the bearing angle under which obstacles are perceived. Second is the time-to-collision deduced from the rate of growth of obstacles in successively perceived images. Inspired by these observations, our model's inputs, i.e. the visual stimuli, are egocentrically perceived obstacles transformed into images of time-derivatives of bearing angles and of times-to-collision. These images are directly computed from the geometries and states of both the static and moving obstacles of the scene. Walkers have simple reactions to these stimuli: they turn to avoid future collisions and decelerate in the case of imminent collisions.

Our contributions are thus the following. We propose a vision-based collision avoidance model for the interactive simulation of crowds of individual humans. We base our approach on cognitive science work on human locomotion. We apply our method to complex situations of interaction: resulting simulations display the emergence of interesting self-organized patterns of walkers at a global scale. We demonstrate our improvements in comparison to previous approaches, with an enhanced emergence of patterns of walkers, improved global efficiency of the walkers traffic and smoother animations.

The remainder of the paper is organized as follows. Section 2 first provides an overview of crowd simulation techniques with particular focus on collision avoidance methods. Following, we present the guiding principles of our approach before describing the proposed model in detail in Section 3. We provide details about its implementation in Section 4. Finally, we illustrate simulation results from several examples and compare our technique with previous work in Section 5. Limitations of our approach and future work are discussed in Section 6, before concluding.

2 Related Work

Virtual Crowds is a wide topic that raises numerous problems including population design, control, simulation and rendering. It was surveyed in recent books [Thalmann and Raupp Musse 2007; Pelechano et al. 2008] and tutorials [Thalmann et al. 2005; Halperin

et al. 2009]. This overview focuses on crowd simulation, whose objective can be restrictively defined as computing global locomotion trajectories to achieve goal-driven collision-free navigation for crowds of walkers.

Several classes of solutions were proposed in the literature. Cellular-automaton approaches [Schadschneider 2001] are used to simulate evacuation scenarios for large crowds: the discrete aspect of resulting trajectories prevent their use for Computer Animation applications. However, grid-based solutions were adapted to meet such requirements [Loscos et al. 2003]. In addition Shao and Terzopoulos [2005] proposed the use of multi-resolution grids to handle large environments. Other techniques consider velocity-fields to guide crowds [Chenney 2004]. This analogy with Physics gave rise to particle-system approaches. Helbing [1995] proposed the social forces model where walkers repulse each other while they are attracted by their goal. The social forces model was later revisited in [Pelechano et al. 2007; Gayle et al. 2009]. Evolved models were proposed using mass-damp-spring systems to compute similar repulsing forces between walkers [Heigeas et al. 2003]. Crowd simulation was also studied as a flowing continuum [Hughes 2003; Treuille et al. 2006] that allows simulating numerous walkers in real-time. Larger crowds were handled using a hybrid continuum-based approach [Narain et al. 2009]. From a general point-of-view, high computation performance is a common point between all of these approaches. Such performance allows simulating large crowds in equally large environments in real-time, which is a crucial need of many interactive applications. Performance is however obtained at the cost of some limitations, such as restricting the total number of goals walkers can have, or using simplistic interaction models that may lower the results realism. Compared to this former set of approaches, our first objective is not to reach a high-performance solution but to simulate local interactions in a realistic manner. By realism, we here mean that we reproduce the human vision-based locomotion control in order to steer walkers in crowds. Synthetic vision raises numerous computations by nature.

Our method can be closely related to rule-based approaches [Reynolds 1999], as well as, geometrically-based local avoidance models approaches [Paris et al. 2007; van den Berg et al. 2008; Kapadia et al. 2009; Pettré et al. 2009; Karamouzas et al. 2009; Guy et al. 2009]. It is generally required to combine local approaches with dedicated techniques in order to reach high-level goals in complex environments [Lamarche and Donikian 2004; Paris et al. 2006; Pettré et al. 2006; Sud et al. 2007]. Nevertheless, geometrically-based avoidance models carefully check the absence of future collisions locally, given the simulation state. This goal is generally achieved by decomposing the reachable velocity-space of each walker into two components: the inadmissible velocity domain and the admissible velocity domain. These domains respectively correspond to velocities leading to collisions and those allowing avoidance. In the opposite, our method makes walkers react to some situations without explicitly computing the admissibility of their motion adaptations. This raises a fundamental question: can explicit collision checks guarantee the absence of residual collisions? We argue the answer is negative. The reason is twofold. First, the admissible velocity domain is computed assuming that the velocity of moving obstacles remains constant. Second, the admissible velocity domain is often made of several independent components, especially in the case of complex interactions - i.e., during simultaneous interactions with several obstacles. Some of these components degenerate in time since moving obstacles may also adapt their own motion. If the current velocity of a given walker is subject to such a degenerative component, switching to another component is required. As a result, traversing the inadmissible velocity domain is required when acceleration is bounded, whereas unbounded accelerations

result into unrealistic motions. Our method does not explicitly check collisions and is not exempt from failure. However, we believe the proposed visual-stimuli/motor-response laws better imitate the most basic level of real human locomotion control.

We previously addressed the question of realism of simulated locomotion trajectories during collision avoidance in [Petré et al. 2009]. We provide a qualitative description of such trajectories: we experimentally show that real humans anticipate avoidance as no further adaptation is required some seconds before walkers pass at close distance. We also show that avoidance is a role-dependent behavior, as the walker passing first makes noticeably less adaptations than the one giving way. We discuss the visual information humans may exploit to be able to achieve avoidance in such a manner. However, we propose a *geometrical model* to reproduce such trajectories calibrated with our experimental dataset. Compared to this work, we address two new problems. First, we address the question of combining interactions. We explore synthetic vision as a solution to implicitly combine them: they are integrated by projection to the perception image, they are filtered when obstacles are invisible, they are weighted by the importance obstacles have in the image. Second, we directly base our motion control laws on the visual information believed to be exploited by real humans.

To the best of our knowledge synthetic-vision based methods were never used to tackle crowd simulation problems, with the exception of Massive software agents [Massive] which are provided with synthetic vision. However, users are responsible for controlling walkers from such an input. Nevertheless, synthetic vision was used to steer a single or few virtual humans [Noser et al. 1995; Kuffner and Latombe 1999; Peters and O’Sullivan 2003] or artificial creatures [Tu and Terzopoulos 1994]. Reynolds’ boids were also recently provided with visual perception abilities [Silva et al. 2009]. Our approach explores a new type of visual stimuli to control locomotion, based on statements from cognitive science. We also improve performance to fit the requirements of interactive crowd simulation. Finally, visual servoing is an active topic in the field of Robotics [Chaumette and Hutchinson 2006]. Major challenges are the process of optic flows acquired with physical systems and the extraction of the relevant information that allows to steer robots. In contrast to this field, we do not process digitally computed images but directly compute the required visual inputs of our model.

3 Vision-based collision avoidance

3.1 Model overview




Figure 2: The bearing angle and its time-derivative, respectively α and $\dot{\alpha}$, allow detecting future collisions. From the perspective of an observer (the walker at the bottom), a collision is predicted when α remains constant in time. **(left)** $\alpha < 0$ and $\dot{\alpha} > 0$: the two walkers will not collide and the observer will give way. **(center)** the bearing angle is constant ($\dot{\alpha} = 0$). The two walkers will collide. **(right)** $\alpha < 0$ and $\dot{\alpha} < 0$: the two walkers will not collide and the observer will pass first.

Humans control their locomotion from their vision [Warren and

Fajen 2004]. According to Cutting and colleagues [Cutting et al. 1995] humans successively answer two questions during interactions with static and moving obstacles: will a collision occur? When will a collision occur? Cutting experimentally observed that these two questions are answered by extracting two indicators from the perceived optic flow:

1. *Will a collision occur?* Humans visually perceive obstacles under a given angle referred to as the *bearing angle* (noted α). A collision is predicted when the time derivative of the bearing angle, $\dot{\alpha}$, is zero (or close to zero because of the body envelopes). This observation is illustrated in Figure 2 from the 3 examples of two walkers displaying converging trajectories.
2. *When will a collision occur?* Humans visually perceive obstacles with given sizes. The rate-of-growth of obstacles in time allow humans to detect obstacles coming toward them when positive. Moreover, the higher the rate, the more imminent the collision. As a result, humans are able to evaluate the time-to-collision (*ttc*).

Therefore, the relevant information necessary to achieve collision-free locomotion according to Cutting is entirely described by the pair $(\dot{\alpha}, ttc)$. It is interesting to note that humans use similar information to intercept mobile targets as described by Tresilian in [Tresilian 1994].

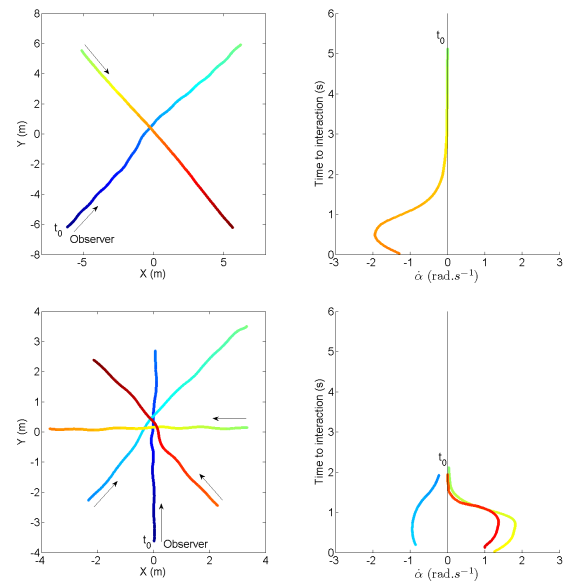


Figure 3: Two examples of real interactions between (top) two walkers and (bottom) four walkers. Motion captured trajectories projected on the ground are shown (plots on the left), as well as in the $(\dot{\alpha}, tti)$ -space (plots on the right), as perceived by one of the participant called ‘observer’. Trajectories are colored in order to enable matching between the two representations.

Figure 3 illustrates Cutting’s theory from 2 examples of real interactions: trajectories are displayed in the horizontal plane as well as in the $(\dot{\alpha}, tti)$ -space, where *tti* is the time-to-interaction. Time-to-interaction is the time remaining before the minimum distance between participants reached, according to current positions and velocities. The notion of time-to-collision *ttc* is generally used in the literature instead of our time-to-interaction *tti*; these two notions are close. By definition *ttc* exists if, and only if, a risk of future collision also exists. Conversely, *tti* exists whatever the relative positions and velocities of moving objects. Also note that

$t t i$ can reach negative values when objects display diverging motions. In the first example, we observe that $\dot{\alpha}$ is initially close to zero whilst $t t i$ decreases: collision is predicted. By turning to the left, the observer solves the interaction: $\dot{\alpha}$ decreases. On the second example, future collision with the observer is predicted for two walkers among the three perceived ones. By turning and decelerating, $\dot{\alpha}$ values are corrected. The impact of motion adaptations on the variations of $(\dot{\alpha}, t t i)$ is not intuitive. However, as a first approximation, turns mainly influence the $\dot{\alpha}$ value, whereas a deceleration mainly changes $t t i$.

The guiding principles of the proposed model - based on Cutting's results - are thus the following. A walker perceives the static and moving obstacles of his environment as a set of points $P = \{p_i\}$ resulting from his synthetic vision. For each perceived point p_i , we compute the bearing angle α_i , its time-derivative $\dot{\alpha}_i$ and the remaining time-to-interaction relatively to the walker $t t i_i$. We deduce the risk of a future collision from $\dot{\alpha}_i$, and the dangerousness of the situation from $t t i_i$. A walker reacts when needed according to two strategies. First, he avoids future collisions by adapting his orientation with anticipation. Second, in the case of an imminent collision, he decelerates until he gets stopped or the interaction is solved. The following sections detail how we put these principles into practice.

3.2 Model inputs




Figure 4: Model's inputs. Any point is perceived under a given bearing angle. The triad $(\alpha_i, \dot{\alpha}_i, t t i_i)$ is deduced from the relative point position and velocity with respect to the walker.

A walker configuration is defined by its position and orientation θ . He has an angular velocity $\dot{\theta}$ and a tangential velocity v . Perceived points $p_i \in P$ may indiscriminately belong to static obstacles - such as walls - or moving ones - such as other walkers. Also note that a single obstacle results in several points with respect to its shape: Figure 6 illustrates how a walker perceives his environment. The variables associated to each $p_i \rightarrow (\alpha_i, \dot{\alpha}_i, t t i_i)$ are deduced from the relative position and velocity of p_i to the walker. We detail their computation in Figure 4 as well as in the following Implementation in Section 4.

3.3 Angular velocity control

As explained in the previous section, a walker detects a risk of future collision when $\dot{\alpha}$ is low and $t t i > 0$. We define the $\dot{\alpha}_i$ threshold τ_1 under which a walker reacts as a function of the perceived

$t t i_i$ as follows:

$$\tau_1(t t i) = \begin{cases} \tau_{1-}(t t i) = a - b \cdot t t i^{-c} & \text{if } \dot{\alpha}_i < 0, \\ \tau_{1+}(t t i) = a + b \cdot t t i^{-c} & \text{otherwise.} \end{cases} \quad (1)$$

where a , b and c are some parameters of the model. These three parameters change a walker's avoidance behavior by adapting his anticipation time, as well as, the security distance he maintains with obstacles. We detail the role of these parameters in the Discussion Section 6. Figure 5 plots the function τ_1 for $a = 0$, $b = 0.6$ and $c = 1.5$. These values were used in the examples shown in Section 5 and were determined by manually fitting τ_1 on a large body of experimental data capturing avoidance between real walkers similar to those shown in Figure 3. Then, the set P_{col} of points $p_i(\dot{\alpha}_i, t t i_i)$ a walker has to react to is defined as follows:

$$p_i \in P_{col} \text{ if } t t i_i > 0 \text{ and } \alpha_i < \tau_1(t t i_i) \quad (2)$$

We now combine the influence of the set of points belonging to P_{col} . For this purpose, we decompose P_{col} into P_+ and P_- , which respectively correspond to points with positive and negative $\dot{\alpha}_i$ values. We then define ϕ_+ and ϕ_- as follows:

$$\phi_+ = \min(\dot{\alpha}_i - \tau_{1+}(t t i_i)), \quad \text{for all } p_i \in P_+ \quad (3)$$

$$\phi_- = \max(\dot{\alpha}_j - \tau_{1-}(t t i_j)), \quad \text{for all } p_j \in P_- \quad (4)$$

At this point, we have identified all interactions requiring walkers to turn to the right to avoid future collision into P_+ , and those asking to turn left into P_- . The required amplitude of a right turn allowing to avoid at once all the interactions provoked by the P_+ set of points directly depends on the amplitude of ϕ_+ (respectively P_- and ϕ_- for a left turn). However, we must ensure walkers do not widely deviate from their goal. For this reason, we now consider the bearing-angle corresponding to the goal α_g , as well as its time-derivative $\dot{\alpha}_g$. Contrarily to obstacles, walkers attempt to intercept their goal, which means that $\dot{\alpha}_g = 0$ is desired. Three cases are then successively considered. Firstly, when $\dot{\alpha}_g$ is small (we arbitrarily choose $|\dot{\alpha}_g| < 0.1 \text{ rad.s}^{-1}$), walkers are currently heading to their goal, hence its influence is neglected. In this case, we simply choose the change of direction which asks for the minimum deviation. $\dot{\theta}$ is controlled as follows:

$$\dot{\theta} = \begin{cases} \phi_+ & \text{if } |\phi_+| < |\phi_-|, \\ \phi_- & \text{otherwise.} \end{cases} \quad (5)$$

Secondly, when $\phi_- < \dot{\alpha}_g < \phi_+$, it cannot be neglected. We choose the change of direction that leads to the smallest deviation from the goal. Then,

$$\dot{\theta} = \begin{cases} \phi_+ & \text{if } |\phi_+ - \dot{\alpha}_g| < |\phi_- - \dot{\alpha}_g|, \\ \phi_- & \text{otherwise.} \end{cases} \quad (6)$$

Thirdly, when $\dot{\alpha}_g < \phi_-$ or $\dot{\alpha}_g > \phi_+$ we choose:

$$\dot{\theta} = \dot{\alpha}_g \quad (7)$$

To avoid unrealistic angular velocities, $\dot{\theta}$ and $\ddot{\theta}$ are finally bounded so that $|\dot{\theta}| < \pi/2 (\text{rad.s}^{-1})$ and $|\ddot{\theta}| < \pi/2 (\text{rad.s}^{-2})$.

3.4 Tangential velocity control

Tangential velocity v is set to the comfort velocity v_{comf} by default. It is only changed in the case of a risk of imminent collision. The imminence of a collision is detected when $t t i_i$ is positive but lower than a threshold τ_2 (we arbitrarily choose $\tau_2 = 3s$). Tangential




Figure 5: τ_1 plot using the following parameter set: $a = 0$, $b = 0.6$ and $c = 1.5$ (cf. Equation (1)). Future collision is detected when $p_i(\hat{\alpha}_i, tti_i)$ is below τ_1 and $tti_i > 0$. The plot also illustrates that the lower the tti_i value, the higher the walker's reaction.

velocity is controlled from the minimum positive tti_{mp} value perceived by the walker. We define P_{pos} as the set of points $p_i \in P_{col}$ for which $tti_i < \tau_2$ and compute tti_{mp} as follows:

$$tti_{mp} = \min(tti_i) \text{ for all } p_i \in P_{pos} \quad (8)$$

Finally, the walker's tangential velocity is controlled as follows:

$$v = \begin{cases} v_{comf} & \text{if } P_{pos} = \emptyset, \\ v_{comf} \cdot (1 - e^{-0.5tti_{mp}^2}) & \text{otherwise.} \end{cases} \quad (9)$$

The position and orientation of the walker are finally updated according to the computed v and $\hat{\theta}$ values, with bounded \dot{v} ($|\dot{v}| < 1m.s^{-2}$).

4 Implementation




Figure 6: Walkers perceive environment obstacles as a set of points $p_i(\hat{\alpha}_i, tti_i)$. The image corresponding to all the perceived $\hat{\alpha}_i$ values of p_i is shown top-left (red are for lowest values). The image corresponding to all the perceived tti_i values of p_i is shown top-right (red are for lowest values). Perception is combined (bottom image) to compute walker reaction. In this example - which corresponds to the circle example, cf. Section 5 - the walker will react to the most red points of the combined perception. In this particular situation, he is likely to follow the walker in front of him on his right.

We implemented our model using OpenGL, shader programming language and CUDA. The algorithm is decomposed into two major stages. First, for each virtual walker:

- Step 1** Set the camera position and orientation at the one of the considered walker (see details below).
 - Step 2** Render to texture environment obstacles using simplified geometries. Compute values $\hat{\alpha}_i$, $\hat{\alpha}_i$ and distance to obstacle d per vertex (Figures 4 and 6).
 - Step 3** Then, using a fragment shader, compute per pixel tti_i , build P_+ and P_- from τ_{1+} and τ_{1-} .
 - Step 4** Copy the resulting texture to the CUDA space and make a parallel reduction to compute ϕ_+ , ϕ_- . The result is stored to an array on the GPU.
- At the end of this first loop, the resulting array is downloaded once to the CPU. Then, for each walker again:
- Step 5** Compute $\hat{\alpha}_g$ and deduce $\hat{\theta}$ and v .
 - Step 6** Update the walker's position accordingly.

Camera Setup Walkers visually perceive their environment through the OpenGL camera set at the first step of our algorithm. The camera's field-of-view has a width of 150° and a height of 80° . The camera position is set to the position of the considered walker at eye level and panning angle is aligned to the walker's motion direction. The tilting angle is set so that the upper clipping plane is horizontal (i.e., the camera is oriented toward the ground with a -40° angle). The resolution is 256×48 pixels.

Simplified Geometries The complexity of the proposed algorithm is dependent on the complexity of the environment (Step 2). Walkers does not need to react to subtle geometrical details of the scene. Simplified bounding geometries can be used as obstacles. In particular, perceived walkers are geometrically simplified as cones of $1.8m$ of height and $0.5m$ of base radius. Cones, similarly to walking humans, are wider at their base than at their top. Real humans can see above others' shoulders: cones better reflect this ability than, for instance, cylinders.

Computation of Inputs The model's inputs are computed as illustrated in Figure 4. In the figure, p_i is one of the perceived points that belongs to a given obstacle o . The relative velocity $\vec{V}_{p_i/w}$ of a perceived point with respect to the considered walker is first deduced:

$$\vec{V}_{p_i/w} = \vec{V}_o - \vec{V}_w \quad (10)$$

where \vec{V}_w is the walker's velocity vector and \vec{V}_o the obstacle's velocity vector the perceived point belongs to. Finally, $\vec{V}_{p_i/w}$ is decomposed into $\vec{V}_{conv_{p_i/w}}$ and $\vec{V}_{orth_{p_i/w}}$ to deduce tti_i and $\hat{\alpha}_i$ (\vec{V}_{conv} is the component of the relative velocity converging to the considered walker, and \vec{V}_{orth} the orthogonal component):

$$\vec{V}_{conv_{p_i/w}} = (\vec{V}_{p_i/w} \cdot \vec{k}) \cdot \vec{k} \quad (11)$$

$$\vec{V}_{orth_{p_i/w}} = \vec{V}_{p_i/w} - \vec{V}_{conv_{p_i/w}} \quad (12)$$

$$tti_i = D \cdot \|\vec{V}_{conv_{p_i/w}}\|^{-1} \quad (13)$$

$$\hat{\alpha}_i = \arctan\left(\frac{\|\vec{V}_{orth_{p_i/w}}\|}{D - \|\vec{V}_{conv_{p_i/w}}\|}\right) \cdot u^{-1} \quad (14)$$

where D is the p_i -walker distance, \vec{k} is the unitary p_i -walker vector, and u the unit of time.

5 Results

5.1 Examples

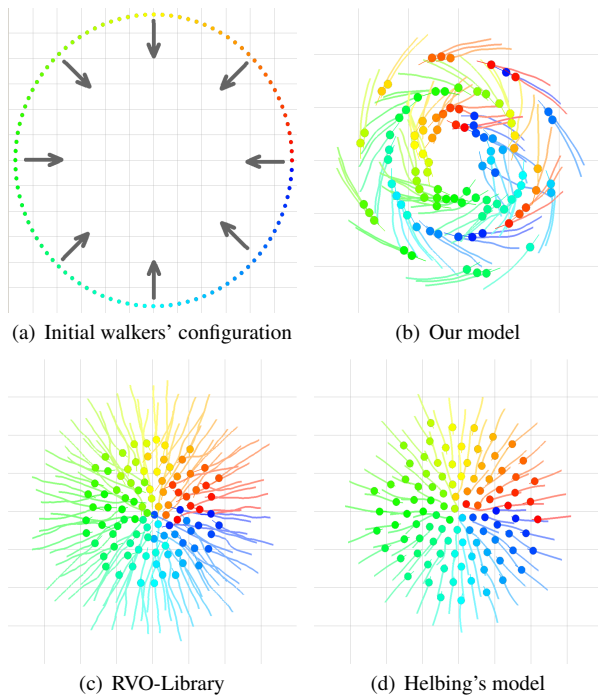


Figure 7: Circle (a) A scene of 100 walkers are initially deployed uniformly along a circle. The walkers goal is to reach the diametrically opposed position. The solution is shown for 3 models (b), (c), (d). Our model (b) is the only one able to provoke the emergence of patterns.

We illustrate our simulation results according to four examples. A comparison with two previously existing techniques is provided for the two first examples in order to illustrate the achieved improvements. We chose:

RVO which is representative of geometrical avoidance models.

Helbing's model which is representative of particle-based approaches. Contrarily to RVO, such models do not take into account anticipation, and interactions are formulated as a function of distance to obstacles.

The available examples are:

Circle: In this example, walkers are initially located along a circle and each one's goal is to reach the diametrically opposed position. In absence of others, each walker would go through the circle passing by its center. The number of interactions occurring in such an example is thus maximized: actually, each walker interacts with all the other ones. The main difficulty raised by this example is that walkers immediately converge to the center of the circle and get stuck there. Such a situation can be efficiently avoided when 'traffic circles' emerge, whilst the center is almost left empty of anyone. Results are shown in Figure 1 and 7.

Group-swap: In this example, walkers are initially separated in two groups. The goal is to swap group positions. Motion is not constrained by static obstacles. A main difficulty raised by this example is to achieve collision avoidance whilst walkers

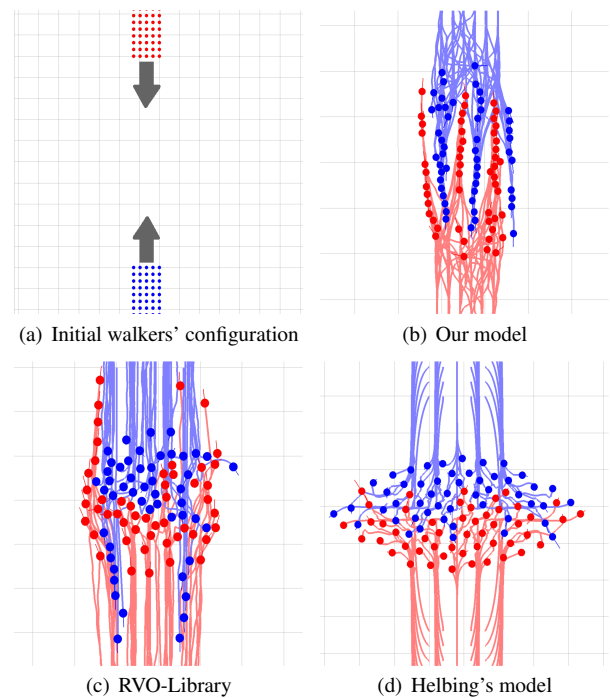


Figure 8: Group-swap A scene with two groups of walkers heading toward each other solved by three different models. In our model (b) distinct lane formations emerge with anticipation. The lane formations in RVO-Library (c) start emerging too late and lead to a congestion. In Helbing's model (d) such a formation does not emerge.

do not excessively deviate from the shortest route in spite of the absence of constraints (e.g., corridor walls). Such a result can be reached only if lane formations emerge. Results are shown in Figure 8.

Pillars: In this example, we increase the difficulty of the group-swap example by adding two rows of pillars in the middle of the scene. We also demonstrate the ability of our model to take in account static obstacles. Results are shown in Figure 9.

Crossing: In this example, two groups of people meet at the intersection of two orthogonal corridors: static obstacles both constraints the motion and prevent walkers to early perceive the ones from the other group. Main difficulties of this example are: first, avoid that one of the two groups get stuck and second, avoid walkers to be excessively deviated along corridors walls. Results are shown in Figure 10.




Figure 9: Pillars This example is identical to the group-swap one, two rows of pillars make the scene more complex. The images show the evolution in time of the simulation starting from the left.

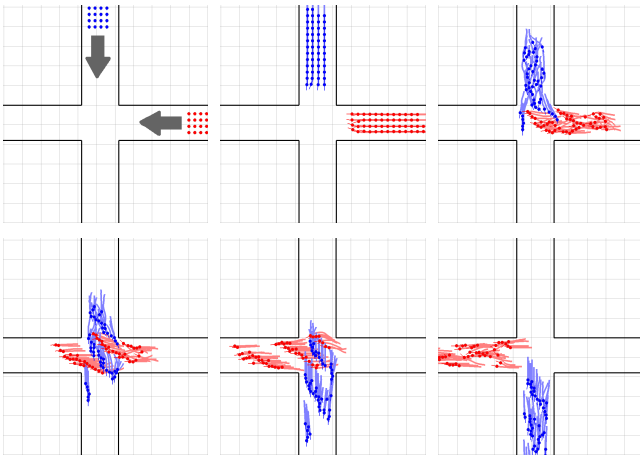


Figure 10: Crossing This example shows two groups meeting at the intersection of orthogonal corridors. The emerging line patterns, whose direction is approximately 45° , allow an efficient global motion (evolution in time is shown from top-left to bottom-right).

All of the displayed examples demonstrate our model ability to let self-organized patterns of walkers emerge from the motion. Emergent patterns allow to efficiently solve the sum of interactions between walkers (cf. Table 1). Improvements compared to previous approaches are perceptible: in identical situations, the walkers *travel time* is lowered using our approach, and the presence of *slow* walkers - with $v < 0.5m.s^{-1}$ and which may affect the overall believability of results - is decreased. In the example of the circle (Figure 7), other techniques concentrate walkers in the center of the scene which lower the efficiency of the circulation. In the case of the group-swap example, Helbing's model fail to find an acceptable solution: groups are widely spread because particles simply repulse each-other.

	circle		group-swap	
	max. travel time	prop. of slow walkers	max. travel time	prop. of slow walkers
our model	53s.	0.97%	55s.	0.74%
Helbing's	90s.	30.4%	71s.	11.0%
RVO	63s.	13.0%	59s.	4.7%

Table 1: The maximum walkers travel time and the proportion of slow walkers are provided for the circle and the group-swap examples, using three different models. The proportion of slow walkers is the mean proportion of time walkers are below $0.5m.s^{-1}$.

Furthermore, a specific feature of our model is the independent control of angular and tangential velocities. Decelerations occur only in case of an imminent collision. The absence of deceleration during anticipated reaction results in smoother trajectories. We believe overall our result is improved compared to previous approaches, especially when virtual humans are animated to follow the generated trajectories. The accompanying video illustrates the quality of synthetic trajectories, emergent self-organized patterns of walkers, as well as final animations.

5.2 Performance

Obtaining reasonable performance is probably the major technical challenge of our approach, due to the synthetic vision technique. We are still able to reach fair results by partly executing our algo-




Figure 11: Performance plot: the computation time for one simulation loop with respect to the number of walkers is measured. The simulation ran on a laptop with an Intel T7800@2.6GHz CPU and a Quadro FX 3600M graphics card. The circle-example situation was used. We detail the total simulation loop time in the rendering and processing plots, which respectively correspond to the time spent during steps 1-3, and 4-6 (the latter includes the GPU-CPU data transfer).

rithm steps on a GPU. Real-time performance, 25 frames per second (fps), is maintained up to 200 walkers (cf. Figure 11, computed on a laptop with an Intel T7800@2.6GHz CPU and a Quadro FX 3600M graphics card). The major bottleneck of our method is the data transfer from the GPU to the CPU (between steps 4 and 5).

Performance can be improved in several ways. Firstly, the camera resolution at step 1 can be lowered: on one hand, the number of perceived points is decreased accordingly and performance is improved; on the other hand, perception accuracy is decreased and may prevent walkers to react with anticipation to partly occluded obstacles. The accompanying video illustrates the impact of lowering the camera resolution. Secondly, we believe that the complete simulation loop can be executed on the GPU (recent approaches demonstrated feasibility [Silva et al. 2009]): on one hand, data-transfer between GPU and CPU is avoided (which represent approximately 30% of the complete simulation loop time); on the other hand, further tasks could be made impossible (e.g., animating virtual walkers). Finally, assuming that each obstacle is represented by a single static or moving point (which is an acceptable assumption for a scene made of walkers only), the model applies without needing to rely on synthetic vision. On one hand, interactions are directly considered between pairs of *moving points*, instead of each walker and a set of *perceived points*. The number of processed interactions is drastically lowered. On the other hand, synthetic vision has many advantages: the visibility of obstacles walkers react to is implicitly checked, obstacles can have any 3D shape, walkers height - which may limit their perception - is taken into account, etc.

6 Discussion

Realism Our results demonstrate the ability of our approach to improve the emergence of self-organized patterns of walkers on several examples. From the standpoint of Computer Animation, our method provides visually appealing results. Interactions are solved more efficiently at the global scale: compared to other approaches, the time required for a walker to reach its goal is noticeably lower using our model (cf. Table 1 and accompanying video for comparisons). We believe the attained efficiency benefits the resulting believability of animations, especially since some locking situations are avoided. It is however still required to quantitatively evaluate

the realism of results. Studies based on spectators' feedback or, better, confrontation with real observations are possible directions for such an evaluation.

High-level behaviors and control Interactions between walkers are today limited to collision avoidance. Locomotion is controlled at the most basic level by visual-stimuli/motor-response laws. A near-future objective is to obtain a higher level of control and to extend simulation abilities. Our first goal is to integrate some new types of interactions, such as following someone, reaching a mobile target, etc. Such interactions can easily be expressed in the $(\dot{\alpha}, tti)$ -space. For instance, following p_i translates into controlling the velocity so that $(\dot{\alpha}_i \rightarrow 0, tti_i \rightarrow cst)$, where cst is a positive constant. Then, our second goal is to combine different types of interactions to further improve the global navigation efficiency by setting mid-term strategies (for instance, temporarily following someone is an efficient strategy to avoid further avoidance interactions) or to allow the simulation of groups inside crowds (e.g., families). We assumed that goals were visible in our examples: a preliminary path planning stage would first be required to achieve navigation in complex environments. Path planners can decompose high-level goals into intermediary way-points that could be successively used as short-term goals in our model. A reactive change of short-term goals according to some external factors (e.g., local population densities) could be of interest: the evaluation of future traffic conditions as well as the route selection process should be deduced from the visually perceived information in order to match our approach philosophy.

Model parameters Model parameters (a, b, c) (cf. Equation (1)) can be adapted for each walker to individualize avoidance behavior with negligible computational overhead. The impact of parameters change on simulations is illustrated in the accompanying video. An intuitive link exists between avoidance behavior and the shape of τ_1 which is completely controlled by (a, b, c) . The higher the peak of τ_1 , the earlier the anticipation. The wider the peak, the stronger the adaptation. Finally, the curvature of τ_1 controls a trade-off between anticipation time and reaction strength: when the maximum curvature is higher, early anticipated reactions remain low whilst they get stronger when tti decreases. The automatic adaptation of parameters with respect to external factors, such as local density of population, may open interesting perspectives.

7 Conclusion

We presented a novel approach to crowd simulation that involves individual walkers avoiding each other. Our main contribution is to steer walkers according to the visual perception they have of their environment. We formulate their collision avoidance behavior as visual-stimuli/motor-response control laws. Compared to previous vision-based approaches, we rely on statements from cognitive science that identified the visual stimuli humans extract from their optic flow to control their locomotion and avoid obstacles. Compared to previous avoidance models, we demonstrate that our approach improves the emergence of self-organized patterns of walkers in crowd simulations. In spite of the computational complexity raised by the synthetic vision technique, we demonstrate the ability of our approach to address complex interaction situations between numerous walkers. Our results are promising and open several future work directions. The first step would be to automatically adapt the model parameters with respect to some external factors. A second direction is to extend our model with new types of interactions. Then, our objective is to add a higher level of control in order to combine several types of interactions and to enable mid-term and long-term navigation strategies. Today, the proposed approach

still results in visually interesting motions that can benefit to many Computer Animation applications. The evaluation of our results by comparing them to real observations and data is now required. Nevertheless, our model is founded on cognitive science work on human locomotion which can open interesting perspectives for realistic simulation purposes.

Acknowledgements

The authors would like to thank Marc Christie for fruitful discussions and valuable comments. This research is funded by the French National Funding Agency for Research (ANR Locanthrope project).

References

- CHAUMETTE, F., AND HUTCHINSON, S. 2006. Visual servo control, part i: Basic approaches and part ii: Advanced approaches. *IEEE Robotics and Automation Magazine* 13(4), 82–90.
- CHENNEY, S. 2004. Flow tiles. In *Proc. 2004 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '04)*, Eurographics Association, Aire-la-Ville, Switzerland, 233–242.
- CUTTING, J. E., VISHTON, P. M., AND BRAREN, P. A. 1995. How we avoid collisions with stationary and moving objects. *Psychological Review* 102(4) (October), 627–651.
- GAYLE, R., MOSS, W., LIN, M. C., AND MANOCHA, D. 2009. Multi-robot coordination using generalized social potential fields. In *Proc. IEEE International Conference on Robotics and Automation (ICRA '09)*, 106–113.
- GUY, S. J., CHHUGANI, J., KIM, C., SATISH, N., LIN, M., MANOCHA, D., AND DUBEY, P. 2009. Clearpath: highly parallel collision avoidance for multi-agent simulation. In *Proc. 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '09)*, ACM, New York, NY, USA, 177–187.
- HALPERIN, C., ANJYO, K., CIOROBA, M., KANYUK, P., REGELOUS, S., YOSHIDA, T., AND SALVATI, M. 2009. Crowd animation: Tools, techniques, and production examples. In *SIGGRAPH Asia '09: ACM SIGGRAPH Asia 2009 Courses*, ACM, New York, NY, USA, 1.
- HEIGEAS, L., LUCIANI, A., THOLLOT, J., AND CASTAGNÉ, N. 2003. A physically-based particle model of emergent crowd behaviors. In *Graphicon 2003*.
- HELBING, D., AND MOLNAR, P. 1995. Social force model for pedestrian dynamics. *Physical Review E* 51, 4282.
- HUGHES, R. L. 2003. The flow of human crowds. *Annual Review of Fluid Mechanics* 35, 169–182.
- KAPADIA, M., SINGH, S., HEWLETT, W., AND FALOUTSOS, P. 2009. Egocentric affordance fields in pedestrian steering. In *Proc. 2009 Symposium on Interactive 3D graphics and games (I3D '09)*, ACM, New York, NY, USA, 215–223.
- KARAMOUZAS, I., HEIL, P., VAN BEEK, P., AND OVERMARS, M. H. 2009. A predictive collision avoidance model for pedestrian simulation. In *Motion in Games*, 41–52.
- KUFFNER, J. J., J., AND LATOMBE, J. C. 1999. Fast synthetic vision, memory, and learning models for virtual humans. In *Proc. Computer Animation*, 118–127.

- LAMARCHE, F., AND DONIKIAN, S. 2004. Crowds of virtual humans : a new approach for real time navigation in complex and structured environments. *Eurographics'04: Computer Graphics Forum 23*, 3 (September), 509–518.
- LOSCOS, C., MARCHAL, D., AND MEYER, A. 2003. Intuitive crowd behaviour in dense urban environments using local laws. *Theory and Practice of Computer Graphics (TPCG'03)*.
- MASSIVE. <http://www.massivesoftware.com>.
- 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*.
- NOSER, H., RENAULT, O., THALMANN, D., AND THALMANN, N. M. 1995. Navigation for digital actors based on synthetic vision, memory, and learning. *Computers & Graphics 19*, 1, 7–19. Computer Graphics Lab., Swiss Federal Inst. of Technol., Lausanne, Switzerland.
- PARIS, S., DONIKIAN, S., AND BONVALET, N. 2006. Environmental abstraction and path planning techniques for realistic crowd simulation. *CASA 2006: Computer Animation and Virtual Worlds 17*, 3-4, 335.
- 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.
- PELECHANO, N., ALLBECK, J. M., AND BADLER, N. I. 2007. Controlling individual agents in high-density crowd simulation. In *SCA '07: Proceedings of the 2007 ACM SIGGRAPH/Eurographics symposium on Computer animation*, 99–108.
- PELECHANO, N., ALLBECK, J., AND BADLER, N. I. 2008. *Virtual Crowds: Methods, Simulation, and Control*. Morgan & Claypool Publishers. (Series Editor: Brian Barsky).
- PETERS, C., AND O'SULLIVAN, C. 2003. Bottom-up visual attention for virtual human animation. In *International Conference on Computer Animation and Social Agents (CASA'03)*, 111–117.
- PETTRÉ, J., CIECHOMSKI, P. D. H., MAĪM, J., YERSIN, B., LAUMOND, J.-P., AND THALMANN, D. 2006. Real-time navigating crowds: scalable simulation and rendering. *Computer Animation and Virtual Worlds 17*, 3-4, 445–455.
- PETTRÉ, J., ONDŘEJ, J., OLIVIER, A.-H., CRETUAL, A., AND DONIKIAN, S. 2009. Experiment-based modeling, simulation and validation of interactions between virtual walkers. In *Proc. 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '09)*, ACM, New York, NY, USA, 189–198.
- 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.
- REYNOLDS, C. W. 1999. Steering behaviors for autonomous characters. In *Game Developers Conference 1999*.
- SCHADSCHNEIDER, A. 2001. Cellular automaton approach to pedestrian dynamicstheory. In *In Pedestrian and Evacuation Dynamics*, 75–85.
- SHAO, W., AND TERZOPOULOS, D. 2005. Autonomous pedestrians. In *Proc. 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation (SCA '05)*, ACM Press, New York, NY, USA, 19–28.
- SILVA, A. R. D., LAGES, W. S., AND CHAIMOWICZ, L. 2009. Boids that see: Using self-occlusion for simulating large groups on gpus. *Comput. Entertain.* 7, 4, 1–20.
- SUD, A., ANDERSEN, E., CURTIS, S., LIN, M., AND MANOCHA, D. 2007. Real-time path planning for virtual agents in dynamic environments. *Proc. IEEE VR 2007*, 91–98.
- THALMANN, D., AND RAUPP MUSSE, S. 2007. *Crowd Simulation*. Springer, London.
- THALMANN, D., KERMEL, L., OPDYKE, W., AND REGELOUS, S. 2005. Crowd and group animation. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Courses*, ACM, New York, NY, USA, 1.
- TRESILIAN, J. R. 1994. Perceptual and motor processes in interceptive timing. *Human Movement Science 13*, 335–373.
- TREUILLE, A., COOPER, S., AND POPOVIĆ, Z. 2006. Continuum crowds. *ACM Transactions on Graphics (SIGGRAPH 2006) 25*, (3).
- TU, X., AND TERZOPOULOS, D. 1994. Artificial fishes: physics, locomotion, perception, behavior. In *SIGGRAPH '94: Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, 43–50.
- 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)*.
- WARREN, W. H., AND FAJEN, B. R. 2004. *Optic Flow and Beyond*. Kluwer (Editors: L. M. Vaina, S. A. Beardsley, and S. Rushton), ch. From optic flow to laws of control, 307–337.