



HAL
open science

Studying and modeling complex interactions for crowd simulation

Julien Bruneau

► **To cite this version:**

Julien Bruneau. Studying and modeling complex interactions for crowd simulation. Graphics [cs.GR]. Université de Rennes, 2016. English. NNT : 2016REN1S119 . tel-01425268v2

HAL Id: tel-01425268

<https://inria.hal.science/tel-01425268v2>

Submitted on 29 Sep 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



THÈSE / UNIVERSITÉ DE RENNES 1
sous le sceau de l'Université Bretagne Loire

pour le grade de
DOCTEUR DE L'UNIVERSITÉ DE RENNES 1

Mention : Informatique

Ecole doctorale Matisse

présentée par

Julien Bruneau

préparée à l'unité de recherche UMR 6074 IRISA
et au centre INRIA - Rennes Bretagne Atlantique

Studying and modeling
complex interactions for
crowd simulation

Thèse soutenue à Rennes
le 30 novembre 2016
devant le jury composé de :

William Warren

Chancellor's Professor, Brown university / rapporteur

Ronan Boulic

Senior Scientist, EPFL / rapporteur

Jean Paul Laumond

Directeur de recherche, LAAS-CNRS / examinateur

François Chaumette

Directeur de recherche, Inria Rennes / examinateur

Anne-Hélène Olivier

Maître de Conférences, M2S-Université de Rennes 2 / exami-
natrice

Julien Pettré

Chargé de recherche, Inria Rennes / directeur de thèse

Contents

List of Figures	vii
List of Tables	xv
1 Introduction	1
1 Problem	3
2 Approach	4
3 Contributions	5
4 Plan	6
2 Background	7
1 Crowd simulation algorithms	7
1.1 Collision avoidance	8
1.1.1 Principal collision avoidance models	8
1.1.2 Extensions to complex situations	11
1.2 Other behaviors	14
1.2.1 Groups	14
1.2.2 Following	15
1.3 Discussion	16
2 Observation of human behaviors and interactions	16
2.1 Real world observations	16
2.2 Controlled experiments	18
2.3 Evaluation	21
2.4 Discussion	24
3 Virtual Reality	25
4 Conclusion	27
3 Simple study case: following model	29
1 Introduction	30
2 Following model	30
2.1 One-dimensional following	31
2.2 Two-dimensional following	34
3 Results	34
3.1 Following distance	35
3.2 Waiting line	35
3.3 Corridor	36
4 Evaluation	37

4.1	Simulation of real situations	38
4.2	Stop-and-go waves	39
4.3	Comparison with other models	41
5	Conclusion	42
4	Planning a sequence of interaction over time	45
1	Introduction	46
2	Effective Avoidance Combination Strategy	47
2.1	Overview	47
2.2	Definitions	48
2.3	Cardinal Interaction Points (CIW)	48
2.4	collision-free Interaction Path	50
2.5	Cost and ranking	51
2.6	Perception	53
2.7	Following case	54
2.8	Extended search	55
3	Results	56
3.1	Performance	56
3.2	Case study	57
3.3	Energy	60
3.4	Limitations	62
4	Conclusion	63
5	VR study: Human behaviors toward groups	65
1	Introduction	66
2	Principle of Minimum Energy: a simulation study	67
2.1	Method	68
2.2	Results	69
2.3	Discussion	70
3	Groups avoidance: a VR-based user study	71
3.1	Objectives	71
3.2	Materials and Methods	72
3.3	Results	73
3.4	Discussion	75
4	Applications	77
4.1	Decision algorithm: Around or through the group?	77
4.2	Integration in RVO2	78
4.3	User-study evaluation	79
5	Conclusion	81
6	Hotspot in human environment perception	83
1	Introduction	84
2	Previous gaze experiments	84
3	Experiment	88
4	Analysis	90
4.1	Fixations computation	90

4.2	Virtual human characteristics	91
4.3	Statistical analysis	93
5	Results	93
6	Discussion	94
7	Conclusion	96
7	Conclusion and future work	97
1	Contributions	97
2	Future Work	99
3	Summary	101
8	Résumé en Français	103
1	Problème	104
2	Approche	105
3	Contributions	106
Appendices		i
A	Collaborations	i
1	Introduction	i
2	Walking with virtual people: Evaluation of locomotion interfaces in dynamic environments	i
2.1	Introduction	i
2.2	Experiments	ii
2.3	Results	iii
2.4	Conclusion	iv
3	Group Modeling: a Unified Velocity-based Approach	iv
3.1	Introduction	iv
3.2	Velocity-based group simulation algorithm	v
3.3	Parameter analysis	vi
3.4	Results	viii
3.5	Conclusion	ix
Bibliography		xxii
Publications list		xxiii

List of Figures

1.1	Virtual crowd used to fill up a large army in the 2014 movie Hercules. This allows to have huge crowds in movies while requiring a manageable number of extras.	1
1.2	Left: player interacting with a virtual crowd in the game Assassin's Creed. Right: Massmotion simulation result of the design for a new Transbay Terminal in San Francisco indicating that early designs of the terminal building contained problem areas from a pedestrian circulation point of view due to insufficient channel widths in what were predicted to be high volume routes.	1
1.3	Picture of a crowd in Central Park showing many different behaviors: a performer blowing giant bubbles, people gathering around him to watch, a medium size group walking together to a wedding, a woman sitting on a stair step in the middle of the way, people walking around, people relaxing around the fountain...	3
2.1	The obstacle avoidance rule on the left and the separation rule on the right, both from Reynolds' rule-based model [Rey99]	9
2.2	On the left, the computation of the velocity obstacle of agent A that leads to a collision with agent B. On the right, the union of all the velocity obstacles to avoid multiple obstacles.	9
2.3	Top: effect of the bearing angle evolution on collision avoidance from [OPOD10]. Bottom: resulting trajectories from Warren and Fajen's model [WF08] compare to real human data.	10
2.4	The Optical flow used in the vision based model [OPOD10] with the risk of collision that is a combination of the derivative of the bearing angle, to check for collision, and the Time To Interaction, to check for the imminence of the danger.	11
2.5	A rule, from the ruled-based algorithm proposed by [ST05], that defines a global strategy to avoid several agents at once in a corridor: pedestrian H searches for a safe direction interval when confronted by oncoming traffic.	12
2.6	Effect of the lookahead mechanism [GNL13], which allows the orange agent to see the approaching crowd and adjust its velocity to avoid it as one entity.	12
2.7	The steering algorithm of the Egocentric Affordance Fields [KSHF09]: (a) The current state of the environment. (b) Static Perception Field indicating low traversability at position of obstacles. (c) Dynamics Perception Field for velocity, $S_{target} = S_0$. (d) Dynamics Perception Field for velocity, $S_{target} = S_0 + \delta V$. (e) Affordance Field indicating path of high fitness to goal.	13
2.8	Different group structure simulated by [QH10].	15

2.9 Pictures from real world observations with A: the train station dataset from [ZWT12], B: the mall dataset from [CLGX12], C: the one year recording dataset from [BK15] 17

2.10 Top: groups observed in Trinity College from [PEC09]. Below: automatically detected line of people going in the same direction from [ZWT12] 18

2.11 Trajectories of people walking through a constrained environment from [BJ03] 19

2.12 Description of the experiment from [OMCP12]. A: Experimental setup. Area is 15mm. Two participants stand at the corners of the area and are synchronously given a start signal. Their task is to walk to the opposite corner. They implicitly start an interaction to avoid any collision. B: Picture taken during experiment. C: Schematic illustration of the minimum predicted distance (MPD) computed at three different times with a motion adaptation occurring between times t2 and t3. 20

2.13 Counter flow study from [DDH+14b]. Top: snapshot of the experiment. Bottom: Trajectories of pedestrians within the channel in each experiment scene. . 21

2.14 Two typical situations used for visual evaluation. On the left, it is a circle with everyone trying to switch places with the one on the opposite side. On the right, it is two flows on a frontal collision course with each other. In both cases, we can see some models having difficulties to produce fluid movement that would look more natural. These comparisons are from [OPOD10] and the "our model" correspond to the vision based model. 22

2.15 Fundamental diagram plotted for different algorithms with optimized parameters to suit different cultures [WJGO+14]. 23

2.16 Steerbench workflow [SKFR09]. 24

2.17 Left: subject wearing HMD and physiological monitoring from the experiment in [MIWB02]. Right: subject in a cave facing a virtual thrower in the handball case study from [BKV+10]. 25

2.18 The evaluation framework, proposed in [COMP13], compares trajectories between an experiment performed in the real world and the same experiment performed in a virtual environment. 26

3.1 The ideal distance between the follower and the leader is described by the difference between the leader's estimated position and the sum of the follower's fixed distance and a dynamic safe distance determined according to the leader's movement. 31

3.2 If there is space beside the leader, the follower accelerates in the perpendicular axis towards that side with space available. 34

3.3 Visualization of the following distance computed for a leader with very stable motion (left) or with jerky motion (right). 35

3.4 Trajectories of people lining up and waiting. Stop-and-go waves are created and propagated when the first walker in the line moves. 36

3.5 Two scenarios. Top: waiting line scenario simulating 1-D following traffic. Center: top view of the periodic corridor simulating 2-D following traffic by using an overtaking rule. The agents in red are walking fast, they will slow down when facing agents blocking their way and try to overtake them. Bottom: perspective view of the periodic corridor. 37

3.6	Trajectories from experiment trial #25 (left) and the simulation results from our model (center) and Lemercier's one (right). Lemercier's model needs some initial motion to work so it has been initialized after 20 s of real motion. In our model, we just need to place the walkers at the initial positions extracted from the real data.	38
3.7	Fundamental diagram computed for real data and their simulated counterpart with different ttr (0.5 s, 0.75 s and 1.0 s).	39
3.8	Waves detection on trial #25 of the real trajectories database. The upper plot shows all the trajectories with the detected minima as black stars and the defined waves propagation as black lines. The lower plot shows the speed of the last person in blue, the smooth speed in red and the detected extrema as red stars.	39
3.9	Relation between waves and density from real data. (a) The wave score according to the density for every real trajectory and (b) the results of three simulations with different densities, showing waves emerging for the one with density around 1.5 ped.m^{-1} , as expected.	40
3.10	Simulations of agents following a leader with a rectangular signal speed with some noise. Reynolds: in Reynolds's model there is no security distance, the agent can be closer to the leader than the desired distance leading to collision risk. Middle: in our model the followers keep close to the leaders with a security margin to prevent getting closer than a fixed distance. Right: in Lemercier's model the following distance increases over time.	41
4.1	In some situations, humans can see a tunnel through a flow. The left figure shows a first person view point where the tunnel is clearly visible. The right figure shows a top view of the situation as well as the set of colliding and collision-free velocities for the red agent. Regular local avoidance models try to find a collision-free speed toward the goal. This means the red agent will choose a velocity from the set of collision-free velocities (green part of the circle), making him slow down and miss the tunnel.	45
4.2	Navigation system architecture with three different levels of decision.	47
4.3	Collision velocity space for one interaction, in red. $\mathbf{v}_{a_{1 2}}$, $\mathbf{v}_{r_{1 2}}$, $\mathbf{v}_{d_{1 2}}$ and $\mathbf{v}_{1 2}$ represent the collision-free velocities that are used to compute the CIW.	49
4.4	Example of collision solved using the four CIW. The collision trajectories are drawn in red and the four trajectories followed to reach the CIW are shown in blue. The position of each CIW is shown together with the position of the other agent when the CIW is reached.	50
4.5	Example showing the construction of a collision-free IP illustrating the different steps taken by the EACS system.	51
4.6	Agent R perception system: First it selects the agents inside the perception area (O, B and G) then for each selected agent it computes the interception speed and selects the n agents with the smallest Minimum Interception Time (MIT) as neighbors (First B then G, O is not selected as it is going away too fast to be intercepted).	53
4.7	Detection of a following case and the two Interaction Waypoints \mathbf{i} and \mathbf{j} computed from the following speed.	54

4.8 1) Without extra search, the algorithm stops at the IS OA and never considers OE. 2) With extra search, the IS OE is created by solving the collision on the IS OA'. 3) Agent's trajectories without extra search, the agent starts to go inside the crowd before backing up and going around. 4) With extra search, the agent is able to directly go around the crowd. 55

4.9 The graphs show the time needed to compute 20 seconds of simulation for the different performance test situations with and without EACS. On the left graph, the results of the three situations: OWF, BF and FCF are presented for different numbers of agents. The right graph presents the results of the SB situation for different densities. 58

4.10 Simulation of an agent (in red) going through a flow, with the flow having a tunnel to facilitate the crossing. Left: simulation including EACS; Right: simulation without EACS 58

4.11 Simulation of a clogged exit, EACS is able to choose other exits that are further away but not clogged. 59

4.12 Simulation of an agent (in red) going through a flow that goes in the opposite direction. The flow has no tunnel but a section of it has a lower density. At the top, we have the simulation with EACS and at the bottom the simulation without EACS. 59

4.13 Left: simulated test where short-term cost leads to bad mid-term cost, agents simulated with EACS on the left and without on the right. Right: the graph quantifies the energy consumption of the red agent (For RVO2 and EACS) as they walk towards the goal (In this situation, it is equivalent to the distance traveled on the x axis). The section of the blue curve circle in red corresponds to RVO2 agents going backwards to avoid the dense group (T=7.5s on the top figure). 60

4.14 Simulation of group avoidance for different group densities. 61

4.15 The different density types used for the energy study from very irregular (left) to uniform (right). 62

4.16 This graph shows the amount of energy preserved by the agents when using EACS to optimize their avoidance strategies compared to EACS-less strategies. 63

5.1 Simulation setting with one agent avoiding a group on an orthogonal or collinear trajectory. On the left, the starting position with movement direction. On the right, example of computed trajectories relative to the group (For the top: group radius=3, interpersonal distance=0.8, for the bottom: group radius=8, interpersonal distance=2) 68

5.2	Energetic cost of paths performed to avoid a group for collinear (Top) and orthogonal (Bottom) relative movements. Graphs on the left and right show respectively the average energy consumed by going around trajectories (Left) or going through trajectories (Right) according to group sizes and group interpersonal distances. Graphs on the middle show the results of the Wilcoxon signed-rank tests which compare the energetic costs for both decisions. Significant differences between both decisions are illustrated in red when going around the group is less costly than going through it, in blue when going around the group is more costly than going through it. When there is no significant difference, the result of the comparison is shown in yellow.	70
5.3	Conditions for the experiment on group avoidance. Top: spacing between the characters of the group was: $\{1.1, 1.4, 1.7, 2, 2.3\}m$. Bottom left: participants avoided a group coming from: left, front or right. Bottom right: the visual appearance of the characters of the group was: ordinary people, soldiers or zombies.	71
5.4	Pictures of the experiment showing participants navigating with a joystick in a VR environment and avoiding virtual agents.	72
5.5	Participants' trajectories relative to the group position (gray circle) with respect to experimental conditions: interpersonal distance (left), direction (top right) and appearance (bottom right). Going around trajectories are displayed in red. Going through trajectories are displayed in blue.	74
5.6	Proportion between going around and going through decisions per distance, over all trials.	75
5.7	Decisions to go through or around amongst the participants and the energy spent in the same condition. For comparison purposes, we only consider decisions for ordinary group.	76
5.8	Proportion between going around and going through decisions per participant, over all trials.	76
5.9	On the left: the situation as it really is, on the right: the situation as it is perceived by the red agent (All the blue agents are seen as one big agent because they are considered as a group while the green ones are too far from each other to be considered as a group)	79
5.10	Example of avoidance decision made by the purple agent. The agent goes through the very dense group in RVO2 while going around it in the modified RVO2.	80
5.11	Experimental setup to evaluate the decision algorithm.	81
6.1	Several results from the experiment performed in [FWK11]. From left to right: the proportion of gaze toward different elements, the probability of a pedestrian being fixated according to the amount of time he has been visible, the proportion of gaze toward pedestrians per time period since their appearance.	85
6.2	Heat maps, from [SvSB ⁺ 13], of a newly built area looked at on two subsequent, static slides in a lab study (top) before the construction of the building as compared to the heat map of the same terminal after construction derived from mobile eye tracking videos recorded at the airport (bottom). Red and green, indicate a high and low number of overall fixations respectively.	86

6.3 Experiment performed in [JMH09]. Left: Illustration of a path around the occluded area (The red circle is the subject, the blue squares are the pedestrians walking in the direction opposite and the gray triangles are pedestrians walking in the same direction. Right: Effect of collision probability on fixations on pedestrians.) 87

6.4 Experiment performed in [KF10]. Left: an example of the field of view recorded by Eye tracker. Right: Relative position of fixated or ignored objects (ignored objects are represented by the cross). 87

6.5 Participant using the joystick to navigate in the virtual environment while his eyes are being tracked by the eye tracking device below the screen (the captor with purple dots on each side). 88

6.6 The generation of the virtual humans' position and speed is done to insure the presence of at least 7 virtual humans in a square around the default player's position at a constant time interval δt 89

6.7 Graphs showing the percentage of fixations toward different rank of MPD, TTCA and distances. 95

A.1 Virtual environment and experimental situation used in this study. The virtual human has several starting positions corresponding to different value of mpd giving the virtual human some advance or delay over the user. Position of both the user and the virtual human are represented for different times: when they can see each other t_{see} and then $t_{0.5}$, t_1 and t_2 which are respectively $0.5s$, $1s$ and $2s$ after t_{see} (if there is no adaptation from the user). A dotted line joins the positions of the character and of the participant at those times for $mpd = 0$ iii

A.2 Participants navigating inside the virtual environment using different locomotion interfaces iii

A.3 (a) The velocity obstacle $VO_{A|B}^\tau$ of a agent B to agent A is a truncated cone in the velocity space. (b) The velocity connection $VC_{A|B}^\tau$ of an agent B to agent A contains the velocities that result in grouping A and B together for at least τ_{min} seconds before time τ . Here $\tau_{min} = 1$. (c) $VC_{A|B}^\tau$ for a rectangular constraint. (d) A situation where not all velocity constraints can be satisfied since $VC_{A|B}^\tau \cap VC_{A|C}^\tau = \emptyset$ (group neighbors B and C are not close to each other). In this case, agent A tries to stay close to both B and C. v

A.4 Simulation results of a group (in blue) going through a crowd of individuals (in green) for different value of n^c . The simulations demonstrate that the group is more and more bound together and is less likely to break into small clusters with higher value of n^c , which is quantitatively confirm by the graph on the right. vi

A.5 Simulations of a group (in blue) going through a crowd of individuals (in green) for different value of preference weights. The more the balance is in favor of grouping the tighter the group's clusters are. Moreover, with a balance in favor of grouping, no blue agent end up alone. vii

A.6	Simulation of a leader (in red) trying to move a whole group (in blue) toward its goal. Different level of n^c were tested showing that an average level of cohesion is require, from the group, to move a whole group. The graph shows the distribution of the agents through the distance from the leader at the end of the simulation.	viii
A.7	Left: a group avoiding a collision with another by either bypassing it or splitting and merging afterward. Right: two formations that is possible to simulate with the proposed algorithm.	viii
A.8	A general scenario that showcase a wide variety of groups, all of which were simulated by the proposed algorithm	ix

List of Tables

6.1	Percent of correctly detected fixation on red virtual humans per participants.	93
-----	--	----

Introduction

1



Figure 1.1 Virtual crowd used to fill up a large army in the 2014 movie Hercules. This allows to have huge crowds in movies while requiring a manageable number of extras.

During this PhD, we worked on crowd simulation. Basically crowd simulation is about computing the movement of virtual crowds so that they visually or quantitatively act like real ones. Such virtual crowd is presented in figure 1.1. Realistic crowd simulation is used in a wide range of application from the entertainment industry, with movies (Figure 1.1) and games (Figure 1.2), to the architecture and safety industry to check the design of future public places for example (Figure 1.2). These applications have greatly participated in the growing interest for such simulations and their evolution. But each application has its own constraints.



Figure 1.2 Left: player interacting with a virtual crowd in the game Assassin's Creed. Right: Massmotion simulation result of the design for a new Transbay Terminal in San Francisco indicating that early designs of the terminal building contained problem areas from a pedestrian circulation point of view due to insufficient channel widths in what were predicted to be high volume routes.

With a constantly improving technology, movies and games offer more and more huge and epic scenes/environments that need to be populated with large crowds. Games need to populate their ever larger environments with virtual humans that should visually act as real ones and interact with players in real time. Movies need to create huge crowds such as armies, zombies or people fleeing a disaster. But unlike games, movies do not need a crowd of virtual humans interacting with a real one in real time. Instead they require even more visually realistic behaviors as the virtual humans are shown on screen with real ones.

While visually realistic crowds are sufficient for the entertainment industry, the architecture and safety industry requires a level of realism even deeper. Crowd simulations are used to check the flow capacity or find ways to improve it when conceiving buildings or restructuring crowded public places such as train-stations or subways. It is also used to simulate disasters and emergency evacuations to be better prepared to face such event or to check the safety measures taken for buildings or other locations. Such applications require virtual crowds that behave as closely as possible to real one, otherwise the conclusion given by the simulation might be counterproductive.

These constraints, that are specific to the different applications, lead to the development of many different solutions. Several algorithms have been conceived so that the virtual crowds are able to reproduce real crowd-like behaviors. There are 2 families of algorithms for crowd simulation. The macroscopic ones consider directly the crowd as one huge entity and replicate its behaviors by modeling it as a fluid that flows through buildings or streets for example. The microscopic ones consider every individuals forming the crowd and focus on replicating their behaviors and interactions between them. Then by simulating several individuals interacting with each other they create virtual crowds the same way real crowds are created and steered by all the individuals interacting with each other. The later requires a good understanding of local interactions between humans, how people's position and movement influence each other's. The work during this PhD has been focused on studying and modeling these local interactions to improve the realism of microscopic simulations.

There is a huge variety of behaviors that can be found in a crowd which varies depending of the type of crowd and the situation. It can go from walking in formation for armies to window shopping in malls. But even in the same crowd there are many different behaviors that can coexist, which is well illustrated by the crowd in Central Park shown in Figure 1.3. The most important one is collision avoidance as it is present in every crowd whatever the situation. Collision avoidance can be seen as a trivial task that we perform every day without even thinking about it. But modeling this behavior is very complex. While it is easy to geometrically solve the collision avoidance problem with one static or moving obstacle, it becomes more and more complex as you add more and more obstacles. It is even more complex as other obstacles can be other people that perform collision avoidance too. It happens in real life that two people try to avoid in the same direction which can lead to awkward moment with both people unsure of what to do next to avoid the collision ([HKK15]). But such events are rare and most of the time people are well synchronized when performing collision avoidance. But the most complex part about modeling collision avoidance is to have a model that will behave similarly to real humans. Indeed, even in simple situations where collision avoidance can be easily solved,



Figure 1.3 Picture of a crowd in Central Park showing many different behaviors: a performer blowing giant bubbles, people gathering around him to watch, a medium size group walking together to a wedding, a woman sitting on a stair step in the middle of the way, people walking around, people relaxing around the fountain...

only a small subset of the solutions is pertinent to model the human behavior. A subset that is very hard to define, as humans are chaotic in nature and won't react the exact same way when confronted twice to the exact same situation.

Many experiments have been performed to study how humans avoid static obstacles or each other's in simple situations. From these studies we have gained a better understanding on how humans avoid obstacles. This has helped create new algorithms or improve old ones to have more realistic simulation. It also allows us to evaluate existing algorithms, to check if the simulation produces trajectories similar to the one performed by real humans.

1 Problem

The problem addressed by this PhD is the following: How to handle complex situations (situations with a great number of various possible interactions) to produce simulation with virtual humans acting similarly to real humans? More specifically, when several interactions are possible, how should they be selected: should all the possible interactions

be performed at once or just a few? Then, once the interactions are selected, how should they be ordered: should they all have the same importance or is there some that should be done in priority? Finally, how should the remaining interactions be combined: should they be treated independently from each other or should some be merged into one bigger interaction?

Thanks to the wide varieties of experiments, done on collision avoidance on simple situations, current algorithms perform well in these situations. But when these simulated situations are scaled to complex situations, current algorithms have difficulties to reproduce realistic behaviors. For example in very high densities, current algorithms show difficulties to achieve the same level of fluidity as in real life. They also have difficulties to reproduce people trajectories snaking in and out of complex crowds. These difficulties seem to imply that there is some difference in the way we handle a collision with one person and the way we handle collisions with lots of them. Studying real data is then needed in complex situations to understand how people combine multiple interactions and what is missing in current algorithms.

The absence of such data on multiple interactions is mainly due to the challenges to perform experiments. Indeed, even for simple situations such as the interaction between two persons, it is hard to standardize the conditions. We can illustrate this challenge in a collision avoidance task between two walkers having perpendicular trajectories. A difficulty will be in that case to provide the situation with a risk of collision since the objective of the experiment is to study the strategies to avoid a collision. Some measures can be performed a priori to improve the odd of an actual risk of collision to appear like having both individuals starting to move at approximately the same time and distance of the crossing point. However, even if there is an actual risk of collision, due to the variability of human motion (not the same walking speed between walkers, not the same reaction time to the starting signal....) it is still difficult to deal with standardized situations. Then, it implies that participants have to perform the task many times to ensure that collision avoidance behavior will be sufficiently recorded to provide robust statistical analysis. Increasing the complexity of the studied situation would only make the experiment even more complicated. It would require to redo the experiment even more times, with more constraints which would be harder to handle efficiently.

2 Approach

To improve algorithms in complex situations, two different approaches have been proposed in this thesis. Interactions between humans can be influenced by many different factors, for example keeping some socially acceptable space with others (large with stranger, small with friends), keeping extra space away from people with unstable walking pace for safety, etc. The number of factors considered, in existing algorithms, are usually limited to spatial factors (position, speed, ...) and sometimes completed by some personal characteristics (preferred walking speed, personal space size, ...). The first approach consisted to add more factors to current algorithms such as psychological or sociological factors. This enables algorithms to simulate trajectories more adapted to the current situation. In the second approach, we keep current collision avoidance algorithms untouched as they

already perform well in simple situations. Instead we add a layer on top of the collision avoidance algorithms. The goal of the new layer is to reduce the complexity of the complex situation. Taking a situation with a dense crowd to avoid, instead of having the collision avoidance algorithm trying to solve all the collision avoidance at once, the new layer analyses the situation, finds a path through the crowd and directs the collision avoidance algorithm to interact with only one or two people of the crowd.

With respect to the complexity previously described to study the navigation behavior as well as the interaction between walkers, we propose to use Virtual Reality (VR) to ease the process. VR is already used to perform experiments and studies in a lot of domains: in sports ([BKV⁺10]), motion control ([FFW07]), perceptual control laws ([WF04]), spatial cognition ([MGvVB98]; [MCRT06]). There are many advantages to use VR for studies as it gives us a perfect control of the virtual environment experienced by the participants. Which means we can control the situation that participants will face, for example we can easily ensure that participants will face a collision risk with a virtual human. This also means that we can set up complex situations, have several instances of the same situation but with one parameter varying (such as the distance between group members if the study involves groups) and introduce every participant to the exact same stimuli. While the experiment process is eased with VR, it is also accelerated. Indeed, with a VR experiment, we only require a VR device (CAVE, HMD ...) and one participant at a time then we launch the simulation in the VR device for the participant. With an experiment in the real world we might have to gather a lot of people at the same time which takes a lot of time and efforts to plan. Moreover, once a sufficient number of people are gathered, we need to set them up according to the studied situation and that is required for every trial if several trials are needed. Despite these interesting features that could help us study many complex interactions, VR is still barely used in the crowd simulation domain.

3 Contributions

The contributions, done during this PhD, can be split into two categories: Crowd simulation algorithms and studies about human navigation behaviors.

The first contribution, on crowd simulation algorithms, was done to design a new algorithm to model following behaviors using additional influencing factors compared to previous one, such as safety. Working on this specific behavior allows us to reduce the simulation to one dimension (every one walking on the same line) and focus on the effect of these extra factors. The resulting trajectories were compared to previous algorithms trajectories as well as trajectories from real humans. The proposed algorithm proved more adaptable than previous one: while it is performing at least as well as previous algorithms, it shows better results in some situations. This work was presented at the MIG 2014 conference [BDP14]. The second contribution is an algorithm allowing virtual humans to plan an efficient path through many static or dynamic obstacles. The algorithm computes different strategies to go through a crowd, by combining collision avoidance through time (one after the other), and selects the most energy efficient one. This planning is done before the usual local collision avoidance done by current algorithm and the result showed an improvement in dealing with several obstacles at once. This work was presented at the SCA 2015 conference [BP15] and was selected for an extension in the CGF journal [BPss].

Finally, a collaboration was done with Zhiguo Ren on his work about group simulation [RCB⁺ss]. He proposed an extension to RVO2 adding new constraints so that agents from the same group would walk while staying together. We assisted on the results analysis, showing that his algorithm was able to simulate a wide variety of grouping behaviors such as small groups walking in specific patterns or huge groups of tourist that can split and go back together.

In parallel to working on locomotion algorithms, some studies on human navigation have been performed. It started with the development of a VR platform to perform experiments on human interactions. This has mainly been focused on the work of Anne-Hélène Olivier to validate the use of VR for navigation experiment [OBKPss, OBCP14]. For this validation, a real world navigation study has been redone in VR and the results have been compared between the VR experiment and the real world one. This study shows that despite some small quantitative differences, the experiment in the real world and the one in VR yield qualitatively similar results thus demonstrating that human locomotion behaviors can be studied using VR. This conclusion leads to the second contribution: using VR to study how people avoid a group of pedestrian. Thanks to VR, we were able to control the size of the group and the distance between the members of the group during the entire task. We were able to define a threshold of distance between group members at which human behavior switches from going around the whole group to going through it. This work was presented at the IEE VR 2015 conference [BOP15]. For the third contribution, we used VR combined with an eye tracker system to gather data about the perception action loop when interacting with other people. In other words, we were focusing on the characteristics of the walkers a human is looking at when navigating in a crowd. This work is still an ongoing work and only preliminary results will be presented in this thesis.

4 Plan

The remainder of the thesis is organized as followed. It starts in chapter 2 with an overview of previous work done to develop crowd simulation algorithms, observe and analyze real data from navigation and interaction tasks, and use Virtual reality to conduct experiments on human behavior. Then the contributions on crowd simulation algorithms are presented. In chapter 3, we present the work on following behavior. In chapter 4, we present the EACS algorithm that plan collision avoidance through time to natural avoiding path through a crowd. After these two chapters, we continue with studies about human navigation behaviors. In chapter 5, we present the study on individual avoidance of groups done in VR. In chapter 6, we present the experiment done to study human gaze during navigation task and the preliminary results of the study. Then the thesis finishes with a conclusion in chapter 7 and a French summary of the thesis in chapter 8. An overview of the work of the two collaborations is given in appendix A.

Background

2

Contents

1	Crowd simulation algorithms	7
1.1	Collision avoidance	8
1.2	Other behaviors	14
1.3	Discussion	16
2	Observation of human behaviors and interactions	16
2.1	Real world observations	16
2.2	Controlled experiments	18
2.3	Evaluation	21
2.4	Discussion	24
3	Virtual Reality	25
4	Conclusion	27

During this PhD, the focus of the work has been on understanding, modeling and simulating human interactions during navigation tasks. The goal was to improve crowd simulation realism: having the virtual crowd behaving as closely as possible as a real crowd (see section 2.3 for comparison methods). This chapter presents the state of the art technology for crowd simulation, the different models and algorithms that have been developed over time. It also presents the various studies performed to improve our understanding of human interactions, the techniques used to perform such study as well as the increasing complexity to study complex behaviors and how to compare simulation results with real data to evaluate the level of realism of a simulation. Finally, as VR has been an important tool during this PhD, a quick overview, of previous work in VR and its use to study some human behaviors, is given.

1 Crowd simulation algorithms

There is a wide variety of algorithms that has been proposed for crowd simulation. A first separation can be done between macroscopic and microscopic models. Macroscopic models aim to simulate a crowd as a whole, for example some approaches consider the crowd as a fluid [Hug03, TCP06]. Doing so, they are able to simulate very dense crowd with good results and small computational cost, but the lack of consideration for the individual behaviors causes a deterioration of the results when simulating less dense situations. On the other hand, microscopic models focus on individual behaviors by simulating each individual and forming a virtual crowd by simulating several of these individuals interacting with each other and with the environment. This PhD has been focused on understanding and modeling interactions between humans in order to improve realism in microscopic

simulation. This is why this overview is focused on microscopic models and does not detail any macroscopic algorithms.

Microscopic models focus on modeling local interactions: how the actions of a human are influenced by its surrounding (walls, trees, other humans ...). There are many different kinds of local interactions. Humans interact between them self by avoiding, following, fleeing, etc. each other. For example, when an individual is confronted to another one that is moving toward him, he can adapt his motion to avoid colliding with the second individual. They also interact with other objects in the environment in a similar way, they can avoid, reach, watch, etc. these objects. Amongst all these local interactions, the collision avoidance is the one that received the most interest. It is the most common interaction as people rarely want to bump into other people or objects. In the first part, we focus on the work done around this collision avoidance interaction and present the principal existing models of collision avoidance. Then, in the second part, we present several models about other interactions. Especially, there are two behaviors that have also received a large amount of interest from the community, namely the following and grouping behaviors. We detail some of the models that focus on the interactions implied by these two behaviors.

1.1 Collision avoidance

Many microscopic models have been proposed to model collision avoidance. We start by presenting the principal models of collision avoidance and detailing some of the algorithms derived from these models. Then, we present some of their limitations and the various extensions done to reduce these limitations.

1.1.1 Principal collision avoidance models

One of the first model of local interactions, the ruled-based model, has been done by Reynolds ([Rey87, Rey99]). It models the interactions by an arbitrary set of rules that define specific actions for specific interactions and state of the environment (positions of the agents). Reynolds defines a separation rule, in [Rey99], to avoid other individuals and an obstacle avoidance rule to avoid static obstacles. These two rules are presented in figure 2.1. The obstacle avoidance rule rotates the agent when an obstacle lies directly in front of it, while the separation rule computes a repulsive force for each nearby agent independently from each other then sum up all these forces to produce the overall steering force.

After the ruled-base model, the social forces model has been proposed by Helbing and Molnár [HM95]. They propose to model individuals as particles subjected to various forces. The motion of the agents is then defined by a sum of forces:

$$\frac{d\vec{v}_\alpha}{dt} = \vec{F}_\alpha(t) + \text{fluctuations}, [\text{HM95}] \quad (2.1)$$

where the agent's velocity \vec{v}_α is computed from various forces $\vec{F}_\alpha(t)$ each defined by an interaction. The collision avoidance interaction is model by a repulsive force, one for each agent around. There are also some algorithms that reduce the disadvantages specific to the two models by mixing social forces and rules [PAB07].




Figure 2.1 The obstacle avoidance rule on the left and the separation rule on the right, both from Reynolds' rule-based model [Rey99]

The algorithms presented so far are based on position and are mainly reactive (move away when too close to a collision). There is another type of algorithms that also considers the evolution of the situation and predicts future collision. Doing so, they become proactive and are able to solve collisions in advance.

The velocity-based model ([PPD07, vdBLM08]), considers the velocity of nearby agents to determine which velocity could lead to a collision. This way, the model is able to solve collision avoidance in advance by defining constraints over the velocity space. For example, the RVO2 algorithm, presented in [vdBGLM11], computes the velocity obstacle $VO_B^A(v_B)$ between agent A and B as the set of velocities v_a for A that will result in a collision with obstacle B moving at velocity v_b . Figure 2.2.1 represents the velocity obstacle $VO_B^A(v_B)$. The current velocity of agent A, V_a , belongs to the velocity obstacle set and the difference between V_a and V_b clearly shows that a collision will happen. These velocity obstacle sets are computed with every other agent then they are all merged into one huge set of velocities that will cause a collision (as presented in Figure 2.2). A velocity outside this set is then chosen to prevent any collision. While these algorithms are adding constraints instead of forces, they still compute individual avoidance strategies for each obstacle and then add them all together.




Figure 2.2 On the left, the computation of the velocity obstacle of agent A that leads to a collision with agent B. On the right, the union of all the velocity obstacles to avoid multiple obstacles.

Warren and Fajen model several interactions by describing the angular acceleration of heading using the bearing angle (the angle at which the target is perceived) and its evolution [WF08]. This way, they model interactions based on variables that can be directly perceived by real humans, which is not the case with the velocity based model. The bearing angle is an important variable influencing human motion, which has been used in other work [OPOD10, SML15]. One important property of the bearing angle is that its evolution gives enough information to judge the collision risk correctly. As illustrated in Figure 2.3-Top, approaching a target with a constant bearing angle leads to a collision. Using this property, Warren and Fajen were able to control the orientation of an agent via a function using distance and bearing angle to intercept or avoid a target. The avoiding function has been modeled and calibrated using experimental data of real humans detouring around a single moving obstacle from [CBW06]. They were able to produce trajectories very close to the mean dominant human paths, these trajectories are presented in Figure 2.3-Bottom. When several targets have to be avoided, each avoiding function for each target is linearly added to control the agent’s angular acceleration of heading.




Figure 2.3 Top: effect of the bearing angle evolution on collision avoidance from [OPOD10]. Bottom: resulting trajectories from Warren and Fajen’s model [WF08] compare to real human data.

The vision based model [OPOD10] exploits the same property of the bearing angle.

But instead of considering the bearing angle of the different objects in the environment, they try to emulate the human behavior by computing collision avoidance strategies based on the optical flow, in imitation of humans and based on cognitive science. They compute the risk of collision for every pixels of the optical flow presented in Figure 2.4 then select a strategy that will globally reduce the risk present on the optical flow. They also compared themselves to previous algorithms showing an improvement in the emergence of self-organized patterns of walkers on several examples.




Figure 2.4 The Optical flow used in the vision based model [OPOD10] with the risk of collision that is a combination of the derivative of the bearing angle, to check for collision, and the Time To Interaction, to check for the imminence of the danger.

These models and algorithms have been constantly evolving to improve the simulated collision avoidance interaction. One of the main issue, with current algorithms, is their difficulties to handle complex situations with many static and dynamic obstacles (including others agents) to avoid. Multiple collision avoidances are mainly handled by integrating all the collision avoidance strategies computed independently for each obstacle. In the social forces model it is a sum of the forces, while in the velocity-based model the velocity constraints are merged to define a limited set of non-colliding velocities. By considering collision avoidance for each obstacle independently of each other and combining them with a simple integration of independent solutions, there is a lack of global strategies to avoid several obstacles efficiently. This leads to some issues appearing when scaled to more complex situations with multiple collision avoidances: agents going through dense group when it is obviously easier to go around, traffic jam appearing at density that real humans are able to handle without too much perturbations, agents' inability to snake through a crowd...

1.1.2 Extensions to complex situations

The limitations of previously presented algorithms are well known and some work has already been done to extend previous algorithms and reduce their limitations. Rule based algorithms can easily be extended with some rules defining some global strategies when facing complex situations with multiple obstacles. This has been done by Shao and Terzopoulos in [ST05]. One rule in particular searches for a gap in the incoming flow

instead of trying to avoid all incoming agent one by one, this rule is presented in Figure 2.5




Figure 2.5 A rule, from the ruled-based algorithm proposed by [ST05], that defines a global strategy to avoid several agents at once in a corridor: pedestrian H searches for a safe direction interval when confronted by oncoming traffic.

The velocity based algorithm RVO2 has been the center of many extensions. Yeh et al. ([YCP⁺08]) added the notion of composite agents: phantom agents that other can perceive and interact with but that do not represent any physical object of the environment. These composite agents can be used to represent some social behaviors. For example by adding a composite agent, that fills the whole exit of a train, agents outside the train, that will interact with it, will no try to enter the train and leave the ones inside exit it. It can also be used to represent several obstacles at once, which represent a new way to combine all the interactions with these obstacles as one big interaction. Golas et al. also merge several obstacles using the lookahead mechanism ([GNL13]). This mechanism allows agents to anticipate the evolution of dense areas and to avoid them by going around. The Figure 2.6 shows how the lookahead applies the new global avoidance strategies to a dense group. Best et al. introduce a similar mechanism with densesense ([BNCM14]), steering the agents toward lower density zone.




Figure 2.6 Effect of the lookahead mechanism [GNL13], which allows the orange agent to see the approaching crowd and adjust its velocity to avoid it as one entity.

Godoy et al. also proposed an extension to the RVO2 algorithm in [GKGG14]. In their extension, they compute several steering strategies for a fixed time period in the future. Then they predict the consequences of each strategy, using RVO2 to predict others' motion according to each strategy, and select the most efficient using an energy criterion. This way, they consider different avoidance strategies by evaluating their impact on all the collision at once instead of doing it independently one by one. The Egocentric Affordance Fields presented in [KSHF09] also checks the impact of a strategy on the future to plan a collision free path using a discrete model. The system is presented in Figure 2.7, the path is computed progressively as the Affordance Fields (circles around the agents containing collision information) are updated with the predicted motion of dynamics obstacles while the path advance in time.




Figure 2.7 The steering algorithm of the Egocentric Affordance Fields [KSHF09]: (a) The current state of the environment. (b) Static Perception Field indicating low traversability at position of obstacles. (c) Dynamics Perception Field for velocity, $S_{target} = S_0$. (d) Dynamics Perception Field for velocity, $S_{target} = S_0 + \delta V$. (e) Affordance Field indicating path of high fitness to goal.

These extensions propose different methods to combine several interactions and compute global strategies to avoid several obstacles efficiently. But most of these extensions are limited to specific situations. In the rule based model, each of the rule defines a global avoidance strategy only for a specific kind of situation and it would require a huge quantity of rules to process all possible situations. Moreover, most of the extensions of RVO2 focus on avoiding dense zones or groups of agents as a whole. Yet sometimes real humans have no other choice then to snake through these dense areas, which they are able to do pretty efficiently. The extension from Godoy et al. and the egocentric Affordance Fields

are able to compute global avoidance strategies in generic situations. But these global avoidance strategies are limited in the immediate vicinity of the agents. In the egocentric Affordance Fields solution, the further one goes from the agent the more details are being lost as the Affordance Fields are growing bigger and bigger. As for Godoy et al., they have to predict the motion of every nearby agent according to a specific strategy. Basically for each considered strategies, they need to run a small crowd simulation which limit greatly the period of time one can consider when evaluating many different strategies.

1.2 Other behaviors

While collision avoidance is the most common and important interaction in a crowd, there are other interactions that remain essential to design a realistic crowd. Reynolds' rules based model [Rey99] was already including some rules beside collision avoidance ones (flee, cohesion, leader following ...). Some behaviors can be specific to a situation, Sung et al. [SGC04] proposed a model that select possible interactions according to the current situation while Pelechano et al. [PAB07] introduced some behaviors specific to emergency evacuation (pushing, falling ...).

There are also some behaviors that are almost as common as collision avoidance. Studies have demonstrated that a regular crowd is composed mostly of groups and that there are only few individuals walking alone. Moussaïd et al. [MPG⁺10], found that 70% of walkers in a shopping street were walking with other people. Grouping interactions are probably the second most important type of interactions in crowds and have received a lot of attention from the community. Another behavior that is very common and has received some attention from the community is the task of following. Indeed, following behavior is used in a lot of different situations: in waiting lines, in dense crowds with no space to overcome the person in front or even in groups where one might have to follow behind another group member. In this subsection, the different work done first on groups then on following will be presented.

1.2.1 Groups

Even before studies measured the importance of groups in crowd, groups were already considered as an important part of crowd simulation as we can notice them every day in crowds. Reynolds was already including some group interactions with flocking rules in his rule based algorithm [Rey87]. Similarly, the group behavior was integrated in the Helbing social forces model by adding several attractive forces [BMdOB03, PV08].

The grouping behavior has been modeled in many different ways. Many define a group as one leader that choose the destination and many followers that try to stay close to the leader [LMM03, YMMT08]. Qiu et al. [QH10] added complexity to such model by defining a leader for each member of the group independently, allowing different kind of groups from the usual one Leader many followers to groups where everyone is following the one in front. Examples of these kinds of group are presented in Figure 2.8. Musse and Thalmann [MT97] also modeled the relation between agents, which will choose to follow other agents dynamically according to their relation and distance to these agents. This allows simulation with groups dynamically forming and changing. Another approach consist of modeling groups as a set of agents walking in specific rigid formations, with some

possible transitions, toward the same destination [TYK⁺09, AMBR⁺12, GD13] much like an army formation or a parade.




Figure 2.8 Different group structure simulated by [QH10].

Recently, much of the work on group behaviors has been focused on small groups of 2 to 6 members. Based on real world data or studies, these contributions reproduce the walking formation of the most common groups that can be seen every day in the streets [LCHL07, PEC09, MPG⁺10]. Moussaïd et al. employed a social forces model-based approach to organize group formations, while Peters and Ennis used discrete formation templates. Unlike the previous rigid formation, these small group formations are dynamic as groups adapt their formation to the surrounding constraints [KO12, HMC⁺12, RY13]: group members walk side by side, in an abreast formation, when there is enough space but will start walking one behind the other, in a river like formation, when necessary.

1.2.2 Following

Following is an interaction present in many different behaviors that are often needed in simulation. In the previous subsection, some group models with agents following one another has been presented. But the following interaction is also present for single individual in waiting line or even when navigating through a crowded corridor where one cannot overtake the people in front. Thus, it is natural to find several contributions focusing on this specific interaction.

There are different methods to model the following behavior, some focus on keeping a specific distance with the leader in front while other matches its speed. There are many variations of these methods, for example the following distance can be adapted to the current speed or a combination of speed and distance can be used. Rio et al. offer a nice overview and comparison of these different methods [RRW14]. For this comparison, they performed several experiments with real humans to evaluate the results of the different method and find which one is closer to the real human behavior. There are also some model derived and calibrated from data of real human following. This is the case of the model proposed by Lemercier et al. [LJK⁺12]. In their model, the follower tries to match the leader's speed and at the same time adjusts its acceleration according to the distance between them. They also demonstrated that the model is able to reproduce the experimentally observed stop-and-go waves when calibrated correctly.

1.3 Discussion

Many different local interaction models have been presented. There is especially a wide variety of models for collision avoidance and even more algorithms. New models and algorithms keep being created to produce more and more realistic simulation. But in order to define which simulation is more realistic, a baseline is needed for comparison. Data on real human interactions is necessary to evaluate the results of a model. It is also useful to detect the limitations of current models and to narrow them to specific kinds of interactions or situations. If we want to continue improving current models and reducing their limitations, we need a deep understanding on human behaviors and interactions and how they handle the situations that current models have difficulties to handle in a human way. This requires thorough studies of real human interactions. The previous work done on such studies is presented in the next section.

2 Observation of human behaviors and interactions

In crowd simulation, observations are very important. Indeed, to be able to create realistic crowd simulations, we need some understanding on how people behave and interact with each other and what are the results of these interactions on the whole crowds. In this section, previous analyses done on crowd observations are presented. We start by considering studies performed on real world observations. Then we present several controlled experiments done to complete and refine real world observations studies. We finish by presenting work that uses such observations and analyses to compare simulated crowds with real ones and evaluate the realism of a crowd simulation.

2.1 Real world observations

Real world observation consists at observing crowds in the real world without any perturbations from the observer (see Figure 2.9). There are many such observations [ZWT12, CLGX12, BK15]. The resulting data have the huge advantage to record human behaviors in their natural habitat without any experimental bias on the observed behaviors. People behave as they would naturally do and the data can capture a wide range of behavior depending for example on the mind state of the walker (people in a hurry, distracted people...) or the social context (walking alone, walking within a group ...).

From such observations, we gain an insight about the global crowd behavior. First, we can analyze patterns that emerge from all these people interacting with each other. For example, groups are an important part of a crowd, we can easily notice them every day. Many studies have been done to determine the different formations of such walking groups (see top of Figure 2.10) and how these formations are adapted to the environment [MPG⁺10, PEC09]. Another example is the formation of lane when two big flows interact with each other (see bottom of Figure 2.10). When such an event happens the space becomes so scarce that people have to follow the one in front going in the same direction, creating line of people going in the same direction. While such patterns are more or less easy to notice visually, they are usually very hard to measure. Some work proposes



Figure 2.9 Pictures from real world observations with A: the train station dataset from [ZWT12], B: the mall dataset from [CLGX12], C: the one year recording dataset from [BK15]

measurements to describe such emerging patterns. For example, Zhou et al. measure the collectiveness of a crowd [ZTZW14]: whether the crowd can be considered as a big entity moving as a whole or as individuals moving independently from each other. This is particularly important to perform objective comparisons, which is greatly needed to evaluate crowd simulator (see section 2.3).

Real world observations are also useful to get some statistical values describing crowds or individuals. For example, these statistical values can be used to trace fundamental diagrams [DDH⁺14a]. These diagrams describe the relation between the density of the crowd and its average speed such as illustrated in Figure 2.15. More specific measurements, such as the distance between people walking together [Cos10], can also be analyzed on real world observations. These statistical values are important to calibrate algorithms by tuning correctly their parameters [BBQ⁺07]. Indeed with bad parameters, such as having the agents' default speed too fast or their desired distance to group members too big, the simulation will never be realistic no matter how good the algorithm is. Finally, trajectories can be extracted from such observations. These trajectories can then be used for comparison (see section 2.3) or integrated into simulations to improve realism [AGS⁺11].

However, using data from real world conditions does not provide any control about the recorded situations and their associated influencing factors. For example, [BK15] studied the same mall corridor for an entire years and show that people behave differently according to the current time or day. This can be explained by the fact that people goals and objectives vary: during rush hours people go to work, during lunch hours people look for a place to eat and during vacation days people are simply shopping or relaxing. There are many influencing factors and unknown variables associated with these data from crowd motion in everyday life conditions. To overpass these limits, authors have




Figure 2.10 Top: groups observed in Trinity College from [PEC09]. Below: automatically detected line of people going in the same direction from [ZWT12]

proposed to develop controlled experiments to study specific situations.

2.2 Controlled experiments

Several experiments have been performed to study specific situations using controlled conditions. In comparison with real world observations, researchers setup the experiment in such a way they get control over the initial conditions as well on some additional knowledge about the participants involved in the experiment (such as their goal or their personal characteristics). Even if this control constrains the observed situation (the specification of the task is fundamental in such an experiment), it allows to isolate a behavior or a specific variable while limiting the effect of uncontrolled factors. A lot of studies have been developed using such controlled situations and concern a wide variety of human behaviors and tasks.

There are experiments with only one human interacting with a static environment ([FW03, BJ03]). These experiments focus on human trajectories performed while avoiding static obstacles or following a specific path constrained by the environment but without any perturbations from other people. Brogan and Johnson did such an experiment to study real human trajectories in different situations: with no turn, with a sharp turn, with a S turn... They were able, for example, to model the variation of the walking trajectory: how and when people decelerate when reaching the goal or having to perform a sharp turn.

Then there are experiments with two humans interacting with each other ([POO⁺09, OMCP12, OMC⁺13, HSK⁺14]). These experiments really focus on the interactions between the two individuals, while some influencing factors, such as the angle of approach, can be constrained by the created environment. The interactions between the two humans usually happen in zones free from all other perturbations from other humans or static obstacles. While making such experiments can be fastidious, one can extract




Figure 2.11 Trajectories of people walking through a constrained environment from [BJ03]

much information from such data. For example, the experiment presented in Figure 2.12, lead to lots of analyses and conclusion on collisions avoidance between two individuals [OMCP12, OMC⁺13]. In their experiment, the researchers had two participants walking on perpendicular trajectories with an initial setup done to maximize the chance of a collision risk to appear. Walls were set to hide the participants from each other at the beginning of their walk so they had time to reach their normal walking speed before interacting. As a tool to analyze the data, the authors introduced the MPD value (Minimum Predicted Distance): the minimum distance the two individuals will be from each other if they keep their current speed, as shown in Figure 2.12.c. With this value, they were able to show that people adapt their trajectories only when there is an actual risk of collision. Moreover, when an adaptation is necessary, it is done in 3 phases: the observation phase (people are evaluating the risk of collision), the adaptation phase (people adapt their trajectory to avoid the collision by increasing the future crossing distance) and the regulation phase (the collision avoidance task is solved and walkers keep a stable future crossing distance). According to authors, this latter phase shows that the collision avoidance task is solved by anticipation. By further analyzing the data, they were able to demonstrate that collision avoidance is asymmetric with the one giving way having to adapt more than the one passing first.

While the more people there are in an experiment the more difficult it is to make, there are also many experiments with several participants. These experiments focus on all kind of situations or factors such as the fundamental diagrams [SSKB05], flows in bottlenecks [DH03, SSW⁺09], walking in line [JARLP12, LJK⁺12, DDH⁺14b], line forming in crossing flow [PCBS11], walking with other people [BRCW12]... The analyses done in such experiments are of the same kind as the one done on real world observations: studying patterns (stop and go wave when walking in line) and statistical values (the flow capacity with different bottlenecks sizes.) An experiment studying bottlenecks from [DDH⁺14b] is presented in Figure 2.13. With this experiment, they were able to measure different statistical values such as pedestrians' velocity, density and flow rate. They also plotted the trajectories of the people, as shown in the Figure, to study patterns such as lane formations. The control over influencing factors allows to perform more precise analysis than the one performed with data from real world observations. Moreover, the extra




Figure 2.12 Description of the experiment from [OMCP12]. A: Experimental setup. Area is $15m$. Two participants stand at the corners of the area and are synchronously given a start signal. Their task is to walk to the opposite corner. They implicitly start an interaction to avoid any collision. B: Picture taken during experiment. C: Schematic illustration of the minimum predicted distance (MPD) computed at three different times with a motion adaptation occurring between times t_2 and t_3 .

knowledge about the participants is also useful for the analysis. For example, Chattaraj et al. were able to find differences in the fundamental diagram between population of different cultures [CSC09] (the differences are shown in Figure 2.15).

While these studies provide detailed and robust results, having many participants increase the difficulties to manage the experiment and to limit the effect of non-studied and uncontrolled factors. Considering two participants, the experiment to study collision avoidance behavior can be designed to increase the probability of a risk of collision to be sure to observe collision avoidance maneuvers with respect to participants' comfort speed. Due to human motion variabilities or the reaction time to starting signals, the risk of collision can not be forced and might not happen during some trials. One key point is to record enough trials to generate enough data with collision avoidance maneuvers. The complexity of the experimental design however increases with the number of participants.




Figure 2.13 Counter flow study from [DDH⁺14b]. Top: snapshot of the experiment. Bottom: Trajectories of pedestrians within the channel in each experiment scene.

2.3 Evaluation

While these data from real world or controlled observations are really useful to improve our understanding about crowd behavior, there are also really useful to evaluate crowd simulators and algorithms. Indeed, a lot of applications for crowd simulations requires a certain level of realism. It is then important to compare the new proposed models to real human behaviors.

There are several ways to evaluate or compare crowd simulations. The easiest one is to perform a visual evaluation. The researcher can check whether or not the simulation looks realistic, whether some strange behaviors happen (collision, unrealistic jam). Some situations, which are difficult to reproduce, are often used for this kind of visual evaluation, such as the circle or bidirectional flows presented in picture 2.14. In some cases, this is all one needs from a simulation: to simply look real (for movies for example). In these cases, a simulation of a recorded real crowd can be done to visually compare the simulation with the real crowd [GCLM12] or the simulation can be evaluated by several people judging whether it looks real or not. This kind of evaluation can even be improved by immersing the people in the simulated crowd using VR technologies [PSAB08, RY13]. Otherwise, researchers can design users studies where naive participants evaluate the realism of the simulation in comparison with other ones. For example, Ahn et al. immersed participants several times into the same virtual environment but with virtual humans control by different algorithms [AWTB12]. By asking participants to score the lifelikeness of the crowd, they demonstrated that their algorithm produces crowds that felt more real than the other algorithms.

This first evaluation, even if it has perceptual relevancy, is subjective. Using the data from real crowd observation, the second way to evaluate the situation is to determine qualitative and quantitative characteristics of the real crowd. Indeed, we know of different emerging patterns that appear in a crowd under specific conditions. By simulating these kinds of conditions we can check whether similar patterns emerge from the simulated crowd [HFV00]. We can thus check if a simulator is capable of reproducing stop and go waves [LJK⁺12], lanes formation [PPD07, MHG⁺09], circular clogging in front of a door




Figure 2.14 Two typical situations used for visual evaluation. On the left, it is a circle with everyone trying to switch places with the one on the opposite side. On the right, it is two flows on a frontal collision course with each other. In both cases, we can see some models having difficulties to produce fluid movement that would look more natural. These comparisons are from [OPOD10] and the "our model" correspond to the vision based model.

[HFV00]... These qualitative comparisons can also be done at the level of the individual. Pettré et al. used their results on a three phases collision avoidance, presented earlier, to check if virtual humans in simulation have similar phases [POO⁺09].

Qualitative comparisons are usually a lot easier to do than quantitative ones but are also more subjective. Thus it is important to perform some quantitative comparisons too. While emerging patterns can be qualitatively defined and visually detected, some measurement of such phenomenon have been proposed [ZTZW14, CARH13]. These values describing emerging patterns can be used to quantitatively compare patterns between simulated crowds and real ones [LJK⁺12]. Crowd analyses also yield many statistical values used to describe a crowd. These statistical values can be used to compare simulated crowds and real ones [BBQ⁺07, VLL⁺14]. One of the most used statistical values is the fundamental diagram. Such diagram has been measured in many different situations which can be simulated to check whether the simulation is able to reproduce the same trend. Such comparison is presented in Figure 2.15. More precise comparison is difficult as crowds are huge entities formed of many individuals and with many factors influencing its behavior.

One solution, for more subtle comparisons, is to compare individual trajectories [BJ03, POO⁺09]. But due to human motion variability, two trajectories can differ for the same situation. Performing a comparison that is too precise would lead to rejecting many trajectories that could have been done by real human. The question to answer is to




Figure 2.15 Fundamental diagram plotted for different algorithms with optimized parameters to suit different cultures [WJGO⁺14].

identify when these differences demonstrate that the evaluated trajectory is not natural. To be able to analyze differences in such precise comparison, we need to find some common characteristics shared by real humans when navigating and interacting with others such as the inertia when turning [BJ03], vorticity (as defined in fluid mechanics) [WJGO⁺14], energy expenditure [GCC⁺10]... Lerner et al. tackle this issue with a huge database, when evaluating simulated trajectories they search in a database for trajectories that are close for situations that are similar [LCSCO09, LCSCO10]. Guy et al. propose to evaluate simulated trajectories by simulating a real situation and, for each simulation step, they would measure the difference in trajectories between the real and simulated one before resetting virtual humans' position and speed to the one of their real counterpart [GvdBL⁺12]. This way, they prevent the snowball effect of the human motion variability that can lead to huge differences in trajectories that are not necessarily significant to an unnatural trajectory.

To summarize, there are many ways to evaluate crowd simulators, each of them with their own advantages and limitations. Then, when performing such an evaluation, it is interesting to use as many evaluation methods as possible. For example, in [GCLM12]

they used visual comparison, checked for emergent patterns, compared the fundamental diagram and compared the path of individuals in one specific situation. Moreover, real world observations clearly show that people have some differences in their behaviors according to their culture, state of mind and situation. Thus, comparing a simulated crowd with one specific set of data only gives an idea on how realistic the simulator is in this kind of situations. To globally evaluate a simulator, many different situations need to be tested, thus the creation of benchmark with a list of situations to test and metrics to compare [SKFR09, WJGO⁺14]. The Steerbench workflow is presented in Figure 2.16 and show that for an evaluation, many situations are simulated and many comparisons methods are used (Several kinds of metrics as well as visual inspection). Wolinski et al. also add an optimization process in their benchmark to ensure that the parameters of a crowd simulator are properly set and suit the simulated situation. As presented in Figure 2.15, this allows to find the proper parameters to simulate different populations.




Figure 2.16 Steerbench workflow [SKFR09].

2.4 Discussion

This section has demonstrated the importance of the observation of human interactions for crowd simulation. These observations improve our understanding of these interactions, which helps create better algorithm to simulate them. They are also used to evaluate the level of realism of a simulation and allow to compare different algorithms. Both types of observations are of equal importance: the global observation from the real world and the

more precise observation from experiments. But due to the difficulties to manage complex situations during experiments, there is a lack of precise observation about complex interactions. These difficulties could be eased by using VR technologies. Virtual environments are far easier to control than real ones. VR is already used to study human behavior in various domains, these studies are presented in the next section.

3 Virtual Reality




Figure 2.17 Left: subject wearing HMD and physiological monitoring from the experiment in [MIWB02]. Right: subject in a cave facing a virtual thrower in the handball case study from [BKV⁺10].

Virtual Reality has been around for a while. The recent technology improvements however have brought a renewed interest for the domain. With new devices such as Head Mounted Display (HMD) or Cave Automatic Virtual Environment (Cave), shown respectively on the left and right in Figure 2.17, immersion (the feeling to be physically present in the virtual environment) is greatly improved. Many studies tried to evaluate the level of immersion that can be reached with these devices, the most common evaluation being the presence questionnaire [WS98, SvdSKvdM01] that tries to determine how much the user feels like physically being in the virtual environment. This evaluation has also been supplemented by physiological measurement, for example physiological reactions when facing a stressful situation in a virtual environment [MIWB02]. Immersion is very important as it has been observed to directly influence participants' behaviors. The higher the immersion, the closest to real world participants' behavior is [SSUS00, GSPR05]. While the most widespread usage is for entertainment, the level of immersion has allowed people to use VR to study human behaviors [BKV⁺10, MGvVB98, MCRT06] or even as a tool to provide health care services [Her05, KAC⁺09] or to teach people [TLvK⁺12].

VR is a powerful tool for perception-action experiments [LBB99]. VR-based experimental platforms allow exposing a population to fully controlled stimuli that can be repeated

from trial to trial with high accuracy. Factors can be isolated and objects manipulations (position, size, orientation, appearance...) are easy to perform. Stimuli can be interactive and adapted to participants' responses. As we have seen in the previous section, doing controlled experiments on complex crowd is very difficult because of all the parameters to control which are impossible to insure once the experiment has started. With virtual reality, we hope to improve the study of complex crowd as much as controlled experiments improved the study of simpler interactions over real world observations.

But before using VR as a platform to conduct experiments on human navigation and interactions in complex situations, the platform itself needs to be evaluated. While the immersion offered by such platform as already been studied in details, we need to make sure that people interact the same way in VR then in real life and that they have the capacity to do it. After all, studies have shown that there are differences in perceived velocity [BSD⁺05] or personal space size [GLRFM08] in VR and walking in VR is performed with increased instability [HBBK07]. It has also been reported in several studies that distances are compressed in VR [LK03, WCCRT04, RVH13] but this effect is reduced after five minutes of continuous visual feedback [MCRTB10].

Validation of the VR platform has already been started, with comparison between real life experiments and their counterpart in VR, as shown in Figure 2.18. Studies about the interaction between one individual and obstacles [FFW07, COMP13] have shown that there are quantitative differences between real world trajectories and the ones performed in virtual reality. Despite these quantitative differences, trajectories had qualitatively similar shapes with some fundamental characteristics of real locomotion. Both studies conclude that VR platform provides a valuable research tool for the study of visually guided locomotion as long as one is careful about the quantitative differences. As for interactions with moving obstacles, comparison has been done on the perception of the risk of collision in virtual reality with cars [SAB07] or virtual humans [OOP⁺10].




Figure 2.18 The evaluation framework, proposed in [COMP13], compares trajectories between an experiment performed in the real world and the same experiment performed in a virtual environment.

There are already a few contributions using VR for crowd simulation. Pelechano et al. used it to immerse participants in a simulation and evaluate the simulation by measuring the presence felt by the participants ([PSAB08]). A similar evaluation method was used in [RY13] to evaluate simulated social group movement and in [AWTB12] to compare different algorithms. Fajen and Warren used it to study the dynamic of steering and

obstacle avoidance, analyzing the trajectories of participants walking freely in an empty room, with a goal and a static obstacle to avoid, through an HMD ([FW03]).

But VR is not yet widely used to study human navigation and interactions despite its numerous advantages. Especially, as seen in the previous section, performing experiments in the real world can be very difficult. Because of these difficulties, there is a lack of precise observations of human interactions during complex situations. This section has demonstrated that VR possesses many features that would ease the experiment process for such complex situations. Using such a tool would allow us to study human interactions in complex situations with more ease and thus gather many useful data to improve crowd simulation.

4 Conclusion

In this chapter, we have presented the state of crowd simulation. There are many existing algorithms to simulate a crowd, and these algorithms can vary significantly. Despite this variety, there is a common problem amongst most of these algorithms: their difficulties to handle complex situations in a realistic way. Most of these algorithms combine multiple interactions in a similar simplistic way causing some unnatural behaviors in complex situations. Some extensions have proposed some methods to combine interactions in specific complex situations and demonstrated improvement in those situations. But there is a lack of a general combining method to improve the results for most complex situations.

This chapter has also demonstrated the usefulness of human observations to improve current algorithms. Such observations could be used to find better ways to combine multiple interactions. While there are some real world observations of complex situations, real world data are not precise enough and need to be completed by some experiments. But due to difficulties in managing complex situations in experiments, such experiments are lacking. These difficulties could be solved by using a VR platform to perform experiments, which has been demonstrated by the many uses of VR to study human behaviors presented in this chapter.

Simple study case: following model

3

Contents

1	Introduction	30
2	Following model	30
2.1	One-dimensional following	31
2.2	Two-dimensional following	34
3	Results	34
3.1	Following distance	35
3.2	Waiting line	35
3.3	Corridor	36
4	Evaluation	37
4.1	Simulation of real situations	38
4.2	Stop-and-go waves	39
4.3	Comparison with other models	41
5	Conclusion	42

1 Introduction

In this chapter, we present the first contribution of the PhD which was done in collaboration with Teofile Dutra. For this first contribution, we worked on following behaviors. In this work, we limited the following behaviors to one target only and the simulations to one dimension. This simplification allows us to focus on how does the follower adapt its motion to the target: the differences in behaviors when following a target that is fast or slow, accelerating or decelerating or having a stable or fluctuating speed. An obvious example would be with cars on a highway as they are required to keep more space with the car in front the faster both car are moving. It is also obvious that we don't follow a drunk man (who has a very random movement) the same way as we follow an old lady (who has a slow and constant movement).

Following someone is to keep a small distance with the moving person (the target) in front and/or to match its speed. Most advanced simulation algorithms suggest that speed matching is controlled by some accelerations, the amplitude of which is modulated by a distance and a density terms [LJK⁺12, RRW14]. In this contribution, we inspect the role of some other factors. We hypothesize that human also explicitly control the following distance from the moving target. This distance results from various aspects of the target's characteristics as well as its motion. We consider that the following distance results from: *physical factors* to avoid contact, *social factors* to keep respectful distances to others, *psychological factors* to take into account reaction times or level of attention, and other factors related to others' *motion perception* to filter out motion fluctuations, that should not be matched, from more long term changes of the target's speed.

The main contribution in this chapter is the design of a numerical model of local following interactions based on an explicit following distance term. One of the main differences to previous approaches is that we use a non-constant distance term to adapt the behavior to the target. In our model, the following distance is a function of the target's motion and some other properties to take into account the previously listed factors. We keep a relatively simple formulation and demonstrate some significant improvements over previous techniques.

The remainder of this chapter is organized as follows. In Section 2, we present our new following model. Section 3 shows the capabilities of our model on some specific situations, and, in Section 4, we evaluate the model with respect to the emergence of patterns observed in real following situations and compare it to real data and previous models. Finally, in Section 5, we draw some conclusions.

2 Following model

This section describes our algorithm to simulate following behaviors. As an input, our technique considers a pair of agents: a follower and a leader. The objective is to compute at each time t the distance $d(t)$ the follower should keep to the leader (Figure 3.1). The distance means distance between agents' centers, modeled as simple circles. When the current distance between agents is much higher than the desired $d(t)$, the follower tries to match its own comfort speed. In the opposite case, the motion is constrained and the

follower's speed adjusted to match the computed distance.




Figure 3.1 The ideal distance between the follower and the leader is described by the difference between the leader's estimated position and the sum of the follower's fixed distance and a dynamic safe distance determined according to the leader's movement.

2.1 One-dimensional following

We consider several factors to compute the following distance between agents' centers. The distance $d(t)$ is decomposed into three terms: d_f , $d_d(t)$ and $d_{err}(t)$.

d_f is a constant term. It results from the sum of the *contact distance* (the sum of the agents' size) and a *personal distance* [Hal66, Hay83]. This distance is directly deduced from simulation parameters.

$d_d(t)$ is our second term. It is a safety distance which increases with respect to the current walking speed of the follower, $v_f(t)$. When walking, people can stop in one step if needed as there is no imbalanced posture in the process. We compute this distance as a reaction distance, i.e., the distance walked by the follower during the time needed to react, ttr , is:

$$d_d(t) = v_f(t) \cdot ttr . \quad (3.1)$$

$d_{err}(t)$ is our third and final term. At this point, we computed an estimation of the motion of the leader out of fluctuations. The distance $d_{err}(t)$ will allow the follower to match the estimated speed with no risk of collision due to the leader's fluctuation of speed. It simulates how a follower perceives a leader's motion. Walking speed can fluctuate over time. For example, children may have a jerky motion with lot of fluctuations. At the opposite, elderly have a more steady walking speed. In real life, followers filter out these fluctuations and avoid matching a constantly changing speed. They also perceive the

intention of their leaders to change their speed, i.e., accelerations that are not due to fluctuations only. To enable a realistic simulation of following behaviors, we decompose the leader's motion into two components: fluctuations, and average *theoretical* motion. To this end, we consider the leader's motion over a short time moving window: $[t - w; t]$. We fit a polynomial $\hat{p}_t(u)$ of degree 2 in a least squares sense to the leader's position $p_l(u)$ over this time windows. Our distance term $d_{err}(t)$ is directly deduced from the observed mean magnitude of the fluctuations:

$$\hat{p}_t(u) \approx p_l(u) \quad (3.2)$$

$$\epsilon(u) = p_l(u) - \hat{p}_t(u) \quad (3.3)$$

$$d_{err} = \bar{\epsilon} + 2\sigma_\epsilon \quad (3.4)$$

where $u : t - w \mapsto t$, and $\bar{\epsilon}$ and σ_ϵ are, respectively, the mean and standard deviation over the residuals $\epsilon(u)$ of the polynomial fit $\hat{p}_t(u)$.

Note that the follower should now follow this estimated motion and not the actual one. To this end, we compute the follower's speed to shadow a virtual position $p_v(t)$:

$$p_v(t) = \hat{p}_t(t + \Delta t) - d(t) \quad (3.5)$$

where

$$d(t) = d_f + d_d(t) + d_{err}(t) \quad (3.6)$$

and Δt is the time step. Then, the follower's shadowing speed can be computed from its current position $p_f(t)$ as follows:

$$v_f(t) = \frac{p_v(t) - p_f(t)}{\Delta t} . \quad (3.7)$$

To simplify the computations, we set the variable $D(t)$ as:

$$D(t) = (\hat{p}_t(t + \Delta t) - p_f(t)) - (d_f + d_{err}(t)) \quad (3.8)$$

$$p_v(t) - p_f(t) = D(t) - d_d(t) , \quad (3.9)$$

which allows us to replace (3.9) in (3.7), resulting in:

$$v_f(t) = \frac{D(t) - d_d(t)}{\Delta t} \quad (3.10)$$

$$\Delta t \cdot v_f(t) = D(t) - v_f(t) \cdot ttr \quad (3.11)$$

$$v_f(t) = \frac{D(t)}{\Delta t + ttr} . \quad (3.12)$$

Finally, we can use this new speed $v_f(t)$ to update the follower's position.

Algorithm 1: Overtaking rule: if the follower has a preferred speed greater than the leader's one, it moves to the closest space available beside the leader, if any.

Data: Current speed $v_f(t)$, Own's preferred velocity $vpref_f$, Leader's preferred velocity $vpref_l$

Result: Add lateral speed to $v_f(t)$

```

1  $v_{lateral} \leftarrow 0$ ;
2 if  $vpref_f > vpref_l$  then
3    $spaceLeft \leftarrow spaceRight \leftarrow 0$ ;
4   if  $hasSpaceToOvertakeLeft()$  then
5      $spaceLeft \leftarrow 1$ ;
6   end
7   if  $hasSpaceToOvertakeRight()$  then
8      $spaceRight \leftarrow -1$ ;
9   end
10  if  $spaceLeft$  and  $spaceRight$  then
11    if  $p_f(t) - p_l(t) > 0$  then
12       $spaceRight \leftarrow 0$ ;
13    else if  $p_f(t) - p_l(t) < 0$  then
14       $spaceLeft \leftarrow 0$ ;
15    else
16      Choose a side randomly;
17    end
18  end
19   $v_{lateral} \leftarrow (spaceLeft + spaceRight) \times 0.5$ ;
20 end
21  $\vec{v}_f(t) \leftarrow (v_f(t), v_{lateral})$ ;
22 return  $\vec{v}_f(t)$ ;

```




Figure 3.2 If there is space beside the leader, the follower accelerates in the perpendicular axis towards that side with space available.

2.2 Two-dimensional following

Despite being a model for simulating 1-D traffic, we can adapt it to handle 2-D traffic as well. In this case, we define separate rules for each axis in our scenario. In a corridor scenario, for example, the following model is responsible for the movement in the corridor-axis, whereas an overtaking rule controls the movement in the perpendicular axis through the use of a lateral speed (see Figure 3.2).

We define a simple overtaking rule as follows: after computing the follower's speed $v_f(t)$, we verify if the follower's preferred speed v_{pref_f} is greater than the leader's preferred speed v_{pref_l} ; if so, we verify if there is enough space to overtake in any of the sides; finally, when there is space, the agent moves towards the chosen side with a lateral speed of 0.5 m.s^{-1} (Figure 3.2). The pseudocode is shown in Algorithm 1, where $p_f(t)$ and $p_l(t)$ are the positions of the follower and the leader respectively at time step t .

3 Results

In this section, we show the results of our model. We start by showing the impact of each distance on the following behaviors. Then, we show trajectories from a waiting line simulation. And, finally, we demonstrate the results obtained in a corridor scenario.




Figure 3.3 Visualization of the following distance computed for a leader with very stable motion (left) or with jerky motion (right).

3.1 Following distance

We have two scenarios to illustrate the effect of each distance on the following behavior. In both scenarios, we simulate one agent following the predetermined motion of a specific leader. First, we use a stable motion: the leader is moving at constant speed and then slows down to stop. This motion is important, it shows the behavior of the model during a basic following situation and its reaction to the leader change of speed, demonstrating how the computed distances prevent any risk of collision. Then, we generate a jerky motion: the leader is moving with an average constant speed and a white noise signal is added to its speed. This motion has been chosen to see the result of the model when facing noisy trajectories and to check that the followers keep a safe distance despite the leaders' speed fluctuations. The trajectories produced are shown in Figure 3.3.

In both scenarios, we can see the dynamic distance d_d changing with respect to the follower's speed. From the stable motion scenario, we can see how this d_d distance absorbs both the delay in realizing that the leader is slowing down and the uncertainty in prediction caused by this change. Without this distance, the agent would have dangerously come close to the leader and would have had to back up to regain a normal distance to him, which would have looked very unnatural.

The error distance d_{err} is almost absent from the first scenario as the motion is very stable and there is no need for it. However, it has a very important impact in the second scenario. We can see that the agent has a motion smoother than the leader's one as it is following a filtered version of its motion. The d_{err} distance ensures that the agent keeps a safe distance to the leader. We can see that the leader's trajectory is rarely below the red zone. Which means that the actual leader's position is rarely closer to the follower than d_{err} from the predicted position.

3.2 Waiting line

The model is used to simulate a waiting line scenario. This scenario has been chosen for several reasons. First, the following distances are well represented in a waiting line as the

followers stop behind their leaders. Moreover, there is a lot of stop-and-go initiated by the people leaving the front of the line and propagated among the people behind. The distances and the stop-and-go propagation are well shown in the companion video¹ and in Figure 3.4, both looking realistic. We can also notice the varieties in behaviors easily set by having different parameters on the model distances among agents as well as clamping speed and acceleration to different values. Finally, this simulation also demonstrates the capacity of our model to simulate specific situations such as the creation of a waiting line from a totally empty line.




Figure 3.4 Trajectories of people lining up and waiting. Stop-and-go waves are created and propagated when the first walker in the line moves.

3.3 Corridor

We simulated traffic in a corridor scenario. This scenario was chosen to illustrate the two-dimensional extension of our model by using the overtaking rule explained in Section 2.2. Our scenario is a periodic corridor with 30 m of length and 9 m of width. The agents traverse the corridor from left to right. When an agent reaches the end of the corridor at the right, it reenters the corridor at the left. The simulation starts with 180 agents stopped. Each agent starts moving according to its leader's motion, if it has a leader. Otherwise, it moves according to its preferred speed which has been randomly set before the simulation starts. The leader of an agent is the agent immediately in front of it. To identify the leader, we consider the periodic aspect of the corridor.

In this scenario, the restriction imposed by the period of the corridor does emerge traffic jams. Agents with low preferred speeds block the traffic holding groups of agents behind them, and these in turn have to adapt their velocities according to the traffic and to try to overtake. The speed variation and propagation as well as the overtaking are shown in the companion video². This scenario is illustrated in Figure 3.5 (center and

bottom).



Figure 3.5 Two scenarios. Top: waiting line scenario simulating 1-D following traffic. Center: top view of the periodic corridor simulating 2-D following traffic by using an overtaking rule. The agents in red are walking fast, they will slow down when facing agents blocking their way and try to overtake them. Bottom: perspective view of the periodic corridor.

4 Evaluation

In this section, we evaluate the model. We used real data composed of several situations which were gathered by [LJK⁺12] and further analyzed in [JacARLP12]. First, we simulate each situation to check that our model produces results close to the real ones. Then, we focus on stop-and-go waves which are often created during following situations. Finally, we also show improvements over the Lemerrier's model [LJK⁺12].

²<http://youtu.be/JbQe19DyUMc>

4.1 Simulation of real situations

The real data consist of trajectories of people walking in-line in a circle way. Different densities have been studied which gives us a good panel of behaviors from stop-and-go waves to unconstrained stable speed. For each experimental trajectory, a corresponding simulation has been performed with similar initial conditions: same circle perimeter, same initial distance between people and same starting speed (0 m.s^{-1}). Moreover, some constraints on agent speed and acceleration have been added to match the participants capacities and the d_f parameter was set according to participants distance to their leader at the beginning.

An example of simulated trajectory compared to real one is presented in Figure 3.6. The parameter was manually tuned to give nice results: $ttr = 0.6 \text{ s}$ and $\omega = 0.5 \text{ s}$, and the simulation was run with $\Delta t = 0.1 \text{ s}$. We obtain a similar wave and in both cases the walkers have traveled around 27.5 m . Moreover, noise has been added to the follower speed in our model. First, it produces curves much more realistic than perfectly smooth ones. But it also has a great impact on the wave formed during the simulation. Waves are more similar to real ones with random noise.




Figure 3.6 Trajectories from experiment trial #25 (left) and the simulation results from our model (center) and Lemerrier’s one (right). Lemerrier’s model needs some initial motion to work so it has been initialized after 20 s of real motion. In our model, we just need to place the walkers at the initial positions extracted from the real data.

Once every trajectory has been simulated, we compute a local fundamental diagram, as described in [JacARLP12], for both set of trajectories: real and simulated. The fundamental diagram is a common tool to compare the average speed according to density. Results are displayed in Figure 3.7. The model produces speeds close to real ones. From the different diagrams, we can notice that ttr influences the speed of the agents. While the value 0.75 s seems to globally fit best the real data, it seems that smaller value of ttr is best for smaller density values.

In conclusion, our model produces following behaviors close to real ones. We find similar patterns in our simulation compared to real data: close average global speed and same kind of waves emergence and propagation. Moreover, the model is able to produce trajectories that are very close to real ones.




Figure 3.7 Fundamental diagram computed for real data and their simulated counterpart with different ttr (0.5 s, 0.75 s and 1.0 s).

4.2 Stop-and-go waves

A very common phenomenon that can emerge from following behaviors is the formation of stop-and-go waves. This type of wave can easily be spotted, we have already shown some in the waiting line simulation (Section 3.2). Quantified data can be extracted from waves for better analysis: the propagation speed from the leader to the followers, the amplitude and how fast are the waves damped or even amplified.




Figure 3.8 Waves detection on trial #25 of the real trajectories database. The upper plot shows all the trajectories with the detected minima as black stars and the defined waves propagation as black lines. The lower plot shows the speed of the last person in blue, the smooth speed in red and the detected extrema as red stars.




Figure 3.9 Relation between waves and density from real data. (a) The wave score according to the density for every real trajectory and (b) the results of three simulations with different densities, showing waves emerging for the one with density around 1.5 ped.m^{-1} , as expected.

An algorithm has been developed to detect waves automatically and to extract quantitative information. It starts by smoothing speed signals and then finding all local minima and maxima for every trajectory. Then, it checks if they are part of a wave and links them together as shown in Figure 3.8. A score is computed as follows: $score = \overline{Ampl} \cdot W/T$ with \overline{Ampl} , the average of the amplitude differences; W , the duration of the wave; and, T , the duration of the simulation. This score allows us to define a specific threshold to separate simple perturbations from waves as it is not always easy to categorize them when waves have very small amplitude or very short duration. This way we make sure that every trajectories are categorized the same way.

This score has been used on real trajectories. By correlating the score to the trajectories, we were able to divide the trajectories into three categories: trajectories with no perturbations at all for score around 0, trajectories with some perturbations for score below 0.12 and trajectories with waves for score above 0.12. In Figure 3.9a, we present the score of each real trajectories according to the density and also show an example of trajectories for each category.

Studying the figure, we can notice that waves appear for a certain level of density only. When density is too low, there is not enough constraints on the walkers' speed to create any perturbation. When density is too high, there is too much constraints on the walkers' speed to allow high enough speed for waves creation. From the real data, waves appear only with density around 1.5 ped.m^{-1} .

We check if the model is able to reproduce the same results: will a density around 1.5 ped.m^{-1} produce a wave while lower and higher produce respectively no perturbation and some perturbations too small to create real waves? We did three simulations with agents starting close to each other with three different densities: $1.9048 \text{ ped.m}^{-1}$, $1.5385 \text{ ped.m}^{-1}$ and 1 ped.m^{-1} . The result is shown in Figure 3.9b and, as we expected, waves

only appear for the density close to 1.5 ped.m^{-1} . Moreover, the wave score for each simulation corresponds to the wave score of real data for similar densities: 0 for the very low density, 0.21538 for the density around 1.5 ped.m^{-1} and 0.063143 for the higher density. Showing that even if no waves are formed for the high density, there is still some perturbations like in the real data.




Figure 3.10 Simulations of agents following a leader with a rectangular signal speed with some noise. Reynolds: in Reynolds’s model there is no security distance, the agent can be closer to the leader than the desired distance leading to collision risk. Middle: in our model the followers keep close to the leaders with a security margin to prevent getting closer than a fixed distance. Right: in Lemercier’s model the following distance increases over time.

Noise has been added to the model speed to produce the three set of trajectories in Figure 3.9b. It has already been confirmed in Section 4.1 that noise helps to produce more realistic waves. It seems that noise has an amplification effect on waves, indeed without noise the waves score from the simulation are all lowered.

In conclusion, our model is able to produce waves starting from a stopped situation. There is no need for any motion initialization as required by some other models (see Section 4.3). Moreover, the density conditions for waves to appear with our model match the ones needed in real life.

4.3 Comparison with other models

Our first comparison will be with Reynolds’ model. Both models are based on a following distance. Unlike our model, Reynolds’ model uses a fixed distance to follow a leader. This distance is set up as a parameter and is equivalent to our fixed distance d_f . An example of trajectories computed with Reynolds’ model is presented in Figure 3.10. The bottom graph shows that the following distance also vary according to the speed. This is due to the speed computation as the agent linearly slows down to 0 m.s^{-1} when closing on its ideal following position. This mechanism has the same effect as our safety distance $d_d(t)$ but in our model we have a direct control over this distance which makes it easier to calibrate it with real data. The main feature that is missing in Reynolds’ model compared to ours is the prediction of the leader’s position and the distance d_{err} to absorb noise. The distance graph in Figure 3.10 shows that the follower in Reynolds’ model can be closer than the fixed distance (set up at 0.6 m) and can cause some collisions (around time 33 s ,

the distance goes down below $0.5m$). The same pick of low distance is present on the distance graph for our model but thanks to $d_{err}(t)$, the actual distance stay above d_f .

Lemercier's model is based on speed matching. The follower will match the leader's speed according to a function that has been calibrated to be close to real human behaviors. But there is one major issue with only using speed matching. Followers will match their leaders' speed no matter how far they are from their leaders. Which means that any simulation done needs to be initialized with all the agents already placed at the right distance between each other. A waiting line simulation as done in Section 3.2 would be impossible with Lemercier's model, the waiting line would need to be filled before running the simulation. Moreover, it needs some speed initialization which renders it unable to simulate the beginning of a motion. When simulating the real data presented in Section 4.1, the model can only start simulating the motion after the wave already appeared, as shown in Figure 3.6.

Even with the right distance initialization, ignoring the distance to the leader causes some issues in the following behaviors. Because of the delay in matching speed and the waves damping effect, the follower does not keep a constant distance to its leader. On the contrary, this distance tends to increase when the leader speed fluctuates. This issue is illustrated in Figure 3.10. The leader switched between $0 m.s^{-1}$ and $1.33 m.s^{-1}$. Looking at the followers' trajectories, we can see that the distance between them and their leaders increases a lot after a while. While other speed matching algorithm can solve this issue by matching the leader's speed at specific distances only ([RRW14]), preventing the distance to become too big or too close; by computing an ideal distance to the leader, our model does not have this kind of issues.

5 Conclusion

We presented a new method to simulate following behaviors which improves the follower's movement anticipation by defining a dynamic ideal safe distance according to the leader's movement. We consider that following distance results from: *physical*, *social* and *psychological factors*. Those factors are incorporated into three terms: first, a constant distance representing a comfortable distance to keep (social and physical constraints); second, a safety distance which allows to avoid a collision if the leader suddenly stops; and, finally, a distance based on the analysis of the leader's motion. This last term allows the follower to adapt its distance to the leader according to the prediction of the its motion.

We demonstrated that the followers in our model are able to adapt their trajectories to the noise present in the leaders' motion. We showed that the model is able to create a waiting line with speed propagation without any initialization. Moreover, we proposed an adaptation to 2-D environments through the use of an overtaking rule. The model was validated against real data and it was shown that the simulated results are close to the real ones when using the same initial conditions. We also proposed a method to detect stop-and-go waves. We used this method to check that the emergence of waves in our model verified the same density conditions as found in the real data. Our model presents some improvements compared to Lemercier's model [LJK⁺12] such as: no initial conditions needed and no increment of distance over time caused by fluctuations in the leader's speed. Finally, the companion video³ shows the visually realistic motion produced

by our model. Thanks to the distance control, the model is able to prevent any weird following distances that would produce silly following behaviors.

The model focuses on the following distance. The follower's speed is chosen to allow it to reach the following distance as fast as possible without passing in front of it. This already produces very good results. It is probable nonetheless that a better control over the speed would lead to even better results. As future work, it would be interesting to integrate a speed matching model to our model. To this end, we would need to carefully define the influence of the distance over the speed as we want to be able to change the following distance while keeping the nice speed properties of the speed matching model. If the influence is too large, it would completely change the speed chosen by the speed matching model, rendering it useless. But if it is too small, the following distance chosen by our model would never be reached.

³<http://youtu.be/JbQe19DyUMc>

Planning a sequence of interaction over time

4




Figure 4.1 In some situations, humans can see a tunnel through a flow. The left figure shows a first person view point where the tunnel is clearly visible. The right figure shows a top view of the situation as well as the set of colliding and collision-free velocities for the red agent. Regular local avoidance models try to find a collision-free speed toward the goal. This means the red agent will choose a velocity from the set of collision-free velocities (green part of the circle), making him slow down and miss the tunnel.

Contents

1	Introduction	46
2	Effective Avoidance Combination Strategy	47
2.1	Overview	47
2.2	Definitions	48
2.3	Cardinal Interaction Points (CIW)	48
2.4	collision-free Interaction Path	50
2.5	Cost and ranking	51
2.6	Perception	53
2.7	Following case	54
2.8	Extended search	55
3	Results	56
3.1	Performance	56
3.2	Case study	57
3.3	Energy	60
3.4	Limitations	62
4	Conclusion	63

1 Introduction

As described in the background chapter 2, solving collision avoidance by considering the surrounding obstacles independently cause some unnatural behaviors when there is many dynamics obstacles. Indeed, real humans are very efficient to snake through individuals or groups when crossing a flow or other dense crowds. An example is shown in Figure 4.1, where a red agent needs to cross a dense flow. There is an empty space in this flow that we call a tunnel. If we look at the first person view point on the left of the figure, the tunnel is clearly visible. A real human would probably detect it and use it to cross the flow with a minimum effort. In a more general way, humans are able to use gap in density to ease their way through dense crowds which is not always possible to do when considering obstacles independently from each others. While some existing solutions have been presented to reduce this problem in the section 1 of chapter 2, there is a lack of general solution that is able to build a global strategies for most situations. In this chapter, we propose a solution that combine collision avoidances through time instead of solving them all at once.

Today, most crowd simulators handle navigation in two steps. First, the long term strategy: the agent searches for a path between its position and its goal in the environment. This step is done by path finding modules. The second step is a short term strategy: while following the path the agent adapts its velocity to avoid any collision with nearby obstacles. This step is done by local avoidance modules. But none of these steps is able to use gap in density or even tunnel such as the one in Figure 4.1. Path finding techniques do not have enough details about the environment and dynamic obstacles to consider tunnel or gap in crowds. They focus on finding paths to the goal through the current static environment, which can be huge, and cannot afford such level of details. Even the one considering dynamics obstacles to avoid dense area, do not have a level of details sufficient. While the path finding step lack details information, local avoidance systems lack the long term considerations. So in this contribution, we explore a third step for navigation in crowd simulators, which looks further into the future than the local avoidance step does, but not as far as the path finder does. Such a step is thus between path finding and local avoidance. It provides a mid-term strategy, allowing agents to perceive tunnel or gap in a crowd and identify those that can be used as pathway through the crowd. In a way, the mid-term strategy defines the "affordances" of different paths through the crowd as define in the theory of affordances by Gibson ([Gib86]).

There exist an infinite number of solutions to avoid a collision with an obstacle. An avoidance strategy can be defined as the nature of adaptations made to the trajectory to perform avoidance. A collision can for example be avoided by adapting the speed or orientation, or a combination of both. The strategy choice has a huge impact on the resulting trajectories since, besides preventing the current collision, it also influences the relative motion with all the other obstacles, and thus changes the way the next potential collisions will be avoided. In this paper, we propose a solution to handle mid-term strategies by building several strategies to navigate through a crowd. These strategies consist of a list of adaptation sequences over time. The less costly strategy is then chosen. The cost of a strategy is a composition of the energy needed to follow the strategy, as well as an extra cost for every collision caused by the strategy. This allows agents to plan their way through a crowd, using empty spaces to navigate while optimizing their effort. We show

the improvement compared to local avoidance systems in terms of energy consumption.

Our contribution is a mid-term motion planning technique for complex scenarios with multiple moving obstacles. Our system, called Effective Avoidance Combination Strategy (EACS), is able to compute an energy-efficient avoidance path made of successive adaptations. To this end, it explores several possible ways of combining interactions, evaluates the energy cost of each possible solution, and selects the most efficient one.

The remainder of this paper is organized as follows. In Section 2.1, we present our system and explain our approach. Section 3 shows the capabilities of our system on some specific situations, and an evaluation in terms of energy gain with respect to previous models. Finally, in Section 4, we draw some conclusions.

2 Effective Avoidance Combination Strategy

2.1 Overview

The objective of the EACS system (EACS stands for Effective Avoidance Combination Strategy) is to enable agents to perform efficient navigation around other moving agents.




Figure 4.2 Navigation system architecture with three different levels of decision.

The typical structure of a crowd simulation framework consists of two levels. As presented in Figure 4.2, there is the path planning for global strategy and the local avoidance for local strategy. The path planning finds a path in the environment which consists of a sequence of positions without any fixed obstacles between two consecutive positions. The local avoidance uses the first position of the path, which is reachable in a straight line from the agent’s position, and computes a speed to reach it without any collision with nearby fixed or moving obstacles. The EACS system stands between these two modules, it takes the reachable first position of the path from the path planning as a mid-term goal and outputs a preferred velocity to the local avoidance system. The

local avoidance system will then follow this preferred velocity while avoiding very close obstacles. Efficient navigation is obtained by setting the relevant avoidance strategy for the few next future interactions with anticipation.

To this end, EACS detects all next interactions on the path to the mid-term goal and computes a sequence of velocity adaptations to successively perform collision avoidance. The variety of possible adaptations to achieve collision-free motion results in different mid-term navigation paths. EACS evaluates and compares their efficiency from the energy point of view and selects the most efficient one to steer agents. In the remainder of this section we start by defining the main concepts used in our system, followed by a description of our technique.

2.2 Definitions

Collision Course: two agents are on a collision course when their position and velocity vectors are set to cause their future distance of closest approach to be below contact distance if the velocity vector is kept constant.

Avoidance Strategy is the nature of the adaptations made to velocity vectors to avoid future collisions. For example, an agent can avoid a collision with another agent by adapting speed (i.e., the norm of the velocity vector) or by turning (i.e., by rotating the velocity vector).

Interaction Waypoints (IW): for two agents on a collision course and performing collision avoidance, the IW are the space-time waypoints by which agents go when avoiding.

Cardinal Interaction Waypoints (CIW) are the 4 IW which result from 4 specific avoidance strategies with the minimum required amount of adaptation: accelerating only, decelerating only, turning left only and turning right only. We note these cardinal interaction points by **a**, **d**, **l** and **r**, respectively.

Interaction Paths (IP) are a sequence of IW ordered by increasing time. An agent that follows an IP will go through all its IW. Between two consecutive IW the agent follows the straight line between the two positions at a constant speed.

Interaction Segments (IS) are a part of an IP composed of two consecutive IW.

Goals (G) are special IW that do not have a time constraint. A goal is a position in space that agents try to reach, regardless of the when reach it. Goals appear only at the end of an IP or IS. When an agent is following an IS between a regular IW and a goal, it simply walks at its preferred speed toward the goal.

2.3 Cardinal Interaction Points (CIW)

In this section, we describe how we compute the 4 CIW. We arbitrarily chose these four points covering the four types of adaptation (accelerating, decelerating, turning right and turning left) as well as the two orders of passage (passing before the other agent or after). Using these 4 CIW yields satisfying results, but other IW could be consider when computing avoidance strategies.

The 4 CIW computation is described for a given agent α_1 on a collision course with agent α_2 . An example is presented in Figure 4.4. Their respective 2-dimensional positions and velocities at current time t_0 are noted $(\mathbf{p}_1, \mathbf{v}_1)$ and $(\mathbf{p}_2, \mathbf{v}_2)$. The CIW (3-dimensional:




Figure 4.3 Collision velocity space for one interaction, in red. $\mathbf{v}_{a_{1|2}}$, $\mathbf{v}_{r_{1|2}}$, $\mathbf{v}_{d_{1|2}}$ and $\mathbf{v}_{l_{1|2}}$ represent the collision-free velocities that are used to compute the CIW.

space and time) are noted: $\mathbf{a}_{1|2}$, $\mathbf{d}_{1|2}$, $\mathbf{l}_{1|2}$ and $\mathbf{r}_{1|2}$. To compute these points, we base our principle on the velocity-obstacle introduced in [FS98] as illustrated in Figure 4.3. The key idea is to compute the 4 admissible velocities which enable collision-free trajectories and which correspond to 4 specific strategies (deceleration, acceleration, left turn, right turn). We recall that the set of admissible velocities for α_1 concerning its interaction with α_2 is noted $AV_{1|2}$ and is defined as follows:

$$AV_{1|2} = \{\mathbf{v}_1 \in V_1 \mid \forall t \in [t_0, t_0 + \tau], \text{dist}_{1,2}(\mathbf{v}_1, t) \geq c\} \quad (4.1)$$

where: V_1 is the set of all the reachable velocities for α_1 , τ is the size of a time window, $\text{dist}_{1,2}(\mathbf{v}, t) = \|(\mathbf{p}_2 + \mathbf{v}_2 t) - (\mathbf{p}_1 + \mathbf{v} t)\|$ and c is the collision distance threshold.

Then, the CIW define trajectories followed by α_1 with specific strategies (velocities): $\mathbf{v}_{a_{1|2}}, \mathbf{v}_{d_{1|2}}, \mathbf{v}_{l_{1|2}}, \mathbf{v}_{r_{1|2}} \in VA_{1|2}$:

- $\mathbf{v}_1 \cdot \mathbf{v}_{a_{1|2}} = 0$, $\|\mathbf{v}_1\| < \|\mathbf{v}_{a_{1|2}}\|$, and $\exists t \in [t_0, t_0 + \tau] \mid \text{dist}_{1,2}(\mathbf{v}_{a_{1|2}}, t) = c$,
- $\mathbf{v}_1 \cdot \mathbf{v}_{d_{1|2}} = 0$, $\|\mathbf{v}_1\| > \|\mathbf{v}_{d_{1|2}}\|$, and $\exists t \in [t_0, t_0 + \tau] \mid \text{dist}_{1,2}(\mathbf{v}_{d_{1|2}}, t) = c$,
- $\det(\mathbf{v}_1, \mathbf{v}_{l_{1|2}}) > 0$, $\|\mathbf{v}_1\| = \|\mathbf{v}_{l_{1|2}}\|$, and $\exists t \in [t_0, t_0 + \tau] \mid \text{dist}_{1,2}(\mathbf{v}_{l_{1|2}}, t) = c$,
- $\det(\mathbf{v}_1, \mathbf{v}_{r_{1|2}}) < 0$, $\|\mathbf{v}_1\| = \|\mathbf{v}_{r_{1|2}}\|$, and $\exists t \in [t_0, t_0 + \tau] \mid \text{dist}_{1,2}(\mathbf{v}_{r_{1|2}}, t) = c$.

The CIW $\mathbf{a}_{1|2}$, $\mathbf{d}_{1|2}$, $\mathbf{l}_{1|2}$ and $\mathbf{r}_{1|2}$ are then:

- $\mathbf{a}_{1|2} = (\mathbf{p}_1 + \mathbf{v}_{a_{1|2}} \cdot \text{time}(\mathbf{v}_{a_{1|2}}), \text{time}(\mathbf{v}_{a_{1|2}}))$,
- $\mathbf{d}_{1|2} = (\mathbf{p}_1 + \mathbf{v}_{d_{1|2}} \cdot \text{time}(\mathbf{v}_{d_{1|2}}), \text{time}(\mathbf{v}_{d_{1|2}}))$,
- $\mathbf{l}_{1|2} = (\mathbf{p}_1 + \mathbf{v}_{l_{1|2}} \cdot \text{time}(\mathbf{v}_{l_{1|2}}), \text{time}(\mathbf{v}_{l_{1|2}}))$,




Figure 4.4 Example of collision solved using the four CIW. The collision trajectories are drawn in red and the four trajectories followed to reach the CIW are shown in blue. The position of each CIW is shown together with the position of the other agent when the CIW is reached.

$$\blacksquare \mathbf{r}_{1|2} = (\mathbf{p}_1 + \mathbf{v}_{\mathbf{r}_{1|2}} \cdot \text{time}(\mathbf{v}_{\mathbf{r}_{1|2}}), \text{time}(\mathbf{v}_{\mathbf{r}_{1|2}})),$$

where $\text{time}(\mathbf{v})$ is the time of closest approach ($ttca$) delayed by a small amount: $\text{time}(\mathbf{v}) = ttca(\mathbf{v}) + \min\left(\frac{c}{v}, \frac{c}{v_2}\right)$; $ttca(\mathbf{v})$ is computed assuming agents α_1, α_2 travel with velocities \mathbf{v}, \mathbf{v}_2 .

This delay allows the two agents to get away from each other to avoid resetting them on a collision course if one of the two agents performs a new manoeuvre too soon (e.g., to solve one next interaction). The example presented in Figure 4.4 clearly shows that once the CIW is reached, the collision risk with the other agent is avoided.

2.4 collision-free Interaction Path

The collision-free Interaction Path construction is presented in Algorithm 2 and illustrated with an example in Figure 4.5. The collision-free IP is built iteratively by the algorithm. It starts with the straightforward IP to the goal (line 3). A collision test is performed with this IP (line 7). If no collision is found, a collision-free IP is detected (line 18). If a collision is found, the IP is discarded and the four CIW are computed and used to build four new IP (line 8). The test is then repeated on these new IP, first on the IS right before the interaction (line 9 to 12) then the IS from the interaction to the goal (line 16) until one IP reaches the goal (line 18).

An example is presented in Figure 4.5. The initial situation is shown in Figure 4.5.A: the blue agent starts at the position \mathbf{s} and has to reach the position \mathbf{g} . As we said the




Figure 4.5 Example showing the construction of a collision-free IP illustrating the different steps taken by the EACS system.

algorithm starts by testing the straight forward IP (Figure 4.5.B), a collision is detected and four new IP are created $sl_{1|3}g$, $sd_{1|3}g$, $sa_{1|3}g$ and $sr_{1|3}g$. Now, these four new IP have to be tested starting with the IS just before the interaction: $sl_{1|3}$, $sd_{1|3}$, $sa_{1|3}$ and $sr_{1|3}$. Figure 4.5.C shows the test for the turning left solution $sl_{1|3}$. A collision is detected on the IS and again four new paths are created. Figure 4.5.D shows the test for the accelerating solution $sa_{1|3}$ and no collision is found. But this is only the first part of the IP, there is still the IS $a_{1|3}g$ to test. This test is done in Figure 4.5.E. A collision is detected and only two paths are created. Both agents are going toward each other and cannot avoid the collision by accelerating. Decelerating will not solve the problem either, but only delay it as the blue agent will have to go toward the goal eventually. The turning left solution is tested on both IS, first on $a_{1|3}l_{1|4}$ (Figure 4.5.F) then on $l_{1|4}g$ (Figure 4.5.G), and no collision is found. The goal is reached and a collision-free IP is built: $sa_{1|3}l_{1|4}g$ (Figure 4.5.H).

2.5 Cost and ranking

On the example presented in Figure 4.5, a collision-free IP $sa_{1|3}l_{1|4}g$ is built. But if the turning right solution for the last interaction had been tested before the turning left one, we would have ended with the collision-free IP $sa_{1|3}r_{1|4}g$. By testing the IS in a different order, we could have also built the collision-free IP $sl_{1|2}g$. In the end, there are many

Algorithm 2: Algorithm used by EACS to build collision-free Interaction Paths

Data: Starting position \mathbf{s} , Goal \mathbf{g}
Result: Build a collision-free IP

```

1 checkList←[];
2 collisionFreeList←[];
3 add( $\mathbf{sg}$ , checkList);
4 pathFound←FALSE;
5 while !pathFound do
6    $\mathbf{bc}$ ←pullSegment(checkList);
7   if collision( $\mathbf{bc}$ ) then
8      $\mathbf{a,d,l,r}$  ← computeCIW( $\mathbf{bc}$ );
9     add( $\mathbf{ba}$ ,checkList);
10    add( $\mathbf{bd}$ ,checkList);
11    add( $\mathbf{bl}$ ,checkList);
12    add( $\mathbf{br}$ ,checkList);
13  else
14    add( $\mathbf{bc}$ , collisionFreeList);
15    if  $\mathbf{c} \neq \mathbf{g}$  then
16      add( $\mathbf{cg}$ , checkList);
17    else
18      pathFound←TRUE
19    end
20  end
21 end
22  $P$ ←buildPath( $\mathbf{s},\mathbf{g},\text{collisionFreeList}$ );
23 return  $P$ ;

```

different IP to go through a crowd. But not all of them would be realistic, so we are looking for one that is efficient and would most likely be picked by a real walker. To that end, we used a cost function to rank the different IP and select a single one among them. We chose to use the energy consumption as the cost function. Real humans often favor the least energy consuming way of performing a task [Zip49]. Moreover it has already been used to improve collision avoidance models [GLM10]. The energy formula, $E = m \int (e_s + e_w |\mathbf{v}|^2) dt$, is from [Whi03]. The values of $e_s = 2.23 J.Kg^{-1}.s^{-1}$ and $e_w = 1.26 J.s.Kg^{-1}.m^{-2}$ minimize the energy consumption per unit of distance for a speed of $|\mathbf{v}| = 1.33 m.s^{-1}$ which will become the preferred speed of the agent. We want to have agents with different preferred speeds. From the energy formula per distance, Equation (4.2), we set a constant distance to look for the speed that minimizes the energy consumption: \mathbf{v}_{pref} (see Equation (4.3)). Then we are able to change the value of e_w for the chosen \mathbf{v}_{pref} (see Equation (4.4)).

$$E = m \left(\frac{e_s}{|\mathbf{v}|} + e_w |\mathbf{v}| \right) d \quad (4.2)$$

$$\frac{dE}{dv} = m \left(\frac{e_s}{|\mathbf{v}|^2} + e_w \right) d = 0 \quad (4.3)$$

$$e_w = \frac{e_s}{|\mathbf{v}_{pref}|^2} \quad (4.4)$$

In some situations, allowing collisions is necessary to find good solutions (see Figure 4.12). So we enabled the search algorithm to consider IP with collisions. To keep a preference for collision-free IP, an extra cost is added for IP with collisions. This extra cost is the equivalent of making a long detour: the energy consumption at preferred speed for a specific distance which is a parameter of the system ($CollisionCost = m(\frac{e_s}{|\mathbf{v}_{pref}|} + e_w|\mathbf{v}_{pref}|)d_{collisionDetour}$). This way, agents will still go around dense areas when these are not too big, and will go through them otherwise.

2.6 Perception

When building an IP, the EACS considers every other agent for each step of the IP. This may lead to a combinatorial explosion. To reduce this risk and improve performance, we limit the number of agents considered by the EACS. This limitation is shared by real humans. For example humans have no eyes on their back and thus cannot perceive most of the people behind them. While their perception is limited, they are still capable of perceiving many people. We can assume that people consider only a fraction of them, that we call neighbors, when performing collision avoidance. In this section we describe a selection process to imitate human limitation and neighbor selection.




Figure 4.6 Agent R perception system: First it selects the agents inside the perception area (O, B and G) then for each selected agent it computes the interception speed and selects the n agents with the smallest Minimum Interception Time (MIT) as neighbors (First B then G, O is not selected as it is going away too fast to be intercepted).

Our selection process is inspired by previous ones (such as the one use in [vdBLM08]) and composed of two steps. The first step represents the physical limitation of human perception and consists of selecting all the agents in an area around the one performing collision avoidance. This area is represented in Figure 4.6 and is composed of two half circles: a big one in front of the agent and a small one behind. In the example from the Figure, Y and P are not selected as neighbors as they are outside of the perception area. The second step selects a limited number of perceived agents using a specific criterion. In our case we define the Minimum Interception Time (MIT) in Equation 4.6 and select

agents with the lowest MIT value. Indeed, the harder it is to intercept another agent, the less likely it is to cause a collision with any IP. This allows us to select fewer agents than with the regular euclidean distance criterion while conserving good behaviors.

$$IT_{1|2} = \{t \in \mathbb{R}^+ \mid \exists \mathbf{v} \in \mathbb{R}^2 : \|\mathbf{v}\| \leq MS_1, \mathbf{p}_1 + \mathbf{v}t = \mathbf{p}_2 + \mathbf{v}_2t\} \quad (4.5)$$

$$MIT_{1|2} = \min(IT_{1|2}) \quad (4.6)$$

In Figure 4.6, B would be selected before G despite G being closest. As for O , it would not be selected as he is going away too fast to be intercepted.

2.7 Following case

When there is someone in front going in the same direction but slower, there is not always enough space to avoid them. In this case, people can try to squeeze to avoid them despite the lack of space or they can match the speed of the person in front to follow them until there is enough space to overtake them or until their trajectories separate. The following behavior is close to collision avoidance as it can be used to prevent a collision. At the same time, it is very different as the collision is never solved by following, only postponed. As said before, people follow until they can properly avoid the person in front of them or the collision risk disappears by itself. When computing the four CIW, the agent follows a velocity that prevents the collision until the collision risk no longer exists. As explained above, when following, the collision risk remains except if one of the agent changes direction. So the agent can still change its speed to match the one of the agent in front and prevent a collision by following, but we need to determine how long it should follow to compute the corresponding IW.




Figure 4.7 Detection of a following case and the two Interaction Waypoints \mathbf{i} and \mathbf{j} computed from the following speed.

Following happens when the tested IS has a collision with an agent that goes in the same direction, this situation is represented in Figure 4.7. Whether the IW \mathbf{h} (the end of the IS) is the goal or just an intermediate IW placed here to avoid another agent, the position and/or the time of the IW is important. This is why it has been decided that when following, two IW should be computed. For the first one \mathbf{j} , the agent will follow until it reaches the previous position he wanted to reach (the position of \mathbf{h}). For the second one \mathbf{i} , the agent will follow for as long as it would have walked with the previous IS \mathbf{kh} .

From the Figure 4.7, we have the velocity of the green agent \mathbf{v} , the IW $\mathbf{k} = (\mathbf{p}_k, t_k)$ and $\mathbf{h} = (\mathbf{p}_h, t_h)$. We then compute the IW $\mathbf{j} = (\mathbf{p}_h, \frac{\mathbf{p}_h - \mathbf{p}_k}{\mathbf{v}} + t_k)$ and $\mathbf{i} = (\mathbf{v} * (t_h - t_k), t_h)$.




Figure 4.8 1) Without extra search, the algorithm stops at the IS OA and never considers OE . 2) With extra search, the IS OE is created by solving the collision on the IS OA' . 3) Agent's trajectories without extra search, the agent starts to go inside the crowd before backing up and going around. 4) With extra search, the agent is able to directly go around the crowd.

2.8 Extended search

The construction of the IP is done step by step. When a collision is found, IS that solve the collision are searched in several directions (by turning or by adapting speed). Once such a segment is found, no other solution is searched in the current direction. If we consider the situation in Figure 4.8, the stop of the search can lead to unnatural behaviors. In this situation, we considered the IP built by the red agent avoiding all the blue agents. The blue agents are all standing in order to simplify the representation of the IS as we do not have to consider their movement through time and collision avoidance solution can only be found by turning. As said, once an IS is found solving a collision avoidance, the search is stopped in the current searched direction. In the current situation, it means that only the IS OA and OB in Figure 4.8.1 will be considered as first IS of the IP. In the end, the agent will be able to only consider a path that goes inside the crowd and exits it again which produces the results presented in Figure 4.8.3.

As we can see on the first person view, an obvious solutions would be to go around the entire crowd directly which means that the agent needs to also consider OE and OF as first IS. In order to extend the search to find these two IS, we added a few extra steps to the algorithm. When an IS free of collision is found, another IS is considered with the same speed for a longer period of time. If no collision is found, then the IS represents a way out of the crowd that does not require to go back on its steps, so the search returns to its normal process. If a collision is found, then we insert the new IS with the collision to the search process to continue searching for more solution in the same direction. In the example in Figure 4.8.2, the IS OA is extended to the segment OA' , then the collision on

OA' is solved by turning left (as turning right would bring us back in the search pattern) and the IS **OE** is created.

As a result, the agent is now able to avoid going inside a crowd if it is only to backtrack on its own path, this fixed behavior is presented in Figure 4.8.4. This extra search has been implemented only for the first IS of the IP. This has the advantage of limiting the extra computation time required by the extra searched paths. This extra search can be extended to the others IS of the IP, but the first interaction is the most important one as it gives the current direction to the local avoidance system and is enough to prevent going inside a dead end situation if there is a better way around.

3 Results

In this section we present the results of our work. The simulations have been done by using our system coupled with RVO2. EACS computed a strategy and then set the preferred velocity required by RVO2. No additional pathfinding algorithms have been used as the studied situations were simple enough and the goal is always reachable from the agents' starting position.

3.1 Performance

The EACS system computes an IP and has to explore many possibilities to find a good one. This exploration leads to a combinatorial explosion as for each constructed IP, a collision detection is done with the others agents and when a collision is found four new IP are created which in turn will need collision detection and might create 4 more IP. Several methods have been used to improve the system's performance:

- **Smart exploration:** an estimation of the final cost of an IP is done to advance the construction of the most promising one in priority.
- **Limited exploration:** when a goal is far away, it is very unlikely that real humans will fully plan collision avoidance to the goal. Moreover, planning too far in the future becomes quickly ineffective as the simple prediction of constant velocity becomes erroneous. The EACS exploration is stopped when the IP reaches a number of consecutive interaction that is too big.
- **Late update:** much like path planning, the interaction planning does not need to be done every step. While the environment in the interaction planning is not static and thus IP need to be regularly updated, IP can be reused for several steps. This allows us to spread the computational cost of all the agents' planning on several step, similarly to what has been done in [SKH⁺11].
- **Neighbor selection:** the perception system presented in section 2.6 reduces the number of agents considered when computing collision detection which greatly improves computation time.

The two methods with the best impact on computation time are the Late update and the Neighbor selection. On a 20s simulation of 400 agents in a four flow crossing, the

simulation took 60.3s to compute without any neighbor selection. With a simple neighbor selection based on the area only, the computation time dropped to 35.4s. When limiting the number of neighbors to 100, the computational time became 3.74s making it possible to simulate larger numbers of agents. On a 20s simulation of 1000 agents in a four flow crossing, the simulation took 35.7s to compute when updating the plan at every step for each agent. When updating the interaction plan every 10 steps, the computational time dropped to 5.19s.

Performance tests have been run with 4 different situations:

- **One-Way Flow (OWF):** one flow with a starting density of $0.9agents.m^{-2}$ moving in one direction with agents having different preferred speeds. The size of the flow is variable (bigger flow, same initial density)
- **Bidirectional Flows (BF):** two symmetrical flows with a starting density of $0.9agents.m^{-2}$ moving in opposite direction with agents having different preferred speeds. The size of the flows is variable (bigger flow, same initial density)
- **Four Crossing Flows (FCF):** four flows with a starting density of $0.9agents.m^{-2}$ crossing each other at the same position with agents having different preferred speeds. The size of the flows is variable (bigger flow, same initial density)
- **Sandbox (SB):** 1000 agents randomly moving in a square, when an agent is close to its current goal a new one is randomly chosen on the square border. The size of the square is variable to simulate different densities.

All the tests have been run on a computer with an Intel®Xeon®Processor E5-1603 (4 cores, 10M Cache, 2.80 GHz, 0.0 GT/s Intel®QPI) and 32GB Memory with the following set of parameters: `updateFrequency=10`, `MaxInteractions=9`. Results are presented in figure 4.9.

As can be seen on the graphs in Figure 4.9, adding the EACS to RVO2 significantly increase the computation time. Nevertheless, thanks to the different mechanisms that reduce the complexity, EACS is able to simulate several thousands of agents in real time (computation time < 20s for 5000 agents). It is true even for complex situations (FCF) as complexity does not influence computational time with large numbers of agents (≈ 5000).

While the size of the simulation is the main factor influencing computation time, density also has a significant impact. This impact is visible on the right graph in Figure 4.9. It seems that computation time increases linearly with density. Indeed, it is harder to plan ahead in high density. Moreover, collisions appear more often causing IP to be smaller. Plans are update every 10 steps except if the agent reaches the end of the current IP. This happens more often with smaller IP.

3.2 Case study

Several situations where designed to test the EACS system and check the improvement it brings to crowd simulation. The first one is the example used in the introduction to show the limitation of the current models. The results are shown in Figure 4.10. As expected, RVO2 alone is unable to use the tunnel to go through the flow. Instead, the red agent goes straight forward and has trouble going through the flow as it spends around 10 seconds




Figure 4.9 The graphs show the time needed to compute 20 seconds of simulation for the different performance test situations with and without EACS. On the left graph, the results of the three situations: OWF, BF and FCF are presented for different numbers of agents. The right graph presents the results of the SB situation for different densities.

in the flow itself. With the EACS system, the red agent is able to take the tunnel which facilitates its navigation greatly: the agent spends only 2 seconds in the flow.




Figure 4.10 Simulation of an agent (in red) going through a flow, with the flow having a tunnel to facilitate the crossing. Left: simulation including EACS; Right: simulation without EACS .

The second one is shown in Figure 4.11 and is very similar to the first one. Instead of going through a flow, agents need to exit or enter a building that has 3 entrances. All the agents are close to one exit which quickly becomes clogged. As for the first situation where the agent was able to detect the tunnel, in this situation, some of the agents detect that the other two exits are empty. Instead of waiting and avoiding many obstacles which can be painful, they decide to go through the two others exits. The extra cost from the detour is largely compensated by all the required velocity adaptations to go through the

clogged exit. This is easily witnessed in real life: in the train station for example, where people will use further exits to avoid the dense crowd of the closest one.




Figure 4.11 Simulation of a clogged exit, EACS is able to choose other exits that are further away but not clogged.

In the third situation shown in Figure 4.12 there is no collision-free solution to cross the incoming flow. Such situations exist in real life where we have to cross a very dense crowd and we require other people to give us some space to go through. In the presented situation, even if there are no collision-free paths through the crowd, there is a weak point where the density is lower. Going through this point would make it easier to go through the flow. EACS is able to detect such weak points in the flow and compute a path through it. If we compare with the trajectory computed without EACS (at the bottom), we can clearly see that it is easier to use the weak point to go through the flow. In the end, the agent controlled by RVO2 spends around 10 seconds in the flow while the one controlled by EACS spends only 2 seconds in it.




Figure 4.12 Simulation of an agent (in red) going through a flow that goes in the opposite direction. The flow has no tunnel but a section of it has a lower density. At the top, we have the simulation with EACS and at the bottom the simulation without EACS.

Another situation is shown in Figure 4.13 and has been designed to show the difference in optimizing the short-term outcome and the mid-term one. If we look at the trajectory

given by RVO2 alone (top right), we can see that the red agent is avoiding by going on the left. If we consider only the first few blue agents, it is the best solution as going on the right would require to go faster and to shift more on the side. This is confirmed by the energy graph, the energy consumption of EACS which goes on the right is higher than the one of RVO2 for the first 3m. But when we look further, we can see that the agent will meet with more people when going to the left. This dense area is so hard to navigate that the agent will even move backwards. Looking at the graph, we can notice the sudden rise in energy consumption for RVO2 around 4m as well as the curve going backwards because the agent walks in the opposite direction of its goal. The curve of EACS energy consumption remains stable which confirms that the EACS strategy is better on the mid-term.




Figure 4.13 Left: simulated test where short-term cost leads to bad mid-term cost, agents simulated with EACS on the left and without on the right. Right: the graph quantifies the energy consumption of the red agent (For RVO2 and EACS) as they walk towards the goal (In this situation, it is equivalent to the distance traveled on the x axis). The section of the blue curve circle in red corresponds to RVO2 agents going backwards to avoid the dense group (T=7.5s on the top figure).

The last situation is about group avoidance. A study, done in parallel of this work and presented in chapter 5, shows that real humans go around dense groups and through sparse ones. We have made several simulations where an agent had to avoid groups of different densities to check that EACS is able to reproduce the real trend. We can see some of them in Figure 4.14, the situation setup is the same as the one in the experiment with real humans. The red agent is able to decide whether to go through the group or around it and reproduces the same trend as that of a real human: it circumvents the two dense groups (Figures 4.14A and 4.14B) and traverses the sparse ones (Figures 4.14C and 4.14D).

3.3 Energy

To check the energy gain of EACS over pure RVO2, we studied the trajectories of agents going through an orthogonal flow. We tried several different conditions to analyze their impact on the results. We performed 2800 simulations: 7 density types x 19 exploration limits for EACS and 1 without EACS x 20 repetitions. For each simulation, a flow was




Figure 4.14 Simulation of group avoidance for different group densities.

randomly generated and 10 agents were randomly placed on the same side of the flow with their goal on the other side. The energy consumption of the 10 agents was computed from their starting position to their goal location and then averaged.

The first condition studied is the density type. For each generated flow, the same dimensions (10mx80m) and number of agents (700) were used. Thus the global density was the same for each flow. The variation among the different conditions was made on the regularity of the density in the flow. These variations are shown in Figure 4.15. We have a uniform flow on the right and the different types of irregularity until the very irregular flow on the left. These irregularities create small dense clusters and leave some space in between to navigate more easily. We expect EACS to be able to use this in-between space and perform better when going through irregular flows.

The second condition studied is the exploration limit for EACS. When EACS is looking for a path it will sequence several interactions over time. The number of interactions can become very big in some situations which can cause performance issues. Moreover, it is doubtful that real humans can consider a high number of future interactions. For this reason, we set a maximum number of interactions when constructing an IP in EACS. When an IP reaches the limit it will not be extended anymore. We expect that a low maximum number of interactions will lead to results close to RVO2. We also expect that a very high number of interactions will perform similarly or worse than a medium number of interactions: the further we plan in the future, the less the extrapolation is relevant.

For the results, for each exploration limit, we compare the average energy consumption with and without EACS. The Figure 4.16 shows this consumption difference in percent. As expected, EACS is performing better for very irregular densities. When the density becomes uniform, the energy consumption with and without EACS becomes similar. If we look at the different results for different maximum numbers of interactions, we can see that having a value higher than 14 does not seem to change the results much. Indeed, optimum results seem to be found for a maximum number of interactions around 4.




Figure 4.15 The different density types used for the energy study from very irregular (left) to uniform (right).

3.4 Limitations

Our system has limitations. The main one is the number of future interactions explored by the system. We showed in the previous section that the system is particularly efficient when used in specific contexts, such as when agents' density is not uniformly distributed in space. We can actually easily guess what type of crowd motions will enable EACS to be efficient. For example, we expected that non-uniform density distributions would open larger efficient interaction paths in the crowd with a great probability that these paths remain valid in spite of all agent adaptations.

More generally, the EACS system is particularly efficient when the initial predictions are good and turn out to be representative of the actual motion of neighbor agents. In the opposite case, predictions turn out to be false, and the energy gain by EACS is poor. We also observe in the energy consumption evaluation of EACS solutions (previous section) that considering interactions too far in the future decreases the quality of the solutions provided by EACS. Note that EACS guided agents toward fewer counterproductive situations than naive short-term solutions.

One solution to this could be to evaluate on-line the quality of the predictions made by EACS. The number of interactions explored in the future could be directly dependent on the quality of this prediction. If predictions are good, agents can explore wider future time windows and more numerous future interactions, and conversely.

Another limitation of EACS is to only consider an energy-related cost function. We are convinced that this criterion is considered by human walkers when setting mid-term strategies, but we are also convinced that other factors play a role. For example, we could check how close the solution paths are to other agents' motion. Additional social distances could be considered in the cost of paths. As another example of a social criterion, passing in front of an agent at a close distance could be penalized in comparison with the strategy




Figure 4.16 This graph shows the amount of energy preserved by the agents when using EACS to optimize their avoidance strategies compared to EACS-less strategies.

of giving way, which could be considered to be more polite.

Finally, EACS explores a subset of all the possible strategies to perform collision avoidance. We make an arbitrary choice to explore adaptations exclusively made of speed or orientations changes, whereas all intermediary solutions could be explored. We are here in a typical trade-off between computation times and quality: we could add more than 4 CIW in the system to consider mixed-adaptation strategies, but this would result in higher computation times.

4 Conclusion

We have modified the usual navigation process of crowd simulation to add some mid-term considerations. To this end we have designed an interaction planner, the Effective Avoidance Combination Strategy that creates interaction paths which enable agents to consider non-straight paths to go through a crowd. As a result, we have shown with specific test cases that the agents are able to select strategies that use density tunnels to facilitate their navigation through a crowd. Moreover, while building the interaction path, EACS evaluates the repercussion of its avoidance choices further in the future. This allows agents to make some extra effort in resolving their current collision so as to simplify future interactions.

With an energy study, we show that by adding EACS to the regular navigation process, agents use less energy to navigate. The study has shown that EACS is especially useful to improve trajectories when facing crowds with irregular densities. The study has also highlighted that too much mid-term planning is not beneficial, which can be explained by the incapacity to correctly extrapolate the other agents' trajectories too far in the future. It seems that planning around 4 interactions ahead yields the best results.

We have shown the improvement in the quality of agents' navigation brought by our

work. The next step is to compare it with real human behaviors and interaction planning. Especially, it would be interesting to study how many interactions ahead real humans plan. Another interesting direction is to work on the evolution of the interaction path. Right now, the IP is computed at a constant interval and each time the computation starts from scratch. Real humans are known to be highly adaptable creatures, in most cases they most likely adapt their plan based on changes in the environment instead of planning a new one from scratch. This would greatly improve performance, but it might also improve results by making the plan more reactive to the environment changes.

VR study: Human behaviors toward groups

5

Contents

1	Introduction	66
2	Principle of Minimum Energy: a simulation study	67
	2.1 Method	68
	2.2 Results	69
	2.3 Discussion	70
3	Groups avoidance: a VR-based user study	71
	3.1 Objectives	71
	3.2 Materials and Methods	72
	3.3 Results	73
	3.4 Discussion	75
4	Applications	77
	4.1 Decision algorithm: Around or through the group?	77
	4.2 Integration in RVO2	78
	4.3 User-study evaluation	79
5	Conclusion	81

1 Introduction

In parallel to the work that has been done on crowd simulation algorithms, some experiments have been performed to better understand human behaviors in complex situations. In this chapter, we present such experiment done to study interactions between individual and groups as well as demonstrate the many advantages of conducting experiment in VR. This experiment follow the validation of the VR platform with moving virtual humans done by Anne-Hélène Olivier, which is presented in appendix A section 2 as part of a collaboration done at the beginning of the PhD.

While there are many studies on how people navigate together as a group, we are missing information on how individuals interact with these groups. In this context, the main objective of this study is to analyze interactions between an individual and groups in order to create new simulation algorithms that correctly handle such situations. More specifically, we focus on the situation of collision avoidance: how do individuals avoid collision with groups of people? Our objective is to understand under which conditions individuals may traverse groups or decide to circumvent them as a whole.

As presented in chapter 2 section 2, most studies performed around group behaviors have been done on real world observation. As it has been explained, performing experiment on groups required to manage several participants at the same time which greatly increase the difficulties to control the experiment. We want to study specific factors and determine the threshold between going through and going around a group, so we need to keep great control over the group. For example, we need to be able to constraint the group members to walk at a specific distance between each other. We first approach this question under the perspective of the Principle of Minimum Energy and used existing crowd simulator to compare the going through and going around strategies. The Principle of Minimum Energy (PME) already used to simulate other navigation mechanism [GCC⁺10], states that humans tend to optimize their trajectory to use as little energy as possible to reach their goals. It is obvious that the circumventing of groups of a large size will represent a considerable loss of energy along the travel path. By contrast, it is also obvious that going through small groups will require several collision avoidance maneuvers whereas it is a small detour to go around. While using the PME principle gives us some insight on which strategy is better, we still need to check what real humans do. But the required constraints over the group members would be near impossible to achieve in the real world with real people. Which is why we turn to VR as it gives us perfect control over the virtual environment including the virtual human forming the groups. We are then able to study the reaction of real humans interacting with specifically parameterized groups.

This study proposes three contributions:

1. We study interactions between individual walkers and groups from a theoretical point of view by applying the PME. We study groups of changing size and density, we compare go around and go through solution paths. Three kinds of situations are drawn. In the group size-density space, we delimit two regions for which go around paths are significantly different on the energetic level from go through paths. In between, go around and go through paths represent the same amount of energy

consumption. We call them inconclusive situations.

2. We use VR to study how real humans behave when avoiding virtual groups of changing size and density. We corroborate our theoretical study: humans behave as predicted by the PME. In addition, we take interest in results obtained for inconclusive situations. We demonstrate that decision changing, between going around and going through, is made according to some individual threshold, and that this threshold can be influenced by secondary factors such as the appearance of the group, or relative direction of motion.
3. We build on the results of both our theoretical and experimental studies. We propose an algorithm to imitate humans in their decision to go around or through a group of people, with few intuitive parameters.

Altogether, these contributions enable us to design the first model of local interactions in crowds (between individual and groups) fully validated through a VR-based study. The chapter is organized as follows. Section 2 presents our study on effort as a main factor to decide on a strategy to avoid a group of people. Section 3 presents a behavioral study conducted in Virtual Reality to confirm our results, and to complete our results with the considering of secondary factors. Finally, Section 4 draws guidelines to integrate our results into crowd simulators, and demonstrates some results based on RVO2, before ending with the conclusion in section 5.

2 Principle of Minimum Energy: a simulation study

An individual walker has two solutions to avoid a group he crosses during his travel: either he goes through it or around it. Going around a group results into a longer travel path whereas going through asks the walker to perform local adaptations maneuvers, i.e., successive changes of speed and orientation. Both strategies raise additional energetic cost along the travel path. How significant are those differences? Do these differences explain human behavior? This section answers the first question.

Group size and density change the length of the path going around the group, as well as the number of local interactions when going through the group. They are both fundamental factors of our study. To evaluate and compare the energy cost of avoidance strategies, we need to have examples of the two kinds of avoidance paths, going around or going through ones, performed under various conditions of group size and density. To face the obvious difficulty to obtain empirical data, we use simulation tools to synthesize data.

Our approach is the following. In simulations, we set situations of interactions between individuals and groups. For each situation, we first simulate the agent going through the group, and then we force him to go around. We record the two kinds of resulting trajectories, evaluate their energetic cost and compare them. Method and Results are detailed below.




Figure 5.1 Simulation setting with one agent avoiding a group on an orthogonal or collinear trajectory. On the left, the starting position with movement direction. On the right, example of computed trajectories relative to the group (For the top: group radius=3, interpersonal distance=0.8, for the bottom: group radius=8, interpersonal distance=2)

2.1 Method

Simulations are based on the RVO2 algorithm [vdBGLM11]. RVO2 considers interaction between individual agents only. Interactions between individuals and groups will naturally result into a going-through strategy. To compare with the going-around strategy, we extended RVO2 to force going-around behaviors as suggested in [YCP⁺08]: the group is simulated as a single 'big' agent with a large radius. The initial setup situation is illustrated in figure 5.1. All agents comfort speed was set to $1.33m.s^{-1}$. The shape of the group was defined as a circle. The positions of agents inside the group are computed randomly according to a Poisson distribution. We considered the following factors values to simulate interactions between the individual agent and the group:

- Group size: the group circle radius varied from $1m$ to $11m$ with a $0.5m$ step (21 different values).
- Group density: Density is expressed as the mean distance between agents. Interpersonal distance between agents varied from $0.75m$ to $3.75m$ with a $0.05m$ step (61 different values).

- Group relative movement: the individual agent and the group had orthogonal or collinear trajectories (initial distance to the crossing point was 40m in both cases).

We studied all the possible combinations of factors values, i.e., 2562 situations. When interpersonal distance is superior to the group size, groups consist of 1 member only. These irrelevant situations were removed from our study (212 removals). For each situation, there can be variations in the energetic cost of avoidance paths depending on the placement of the agents in the group, especially when going through given the exact interactions occurring there. We averaged this cost by calculating the 50 different distributions for agents' positions in group and by simulating one go-around path and one go-through path for each of the 50 distributions.

In the end, our simulations resulted into two sets of 117500 paths each (going around vs. going through sets, made of 50 paths for 2350 situations each). The energetic cost evaluation was performed for each path based on the equation proposed by [Whi03]:

$$E = m \int (e_s + e_w |v|^2) dt \quad (5.1)$$

where m is the mass of a person, v the instantaneous speed and e_s and e_w are constants. The energy is computed only to be compared with each other to find the trajectory that minimize the energy consumption, the mass has no influence on such computation and has been set to $m = 1kg$. The two constants have been set to $e_s = 2.23J.Kg^{-1}.s^{-1}$ and $e_w = 1.26J.s.Kg^{-1}.m^{-2}$ which are values given by [Whi03] to represent the average human.

We used the Wilcoxon signed-rank tests to compare sets of going around and going through paths computed under identical conditions.


2.2 Results

Figure 5.2 shows the average energetic cost of simulated paths depending on group size and interpersonal distance. Left and right plots are respectively for go around and go through paths. Collinear (top) and orthogonal (bottom) relative movements are displayed separately. The figure also summarizes the statistical results (middle) from the comparison between the energy spent with go around or go through strategies.

The energy graphs highlight how energetic cost varies with respect to group size and interpersonal distance. The cost of going around increases with the group size. This was expected because larger groups result in longer avoidance paths. The cost of going through decreases with an increase of the group interpersonal distance. This was also expected as the decrease of distance result in denser groups and more numerous people to be locally avoided one after the other. Note that an increase of the group size also results into more local avoidances and in an increase of the path cost, but in smaller amounts.

The comparison test results between go through and go around paths is reported in figure 5.2. Three cases can be distinguished. The red area corresponds to when going around is significantly less costly than going through. Conversely, the blue area corresponds to when going through is significantly less costly than going around. Finally, the yellow zone corresponds to a transition zone from the red to the blue: there, no significant difference in energy between going through and around is found.

Collinear



Orthogonal




Figure 5.2 Energetic cost of paths performed to avoid a group for collinear (Top) and orthogonal (Bottom) relative movements. Graphs on the left and right show respectively the average energy consumed by going around trajectories (Left) or going through trajectories (Right) according to group sizes and group interpersonal distances. Graphs on the middle show the results of the Wilcoxon signed-rank tests which compare the energetic costs for both decisions. Significant differences between both decisions are illustrated in red when going around the group is less costly than going through it, in blue when going around the group is more costly than going through it. When there is no significant difference, the result of the comparison is shown in yellow.

2.3 Discussion

Our results show that, considering the energetic cost of collision avoidance paths, the PME can suggest clear strategies: walkers have strong interest in going through large and sparse groups, and conversely, going around small and dense groups of people. However, the PME is unable to determine a better strategy for some combinations of group size and interpersonal distance. These are called inconclusive situations.

In the next section, we propose a study with a twofold objective. Using VR, we observe interaction between real individuals and groups of people. We first check that users behave as predicted by the PME. We expect that users will choose their strategy to go through or around a group in accordance to PME results for situations picked in the red or the blue areas of figure 5.2. Second, we examine the role of secondary factors in inconclusive situations.

3 Groups avoidance: a VR-based user study

3.1 Objectives

In the previous section we showed that PME is a relevant decision criterion for individual to set avoidance strategies during avoidance with groups. However, several questions remain: do humans behave as dictated by the PME? How do they make a transition in the inconclusive situations (where PME cannot distinguish one best strategy)? Do some non-physical factors influence avoidance strategies?

An obvious way to answer these questions is to replicate some of the previously studied situations. It is not trivial to get a group of people to walk in a perfectly controlled way, or to accurately replicate similar situations for several subjects. For these reasons, we decided to conduct our study using an immersive virtual reality platform: individual subjects interact with groups of virtual agents. We picked several situations considered in the PME study, some where going around is the best strategy, some where going through is the best one, and finally some where PME cannot determine a best option. For the two first kinds of situations, we expect subjects to behave as predicted by our previous study. We inspect how there is a sudden or a smooth change in their strategy for the latter kind of situations. Beyond physical factors, we also consider the effect of a social factor. We consider the role of the visual appearance of the group on the avoidance decision. Two hypotheses are tested: humans try not to go between people that have a visually perceptible social link and they try to stay out of groups that have a repulsing aspect. Our baseline condition is groups with ordinary aspect: ordinary groups are made of virtual men and women with individual appearance each (cloth, hair, skin). To give groups a repulsing appearance, we composed them with virtual zombies (see fig.5.3). Finally, to visually exhibit social links in groups, we composed them with soldiers.



Figure 5.3 Conditions for the experiment on group avoidance. **Top:** spacing between the characters of the group was: $\{1.1, 1.4, 1.7, 2, 2.3\}m$. **Bottom left:** participants avoided a group coming from: left, front or right. **Bottom right:** the visual appearance of the characters of the group was: ordinary people, soldiers or zombies.



Figure 5.4 Pictures of the experiment showing participants navigating with a joystick in a VR environment and avoiding virtual agents.

3.2 Materials and Methods

Participants 13 people (1 woman, 12 men) volunteered for this experiment. They were $28.4 (\pm 7.5)$ years old (range: 23 to 52). They were naive with respect to the purpose of the experiment. All had normal or corrected-to-normal vision. They gave written and informed consent and the study conformed to the declaration of Helsinki.

Apparatus Experiments took place in the 4-screen Computer Assisted Virtual Environment (CAVE) as illustrated on figure 5.4, equipped with 13 projectors, 15MPixels resolution in total, 9m large, 3m high and 3m deep. 3D environment display and character animation were designed in the Unity game engine. Multi-surface rendering was performed by the MiddleVR plugin. Active stereoscopy was achieved with Volfony ActiveEyes Pro Radiofrequency wearable glasses. Glasses were tracked by an ART tracking system made up of 16 cameras. Participants were standing in the CAVE and were immersed in a first person perspective in a street-like virtual environment. They were given a joystick to control their virtual motion. The longitudinal axis of the joystick controls speed linearly from $0.8m.s^{-1}$ to $2.0m.s^{-1}$. The lateral axis control the angular rotation speed linearly from $-25deg.s^{-1}$ to $25deg.s^{-1}$. When the joystick is in rest position (nobody is touching it), the speed is $1.4m.s^{-1}$ and the angular rotation speed is $0deg.s^{-1}$.

Task Participants were asked to go to a visible target (a gate in the wall in front of them, see figure 5.3). They were asked to move like they would have done in real conditions, especially they should try to avoid any collision with the virtual characters walking in the environment. The task was described to participants using slides with images and text. First, they trained with the joystick in a dedicated environment. Then, they performed the experimental task 6 times before we began recording data. At the end of the experiment, participants filled a questionnaire to report their feedback.

Conditions We fixed the virtual group radius to 3m. We reuse two factors (interpersonal distance and group relative movement) from the PME experiment. We introduced a new factor in this experiment on visual appearance. Factors are illustrated in figure 5.3. We considered the following conditions:

- **Interpersonal Distance:** the distances between characters in groups were: $\{1.1, 1.4, 1.7, 2, 2.3\}m$. With a group of 3m radius, 1.1 and 2.3 corresponds to situations where PME has a well defined optimized strategy while 1.4, 1.7 and 2 are distances

from the inconclusive zone.

- **Group relative movement:** participants interacted with groups coming from the front (Collinear motion), or from the left or the right (Orthogonal motion). We did not distinguish the left and right conditions in our analysis. There were as many repetitions from the front as from the left plus the right.
- **Appearance of the group:** we used various 3D shapes and textures for virtual characters. We studied groups made of ordinary people, zombies, and soldiers. They were all animated by a single walking cycle and soldiers had a synchronized walk.

In this study, we selected only one condition of group radius to avoid a huge number of repetitions for participants. The relative position of characters in the group was computed according to a homogeneous Poisson distribution.

Plan There were 4 repetitions for each condition, i.e., 120 trials in total. They were presented in a randomized order. We recorded the trajectories performed by participants in the virtual environment.

Analysis Dependent variables were the percentages of trials for which real humans went around or through the group of virtual walkers. Each dependent variable was analyzed into separate 3-way analysis of variance (ANOVA) with repeated measures on the following factors: interpersonal distance, character appearance and group relative movement. Greenhouse-Geisser adjustments to the degrees of freedom were applied when appropriate, to avoid any violation of the sphericity assumption. The effect size was computed using partial eta squared (η_p^2). When appropriate, Bonferroni post-hoc tests were used to further analyze significant effects between conditions. We also compared going around and going through strategies using Wilcoxon signed-rank test to evaluate the preferred strategy across participants.

3.3 Results

Figure 5.5 illustrates all the trajectories performed in the virtual environment by the participants, going towards the gate while avoiding collisions with a group of walkers. Virtual trajectories are grouped by conditions of interpersonal distances, relative motion, and appearance. Trajectories are plotted relatively to the moving group position to easily distinguish going through from going around paths. From these trajectories, we can already notice a switch in strategy from 1.1m interpersonal distance to 2.3m which corroborates the PME results. There are visually no real differences between collinear or orthogonal relative movement or among the different appearances.

The quantitative results concerning the decisions made to go around or through the group during the interactions are summarized in figure 5.6. Results showed that users performed more going around trajectories when the direction of the group motion was orthogonal to their own motion (54.1% of going around) than when it was collinear (44.2% of going around) ($F(1, 12) = 7.33, p < 0.02, \eta_p^2 = 0.38$). Participants also avoided groups of soldiers (54.4%) using more going around decisions than for zombies, and ordinary people (respectively 49% and 44%, $F(1.3, 15.5) = 8.22, p < 0.01, \eta_p^2 = 0.41$). Note that the difference between zombies and ordinary people is not significant. Finally, interpersonal distance had an effect on the user decisions ($F(1.8, 21.5) = 26.46, p < 0.0001, \eta_p^2 = 0.69$).




Figure 5.5 Participants' trajectories relative to the group position (gray circle) with respect to experimental conditions: interpersonal distance (left), direction (top right) and appearance (bottom right). Going around trajectories are displayed in red. Going through trajectories are displayed in blue.




Figure 5.6 Proportion between going around and going through decisions per distance, over all trials.

The percentage of going around trajectories w.r.t. distance was: 85,2% for 1.1m, 66.9% for 1.4m, 48.7% for 1.7m, 25% for 2m, and 19.9% for 2.3m. The behavior was not significantly different between 1.1m and 1.4m nor between 1.4m and 1.7m.

3.4 Discussion

As a main result to this experiment, we corroborate the strategies predicted by the PME is representative of the observed users' behavior. On average, whatever the other factors values, sparse groups are traversed by users in large proportions, and conversely, dense groups are avoided as a whole.

Subjects operate a smooth transition between those extreme groups density. Figure 5.7 gives a parallel view on the results of our 2 studies, we can see similar trends with the best strategy switching from going around to going through when the interpersonal distance grow bigger. On average, strategies are equally set when the group interpersonal distance is 1.7m. This situations as well as the surrounding ones (1.4m and 2m) belong

to the inconclusive set of situations according to the PME.




Figure 5.7 Decisions to go through or around amongst the participants and the energy spent in the same condition. For comparison purposes, we only consider decisions for ordinary group.

For intermediary density values, we revealed the importance of relative motion: effect is more important than what is predicted by the PME study. We believe that, even if situations are similar on the energetic point of view, the situation is visually very different for users. The difference in the group motion perception may explain this difference in the avoidance decisions set by users. Out of inconclusive conditions, when the difference of energy consumption is significant (according to the PME study), the difference concerning the relative proportion in users' decisions is very clear (order of magnitude 80 vs. 20 % for extreme conditions of interpersonal distances).

Individual preferences were not a studied factor. Nevertheless, figure 5.8 shows a variability in decision amongst participants. Such variation amongst people is common [Cos10, Hay83] and is even stronger between culture [SBK⁺10, Bea04].




Figure 5.8 Proportion between going around and going through decisions per participant, over all trials.

As expected, we could influence the strategy set by users by playing on the visual appearance of groups. The result is clear for groups of soldiers, where a tendency to go around more often is observed. However, we did not obtain expected reactions to groups of zombies: they did not have any effect on participants in comparison with ordinary people. Zombies are not real, which might have hindered the participants' capacity to experience it as a real situation.

Finally, trajectories can also be a bias as the input device (a joystick in our experiment) is far from the actual walking action. This bias has been studied by [COMP13]. But in our experiment we analyze the decision made and not the trajectory shape. Such a decision is mainly based on perception and is done pretty soon during the avoidance which should minimize the bias.

4 Applications

Sections 3 and 4 provide us theoretical and experimental results on individual strategies and decisions to avoid a group of walkers. Based on these results, we propose an algorithm for group avoidance decision that selects the strategy that best fits the human one. This section presents the algorithm as well as an example of integration to the RVO2 model ([vdBGLM11]).

4.1 Decision algorithm: Around or through the group?

The PME study has delimited three zones where the best strategy is well defined. The first zone concerns very small groups, below $2m$ radius, that are so small that going around is always a good strategy. For the groups that are bigger, there is a zone with interpersonal distance below $1.35m$ where going around is better and a zone with interpersonal distance above $2.05m$ where going through is better. Between these two zones, the user case study has shown that the decision rate goes from a majority of people going around the group to a majority of people going through the group between $1.1m$ and $2.3m$ interpersonal distance. Figure 5.6 show that this evolution is linear and the user case study has highlighted personal preferences as a reason for this linear evolution. Moreover, the group's appearance and the relative movement have both demonstrated some influence over this decision.

From these results, we developed the algorithm 3. This algorithm takes as input both an agent and a group and returns a Boolean answering the question whether the agent should cross the group or go around. If the group size is less than $2m$ than the strategy is to go around as hinted by the PME study. Otherwise, a specific threshold is computed according to the agent own preference ρ_a , the group appearance score G_{vu} and the relative motion score G_θ . Both appearance and relative score measure the effect of these factors on the threshold. For example, a group of soldier will have the highest effect on the threshold with an appearance score of 1. The relative motion score goes from 0 to 1 with 0 being collinear and 1 being orthogonal. For the three parameters: the personal preference ρ_a , the appearance maximum effect C_{vu} and the relative movement maximum effect C_θ , we propose default values based on our best guest from our previous results. To represent the variability in decisions between people and the linear progression seen in figure 5.6,

Algorithm 3: Function that select a strategy to avoid a specific group for a specific agent

```

Data:  $agentID, groupID$ 
Result: true if the agent should cross the group, false otherwise
  // test
1 if  $groupSize(groupID) < 2$  then
2   | return false;
3 end
  // Get the personal preference of the agent
4  $\rho_a \leftarrow getAgentThreshold(agentID)$ ;
  // Get the visual factor of the group
5  $G_{vu} \leftarrow getGroupVisualFactor(groupID)$ ; //  $G_{vu} \in [0 : 1]$ 
  // Compute the relative movement factor
6  $\vec{dir}_a \leftarrow getAgentDirection(agentID)$ ;
7  $\vec{dir}_g \leftarrow getGroupDirection(groupID)$ ;
8  $G_\theta \leftarrow 1 - \left| \vec{dir}_a \cdot \vec{dir}_g \right|$ ; //  $G_\theta \in [0 : 1]$ 
  // Compute threshold and choose strategy
9  $threshold \leftarrow \rho_a + G_{vu} \cdot C_{vu} + G_\theta \cdot C_\theta$ ;
10 if  $groupInterDistance(groupID) < threshold$  then
11   | return false;
12 end
13 return true;

```

the $\rho_{default}$ was set by a uniform distribution from $1.3m$ to $2.1m$. We set $C_{vu} = 0.15m$, by considering that soldiers have the highest effect from group appearance, and $C_\theta = 0.06m$.

4.2 Integration in RVO2

The algorithm we proposed determines individual decisions to avoid a group of walkers depending on several factors. It appeared to be easy to combine it with existing crowd simulation models to adapt their avoidance behavior to groups. To prove the feasibility of such an adaptation, we proposed to combine our decision algorithm with RVO2 model. This model deals with local interactions and was not designed to consider groups as a whole. The objective here was then to define in the simulation which agents should be merged into a group to be avoided by going around it. The first adaptation was to include a group criterion. It was based on [YCP+08]. When agents are identified as part of a group, all of them are replaced by a single 'big' proxy agent, as presented in figure 5.9. The proxy agent current speed was set as the average speed of all agents of the group. In that way, the group is avoided as a whole and agents of the group are not considered individually with succession of local avoidances: the resulting avoidance trajectories will go around the whole group instead of going through the group around each individual agent.

To improve the trajectories performed to avoid a group (by going around the group), we also proposed another adaptation. Indeed, when avoiding large obstacles, the resulting




Figure 5.9 On the left: the situation as it really is, on the right: the situation as it is perceived by the red agent (All the blue agents are seen as one big agent because they are considered as a group while the green ones are too far from each other to be considered as a group)

trajectory can be jerky if the adaptation is performed too late. In such a case, walkers need to adapt their motion earlier to go around the group while keeping a smooth trajectory. To consider this point, we adjusted the timeHorizon parameters of RVO2 which defined how soon an agent should avoid other agents. We defined it as a factor of the size of the agent or proxy agent to avoid, allowing sooner avoidance when interacting with larger groups.

Results of the combination of RVO2 with the decision algorithm are presented in the companion video. An example of RVO2 behavior after its combination with our decision algorithm is illustrated on figure 5.10. On the left, the purple agent goes through the group, only considering local interactions. After combining RVO2 with the decision algorithm (on the right), the purple agent avoids all the red agents as part of a whole group instead of trying to go through it.

4.3 User-study evaluation

We integrated our decision algorithm in RVO2 simulation model. The objective of this study is to check if our decision algorithm is able to improve the visual quality of crowd simulation results. To this end, we proposed a subjective evaluation based on a perceptual user study: we asked some participants to compare simulation results from RVO2 only with ones from RVO2 extended with our algorithm. 10 male participants volunteered for this experiment. They were $28.3 (\pm 4.7)$ years old (range: 24 to 40). They were naive with respect to the purpose of the experiment. All had normal or corrected-to-normal vision. They gave written and informed consent and the study conformed to the declaration of Helsinki.

We prepared 4 situations (20s each) where one single agent walks counter flow in front of a crowd of walking people. This crowd was made up of individuals and groups of various densities. For each situation, agent motion was simulated both with RVO2 and with RVO2 combined with the decision algorithm, creating 2 video stimuli. Participants were seated in front of a desk (cf. figure 5.11). Stimuli were displayed on a 24inches screen.




Figure 5.10 Example of avoidance decision made by the purple agent. The agent goes through the very dense group in RVO2 while going around it in the modified RVO2.

Participants were asked to observe successively the 2 video stimuli and to evaluate which one complies the most with the decision they would have made as a single walker avoiding other people in the same conditions. The 2 video were made from a third person view point to give participants a good view of the situation as the computer screen limits greatly the field of view of a first person view point. In total, they repeated this task 20 times (4 situations*5 repetitions). Stimuli were presented in a randomized order. We compared the percentage of answers supporting RVO2 with respect to the percentage of answers supporting RVO2 combined with the decision algorithm using a Wilcoxon signed-rank test. Results showed that RVO2 combined with the decision algorithm was perceived as more compliant with decisions a human would have used to avoid people than RVO2 alone (RVO2+Decision algorithm= $77.5\% \pm 17.36$, RVO2= $22.5\% \pm 17.36$, $Z = 2.66$, $p < 0.01$). The participants were asked which criteria they used to choose between the two video stimuli. Some answers are transcribed below; globally people chose what seems to be the most natural motion for them. Trying to go through a dense group is the main reason that breaks the naturalness of the motion, according to the commentary 3 below, splitting small group is worse than having weird trajectories.

1. I chose according to what the path I would have selected for my own motion.
2. Avoiding most of dense crowds seems more natural.
3. Splitting small groups is a deal-breaker (happened in some cases where trajectories were actually better than other solution)
4. My criteria would change according to the density (homogeneous or heterogeneous). In homogeneous density, the behavior "I go directly toward the goal" seems more realistic as there were no visible group to break. In heterogeneous density, I preferred the other behavior that seems to preserve the groups from the other agents.

This result is promising and shows that including our decision algorithm to existing crowd simulator will improve their level of realism. The commentary 4 highlighted a difference for homogeneous and heterogeneous density. In higher density situation, people get closer to each other leading to smaller personal space and less visible groups. It is fair to say that this could greatly influence the group avoidance strategy. The experiment done in this paper was done with no other agent then the group members to be able to consider only the interaction with the group. But experiments in high density are also needed to check how the environment around the groups influences this group avoiding behaviors. Beyond that, our perception study shows that spectators perceive improvements on simulation results, and find them more natural. Some people find that the reaction facing a group can be even more important than the avoidance trajectory itself for realism. This demonstrates the need for improving the ability of crowd simulators to deal with more complex situations of interactions.




Figure 5.11 Experimental setup to evaluate the decision algorithm.

5 Conclusion

In this paper, we study interaction between individuals and groups; we focus on strategies set by individuals to go through or around a group of people. We extend the use of the Principle of Minimum Energy from choosing local speed variation to selecting global avoidance strategy. We confirm the influence of the energy factor, in this specific situation, and complete its results with an experiment where real individuals were interacting with different groups. We also highlight, with the experiment, two secondary factors that influence the chosen strategy: the relative motion and the visual aspect of the group. We propose an algorithm to simulate realistic avoidance strategies between individuals or groups. We also integrate our results in an existing crowd simulator and demonstrate that spectators can perceive some improvement in the crowd animation.

This work opens various perspectives. We propose to use VR as a new method to gather observation on crowd behaviors and local interactions in crowds. We show its strength and how it can help to better understand complex interactions that are hard to study using regular observation gathering method: only one subject is required to observe

behaviors in crowds, we have no technical difficulty with tracking, we can perfectly control complex experimental conditions. Especially, crowd simulation has only considered interactions between individuals for too long, real situations of interactions in crowds are often much more complex. We focused on interactions between individuals and groups. Our results also reveal how humans process and simplify interactions, e.g., by considering a group of many as a whole. We need to continue performing experiments on individual behaviors in crowds: simulators would widely benefit from conclusions about the way people in crowds filter, select and combine the numerous interactions they have. We expect large improvement of microscopic crowd simulation algorithm in the near future thanks to next iteration of experiments to come based on our VR-based experimental facility.

Hotspot in human environment perception

6

Contents

1	Introduction	84
2	Previous gaze experiments	84
3	Experiment	88
4	Analysis	90
	4.1 Fixations computation	90
	4.2 Virtual human characteristics	91
	4.3 Statistical analysis	93
5	Results	93
6	Discussion	94
7	Conclusion	96

1 Introduction

When navigating through a crowded street, there are many possible targets to interact with and many possible interactions to perform. An individual attention is limited, it is unlikely that humans could get all the required information from each target and then solve all these possible interactions at once. They must focus on a few elements of the environment and select a limited number of interactions to perform. However, there are few results explaining how humans select the information about their surrounding. We can then question this selection process and wonder how humans select the interactions to solve in priority.

Optic flow was shown to be one of the main sources of information used by humans to gather data about the surrounding environment and assess the situation during navigation tasks. The importance of the optical flow during human interactions has already been discussed in previous work [CVB95, Ber97, FRT09]. Some algorithms even base their avoidance computation on optical information (see chapter 2 section 1.1.1). Many relevant data are extracted from it, such as the derivative of the bearing-angle [CVB95] or the time to collision [CL88]. We hypothesize that people look toward one of the most important target they chose to interact with (whether it is an object of the environment or another person), as they require precise information about the target and its surrounding for the interaction.

To gain some insights on this selection process, we propose to study the human gaze during a navigation task in a virtual environment. For this study, an experiment has been performed in a virtual environment using a computer screen and a joystick. During the experience, the participant had to navigate through a crowd of virtual humans, giving him many possible targets to look at. The participant's gaze was studied using an eye tracking device. Eye tracking devices can be used to record the human's gaze direction and information about what people are looking at can be extracted ([PB06, ANH10, HNM12]). This experiment has been performed at the end of the PhD and the data are currently still being analyzed.

In this chapter, we start by presenting some previous gaze studies about navigation tasks in section 2. Then, we present the experiment in section 3. The data analysis, and the method used to compute the gaze fixation, are explained in section 4. Then some preliminary results are presented in section 5 followed by a discussion about these results and future analysis in section 6. Finally, a conclusion is given in section 7.

2 Previous gaze experiments

Eye tracking devices have already been used to study human gaze during navigation task. Most of these studies have been performed to analyze the visual cues used by humans to find a path in the environment [ADKB12, WHBK12, SvSB⁺13]. Many of them are done in laboratory settings with participants looking at pictures or videos. Some evaluations have been performed to compare results between laboratory experiments and real world ones.

Foulsham et al. conducted such an evaluation [FWK11]. They had participants walking to a coffee place while wearing a portable eye tracking device. A week later the same participants returned to perform a lab experiment. During the lab experiment, participants were presented with a video of another participant walking to the coffee, taken from the perspective of the walker. They were asked to imagine that they were walking the route and to behave as they normally would have done. The results of the comparison showed both similarities and differences in participants' gaze (see Figure 6.1). Many of the differences could be explained by the fact that participants watching the video were not active in the navigation. For example, there is a huge difference in the path fixation, but there is no real need for participants watching the video to look at the path since they are not the one choosing where to go. The authors concluded: "The results suggested both similarities and differences, with the differences being attributable to the freedom of immersed participants to move their head and their body, to their engagement in an active task, and to the social constraints of having other real people nearby".




Figure 6.1 Several results from the experiment performed in [FWK11]. From left to right: the proportion of gaze toward different elements, the probability of a pedestrian being fixated according to the amount of time he has been visible, the proportion of gaze toward pedestrians per time period since their appearance.

Schwarzkopf et al. did a similar comparison while studying the gaze of participants having to find a gate in an airport [SvSB⁺13]. This experiment has been performed in a lab with a stationary eye tracking device and static images of the airport indoor environment (Figure 6.2-Top). The experiment has been performed again with a mobile eye tracking in the real airport (Figure 6.2-Bottom). Results have been compared to validate the lab experiment. As illustrated by Figure 6.2, there are many similarities between the recorded gaze in the lab and in the real airport: participants seem to look at the same zone (the sign and the area next to the elevator). But there are some differences too, as the sign from the elevator scene was ignored during the experiment in the real airport. Several reasons were given by the authors, the most important being that, in the real airport, the sign was not exactly at the same place and was quickly passed by the participants before they could notice it.

These studies show that despite some differences in gaze between lab experiments and real world ones, there are many similarities and results from the lab experiments are still relevant while not entirely complete. But most of the differences seem to be caused by the lack of engagement of the participants in the active task as they mostly watch pictures or videos. This means that these differences could probably be significantly reduced by having the participants actively navigate through a virtual environment instead of using




Figure 6.2 Heat maps, from [SvSB+13], of a newly built area looked at on two subsequent, static slides in a lab study (top) before the construction of the building as compared to the heat map of the same terminal after construction derived from mobile eye tracking videos recorded at the airport (bottom). Red and green, indicate a high and low number of overall fixations respectively.

pictures and videos.

Amongst the gaze studies done during navigation tasks, some focused on studying humans interactions with the environment [CPA09] or with other people [JMH09, KF10, FUCH15]. Jovancevic-Misic and Hayhoe place participants in a circular corridor with other pedestrians (see Figure 6.3-Left). Three pedestrians were walking in an opposite direction and had different roles: the Rogue who would go on a collision path every time he approached the participant, the Risky that would also go on a collision path but only half the time and the Safe that would never collide. During the experiment, the gaze of the participant was recorded with a portable eye tracker. Their results showed an increase in fixation toward the pedestrians that would be on a collision course more often (Figure 6.3-Right). Participants learned quickly to recognize these dangerous pedestrians and looked at them when performing collision avoidance maneuvers.

Kitazawa and Fujiyama performed a study of the human gaze during navigation while interacting with obstacles and other pedestrians (see Figure 6.4-Left). During the experiment, up to 4 participants had to walk on a platform, with or without static obstacles on the way, while their gazes were tracked by a head-mounted tracking device. While it was expected that most of the fixations would appear on static obstacles and pedestrians, results showed many fixations on the platform surface. Walking safely on the platform




Figure 6.3 Experiment performed in [JMH09]. Left: Illustration of a path around the occluded area (The red circle is the subject, the blue squares are the pedestrians walking in the direction opposite and the gray triangles are pedestrians walking in the same direction.) Right: Effect of collision probability on fixations on pedestrians.)

may represent an immediate issue, while pedestrians interact with static and dynamic obstacles, they also interact with the environment itself, watching for environmental hazards (e.g. a gap in the surface). This should be taken into account, when performing an experiment, to reduce its effect when other factors are studied. They have also gathered very interesting data about the shape of the perception (see Figure 6.4-Right). Participants seldom fixated objects with more than 45 degrees angle with the walking trajectory of the observer, giving higher priority for the front area. This result confirmed that using a computer screen, which limits peripheral vision, will have an impact on the quality of the experiment, but this impact will be limited as the front area remains the most important fixation zone.




Figure 6.4 Experiment performed in [KF10]. Left: an example of the field of view recorded by Eye tracker. Right: Relative position of fixated or ignored objects (ignored objects are represented by the cross).

All these studies on interactions are limited to a small number of objects (static obstacles and pedestrians) to interact with or to fixate. To ensure the recorded fixations are not selected by defaults, participants need to be given a large choice in objects to fixate or interact with. In the best of our knowledge, such experiment has not yet been done. But the experiments presented in this section have shown that using VR and a computer screen should have a limited impact on the quality of the results. While performing the experiment in the real world, or even with a more immersing device such as a CAVE, would yield less biased results, the results obtained with the computer screen should already give us plenty of reliable information on human gaze activities during navigation.

3 Experiment



Figure 6.5 Participant using the joystick to navigate in the virtual environment while his eyes are being tracked by the eye tracking device below the screen (the captor with purple dots on each side).

Participants 19 people (6 women, 14 men) volunteered for this experiment. They were 25 (± 4) years old (range: 24 to 40). They were naive with respect to the purpose of the experiment. All had normal or corrected-to-normal vision. They gave written and informed consent and the study is conformed to the declaration of Helsinki.

Apparatus The experiment took place on a 24inch screen ($53.1 \times 29.9\text{cm}$) as illustrated in Figure 6.5. A desktop eye tracker device was positioned below the screen. The eye tracker used is an Eye Tribe with an accuracy of $0.5 - 1^\circ$ and a precision of 0.1° (RMS). Participants were seated facing the screen at a distance comprise by the eye tracker operating distance (between 45 and 75cm). They were given a joystick to control their virtual motion. The longitudinal axis of the joystick controls speed linearly from 0.8m.s^{-1} to 2.0m.s^{-1} . The lateral axis controls the angular rotation speed linearly from -25deg.s^{-1} to 25deg.s^{-1} . When the joystick is in rest position (nobody is touching it), the speed is 1.33m.s^{-1} and the angular rotation speed is 0deg.s^{-1} .

Task Participants were asked to go toward a visible target (a light up in the sky in front of them). They were asked to move like they would have done in real conditions; especially they should try to avoid any collision with virtual humans walking in the environment. First, they trained with the joystick navigating in the environment until they were satisfied with their virtual motion.

Scene Participants had to navigate inside a virtual environment. The virtual environment is populated by several virtual humans. Smog has also been added to the environment to limit participants' fields of view to 24m ahead and hide virtual humans that are further. Virtual humans walk straight forward at constant speed but still avoid collision thanks to the RVO2 collision avoidance algorithm. As we do not want the virtual humans to do all the avoiding, their avoidance reaction is parameterized to be late (time-

Horizon=5). They are spawn randomly along the default path of the player (controlled by the participants) if no adaptation is made by the participant. To achieve that, we build a rectangle ($20 \times 30m$) around the default position of the player at different time (every $2.5s$) and randomly spawn 7 virtual humans so that they will reach a position in this rectangle at the same time (as presented in Figure 6.6). While participants may adapt their trajectory to interact with the virtual humans, the trials are short enough so that the deviation in the trajectory should remain within the built rectangles. All virtual humans share the following properties:

- walking speed is $1.33m.s^{-1}$, which corresponds to a comfort speed
- Their direction is the direction of the player at the start rotated randomly between 90° and 270° . That way, virtual humans would come only from the front or the side of the player but not from behind.
- There is no collision between any of the virtual humans but they may still collide with the default player trajectory.




Figure 6.6 The generation of the virtual humans' position and speed is done to insure the presence of at least 7 virtual humans in a square around the default player's position at a constant time interval δt .

Conditions Four experimental conditions were designed:

- The default condition as presented above.
- The speed variation condition with 3 different type of virtual humans: the slow one ($0.7m.s^{-1}$), the normal one ($1.33m.s^{-1}$) and the fast one ($2.1m.s^{-1}$). With this condition, we want to study the influence of speed on the gaze: is the participant's gaze more attracted by faster moving object?
- The sudden rotation condition with visible virtual humans randomly selected to perform a sudden rotation in front of the participant with a random period comprise between 2 and 8s. The sudden rotation consisted of a $0.4s$ rotation, randomly distributed between -90° and 90° . These rotations could cause or prevent a collision with the participant but not necessarily. With this condition, we want to study the effect of unexpected changes in the visual flow: is the participant's gaze attracted to perturbation in the visual flow?

- The eye tracking test condition used to check the precision of the eye tracking device and the data processing. This condition is similar to the default one with two exceptions: the participant is unable to adapt the player motion and virtual humans have no risk of collision with the player. Visible virtual humans are randomly selected to turn red for a small duration (there can be only one red virtual human at the same time). The participant is asked to fix the red virtual human when one is present and to look straight toward the light when none is present.

Plan There were two repetitions of default, speed variation and sudden rotation conditions. Each repetition lasts approximately 2 minutes. They were presented in a randomized order after the training. Then once all the repetitions of the 3 conditions were performed, participants were presented with the eye tracking test condition which lasts 1 minute. We recorded the trajectories performed by participants and virtual humans in the virtual environment as well as the participants' eye position and direction in the virtual environment. We also recorded which virtual humans are visible on the screen.

4 Analysis

4.1 Fixations computation

From all the trajectories and the line of sight of the participants, we still have to compute fixations: which virtual human is the participant looking at if he is looking at a virtual human at all. Eye tracker technology has already been used for various experiments. Fixation is the most important measure used during these experiments, thus there is already some methods to compute the fixation from the raw data of the tracker [SG00]. But these experiments rarely consider moving objects as fixation points and thus their methods to compute fixations are not adapted to our moving virtual humans and need to be modified. For our purpose, the fixations are computed using the 4 following steps:

1. first selection and punctual score: this step determines which objects are in the line of sight of the participant for each time step,
2. continuity score: this step prioritizes the objects that are followed by the participant's gaze for a longer period,
3. selection of best score: this step selects the most likely fixation,
4. filter fixation: this step invalids fixation or no fixation periods which last less than 0.2s (According to Salvucci and Goldberg, fixations are rarely less than 100ms often in the range of 200 – 400ms).

The first step is to select possible target points and compute a first punctual score for each time step. The selection is done amongst all the visible virtual human positions (the one on screen not hidden by the smog) and a position representing the goal. The position representing the goal is set at the fixed distance of 24m from the player (max visibility distance in the smog) in the direction of the goal. The score is computed as a ratio of the distance between the position of the possible target and the line of sight. All that are

too far from the line of sight are removed from the possible target points (above 1m) for the other the score is: $PS(A, t) = \frac{4}{3 \times \min(1 - \text{abs}(\text{LoSdist}(A, t)), 0.75)}$ with A a possible target, t the current time and $\text{LoSdist}(A, t)$ the distance between the position of A and the line of sight of the participant at time t . This gives us scores between 0 and 1 with 1 being the more likely to be the target of the participant sight than 0 (this is a score and not a probability, the sum can be over 1).

Once this is done, a continuity score is computed for each time step, only for possible target points (those with punctual scores above 0). The algorithm for the continuity score computation is presented in algorithm 4. It starts by creating possible fixations for every target point: a set of non zero punctual scores consecutive in time with gap between them smaller than 0.2s (lines 2 to 18). For each possible fixation, the punctual score at each time step is replaced by a continuity score. The continuity score is composed of the punctual score, a variation penalty and a duration boost. The variation penalty checks whether the distance between the target points and the line of sight is stable for the duration of the fixation. For each time step of the fixation, the distance to the line of sight is compared to the median distance of the entire fixation duration and the continuity score is decreased accordingly (line 22). The duration boost increases the score by a factor of the duration of the possible fixation (line 23). As all the possible fixation points are moving, this continuity score allows to check if the current possible target is just on the way of the eye movement or have really been followed by the eyes.

Then for each time step, the possible fixation point with the highest score is elected as the fixation. Moreover, fixations can also be separated from random noises and eye movement by its duration. So all the fixations that last less than 0.2s are removed. Moreover, all lack of fixation of less than 0.2s, surrounded by two fixations on the same target, are also removed and the two fixations merged into one.

4.2 Virtual human characteristics

During the experiment, many virtual humans are visible at the same time. Participants only select one of them to fixate at a time. We want to determine what are the criteria influencing this selection. For this, we need to check the characteristics of the fixated virtual human compare to the other visible ones, that have been discarded by the participant for the time being. The considered characteristics are the following:

- **MPD rank:** the Minimum Predicted Distance measures the risk of collision, the virtual human of rank one being the one with the higher risk of collision with the participant (MPD closer to zero).
- **TTCA rank:** the Time To Closest Approach measures the imminence of the crossing (and of the collision if there is a risk), the virtual human of rank one will be the first one to be crossed by the participant (lowest TTCA value).
- **Distance rank:** this is a distance ranking with rank one being the closest virtual human from the participant.
- **Speed:** this measures the speed of the virtual human and is used only for the speed variation condition.

Algorithm 4: The continuity score computation

Data: Punctual score array PS , Distance from Line Of Sight array $LoSdist$, All simulation targets O

Result: Continuity score array CS

```
1 forall  $o \in O$  do
2   Sequences $\leftarrow$ [];
3   NumOfSequences $\leftarrow$ 1;
4   SequenceIdx $\leftarrow$ 1;
5   GapSize $\leftarrow$ 10;
6   for  $t = [0 : endSimTime]$  do
7     if  $PS(o, t)$  is not null then
8       Sequences(NumOfSequences, SequenceIdx) $\leftarrow$  $t$ ;
9       SequenceIdx $\leftarrow$ 2;
10      GapSize $\leftarrow$ 0;
11    else
12      if  $GapSize < 0.2$  and  $GapSize + timeStep \geq 0.2$  then
13        NumOfSequences $\leftarrow$ NumOfSequences+1;
14        SequenceIdx $\leftarrow$ 1;
15      end
16      GapSize $\leftarrow$ GapSize+timeStep;
17    end
18  end
19  for  $s = [1 : NumOfSequences]$  do
20    medianValue $\leftarrow$ median( $LoSdist(o, Sequences(s, 1) : Sequences(s, 2))$ );
21    localScore $\leftarrow$  $PS(o, Sequences(s, 1) : Sequences(s, 2))$ ;
22    variationPenalty $\leftarrow$  $\frac{1}{1 + 2 \times \text{abs}(\text{medianValue} - LoSdist(o, Sequences(s, 1) : Sequences(s, 2)))}$ ;
23    durationBoost $\leftarrow$   $1 + (\sum(\text{variationPenalty})) \times 0.05$ ;
24     $CS(o, Sequences(s, 1) : Sequences(s, 2)) \leftarrow$  durationBoost  $\times$  localScore  $\times$ 
      variationMalus $^T$ ;
25  end
26 end
27 return  $CS$ ;
```

- **TFSR**: the Time From Sudden Rotation measures the time from the last sudden rotation of the virtual human and is used only for the sudden rotation condition.
- **Rotation amplitude**: this measures the rotation angle during the last sudden rotation.

We then computed the percentage of fixations corresponding to each of these presented characteristics. For example, with the MPD we can check whether the virtual human, posing the greatest risk of collision, is being look at more often than the other.

4.3 Statistical analysis

Statistics were performed using sigmastat software. The percentage of fixations associated to the virtual human characteristics were defined as the dependent variables of the statistical analysis. A Kolmogorov-Smirnov test has been performed to compare the data distribution to a normal distribution. Because all of the data did not follow a normal distribution, we used a non parametric Friedman test to determine whether the ranking of virtual human characteristics has an influence on the percentage of fixations. There to highlight which ranks differ from each others concerning the number of fixations, a Student-Newman-Keuls post hoc test has been used.

5 Results

This section presents the preliminary results of the data analysis.

The first result is the validation of the fixation computation. The eye tracking test condition has been done for this purpose. In this condition, we know what participants are looking at and can check if that corresponds to the computed fixation point. It also allows us to tweak the algorithm and its parameters to improve the computation of fixation. Each time a red virtual human appears, we wait 0.8s to give time to the participant to see the red virtual human and to fixate it, then we check at every time step whether the computed fixation corresponds to the red virtual human. This gives us the percent of good fixation detection that is shown in table 6.1 for every participants, with an average of 77% of good detection. This also shows that there might have been an issue during the eye tracking with participants 2 and 16. The same computation was done when no red virtual human was visible and participants had to fix the goal (the light). The results for the goal fixation are not detailed as they are dispersed between 0% and 78%. We probably did not emphasize enough during the instructions that participants had to fix the goal when there were no visible red virtual humans.

Participant	1	2	3	4	5	6	7	8	9	10
Red virtual human detection (%)	76	38	92	97	77	91	53	99	100	100

Participant	11	12	13	14	15	16	17	18	19
Red virtual human detection (%)	80	78	81	74	84	1	88	77	83

Table 6.1 Percent of correctly detected fixation on red virtual humans per participants.

With the fixation computation validated, we can now present our preliminary results about the effect of the rank of MPD, TTCAA and distance on the percentage of fixations. The results of this analysis are presented in Figure 6.7 and detailed below:

Effect of the risk of collision (MPD): The percentage of fixations is influenced by the MPD ranking ($\chi^2(18) = 315.5; p < 0.001$). The post hoc test reveals significant differences between ranks ($p < 0.05$):

1>2>3>4>5>6>7-8>9-10>11>12>13-14>15-16>17-18-20.

Effect of the imminence of the interaction (TTCA): There is also an effect of the TTCA on the percentage of fixations ($\chi^2(19) = 337.1; p < 0.001$). No significant differences has been found between the rank 2, 3, 4 and 5 which are significantly more gaze upon than the other ranks. The complete order is the following:

2-3-4-5>1-6>7>8>9>10-11>12>13>14>15>16-17-18-19-20

Effect of the distance: The distance also influences the percentage of fixations ($\chi^2(20) = 349.780; p < 0.001$). This time, rank 2 is significantly the most gaze upon rank with the complete order being:

2>1>3>4>5>6>7>8>9>10>11>12>13>14>15-16-17-18-20

These results show that human gaze, during navigation, seems to be attracted by dangerous individuals: the one with the highest risk of collision. When looking at the imminence of the closest approach, while the gaze is attracted by individuals that will soon be close to us, the most imminent one is less gaze upon than the following ranks. This could be explained by the fact that humans avoid collision by anticipation (which has been discussed in chapter 2 section 2.2). Thus the individual with the minimum time before closest approach should most of the time already been avoided. This requires more analysis, like analyzing the TTCA ranking only amongst the virtual humans presenting an actual risk of collision.

6 Discussion

The preliminary results presented in the previous section are promising. It seems that MPD and TTCA have a strong influence over whom the participants' gaze is attracted. The MPD represents a danger (the collision risk) and the TTCA represents the imminence of this danger. We believe that humans' gaze is attracted to imminent danger in priority. Thus, a human being would not look at the person with the higher risk of collision or the person that will cross his path the sooner, but will select a person with a specific balance over these two factors: a person with a high risk of collision that is going to happen soon. To test this hypothesis, more analysis is required. A heat map could be analyzed between the two variables TTCA and MPD to check whether a specific balance of the two variables attract more the gaze of the participants. But such analysis needs to be carefully performed as participants probably face only a few number of virtual humans that are of rank one in MPD and TTCA (the virtual humans with the lowest TTCA have probably already been avoided by the participant, or at least been looked at). We could also redo the analysis after removing those virtual humans that have already been looked at. This would allow us to check whether the fixated one are the most dangerous amongst the ones that have not yet been studied by the participant.

But before using these results to improve collision avoidance algorithm, we need to




Figure 6.7 Graphs showing the percentage of fixations toward different rank of MPD, TTCA and distances.

validate our main hypothesis: people look toward one of the most important target they chose to interact with. For this validation, a deeper analysis is required to check whether fixations are correlated to avoidance maneuvers (the main interaction present in the experiment). This could be achieved by measuring the evolution of the MPD of a fixated target around the fixation time. If the MPD before the fixation reveals a collision and its value rises above the collision level during or after the fixation, it would mean that the participant did avoid the fixated target. It would also be interesting to study whether participants keep fixating their target during avoidance to make sure that the avoidance is well done or if they start looking for their next target directly after assessing the collision risk and the avoidance strategy.

If the hypothesis is confirmed, then the results can be used to improve the avoidance selection done in current collision avoidance algorithms. By using the same selection criteria as human do, collision avoidance could be focused on the right collision to solve which could in turn improve the results, making the simulation more realistic. Moreover, the selection could be decrease to the few relevant surrounding virtual humans which

should ease the avoidance process and improve the computation time.

7 Conclusion

In this chapter, we have presented an experiment recording the gaze direction of participants while they were navigating in an environment with several virtual humans to avoid. The preliminary results show an interesting trend: more gaze toward dangerous individuals (the one with high collision risk that might be imminent). A deeper analysis is required to confirm and precise this trend. We also need to check other factors, such as speed differences and sudden rotations, which might influence the fixation selection. Finally, we need to confirm that participants are interacting with their fixated targets. But we have hopes that these data will yield many useful results to update the selection process during collision avoidance in crowd simulation.

Conclusion and future work

7

This PhD main focus has been to improve crowd simulation in complex situations: to reduce the number of unnatural behaviors from the virtual humans in situations with many different individuals to interact with. The complexity of these situations is twofold: there are many possible interactions that can be performed at the same time and there can be huge differences amongst the possible targets that would greatly influence the way the interactions are handled. Our work resulted in two types of contributions: one was about working on the algorithms to simulate crowds, and the other was about performing studies to better understand human interactions. The work done on the algorithms has been focused of expanding the type of situations that can be correctly simulated while not deteriorating results on situations that are currently well simulated. This has been done by improving the way many interactions and different targets are handled. The studies have also been focused on these complex situations and how humans handle this complexity with such ease. The goal was to better understand how humans select the people to interact with and then how the interactions are combined.

1 Contributions

Algorithms: The first contribution has been to explore some factors that influence people interactions with others. The considered factors were physical factors to avoid contact, social factors to keep respectful distances to others, psychological factors to take into account the reaction time or the level of attention, and other factors related to others' motion perception to filter out irrelevant motion fluctuations. An algorithm has been implemented that adapts the following distance in relation to these additional factors. The results were evaluated with data from real people following and compared to previous solutions. The proposed algorithm proved to be able to produce results at least as good as previous algorithms in the most common situations. But thanks to the additional factors, it was able to adapt to other kinds of situations (following a drunk person, an old lady, a runner ...) and gave better results than previous algorithms.

The second contribution has been focused on the combination of collision avoidance with several obstacles. We proposed a new way to combine these interactions by ordering them through time, planning one collision avoidance after another. This planning process is done before the current local avoidances, which are done by current algorithms but reduced to the immediate surrounding. An implementation of this solution has been proposed: the EACS algorithm. Results with EACS show an improvement in complex situations compared to results without the planning process. The planning process allows agents to select more efficient strategies to avoid several individuals and prevent many unnatural behaviors.

Finally, there has been a collaboration with Zhiguo Ren on his work about group behaviors. A modification of RVO2 has been proposed to incorporate group behavior in the simulation with agents trying to stay close to the one they are related to. Results have been analyzed to demonstrate the wide varieties of groups that can be simulated by the algorithm and its capacity to simulate complex scenarios with many individuals and different kind of groups navigating and interacting with each other.

Studies: The first contribution was a collaboration with Anne-Hélène Olivier on her work to validate a VR platform for navigation studies with walking virtual humans. This work yields similar results than previous validation work done in static virtual environments: there are some quantitative differences between trajectories made in VR and the one made in the real world, but qualitatively they are very similar as humans make the same decisions. It also determined which locomotion devices yielded the best results amongst the one we had at our disposition but the proposed framework can be used to evaluate other devices. The validation of the VR platform allowed us to use it as a tool to study human interactions and led to the second contribution.

The second contribution has been a study on how do individuals avoid groups. This study was done using VR with participants interacting with groups of virtual humans. This study confirms that humans can interact either with the group as a whole big entity or with the individuals forming the group. Thanks to VR, the group's parameters were perfectly controlled which allowed us to determine when humans consider groups as one big entity or as several individuals. This knowledge was then used to improve the way groups are avoided in crowd simulation. The study also demonstrated that humans do not always combine interactions independently from each other and that VR can be a very important tool to study human interactions

At the end of the PhD, a study has been started about real human gaze during navigation task but the data analysis is still an ongoing work. An experiment has been done with participants navigating in a virtual environment, filled with many virtual humans to interact with, while an eye tracker was recording the direction of their gaze. Preliminary results hint that real humans fixate virtual humans in a non-random pattern. Additional analysis is required to study this selection process.

All the studies done during this PhD have been performed using VR technology. These studies have proven that VR can be a very important tool to present participants with complex situations to interact with. Complex situations are fast and easy to setup in a virtual environment and VR gives us enough control to prevent any variability between trials or to limit it for studied factors only. Data analysis is also simplified and accelerated when using VR as the data is less noisy and does not require complex reconstruction, which is often the case in laboratory settings. While the use of VR introduces some biases, these biases are known and can be circumvented during a study as it has been done during this PhD.

Both: Crowd simulation has been in constant evolution. This constant evolution led to current algorithms being able to handle simple situations with realistic results. The contributions done during this PhD are in the continuity of this constant evolution. They have been done to extend the range of situations that can be realistically simulated with: to be able to handle more interactions, combine different interactions, interact with

very different targets... On one side, new algorithms have been designed to handle the complexity of this new range of situations. On the other side, studies have been performed to improve our knowledge of such complex situations and to give us some reference data to check simulation results in these complex situations. Both algorithms and studies have been done in parallel, the experience gained during the studies was a huge help when designing the algorithms, and some of the trends found during the studies confirmed an improvement with the proposed algorithms in some situations (Virtual humans with EACS tend to go around dense groups as it has been demonstrated by the group study). But there has not been enough time to perform a precise and complete evaluation or calibration using the results from our studies. Evaluation is difficult in simple situations and becomes harder in complex situations as the number of involved factors increases. Finally, time is limited during a PhD and there are still many factors and processes to consider to handle all the various complex situations.

2 Future Work

The next step of our work on the following behaviors would be to implement the same kind of adaptability to other interactions like collision avoidance. While some of the used factors for the following behaviors might be very specific to following, the other factors could probably be generalized. For example, checking the stability of another motion to follow from further back if they have a random motion (drunk people) could easily be generalized by trying to stay further from these people for any interactions (when avoiding him, talking to him, ...). Wolinsky et al. did a similar work on collision avoidance considering some common and different factors.

About the combination of interactions and the EACS algorithm, there are many perspectives for future research work. While we do not consider collision avoidance independently for each obstacles (When avoiding an obstacle causes another collision, strategies are made to either widen the avoidance to avoid both at once or plan 2 adaptations to avoid one obstacle after the other), obstacles are still perceived independently and treated one by one. The study about individuals avoiding groups clearly demonstrates that some groups can be perceived and interacted with as one big entity instead of several moving obstacles. This perception modification has been done in some RVO2 extensions already. Such modifications could be added to the EACS algorithm too, which could improve results and would greatly reduce the computation complexity.

The quality of the planning done by EACS greatly depends on the prediction of others' motion. When an agent plans its way through a stable crowd, results are far better than when the agent is facing others that keep changing their motion. Indeed, prediction is done with a constant current speed and the planning quickly becomes erroneous when the motions prediction of others' motion is wrong. Wolinsky et al. proposed a better prediction system based on probabilities. Using such prediction system could greatly improve EACS results dealing with more complex motions. This would require the planning process to be adapted to handle probabilities and might increase the number of considered plans.

Finally, an evaluation and a calibration, with real data to compare to, are also required for the EACS algorithm. This algorithm has been done to compute more efficient strategies when navigating through many obstacles. Data about such complex situations are hard to come by, as lab experiments are difficult to do on such complex situations. But we have demonstrated the power of VR to study such complex situations, thus the VR platform could be used to evaluate the EACS algorithm. For example, we could study humans' strategies when facing a crowd with more or less obvious density gaps to exploit. This study would tell us how much humans really exploit these gaps and would give us some data to calibrate EACS's parameters.

In a more general way, the VR platform is a great tool to study interactions in complex situations. Many more studies are needed, complex situations have many factors that can influence the interactions and that still need to be studied. For example, when choosing between a direct path to the goal and a detour, what is the influence of the crowd density in the direct path compare to the length of the empty detour? Or during a collision avoidance, is the perform adaptation influence by secondary targets to insure that the avoidance will not cause another collision with one of the secondary targets? Moreover, a bigger reference database is required to be able to evaluate crowd simulation algorithms on different complex situations. VR offers a quick and easy way to perform these studies and build this reference database compared to real experiment.

However, there is some biases introduce by VR that should not be overlook. First, when using a joystick as a locomotion interface, there are some quantitative differences with real trajectories. Thus, it should be more relevant to design experiments that do not require a precise quantitative analysis of trajectories in such experimental conditions. For example, it can allow studying human decision during locomotion, such as choosing to go toward less dense area or to stop to let the people pass first. Of course, other interfaces that are closer to real walking, such as omnidirectional treadmill or even real walking itself, might produce trajectories close enough to real ones to study the trajectories themselves, but such interfaces still need to be evaluated. These interfaces have not been used during this PhD because the omnidirectional treadmill is not easy to come by and real walking requires a lot of space even with redirected walking techniques. Secondly, even when the trajectories are set aside from the analysis, some biases remains (such as the distance perception). The evaluations of the VR platform have shown these biases to be small enough to produce results roughly similar to real experiments. Nevertheless, these biases have to be taken into account during the design of the experiment as well as the data analysis. For example, the studied situations should not be too ambiguous as the perception biases would most probably reinforce the ambiguity: in a situation with no collision but a crossing distance very close to the collision threshold, it could be perceived as a non-collision situation in real life and as a collision situation in VR. Also during the analysis, the fact that personal space in VR tend to be a bit bigger is important as it means that people will probably require a bit more space to navigate through a crowd than in real life.

In the end, real experiments are still necessary to understand the biases from VR, the same way that real world observations are necessary to understand the biases from experiments in laboratory. But as long as the biases from VR are well understood and taken into account, results from VR experiments will greatly improve our knowledge of

human interactions in complex situations. Finally, other new technologies can also be used to increase the gathered information during experiment. This is what has been started with the gaze experiments and shall be continued with other experiments.

3 Summary

This PhD has been focused on human interactions in complex situations whether by studying such interactions or by proposing algorithms to simulate them. While also exploring how to adapt an interaction to the current situation and target, in order to simulate a wider range of situations, the focus has been on the combination of interactions. An algorithm has been proposed to improve the way collision avoidance are combined which improved results in complex situations. Meanwhile, experiments have been performed, highlighting the importance of the combination of interactions, and have started to give us some insight on how humans perform such combination (for example, by merging several collision avoidances into a big one). VR has proven to be a very useful tool for such experiments and should be used more and more to complete regular observations and improve our knowledge on human interactions.

Résumé en Français

8

Cette thèse a été effectuée sur la simulation de foule. La simulation de foule consiste à créer des foules virtuelles qui agissent visuellement ou quantitativement comme une vraie (Par exemple, les foules présentes dans les jeux vidéo). Des algorithmes contrôlant les foules virtuelles sont conçus pour reproduire le comportement des foules réelles. Il y a deux familles d'algorithmes pour la simulation de foule. Les algorithmes macroscopiques qui considèrent la foule comme une unique entité et reproduisent son comportement en la modélisant comme un fluide, qui se répand dans un bâtiment ou une rue, par exemple. A l'inverse, les algorithmes microscopiques considèrent chaque individu formant la foule et s'appliquent à reproduire leurs comportements et leurs interactions. Ainsi, en simulant un grand nombre de ces individus interagissant entre eux, ils forment une foule virtuelle dont le comportement dépend de celui des individus, de la même manière que le comportement d'une vraie foule dépend des êtres humains la formant. Afin de reproduire les comportements des individus, on a besoin d'une bonne compréhension des comportements et des interactions entre les êtres humains : Comment est-ce que la position et le mouvement, des personnes autour de nous, influencent notre propre mouvement ? Le travail effectué durant cette thèse c'est focalisé sur l'étude et la modélisation de ces interactions afin d'améliorer le réalisme des simulations microscopiques.

Les simulations de foule réaliste sont utilisées pour de nombreuses applications allant du divertissement, avec le cinéma et les jeux vidéo, à l'architecture, pour vérifier la conception des lieux publics par exemple. Ces applications ont grandement participé à l'intérêt grandissant porté à la simulation de foule ainsi qu'à son évolution. Mais chaque application possède ses propres contraintes.

Grâce aux constantes avancées technologiques, les films et jeux vidéo offrent des environnements et des scènes de plus en plus épiques qui ont besoin d'être peuplés avec des foules de plus en plus grandes. Les jeux ont besoin de peupler leurs larges environnements avec des humains virtuels qui doivent agir comme des vrais et interagir avec le joueur en temps réel. Les films ont aussi besoin d'êtres humains virtuels pour créer des foules immenses telles que des armées, des zombies ou toute une population fuyant un désastre. Contrairement aux jeux vidéo, les films n'ont pas besoin d'êtres humains virtuels interagissant en temps réel avec les acteurs. Par contre, ils ont besoin d'un niveau de réalisme visuel accru puisque les êtres humains virtuels seront montrés à côté de véritables êtres humains ce qui accentuera les défauts de la simulation.

Bien qu'une simulation visuellement réaliste soit suffisante pour l'industrie du spectacle, un niveau bien plus élevé de réalisme est nécessaire pour les applications liées à l'architecture. La simulation de foule est utilisée pour évaluer la circulation des personnes et tester des modifications lors de la conception d'un bâtiment ou la modification d'un

lieu public tel qu'une gare ou une station de métro. La simulation est aussi utilisée pour simuler des évacuations en situation d'urgence afin de mieux se préparer à ce genre d'éventualité et d'évaluer les règles de sécurité. Pour ce genre d'utilisation, le simulateur de foule doit produire un résultat aussi proche de la réalité que possible sinon les conclusions tirées de telles simulations risqueraient d'être contreproductif.

Les comportements pouvant être présents dans une foule sont extrêmement variés et dépendent de la situation ainsi que du type de foule. Cela peut aller de la marche en formation, chez les militaires, au lèche-vitrines, dans les centres commerciaux. Rien que dans une même foule, de nombreux comportements différents peuvent coexister. Le plus important étant l'évitement de collision qui est présent dans tous les types de foule. L'évitement de collision peut être perçu comme une tâche triviale que l'on effectue tous les jours sans même y penser. Mais reproduire ce comportement est extrêmement complexe. Il est relativement facile de résoudre le problème de l'évitement d'un obstacle statique ou mobile, mais la complexité du problème augmente avec le nombre d'obstacles à éviter. La complexité augmente encore plus quand ces obstacles sont d'autres êtres humains qui essaient, eux aussi, d'éviter les collisions. D'ailleurs, il arrive dans le monde réel que deux personnes essaient de s'éviter dans la même direction ce qui, en général, conduit à une situation où les deux personnes ne savent plus trop comment éviter l'autre. Mais ce genre d'évènements reste rare et la plus part du temps les gens sont bien synchronisés lors de leurs évitements de collision. Malgré cela, trouver une solution pour résoudre le problème d'évitement de collision n'est pas la partie la plus compliquée. Ce qui est le plus compliqué, c'est de s'assurer que la solution utilisée soit proche de celle utilisée par les hommes eux-mêmes. En effet, même dans les situations simples où des solutions pour éviter les collisions sont facilement calculables, il n'existe qu'un ensemble extrêmement réduit de solutions permettant de reproduire l'évitement de collision d'un véritable être humain. Ensemble qui est difficile à définir dû à la nature chaotique de l'être humain qui ne réagira pas exactement de la même manière même s'il est confronté deux fois à la même situation.

De nombreuses expériences ont été effectuées pour étudier comment un être humain évite les obstacles statiques et les autres êtres humains. Ces études nous permettent d'améliorer notre compréhension de ce phénomène. Grâce à cela, de nombreux algorithmes d'évitement ont vu le jour et d'autres ont pu être améliorés pour simuler des foules de manière plus réaliste. Cela a aussi permis d'évaluer les algorithmes existants pour mesurer à quel point la simulation est proche de la situation réelle avec de véritables êtres humains.

1 Problème

Grâce aux nombreuses expériences, sur l'évitement de collision, effectuées dans des situations élémentaires, les algorithmes actuels sont capables de simuler ces situations de manière satisfaisante. Mais lorsque la complexité des situations simulées augmente, des difficultés à reproduire des comportements réalistes apparaissent. Par exemple, les algorithmes actuels ont beaucoup de difficulté à atteindre le même niveau de fluidité que dans la vraie vie lorsqu'ils simulent des foules à très haute densité. Ils peinent aussi à reproduire la manière dont les êtres humains se faufilent entre les gens pour traverser

une foule. Ces difficultés semblent démontrer qu'il existe une différence non négligeable entre l'évitement d'une personne et l'évitement de nombreuses personnes. Pour comprendre comment les êtres humains combinent ces évitements multiples et ce qui manque aux algorithmes actuels, il est nécessaire d'effectuer des études sur ces situations plus complexes.

L'absence de données, sur ces interactions multiples, est principalement due à la difficulté de mener des expériences sur ce genre de situations. En effet, rien que sur des situations plus simple tel qu'une interaction entre seulement deux personnes, il est difficile de standardiser les conditions. Ce problème peut être illustré avec l'évitement de collision entre deux personnes qui se croisent perpendiculairement. Une des difficultés, dans ce cas-là, est de forcer le risque de collision entre les deux individus pour pouvoir étudier le comportement d'évitement. Des mesures peuvent être prises pour augmenter les chances d'obtenir un réel risque de collision comme, par exemple, s'assurer que les deux individus soient positionnés à la même distance du point de croisement et partent en même temps. Malgré tout, même si un réel risque de collision est présent au début, à cause des variations chez les êtres humains (différentes vitesses de marche, différents temps de réaction au départ, ...) qui peuvent modifier les conditions de départ, il reste difficile de traiter des conditions standardisées. Il est alors nécessaire d'effectuer cette tâche d'évitement un grand nombre de fois afin de s'assurer que l'on ait assez de données sur de véritable évitement de collision pour produire une analyses statistique robuste de l'interaction. Ajouter de la complexité à la situation ne fait qu'accroître ce problème et nécessiteras encore plus d'itérations, itérations étant déjà plus compliquées à mettre en place.

2 Approche

Pour améliorer les algorithmes afin de gérer correctement les situations plus complexes, deux approches furent employées pendant cette thèse. La première a consisté à considérer plus de facteurs en compte, tel que des facteurs psychologiques ou sociologiques. Par exemple, respecter une distance socialement acceptable avec les autres (plus éloigné avec des étrangers, plus proche avec des amis), garder de plus grande distance avec ceux qui ont une marche instable par sécurité (une personne saoule), etc. Cela permet aux algorithmes de s'adapter à un plus grand nombre de situations différentes. Une deuxième approche fut explorée : garder les algorithmes d'évitement de collision tel quels, puisqu'ils permettent d'obtenir de bon résultats dans des situations simples, et ajouter un système par-dessus pour réduire la complexité des situations qui sont difficilement traitées par les algorithmes actuels. Dans le cas d'une foule assez dense à traverser, au lieu d'avoir l'algorithme d'évitement de collision qui essaie de trouver une solution en considérant tous les individus autour, le nouveau système analyse la situation pour déterminer un chemin au travers de la foule et limite l'algorithme d'évitement de collision à une petite sélection d'individus seulement (un ou deux).

A propos de la complexité d'effectuer des études sur la navigation et les interactions des êtres humains, décrite dans la section précédente, on propose d'utiliser la Réalité Virtuelle (VR) pour faciliter le processus. La VR possède de nombreux avantages pour mener à bien ce type d'expériences, notamment on a le control parfait de l'environnement virtuel

dans lequel le participant évolue. Ce qui nous donne le control sur la situation que le participant rencontre, par exemple on peut facilement forcer un risque de collision avec un humain virtuel. De plus, on peut mettre en place une situation complexe, la répéter en ne changeant qu'un paramètre de la situation et représenter exactement les mêmes situations à d'autres participants. Ainsi, effectuer une expérience en VR est bien plus simple et rapide que dans le monde réel. En effet, pour mener une expérience en VR on a besoin de la plateforme de VR (CAVE, HMD, ...) et d'un participant à la fois. Pour chaque situation à tester, il suffit d'exécuter la simulation correspondante dans le système de VR. Alors que pour une expérience hors VR il faudrait rassembler beaucoup de personnes à la fois ce qui demande beaucoup de temps et d'efforts à planifier. De plus, pour chaque situation étudiée il faudrait placer les participants correctement et leur fournir les bonnes instructions. Grâce à ces avantages, la VR permet aux chercheurs d'effectuer des expériences dans les domaines du sport, du control du mouvement, de la perception, de la représentation dans l'espace et, dans notre cas, de la simulation de foule.

3 Contributions

Les contributions apportées par cette thèse peuvent être séparées en deux catégories : les algorithmes de simulation de foule et les études sur la navigation des êtres humains.

Au niveau des algorithmes de simulation, la première contribution s'est limitée au comportement de suivi, ce qui a permis de travailler sur des simulations simplifiées à une dimension. Un nouvel algorithme fut proposé incluant de nouveaux facteurs, par rapport aux algorithmes précédents, influençant le résultat. Les trajectoires ainsi produites furent comparées à celles des algorithmes précédant ainsi qu'à celles de véritables êtres humains. Cela a permis de démontrer la robustesse de l'algorithme proposé : des résultats au moins aussi bien que les précédents algorithmes et meilleurs dans certains cas qui ne sont pas bien gérés par les algorithmes précédents. Cette contribution a été présentée à la conférence MIG 2014. La deuxième contribution a été de proposer une étape supplémentaire dans le workflow habituel des simulateurs de foule afin de trouver des chemins plus naturels au travers d'une foule. Une implémentation fut effectuée incluant cette étape supplémentaire et les résultats ont démontré une amélioration dans la gestion de multiples évitements de collision. Ce travail a été présenté à la conférence SCA 2015 et a été sélectionné pour une extension dans le journal CGF. Enfin, une collaboration fut effectuée avec Zhiguo Ren sur son travail à propos de la simulation du comportement de groupe. Il a proposé une extension à RVO2 ajoutant une nouvelle contraintes forçant les agents d'un même groupe à rester près les uns des autres. L'assistance apportée, dans le cadre de cette thèse, fut principalement concentrée au niveau de l'analyses des résultats afin de démontrer que l'extension était capable de simuler une grande variété de groupes tels qu'un petit groupe marchant en formation ou un énorme groupe de touristes qui peut facilement se séparer et se rassembler.

En parallèle des travaux sur les algorithmes de simulation, des études sur la navigation humaine furent menées. La première contribution a été de collaborer avec Anne-Hélène Olivier sur ces travaux de validation de la plateforme VR pour étudier le comportement humain. Pour cette validation, une expérience menée dans le monde réel fut reproduit en VR afin de comparer les résultats entre les deux. Cette étude a démontré que, malgré

quelques différences quantitatives, les résultats obtenus en VR et dans le monde réel sont qualitativement similaires. Il est donc possible d'utiliser la VR pour étudier la navigation humaine. Cette conclusion a conduit à la seconde contribution : l'utilisation de la VR pour étudier comment un individu évite un groupe de personnes. Grâce à la VR, les paramètres du groupe ont pu être parfaitement contrôlés pendant toute la durée de l'expérience (notamment la taille du groupe et la distance entre ces membres). Ainsi, on a pu déterminer un seuil de distance, entre les membres du groupe, à partir duquel la stratégie d'évitement passe de contourner le groupe à traverser le groupe. Cette étude fut présentée à la conférence IEE VR 2015. Pour la troisième contribution, la VR fut combinée avec un système de suivi du regard afin de mieux comprendre la boucle perception/action lorsque qu'un individu interagi avec d'autres personnes. En particulier, on s'est intéressé aux particularités du marcheur regardé par un humain pendant qu'il évolue dans une foule. Cette expérience a été effectuée à la fin de la thèse et les données sont toujours en cours d'analyses mais les premiers résultats sont encourageants.

Appendices

Collaborations

A

1 Introduction

During this PhD, there have been some opportunities to work with other people and to contribute to their work. When the PhD started, Anne-Hélène Olivier was working on the development of a VR platform to perform experiments on human interactions. An evaluation of the platform, with moving virtual humans, had just begun and the opportunity arose to help on this evaluation. This has been a good opportunity to get familiarize with the platform itself, to learn about the advantages and weaknesses of the platform and what kind of experiment would be the best suited for the platform. Thanks to this evaluation, we were able to use the platform to study human interactions. Especially, it led to the study on individual avoidance of group, presented in chapter 5.

Another collaboration happened near the end of the PhD. This collaboration has been done with Zhiguo Ren on its work about group modeling. Both the experiences from the study on individual avoidance of group and from combining multiple collision avoidances with the EACS algorithm (chapter 4) proved to be useful during this collaboration. Moreover, this was a good way to work on a different kind of interaction combination: combining two very different interactions (staying close to group members while avoiding collision).

An overview of both these contributions is presented in this appendix. The evaluation of the VR platform is presented first in section 2. It is followed by the presentation of the work on group modeling in section 3.

2 Walking with virtual people: Evaluation of locomotion interfaces in dynamic environments

2.1 Introduction

VR can be a great tool to perform experiment on human interactions during locomotion tasks. But when doing an experiment, there might be some intrinsic biases caused by VR. Before being able to use it to study human behaviors in complex navigation tasks, we need to ensure that the gathered data are exploitable. Some work, presented in chapter 2 section 3, has already started to evaluate VR platforms to insure that humans act in VR as they would in real life. The work that is presented in this chapter aims to continue this evaluation. Its main purpose is to investigate whether small scale interactions between a real user and a virtual human, during a locomotion task, are preserved in VR compared to interactions in real conditions.

A lot of research has been performed in Biomechanics to describe the collision avoidance interaction, which can be used as a reference for evaluation purposes. More specifically, it has been shown that humans accurately perceive others' motion and risk of collision, allowing them to perform anticipated adaptations of their trajectories. As a result, realistic interactions, between a real human and a virtual human, first require that the relative motion with the virtual human is correctly perceived, and the risk of future collision detected. Second, it is required that the employed locomotion interface does not induce biases in the avoidance trajectories.

The contributions of this work are as follows: 1) we experimentally evaluate the ability of a real user to accurately perceive the risk of collision with a moving character. We show that the level of perception is accurate enough to expect realistic interactions 2) we experimentally compare various basic locomotion techniques and show how they influence the virtually formed locomotion trajectories. 3) we provide recommendations for the design of VR platforms to perform human locomotion studies in the context of interactions.

The collaboration has been mainly focused on helping to prepare and perform the experiments. Two experiments were performed, the first one evaluating perceptual biases and the second one evaluating biases from different locomotion interfaces. These two experiments are presented in this section. Then a quick overview of the results is given. For more details about the experiments and the results, refer to the paper [OBKPss].

2.2 Experiments

The experiments were performed in a 4-screen CAVE. For both experiments, participants were immersed in the same virtual environment, presented in Figure A.1, facing a situation similar to the experiment done in the real world from [OMCP12]. In this situation, the participant moves toward a goal while a virtual human crosses his path perpendicularly from the left or the right. This environment is composed of four walls used to control the time at which the participant sees the virtual human and starts interacting with it (see Figure A.1).

For the first experiment, the participant passively moves in the virtual environment along a straight trajectory. The future distance of closest approach (e.g., mpd) with the virtual human is accurately controlled by delaying or advancing the character on its trajectory at initialization. The participant is only able to perceive the beginning of the virtual human's motion. From the time he can see the virtual human (after they are no longer hidden from each other by the wall), we consider 3 cutoff times (presented in Figure A.1): 0.5s, 1s and 2s. After the cutoff time, visual information is stopped by displaying a black screen. Thus the participant has a limited time to assess the situation. Note that cutoff times were chosen based on a real experiment which showed that adaptations may start 0.5s after the beginning of the interaction (i.e., when participants are able to see each other) [OMCP12]. To verify if the collision risk was correctly perceived, the participant was then asked about it: "Will you collide with the virtual character?" "Will you pass first or give way?". For each of these questions, they had to rate their level of confidence on a 7 point Likert scale (1: not at all confident to 7: very confident). They were trained to the task 6 times before recording their answers.

For the second experiment, the participant is able to adapt his own motion to avoid collision with the virtual human when he deems it necessary. The performed adaptation




Figure A.1 Virtual environment and experimental situation used in this study. The virtual human has several starting positions corresponding to different value of mpd giving the virtual human some advance or delay over the user. Position of both the user and the virtual human are represented for different times: when they can see each other t_{see} and then $t_{0.5}$, t_1 and t_2 which are respectively $0.5s$, $1s$ and $2s$ after t_{see} (if there is no adaptation from the user). A dotted line joins the positions of the character and of the participant at those times for $mpd = 0$.

is evaluated, compared to the real one perform in the real world. Several locomotion interfaces were tested using either a joystick or a whole-body locomotion metaphor. In Figure A.2, participants are shown using different locomotion interfaces to interact with the virtual human.




Figure A.2 Participants navigating inside the virtual environment using different locomotion interfaces

2.3 Results

Results from the first experiment showed that on average participants correctly estimated the condition of interaction with the virtual human, i.e., the risk of collision as well

as the crossing order. Participants were able to identify the order of passage very early, there was no effect of cutoff time. Also, the level of confidence was correlated with answers accuracy, both for order and risk of collision. Participants were conscious about their level of accuracy which is an important point. Indeed, their reaction will probably be delayed if they do not perceive the situation in a clear enough way, just as in reality. Collision estimation was not symmetrical between passing first or behind: many estimated a collision when passing first without collision but close to the virtual human and estimated no collision when colliding slightly with the virtual human passing behind. These observations have been interpreted as a shift in space made by the participant due to a gap in the perception of their envelope in the virtual world, as if they feel like they are ahead of their actual virtual position.

The study of the trajectories, produced during the second experiment, reveals some quantitative differences compared to the one from the real world. In VR participants adapt their locomotion for lower risk of future collision (when the absence of adaptation would lead to a collision free motion), and perform larger adaptations which result into larger passage distance. We interpret that this difference in the adaptation threshold could be a consequence of the compression of distance perception in VR. Despite these quantitative differences, both set of trajectories were qualitatively similar, sharing the same structure: a small period of collision estimation followed by an adaptation reducing the risk of collision and finishing with a regulation phase recovering a default speed toward the goal before crossing the virtual human.

2.4 Conclusion

In this work, we studied and discussed the use of Virtual Reality to observe human interactions during a locomotion task. Results similar to previous study done with static obstacles were obtained: trajectories done in VR show some quantitative differences with the ones done in the real world but are qualitatively very similar. This offers opportunities for qualitative analysis of such human behaviors and opens several perspectives (for example, the study presented in chapter 5).

3 Group Modeling: a Unified Velocity-based Approach

3.1 Introduction

As presented in chapter 2 section 1.2.1, some work has already been done to include group behaviors in crowd simulation. But groups in crowds may exist under many forms. Groups are of different sizes: a family of four, ten friends, fifty tourists, one thousand in a march, etc. Groups are structured by different internal relations: followers with a leader, tourists made of families with a guide, schoolmates with a set of teachers, etc. Groups are guided by different goals and constraints: armies have specific formations, parents keep contact with their kids, friends chat together, etc. It is still a challenge to model groups in all these dimensions. Previous solutions do not generalize to a large variety of groups. The goal of this work is to offer a solution that is able to simulate these different kinds of group. We implement this solution based on the velocity-obstacle formalism. Agents

select their speed according to their goal with the regular constraint of not colliding with other neighbor agents and the extra constraint of staying at a bounded distance to some other agents belonging to the same group.

This collaboration has been mainly focused on the results part. In the results, we had to demonstrate that the proposed solution was able to simulate the different kinds of group and explain how each kind of group could be simulated. For that, an extended study of the different parameters of the algorithm and their influence over the groups in the simulation had to be done. This study is presented in this section, with some example of the kinds of group that can be simulated with this solution. But first, an overview of the algorithm itself is given to understand the role of some important parameters. For a more details explanation see the article [RCB⁺ss].

3.2 Velocity-based group simulation algorithm




Figure A.3 (a) The velocity obstacle $VO_{A|B}^\tau$ of an agent B to agent A is a truncated cone in the velocity space. (b) The velocity connection $VC_{A|B}^\tau$ of an agent B to agent A contains the velocities that result in grouping A and B together for at least τ_{min} seconds before time τ . Here $\tau_{min} = 1$. (c) $VC_{A|B}^\tau$ for a rectangular constraint. (d) A situation where not all velocity constraints can be satisfied since $VC_{A|B}^\tau \cap VC_{A|C}^\tau = \emptyset$ (group neighbors B and C are not close to each other). In this case, agent A tries to stay close to both B and C.

The RVO2 algorithm computes a set of velocity that will lead to a collision with an obstacle. This is done by computing all the velocities that will lead an agent under a specific distance of the obstacle, as presented in Figure A.3.a or explained in more details in chapter 2 section 1.1.1. For collision avoidance, the distance is set as the minimum distance before collision. For grouping, the computation is exactly the same as group members try to stay under a specific distance from each other. The computed velocity set can be adapted to reach the distance under a specific amount of time or to adapt the distance to the position (to be inside a rectangle center on the targeted group member) as shown in Figure A.3 (b and c). Then by combining the constraints: the collision avoidance, the grouping and the goal reaching, we can select a velocity that will be collision free while remaining close to some of the group members and in the direction of the goal (Figure A.3.d). As perfect solutions rarely exist, a balance between the three constraints is needed and preferences are to be set between avoiding collisions, staying with influencing agents and following the desired velocity. This is done via three weights: w_{vo} for collision avoidance, w_g for grouping and w_v for goal reaching.

Groups are created through a relation matrix. The relation matrix defines the desire for an agent to stay close to another one. This relation is set between 0 (no desire to stay

close to that agent) and 1 (very high desire to stay close to that agent). This relation matrix is set by users before the start of the simulation. Conceptually, the relation matrix is a weighted directed graph describing logical (often asymmetric) relationships between agents. This allows the construction of different kinds of group. During the simulation, agents will try to stay close to a limited number of agents (the number of connection: n^c), these agents are selected according to a weighted distance $d_{ij}^{rel} = d_{ij}^{euc} f(e_{ij})$. (d_{ij}^{euc} the euclidean distance between agent i and j , $f(e_{ij})$ the relation influence between i and j)

3.3 Parameter analysis

There are three parameters that have a huge impact on the simulated groups: the weight w_g for grouping, the weight w_v for goal reaching (the weight for collision avoidance is fixed to a value far above the two others to insure there is no collision) and the number of connection n^c . The number of connection has a direct effect on the size of the cluster of agent form by the group and thus on its cohesion. With high values of n^c , the group will be divided into a few very big clusters, if n^c is sufficiently high the whole group will form one big cluster and remain as such. With low values of n^c , the group will be broken into many small clusters. This phenomenon can be seen in Figure A.4, simulations, with one big group (in blue) going through a crowd of individuals (in green), are shown for several values of n^c . With $n^c = 1$ the groups split into many small subgroup while with $n^c = 45$ the whole group stay together. One side effect of having big clusters of agents moving together is the average speed dropping. All the pictures in Figure A.4 are taken after the same amount of time in the simulation. The smaller n^c is the further blue agents are toward their goal (on the right).




Figure A.4 Simulation results of a group (in blue) going through a crowd of individuals (in green) for different value of n^c . The simulations demonstrate that the group is more and more bound together and is less likely to break into small clusters with higher value of n^c , which is quantitatively confirm by the graph on the right.

While n^c has a direct impact on the cohesion of the group and the size of clusters, the balance, between the weight w_g and w_v , impacts the cohesion of the clusters them self (and of the group if there is only one cluster). Some results with different weight are presented in Figure A.5. With $W_g = 0$ and $W_v = 0$, the blue agents are scattered in the environment as they have no desire to stay together or even go toward a same goal. When W_v is set to 0.2, the blue agents start moving toward a same goal but are still scattered

along the way. With both weight set at 0.2, there is no longer blue agents left behind by themselves. Increasing W_g to 0.5, the clusters of blue agents become more dense and tight.




Figure A.5 Simulations of a group (in blue) going through a crowd of individuals (in green) for different value of preference weights. The more the balance is in favor of grouping the tighter the group's clusters are. Moreover, with a balance in favor of grouping, no blue agent end up alone.

Finally, the relation matrix has also a huge impact on the resulting group. But this impact is more straight forward:

- if one wants a regular group, every agent will have the same relation with each other,
- if one wants a leader, one defines an agent with whom all the other agents will have a greater relation with,
- if one wants a group composed of subgroup, one increases the relation between the members of the subgroup and keep a lower relation with the other members of the group,
- ...

A group with one leader has been simulated to check how many agents one leader could lead toward the goal. The results are shown in Figure A.6, depending on the cohesion of the group (set with the parameter n^c), the leader is able to move more or less agents toward the goal. If the group is too cohesive, one leader is unable to move anyone as all the blue agent are glued together (see simulation with $n^c = 45$). If the group is not cohesive enough, the leader is able to grab some of the blue agents but many will stay behind (see simulation with $n^c = 1$). For the leader to be able to move all the agents of the group toward the goal, the group has to be at a specific level of cohesion which is $n^c = 4$ in the simulated situation.




Figure A.6 Simulation of a leader (in red) trying to move a whole group (in blue) toward its goal. Different level of n^c were tested showing that an average level of cohesion is require, from the group, to move a whole group. The graph shows the distribution of the agents through the distance from the leader at the end of the simulation.

3.4 Results

With the proposed algorithm and the knowledge of the parameters effect on the simulation, many kinds of grouping behaviors can be simulated. Playing with the level of cohesion of a group, it can be made to be more prone to split and merge afterward when avoiding or to stay together and bypass the obstacle, this is demonstrated in Figure A.7 on the left side. On the right side of the Figure, some group formations are shown which are set by adding geometrical constraints when computing the velocity connection (see Figure A.3.c). These formations can easily be changed during the simulation which would enables a user or a control algorithm to adapt the formation to the environment.




Figure A.7 Left: a group avoiding a collision with another by either bypassing it or splitting and merging afterward. Right: two formations that is possible to simulate with the proposed algorithm.

Only of few examples of the possible kinds of group is presented in this chapter, many more are presented in the paper [RCB⁺ss] and in the companion video. The Figure A.8 gives a quick glance at a more general scenario showing many different kinds of group.



Figure A.8 A general scenario that showcase a wide variety of groups, all of which were simulated by the proposed algorithm

3.5 Conclusion

In this collaboration, a microscopic algorithm to simulate groups has been presented. The algorithm is an extension of the RVO2 algorithm for collision avoidance and combine both interactions (collision avoidance and grouping) directly when selecting a current speed for the agents. Its ability to handle a wide variety of groups have been demonstrated and is done by tuning the parameters. Some are quite intuitive: the relation matrix or the geometrical constraints to the velocity connection. Others have been analyzed to understand their effects on the simulation: the number of connection or the preference weights.

Bibliography

- [ADKB12] Nicolas E Andersen, Louisa Dahmani, Kyoko Konishi, and Véronique D Bohbot. Eye tracking, strategies, and sex differences in virtual navigation. *Neurobiology of learning and memory*, 97(1):81–89, 2012. [84](#)
- [AGS⁺11] Junghyun Ahn, Stéphane Gobron, Quentin Silvestre, Horesh Ben Shitrit, Mirko Raca, Julien Pettré, Daniel Thalmann, Pascal Fua, and Ronan Boulic. *Long Term Real Trajectory Reuse through Region Goal Satisfaction*, pages 412–423. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011. [17](#)
- [AMBR⁺12] Javier Alonso-Mora, Andreas Breitenmoser, Martin Ruffli, Roland Siegwart, and Paul Beardsley. Image and animation display with multiple mobile robots. *The International Journal of Robotics Research*, 31(6):753–773, 2012. [15](#)
- [ANH10] Richard Andersson, Marcus Nyström, and Kenneth Holmqvist. Sampling frequency and eye-tracking measures: how speed affects durations, latencies, and more. *Journal of Eye Movement Research*, 3(3):6, 2010. [84](#)
- [AWTB12] Junghyun Ahn, Nan Wang, Daniel Thalmann, and Ronan Boulic. Within-crowd immersive evaluation of collision avoidance behaviors. In *Proceedings of the 11th ACM SIGGRAPH International Conference on Virtual-Reality Continuum and Its Applications in Industry*, VRCAI '12, pages 231–238, New York, NY, USA, 2012. ACM. [21](#), [26](#)
- [BBQ⁺07] J. L. Berrou, J. Beecham, P. Quaglia, M. A. Kagarlis, and A. Gerodimos. *Calibration and validation of the Legion simulation model using empirical data*, pages 167–181. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007. [17](#), [22](#)
- [BDP14] Julien Bruneau, Teófilo B. Dutra, and Julien Pettré. Following behaviors: A model for computing following distances based on prediction. In *Proceedings of the Seventh International Conference on Motion in Games*, MIG '14, pages 17–24, New York, NY, USA, 2014. ACM. [5](#)
- [Bea04] Catherine Beaulieu. Intercultural study of personal space: A case study. *Journal of Applied Social Psychology*, 34(4):794–805, 2004. [76](#)
- [Ber97] Alain Berthoz. *Sens du mouvement (Le)*. Odile Jacob, 1997. [84](#)

- [BJ03] D.C. Brogan and N.L. Johnson. Realistic human walking paths. In *Computer Animation and Social Agents, 2003. 16th International Conference on*, pages 94–101, 2003. [viii](#), [18](#), [19](#), [22](#), [23](#)
- [BK15] D. Bršćić and T. Kanda. Changes in usage of an indoor public space: Analysis of one year of person tracking. *IEEE Transactions on Human-Machine Systems*, 45(2):228–237, April 2015. [viii](#), [16](#), [17](#)
- [BKV⁺10] B. Bideau, R. Kulpa, N. Vignais, S. Brault, F. Multon, and C. Craig. Using virtual reality to analyze sports performance. *Computer Graphics and Applications, IEEE*, 30(2):14–21, 2010. [viii](#), [5](#), [25](#)
- [BMdOB03] Adriana Braun, Soraia Raupp Musse, Luiz Paulo Luna de Oliveira, and Bardo EJ Bodmann. Modeling individual behaviors in crowd simulation. In *Computer Animation and Social Agents, 2003. 16th International Conference on*, pages 143–148. IEEE, 2003. [14](#)
- [BNCM14] A. Best, S. Narang, S. Curtis, and D. Manocha. Densesense: Interactive crowd simulation using density-dependent filters. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA '14, pages 97–102, Aire-la-Ville, Switzerland, Switzerland, 2014. Eurographics Association. [12](#)
- [BOP15] Julien Bruneau, Anne-Hélène Olivier, and Julien Pettré. Going through, going around: A study on individual avoidance of groups. *IEEE Transactions on Visualization and Computer Graphics*, 21(4):520–528, April 2015. [6](#)
- [BP15] Julien Bruneau and Julien Pettré. Energy-efficient mid-term strategies for collision avoidance in crowd simulation. In *Proceedings of the 14th ACM SIGGRAPH / Eurographics Symposium on Computer Animation*, SCA '15, pages 119–127, New York, NY, USA, 2015. ACM. [5](#)
- [BPss] Julien Bruneau and Julien Pettré. Eacs: Effective avoidance combination strategy. *Comp. Graph. Forum*, in press. [5](#)
- [BRCW12] Stéphane Bonneaud, Kevin Rio, Pierre Chevaillier, and William H. Warren. *Accounting for Patterns of Collective Behavior in Crowd Locomotor Dynamics for Realistic Simulations*, pages 1–11. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012. [19](#)
- [BSD⁺05] T Banton, J Stefanucci, F Durgin, A Fass, and D Proffitt. The perception of walking speed in a virtual environment. *Presence*, 14(4):394–406, Aug 2005. [26](#)
- [CARH13] Julien Cividini, Cecile Appert-Rolland, and Hendrik-Jan Hilhorst. Diagonal patterns and chevron effect in intersecting traffic flows. *EPL (Europhysics Letters)*, 102(2):20002, 2013. [22](#)

- [CBW06] Jonathan A Cohen, Hugo Bruggeman, and William H Warren. Combining moving targets and moving obstacles in a locomotion model. *Journal of Vision*, 6(6):135–135, 2006. [10](#)
- [CL88] V Cavallo and M Laurent. Visual information and skill level in time-to-collision estimation. *Perception*, 17(5):623–632, 1988. [84](#)
- [CLGX12] Ke Chen, Chen Change Loy, Shaogang Gong, and Tao Xiang. Feature mining for localised crowd counting. In *In British Machine Vision Conference*, 2012. [viii](#), [16](#), [17](#)
- [COMP13] G. Cirio, A. Olivier, M. Marchal, and J. Pettre. Kinematic evaluation of virtual walking trajectories. *Visualization and Computer Graphics, IEEE Transactions on*, 19(4):671–680, 2013. [viii](#), [26](#), [77](#)
- [Cos10] Marco Costa. Interpersonal distances in group walking. *Journal of Non-verbal Behavior*, 34(1):15–26, 2010. [17](#), [76](#)
- [CPA09] Michael E Cinelli, Aftab E Patla, and Fran Allard. Behaviour and gaze analyses during a goal-directed locomotor task. *The Quarterly Journal of Experimental Psychology*, 62(3):483–499, 2009. [86](#)
- [CSC09] Ujjal Chattaraj, Armin Seyfried, and Partha Chakroborty. Comparison of pedestrian fundamental diagram across cultures. In *Advances in Complex Systems*, pages 393–405, 2009. [20](#)
- [CVB95] J. E. Cutting, P. M. Vishton, and P. A. Braren. How we avoid collisions with stationary and with moving obstacles. *Psychological Review*, 102:627–651, 1995. [84](#)
- [DDH⁺14a] Winnie Daamen, Dorine C. Duives, Serge P. Hoogendoorn, Alessandro Corbetta, Luca Bruno, Adrian Muntean, and Federico Toschi. The netherlands high statistics measurements of pedestrian dynamics. *Transportation Research Procedia*, 2:96–104, 2014. [17](#)
- [DDH⁺14b] Winnie Daamen, Dorine C. Duives, Serge P. Hoogendoorn, Xiao dong Liu, Wei guo Song, and Wei Lv. Empirical data for pedestrian counterflow through bottlenecks in the channel. *Transportation Research Procedia*, 2:34–42, 2014. [viii](#), [19](#), [21](#)
- [DH03] W Daamen and SP Hoogendoorn. Qualitative results from pedestrian laboratory experiments. 2003. [19](#)
- [FFW07] Philip W. Fink, Patrick S. Foo, and William H. Warren. Obstacle avoidance during walking in real and virtual environments. *ACM Trans. Appl. Percept.*, 4(1), January 2007. [5](#), [26](#)
- [FRT09] Brett R Fajen, Michael A Riley, and Michael T Turvey. Information, affordances, and the control of action in sport. *International Journal of Sport Psychology*, 40(1):79, 2009. [84](#)

- [FS98] Paolo Fiorini and Zvi Shillert. Motion planning in dynamic environments using velocity obstacles. *International Journal of Robotics Research*, 17:760–772, 1998. [49](#)
- [FUCH15] S Fotios, J Uttley, C Cheal, and N Hara. Using eye-tracking to identify pedestrians’ critical visual tasks, part 1. dual task approach. *Lighting Research and Technology*, 47(2):133–148, 2015. [86](#)
- [FW03] Brett R. Fajen and William H. Warren. Behavioral dynamics of steering, obstacle avoidance, and route selection. *Journal of Experimental Psychology: Human Perception Performance*, 2003. [18](#), [27](#)
- [FWK11] Tom Foulsham, Esther Walker, and Alan Kingstone. The where, what and when of gaze allocation in the lab and the natural environment. *Vision research*, 51(17):1920–1931, 2011. [xi](#), [85](#)
- [GCC⁺10] Stephen J. Guy, Jatin Chhugani, Sean Curtis, Pradeep Dubey, Ming Lin, and Dinesh Manocha. Pedestrians: A least-effort approach to crowd simulation. In *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA ’10, pages 119–128, Aire-la-Ville, Switzerland, Switzerland, 2010. Eurographics Association. [23](#), [66](#)
- [GCLM12] Stephen J. Guy, Sean Curtis, Ming C. Lin, and Dinesh Manocha. Least-effort trajectories lead to emergent crowd behaviors. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, 85(1), 1 2012. [21](#), [23](#)
- [GD13] Q. Gu and Z. Deng. Generating freestyle group formations in agent-based crowd simulations. *IEEE Computer Graphics and Applications*, 33(1):20–31, Jan 2013. [15](#)
- [Gib86] James J Gibson. *The ecological approach to visual perception*. Hillsdale, NJ: Erlbaum, 1986. [46](#)
- [GKGG14] J. Godoy, I. Karamouzas, S. J. Guy, and M. Gini. Anytime navigation with progressive hindsight optimization. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 730–735, Sept 2014. [13](#)
- [GLM10] Stephen J. Guy, Ming C. Lin, and Dinesh Manocha. Modeling collision avoidance behavior for virtual humans. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 2 - Volume 2*, AAMAS ’10, pages 575–582, Richland, SC, 2010. International Foundation for Autonomous Agents and Multiagent Systems. [52](#)
- [GLRFM08] Martin Gérin-Lajoie, Carol L. Richards, Joyce Fung, and Bradford J. McFadyen. Characteristics of personal space during obstacle circumvention in physical and virtual environments. *Gait & Posture*, 27(2):239–247, 2008. [26](#)

- [GNL13] Abhinav Golas, Rahul Narain, and Ming Lin. Hybrid long-range collision avoidance for crowd simulation. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, I3D '13*, pages 29–36, New York, NY, USA, 2013. ACM. [vii](#), [12](#)
- [GSPR05] Maia Garau, Mel Slater, David-Paul Pertaub, and Sharif Razzaque. The responses of people to virtual humans in an immersive virtual environment. *Presence: Teleoper. Virtual Environ.*, 14(1):104–116, February 2005. [25](#)
- [GvdBL⁺12] Stephen J. Guy, Jur van den Berg, Wenxi Liu, Rynson Lau, Ming C. Lin, and Dinesh Manocha. A statistical similarity measure for aggregate crowd dynamics. *ACM Trans. Graph.*, 31(6):190:1–190:11, November 2012. [23](#)
- [Hal66] E.T. Hall. *The hidden dimension*. Doubleday Anchor Books. Doubleday, 1966. [31](#)
- [Hay83] Leslie A Hayduk. Personal space: Where we now stand. *Psychological bulletin*, 94(2):293, 1983. [31](#), [76](#)
- [HBBK07] John H Hollman, Robert H Brey, Tami J Bang, and Kenton R Kaufman. Does walking in a virtual environment induce unstable gait?: An examination of vertical ground reaction forces. *Gait & Posture*, 26(2):289–294, 2007. [26](#)
- [Her05] Bruno Herbelin. *Virtual reality exposure therapy for social phobia*. PhD thesis, Ecole Polytechnique Federale de Lausanne, 2005. [25](#)
- [HFV00] Dirk Helbing, Illés Farkas, and Tamas Vicsek. Simulating dynamical features of escape panic. *Nature*, 407(6803):487–490, 2000. [21](#), [22](#)
- [HKK15] Motoyasu Honma, Shinichi Koyama, and Mitsuru Kawamura. Hesitant avoidance while walking: an error of social behavior generated by mutual interaction. *Frontiers in psychology*, 6:1013, 2015. [2](#)
- [HM95] D. Helbing and P. Molnar. Social force model for pedestrian dynamics. *Phys. Rev. E*, 51(5):4282–4286, May 1995. [8](#)
- [HMC⁺12] Rafael Hocevar, Fernando Marson, Vinícius Cassol, Henry Braun, Rafael Bidarra, and Soraia R. Musse. From their environment to their behavior: A procedural approach to model groups of virtual agents. In *Proceedings of the 12th International Conference on Intelligent Virtual Agents, IVA'12*, pages 370–376, Berlin, Heidelberg, 2012. Springer-Verlag. [15](#)
- [HNM12] Kenneth Holmqvist, Marcus Nyström, and Fiona Mulvey. Eye tracker data quality: What it is and how to measure it. In *Proceedings of the Symposium on Eye Tracking Research and Applications, ETRA '12*, pages 45–52, New York, NY, USA, 2012. ACM. [84](#)

- [HSK⁺14] Markus Huber, Yi-Huang Su, Melanie Krüger, Katrin Faschian, Stefan Glasauer, and Joachim Hermsdörfer. Adjustments of speed and path when avoiding collisions with another pedestrian. *PloS one*, 9(2):e89589, 2014. [18](#)
- [Hug03] Roger L. Hughes. The flow of human crowds. *Annual Review of Fluid Mechanics*, 35(1):169–182, 2003. [7](#)
- [JacARLP12] Asja Jelić, Cécile Appert-Rolland, Samuel Lemerrier, and Julien Pettré. Properties of pedestrians walking in line: Fundamental diagrams. *Phys. Rev. E*, 85:036111, Mar 2012. [37](#), [38](#)
- [JARLP12] Asja Jelić, Cécile Appert-Rolland, Samuel Lemerrier, and Julien Pettré. Properties of pedestrians walking in line. ii. stepping behavior. *Phys. Rev. E*, 86:046111, Oct 2012. [19](#)
- [JMH09] Jelena Jovancevic-Misic and Mary Hayhoe. Adaptive gaze control in natural environments. *The Journal of Neuroscience*, 29(19):6234–6238, 2009. [xii](#), [86](#), [87](#)
- [KAC⁺09] Joung H. Kwon, Chalmers Alan, Silvester Czanner, Gabriela Czanner, and John Powell. A study of visual perception: Social anxiety and virtual realism. In *Proceedings of the 25th Spring Conference on Computer Graphics, SCCG '09*, pages 167–172, New York, NY, USA, 2009. ACM. [25](#)
- [KF10] Kay Kitazawa and Taku Fujiyama. Pedestrian vision and collision avoidance behavior: Investigation of the information process space of pedestrians using an eye tracker. In *Pedestrian and evacuation dynamics 2008*, pages 95–108. Springer, 2010. [xii](#), [86](#), [87](#)
- [KO12] I. Karamouzas and M. Overmars. Simulating and evaluating the local behavior of small pedestrian groups. *Visualization and Computer Graphics, IEEE Transactions on*, 18(3):394–406, 2012. [15](#)
- [KSHF09] Mubbasir Kapadia, Shawn Singh, William Hewlett, and Petros Faloutsos. Egocentric affordance fields in pedestrian steering. In *Proceedings of the 2009 Symposium on Interactive 3D Graphics and Games, I3D '09*, pages 215–223, New York, NY, USA, 2009. ACM. [vii](#), [13](#)
- [LBB99] Jack M. Loomis, James J. Blascovich, and Andrew C. Beall. Immersive virtual environment technology as a basic research tool in psychology. *Behavior Research Methods, Instruments, & Computers*, 31(4):557–564, 1999. [25](#)
- [LCHL07] Kang Hoon Lee, Myung Geol Choi, Qyoun Hong, and Jehee Lee. Group behavior from video: A data-driven approach to crowd simulation. In *Proceedings of the 2007 ACM SIGGRAPH/Eurographics Symposium on Computer Animation, SCA '07*, pages 109–118, Aire-la-Ville, Switzerland, Switzerland, 2007. Eurographics Association. [15](#)

- [LCSCO09] Alon Lerner, Yiorgos Chrysanthou, Ariel Shamir, and Daniel Cohen-Or. *Data Driven Evaluation of Crowds*, pages 75–83. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. 23
- [LCSCO10] Alon Lerner, Yiorgos Chrysanthou, Ariel Shamir, and Daniel Cohen-Or. Context-dependent crowd evaluation. *Computer Graphics Forum*, 29(7):2197–2206, 2010. 23
- [LJK⁺12] S. Lemercier, A. Jelic, R. Kulpa, J. Hua, J. Fehrenbach, P. Degond, C. Appert-Rolland, S. Donikian, and J. Pettré;. Realistic following behaviors for crowd simulation. *Comp. Graph. Forum*, 31(2pt2):489–498, May 2012. 15, 19, 21, 22, 30, 37, 42
- [LK03] Jack Loomis and Joshua Knapp. Visual perception of egocentric distance in real and virtual environments. *Virtual and adaptive environments: Applications, implications, and human performance issues*, pages 21–46, 2003. 26
- [LMM03] C. Loscos, D. Marchal, and A. Meyer. Intuitive crowd behavior in dense urban environments using local laws. In *Theory and Practice of Computer Graphics, 2003. Proceedings*, pages 122–129, June 2003. 14
- [MCRT06] Betty J Mohler, Sarah H Creem-Regehr, and William B Thompson. The influence of feedback on egocentric distance judgments in real and virtual environments. In *Proceedings of the 3rd symposium on Applied perception in graphics and visualization*, pages 9–14. ACM, 2006. 5, 25
- [MCRTB10] Betty J. Mohler, Sarah H. Creem-Regehr, William B. Thompson, and Heinrich H. Bühlhoff. The effect of viewing a self-avatar on distance judgments in an hmd-based virtual environment. *Presence: Teleoper. Virtual Environ.*, 19(3):230–242, June 2010. 26
- [MGvVB98] Hanspeter A Mallot, Sabine Gillner, Hendrik AHC van Veen, and Heinrich H Bühlhoff. Behavioral experiments in spatial cognition using virtual reality. In *Spatial cognition*, pages 447–467. Springer, 1998. 5, 25
- [MHG⁺09] Mehdi Moussaïd, Dirk Helbing, Simon Garnier, Anders Johansson, Maud Combe, and Guy Theraulaz. Experimental study of the behavioural mechanisms underlying self-organization in human crowds. *Proceedings of the Royal Society B: Biological Sciences*, 276(1668):2755–2762, August 2009. 21
- [MIWB02] Michael Meehan, Brent Insko, Mary Whitton, and Frederick P. Brooks, Jr. Physiological measures of presence in stressful virtual environments. *ACM Trans. Graph.*, 21(3):645–652, July 2002. viii, 25
- [MPG⁺10] Mehdi Moussaïd, Niriaska Perozo, Simon Garnier, Dirk Helbing, and Guy Theraulaz. The walking behaviour of pedestrian social groups and its impact on crowd dynamics. *PloS one*, 5(4):e10047, 2010. 14, 15, 16

- [MT97] S. R. Musse and D. Thalmann. A model of human crowd behavior: Group inter-relationship and collision detection analysis. In *Workshop Computer Animation and Simulation of Eurographics*, pages 39–52, 1997. 14
- [OBBCP14] Anne-Hélène Olivier, Julien Bruneau, Gabriel Cirio, and Julien Pettré. A virtual reality platform to study crowd behaviors. *Transportation Research Procedia*, 2:114–122, 2014. 6
- [OBKPss] Anne-Hélène Olivier, Julien Bruneau, Richard Kulpa, and Julien Pettré. Walking with virtual people: Evaluation of locomotion interfaces in dynamic environments. *Visualization and Computer Graphics, IEEE Transactions on*, in press. 6, ii
- [OMC⁺13] Anne-Hélène Olivier, Antoine Marin, Armel Crétual, Alain Berthoz, and Julien Pettré. Collision avoidance between two walkers: Role-dependent strategies. *Gait & Posture*, 38(4):751–756, 2013. 18, 19
- [OMCP12] Anne-Hélène Olivier, Antoine Marin, Armel Crétual, and Julien Pettré. Minimal predicted distance: A common metric for collision avoidance during pairwise interactions between walkers. *Gait & Posture*, 36(3):399–404, 2012. viii, 18, 19, 20, ii
- [OOP⁺10] Anne-Hélène Olivier, Jan Ondřej, Julien Pettré, Richard Kulpa, and Armel Crétual. Interaction between real and virtual humans during walking: Perceptual evaluation of a simple device. In *Proceedings of the 7th Symposium on Applied Perception in Graphics and Visualization, APGV '10*, pages 117–124, New York, NY, USA, 2010. ACM. 26
- [OPOD10] Jan Ondřej, Julien Pettré, Anne-Hélène Olivier, and Stéphane Donikian. A synthetic-vision based steering approach for crowd simulation. *ACM Trans. Graph.*, 29(4):123:1–123:9, July 2010. vii, viii, 10, 11, 22
- [PAB07] N. Pelechano, J. M. Allbeck, and N. I. Badler. Controlling individual agents in high-density crowd simulation. In *Proceedings of the 2007 ACM SIGGRAPH/Eurographics Symposium on Computer Animation, SCA '07*, pages 99–108, Aire-la-Ville, Switzerland, Switzerland, 2007. Eurographics Association. 8, 14
- [PB06] Alex Poole and Linden J Ball. Eye tracking in hci and usability research. *Encyclopedia of human computer interaction*, 1:211–219, 2006. 84
- [PCBS11] Matthias Plaue, Minjie Chen, Günter Bärwolff, and Hartmut Schwandt. Trajectory extraction and density analysis of intersecting pedestrian flows from video recordings. In *Proceedings of the 2011 ISPRS Conference on Photogrammetric Image Analysis, PIA'11*, pages 285–296, Berlin, Heidelberg, 2011. Springer-Verlag. 19
- [PEC09] C. Peters, C. Ennis, and O'Sullivan C. Modeling groups of plausible virtual pedestrians. *Computer Graphics and Applications, IEEE*, 29(4):54–63, 2009. viii, 15, 16, 18

- [POO⁺09] Julien Pettré, Jan Ondřej, Anne-Hélène Olivier, Armel Cretual, and Stéphane Donikian. Experiment-based modeling, simulation and validation of interactions between virtual walkers. In *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA '09, pages 189–198, New York, NY, USA, 2009. ACM. [18](#), [22](#)
- [PPD07] Sébastien Paris, Julien Pettré, and Stéphane Donikian. Pedestrian reactive navigation for crowd simulation: a predictive approach. In *Computer Graphics Forum*, volume 26, pages 665–674. Wiley Online Library, 2007. [9](#), [21](#)
- [PSAB08] Nuria Pelechano, Catherine Stocker, Jan Allbeck, and Norman Badler. Being a part of the crowd: Towards validating vr crowds using presence. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '08, pages 136–142, Richland, SC, 2008. International Foundation for Autonomous Agents and Multiagent Systems. [21](#), [26](#)
- [PV08] Claudio Pedica and Hannes Vilhjálmsson. Social perception and steering for online avatars. In *Intelligent Virtual Agents*, pages 104–116. Springer, 2008. [14](#)
- [QH10] Fasheng Qiu and Xiaolin Hu. Modeling group structures in pedestrian crowd simulation. *Simulation Modelling Practice and Theory*, 18(2):190–205, 2010. [vii](#), [14](#), [15](#)
- [RCB⁺ss] Zhiguo Ren, Panayotis Charalambous, Julien Bruneau, Qunsheng Peng, and Julien Pettré. Group modeling: a unified velocity-based approach. *Comp. Graph. Forum*, in press. [6](#), [v](#), [viii](#)
- [Rey87] Craig W. Reynolds. Flocks, herds and schools: A distributed behavioral model. *SIGGRAPH Comput. Graph.*, 21(4):25–34, August 1987. [8](#), [14](#)
- [Rey99] Craig Reynolds. Steering behaviors for autonomous characters. In *Game Developers Conference 1999*, pages 763–782, San Francisco, California, 1999. Miller Freeman Game Group. [vii](#), [8](#), [9](#), [14](#)
- [RRW14] Kevin W. Rio, Christopher K. Rhea, and William H. Warren. Follow the leader: Visual control of speed in pedestrian following. *Journal of Vision*, 14(2), 2014. [15](#), [30](#), [42](#)
- [RVH13] Rebekka S. Renner, Boris M. Velichkovsky, and Jens R. Helmert. The perception of egocentric distances in virtual environments - a review. *ACM Comput. Surv.*, 46(2):23:1–23:40, December 2013. [26](#)
- [RY13] Francisco Arturo Rojas and Hyun Seung Yang. Immersive human-in-the-loop hmd evaluation of dynamic group behavior in a pedestrian crowd simulation that uses group agent-based steering. In *Proceedings of the*

- 12th ACM SIGGRAPH International Conference on Virtual-Reality Continuum and Its Applications in Industry*, pages 31–40. ACM, 2013. [15](#), [21](#), [26](#)
- [SAB07] A. Elizabeth Seward, Daniel H. Ashmead, and Bobby Bodenheimer. Using virtual environments to assess time-to-contact judgments from pedestrian viewpoints. *ACM Trans. Appl. Percept.*, 4(3), November 2007. [26](#)
- [SBK⁺10] Armin Seyfried, Maik Boltes, Jens Kähler, Wolfram Klingsch, Andrea Portz, Tobias Rupprecht, Andreas Schadschneider, Bernhard Steffen, and Andreas Winkens. *Enhanced Empirical Data for the Fundamental Diagram and the Flow Through Bottlenecks*, pages 145–156. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010. [76](#)
- [SG00] Dario D. Salvucci and Joseph H. Goldberg. Identifying fixations and saccades in eye-tracking protocols. In *Proceedings of the 2000 Symposium on Eye Tracking Research & Applications*, ETRA '00, pages 71–78, New York, NY, USA, 2000. ACM. [90](#)
- [SGC04] Mankyu Sung, Michael Gleicher, and Stephen Chenney. Scalable behaviors for crowd simulation. *Computer Graphics Forum*, 23(3):519–528, 2004. [14](#)
- [SKFR09] Shawn Singh, Mubbasir Kapadia, Petros Faloutsos, and Glenn Reinman. Steerbench: a benchmark suite for evaluating steering behaviors. *Computer Animation and Virtual Worlds*, 20(5-6):533–548, 2009. [viii](#), [24](#)
- [SKH⁺11] Shawn Singh, Mubbasir Kapadia, Billy Hewlett, Glenn Reinman, and Petros Faloutsos. A modular framework for adaptive agent-based steering. In *Symposium on Interactive 3D Graphics and Games*, I3D '11, pages 141–150 PAGE@9, New York, NY, USA, 2011. ACM. [56](#)
- [SML15] Manish Sreenivasa, Katja Mombaur, and Jean-Paul Laumond. Walking paths to and from a goal differ: On the role of bearing angle in the formation of human locomotion paths. *PLoS ONE*, 10(4):1–16, 04 2015. [10](#)
- [SSKB05] Armin Seyfried, Bernhard Steffen, Wolfram Klingsch, and Maik Boltes. The fundamental diagram of pedestrian movement revisited. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(10):P10002, 2005. [19](#)
- [SSUS00] Mel Slater, Amela Sadagic, Martin Usoh, and Ralph Schroeder. Small-group behavior in a virtual and real environment: A comparative study. *Presence: Teleoper. Virtual Environ.*, 9(1):37–51, February 2000. [25](#)
- [SSW⁺09] Armin Seyfried, Bernhard Steffen, Andreas Winkens, Tobias Rupprecht, Maik Boltes, and Wolfram Klingsch. *Empirical Data for Pedestrian Flow Through Bottlenecks*, pages 189–199. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. [19](#)

- [ST05] Wei Shao and Demetri Terzopoulos. Autonomous pedestrians. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA '05, pages 19–28, New York, NY, USA, 2005. ACM. [vii](#), [11](#), [12](#)
- [SvdSKvdM01] Martijn J. Schuemie, Peter van der Straaten, Merel Krijn, and Charles A. van der Mast. Research on presence in virtual reality: A survey. *Cyberpsychology and Behavior*, 4(2):183–201, 2001. [25](#)
- [SvSB⁺13] Sarah Schwarzkopf, Rul von Stülpnagel, Simon J BÜchner, Lars Konieczny, G Kallert, and C Hölscher. What lab eye-tracking tells us about wayfinding. a comparison of stationary and mobile eye-tracking in a large building scenario. In *1st Intl. workshop on eye tracking for spatial research, ET4S*, 2013. [xi](#), [84](#), [85](#), [86](#)
- [TCP06] Adrien Treuille, Seth Cooper, and Zoran Popović. Continuum crowds. *ACM Trans. Graph.*, 25(3):1160–1168, July 2006. [7](#)
- [TLvK⁺12] S. Thovuttikul, D. Lala, N. van Kleef, Y. Ohmoto, and T. Nishida. Comparing people’s preference on culture-dependent queuing behaviors in a simulated crowd. In *Cognitive Informatics Cognitive Computing (ICCI*CC), 2012 IEEE 11th International Conference on*, pages 153–162, 2012. [25](#)
- [TYK⁺09] Shigeo Takahashi, Kenichi Yoshida, Taesoo Kwon, Kang Hoon Lee, Jeehee Lee, and Sung Yong Shin. Spectral-based group formation control. In *Computer Graphics Forum*, volume 28, pages 639–648. Wiley Online Library, 2009. [15](#)
- [vdBGLM11] Jur van den Berg, Stephen J. Guy, Ming Lin, and Dinesh Manocha. Reciprocal n-body collision avoidance. In Cédric Pradalier, Roland Siegwart, and Gerhard Hirzinger, editors, *Robotics Research*, volume 70 of *Springer Tracts in Advanced Robotics*, pages 3–19. Springer Berlin Heidelberg, 2011. [9](#), [68](#), [77](#)
- [vdBLM08] J. van den Berg, Ming Lin, and D. Manocha. Reciprocal velocity obstacles for real-time multi-agent navigation. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, pages 1928–1935, 2008. [9](#), [53](#)
- [VLL⁺14] Vaisagh Viswanathan, Chong Eu Lee, Michael Harold Lees, Siew Ann Cheong, and Peter M.A. Sloot. Quantitative comparison between crowd models for evacuation planning and evaluation. *The European Physical Journal B*, 87(2):1–11, 2014. [22](#)
- [WCCRT04] Peter Willemsen, Mark B. Colton, Sarah H. Creem-Regehr, and William B. Thompson. The effects of head-mounted display mechanics on distance judgments in virtual environments. In *Proceedings of the 1st Symposium on Applied Perception in Graphics and Visualization, APGV '04*, pages 35–38, New York, NY, USA, 2004. ACM. [26](#)

- [WF04] W. H. Warren and B. R. Fajen. *Optic Flow and Beyond*, chapter From optic flow to laws of control, pages 307–337. Kluwer (Editors: L. M. Vaina, S. A. Beardsley, and S. Rushton), 2004. [5](#)
- [WF08] William H. Warren and Brett R. Fajen. *Behavioral Dynamics of Visually Guided Locomotion*, pages 45–75. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008. [vii](#), [10](#)
- [WHBK12] Jan M Wiener, Christoph Hölscher, Simon Büchner, and Lars Konieczny. Gaze behaviour during space perception and spatial decision making. *Psychological research*, 76(6):713–729, 2012. [84](#)
- [Whi03] Michael W Whittle. *Gait analysis: an introduction*. Elsevier, 2003. [52](#), [69](#)
- [WJGO⁺14] D. Wolinski, S. J. Guy, A.-H. Olivier, M. Lin, D. Manocha, and J. Pettré. Parameter estimation and comparative evaluation of crowd simulations. *Computer Graphics Forum*, 33(2):303–312, 2014. [viii](#), [23](#), [24](#)
- [WS98] Bob G. Witmer and Michael J. Singer. Measuring presence in virtual environments: A presence questionnaire. *Presence: Teleoper. Virtual Environ.*, 7(3):225–240, June 1998. [25](#)
- [YCP⁺08] H. Yeh, S. Curtis, S. Patil, J. van den Berg, D. Manocha, and M. Lin. Composite agents. In *Proceedings of the 2008 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA '08, pages 39–47, Aire-la-Ville, Switzerland, Switzerland, 2008. Eurographics Association. [12](#), [68](#), [78](#)
- [YMMT08] Barbara Yersin, Jonathan Maïm, Fiorenzo Morini, and Daniel Thalmann. Real-time crowd motion planning. *The Visual Computer*, 24(10):859–870, 2008. [14](#)
- [Zip49] G. K. Zipf. *Human Behaviour and the Principle of Least-Effort*. Addison-Wesley, Cambridge MA, 1949. [52](#)
- [ZTZW14] Bolei Zhou, Xiaoou Tang, Hepeng Zhang, and Xiaogang Wang. Measuring crowd collectiveness. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(8):1586–1599, August 2014. [17](#), [22](#)
- [ZWT12] B. Zhou, X. Wang, and X. Tang. Understanding collective crowd behaviors: Learning a mixture model of dynamic pedestrian-agents. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2871–2878, June 2012. [viii](#), [16](#), [17](#), [18](#)

Publications list

In press

- Julien Bruneau and Julien Pettré. EACS: Effective avoidance combination strategy. *Comp. Graph. Forum*, in press.
- Zhiguo Ren, Panayotis Charalambous, Julien Bruneau, Qunsheng Peng, and Julien Pettré. Group modeling: a unified velocity-based approach. *Comp. Graph. Forum*, in press.
- Anne-Hélène Olivier, Julien Bruneau, Richard Kulpa, and Julien Pettré. Walking with virtual people: Evaluation of locomotion interfaces in dynamic environments. *Visualization and Computer Graphics*, IEEE Transactions on, in press.

2015

- Julien Bruneau and Julien Pettré. Energy-efficient mid-term strategies for collision avoidance in crowd simulation. In *Proceedings of the 14th ACM SIGGRAPH / Eurographics Symposium on Computer Animation, SCA '15*, ACM, pages 119-127, New York, NY, USA, 2015.
- Julien Bruneau, Anne-Hélène Olivier, and Julien Pettré. Going through, going around: A study on individual avoidance of groups. *IEEE Transactions on Visualization and Computer Graphics*, 21(4):520-528, April 2015.

2014

- Julien Bruneau, Teófilo B. Dutra, and Julien Pettré. Following behaviors: A model for computing following distances based on prediction. In *Proceedings of the Seventh International Conference on Motion in Games, MIG '14*, ACM, pages 17-24, New York, NY, USA, 2014.
- Anne-Hélène Olivier, Julien Bruneau, Gabriel Cirio, and Julien Pettré. A virtual reality platform to study crowd behaviors. *Transportation Research Procedia*, 2:114-122, 2014.