Bar-Ilan University
Department of Computer Science

Modeling Crowd Behavior Based On Social Comparison Theory

by

Natalie Fridman

Advisor: Dr. Gal Kaminka

Submitted in partial fulfillment of the requirements for the Master's degree
in the department of Computer Science

Ramat-Gan, Israel
May 2007

Abstract

Modeling crowd behavior is an important challenge for cognitive modelers. Models of crowd behavior facilitate analysis and prediction of human group behavior, where people are close geographically or logically states, and are affected by each other's presence and actions. Existing models of crowd behavior, in a variety of fields, leave many open challenges. In particular, psychology models often offer only qualitative description, and do not easily permit algorithmic replication, while computer science models are often simplistic, treating agents as simple deterministic particles. We propose a novel model of crowd behavior, based on Festinger's Social Comparison Theory (SCT), a social psychology theory known and expanded since the early 1950's. We propose a concrete algorithmic framework for SCT, and evaluate its implementations in several crowd behavior scenarios such as pedestrian movement, gathering and imitational behavior. We show that our SCT model produces improved results compared to base models from the literature. We describe two possible implementations of SCT process in an architectural level. The first, which seems to follow directly from Festinger's Social Comparison theory, treats the SCT process as an uncertainty-resolution method. The second, takes a different approach, in which an SCT process is constantly active, in parallel to any problem solving activity. We present the implementation of these approaches in Soar cognitive architecture. Moreover, we examine these approaches in the context of crowd behavior simulations and argue that surprisingly, it is the second approach which is correct.

Acknowledgments

It is very difficult to put into words an appropriate acknowledgments that reflect all my appreciation and gratitude to my colleagues, family and friends.

First, I would like to express my deepest gratitude to my advisor Gal A. Kaminka, for his constant support, outstanding guidance and excellent advice throughout this research. He has been everything that one could want in an advisor and I consider myself very lucky and most honored to have been one of his students. Gal has always provided me with direction and while giving a total freedom to explore new ideas, he always managed to keep me in focus. Gal's continuous support was not only on scientific side but also personal. His advises went from how to make a good research to what is really important in life; from how to write a good paper to where to rent an apartment. I admire his energy, enthusiasm and generosity. I am really looking forward to keep working with him.

I am very grateful to Prof. Shaul Fox for providing a consultation on psychological and sociological themes involving this research. Although we only had a few meetings, they were fruitful and helpful.

Many thanks to the Maverick group for their friendly, supportive and mostly fun atmosphere. Especially, I would like to thank to Nirom Cohen-Nov Slapak and Ari Yakir for their great collaboration and for the interesting and fruitful discussions. It was always a pleasure to work with you.

My deepest appreciations to my parents for their love, support and encouragement throughout my entire life. Thank you for providing me a warm home and a clean head for study by always preventing me from worry about financial issues. I am also very grateful to my brother Alex for all his love and support. Despite his busy schedule, he always try to stay updated with my research progression, even if it evolved a phone call in the middle of the night.

Most importantly, my deepest thanks to the most wonderful person in my life, my closest friend and my husband, Alex. Thank you for all your love, for being a patience with me and for tolerating my long working hours. His support and encouragement were a large part of what kept me going. Alex has always been a great source of strength for me and most importantly, he makes my life more enjoyable.

Finally, I would like to express my gratitude to Israeli Ministry of Defense, that provided partial support for this research.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Modeling crowd behavior is an important challenge for cognitive modelers. Models of crowd behavior facilitate analysis and prediction of the behavior of groups of people, who are in close geographical or logical states, and are affected by each other's presence and actions. Accurate models of crowd behavior are sought in training simulations [30], safety decision-support systems [5], traffic management [14, 27], business and organizational science.

Existing models of crowd behavior, in a variety of fields, leave many open challenges. In social sciences and psychology, models often offer only qualitative description, and do not easily permit algorithmic replication. In computer science, models are often simplistic, and typically not tied to specific cognitive science theories or data. Moreover, existing computer science models often focus only on a specific phenomenon (e.g., flocking, pedestrian movement), and thus must be switched depending on the goals of the simulation.

We propose a novel model of crowd behavior, based on Social Comparison Theory (*SCT*) [10], a social psychology theory that has been continuously evolving since the 1950s. The key ideas in this theory is that humans, lacking objective means to evaluate their state, compare themselves to others that are similar. Similarity, in SCT, is very loosely defined—indeed much of the literature on SCT addresses with exploring different ways in which humans judge similarity.

While inspired by SCT, we remain deeply grounded in computer science; we propose a concrete algorithmic framework for SCT, and evaluate its implementations in several crowd behavior scenarios. We quantitatively compare the performance of SCT crowd behavior model with popular models in the literature, and show that SCT generates behavior more in-tune with human crowd behavior. Moreover, unlike many previous models, SCT generalizes across social phenomena. In particular, we evaluate the use of SCT model in generation of pedestrian movement, gathering and imitational behavior.

In pedestrian movement generation, the SCT model accounts for group formation in pedestrians that are inter-related, a phenomenon unaccounted for by previous models; and where previous techniques apply, SCT shows improved results. In addition the SCT model accounts for group behavior in the presence of obstacles, modeling the selection of group members to bypass obstacles in the same direction as other members of the group.

In gathering towards a target location, the SCT model accounts for successful gathering where participants are uninformed about the location. We show that the SCT model results in a middle-ground between individual behavior (where each agent is only aware of itself, and independently searches for the target location), and a full-knowledge model (where all agents know the target location). We additionally show that we can improve these results by adding a relatively small percentage of agents that know the target location, to the general agent population. This shows that in principle, social comparison can compensate for lack of objective knowledge by the agent, as predicted by the psychology theory.

In the context of imitational behavior, the SCT model was evaluated in studies with human subjects. The model was used to control the behavior of agents in a 3D virtual environment. The subjects ranked SCT to be a middle-ground between completely individual behavior, and perfect synchronized ("solider-like") behavior. Independently, human subjects gave similar rankings to short clips showing human crowds. While the similarity in the results is no proof that social comparison the best model for modeling human behavior, it is certainly encouraging in that the SCT is at least shown to be compatible with such behavior.

Finally, we discuss SCT as an hypothesized cognitive process. We describe two possible implementations of SCT process in an architectural level. The first, which seems to follow directly from Festinger's Social Comparison theory, treats the SCT process as an uncertainty-resolution method, i.e., as a weak (read: general) problem-solving method, which is *social*. The second, which is less intuitive, takes a different approach, in which an SCT process is constantly active, in parallel to any problem solving activity. Such a view necessitates the agents to be constantly aware of others around them, and seems to require greater computational cost. We present the implementation of these approaches in the Soar cognitive architecture. Moreover, we examine these approaches in the context of crowd behavior simulations and conducted a set of experiments that contrast these approaches.

Based on the results from experiments, we argue that surprisingly, it is the second approach which is correct, i.e., that SCT is constantly active, regardless of the problem-solving activity of the agent. This conclusion raises questions as to the role of social reasoning in cognitive architectures and the mind. In particular, it leads to the conclusion that modeling of other agents, which is a precursor to social comparison, occurs as a fundamental cognitive process at an architectural level, rather than as a part of problem solving.

Chapter 2

Background and Motivation

Social psychology literature provides several views on the emergence of crowds and the mechanisms underlying its behaviors. These views can inspire computational models, but are unfortunately too abstract to be used algorithmically. In contrast, computational crowd models tend to be simplistic and focus on specific crowd behaviors (e.g, flocking). A common theme in all of them is the generation of behavior from the aggregation of many local rules of interaction, e.g., [27, 25, 26, 20]. However, these models have rarely, been validated against human (or animal) data. Indeed, there is generally limited quantitative data on the behavior of human crowds at a resolution which permits accurate modeling. The exception is the formation of lanes (in opposing directions) in human pedestrian movements and evacuation behavior [9, 20], which have been extensively investigated and for which specific performance measures are well defined (reduced lane changes, flow, time between alarm and last person that leave the building etc.).

Social psychology. A phenomenon observed with crowds, and discovered early in crowd behavior research is that people in crowds act similar to one another, often acting in a coordinated fashion, as if governed by a single mind [21, 2]. However, this coordination is achieved with little or no verbal communication.

Le Bon [21] emphasized a view of crowd behaviors as controlled by a "Collective Mind", and observed that an individual who becomes a part of the crowd is transformed into becoming identical to the others in the crowd. Le Bon noted that individuals seem to lose their individuality (in terms of personality and thought) when becoming part of a crowd. Le Bon explains the homogeneous behavior of a crowd by two processes: (i) *Imitation*, where people in crowds imitate each other; and (ii) *Contagion*, where people in a crowd behave very differently from the way they usually do, individually.

Freud [11] expanded on Le Bon by theorizing that individuals in the crowd identify with the leader and with each other, and therefore behave as one. As a corollary, crowd behavior can be controlled by the leader, as the individuals imitate that person.

Another phenomenon that was addressed by researchers is what makes an individual be part of a crowd. According to Allport's theory, individuals become a part of the crowd behavior when they have a "common stimulus" with people inside the crowd; for example, a common cause [2]. Allport agrees with Le Bon [21] about the homogeneous behavior of the crowd, but his explanation of this phenomenon is that similar people act in similar ways; otherwise they would not be a part of the same group. Thus, according to Allport, "the individual in the crowd behaves just as he would behave alone, only more so."

We based our work on social comparison theory [10], which (to the best of our knowledge) has never been applied to modeling crowd behavior. Nevertheless, as we show in the next section, key elements of the theory are at the very least compatible with those theories discussed above.

Computational models. Work on modeling collective behavior has been carried out in other branches of science, in particular for modeling and simulation. Inspired by different science fields, researchers are developing models for simulation of collective behavior.

Reynolds [25] simulated bird flocking using simple, individual-local rules, which interacted to create coherent collective movement. There are only three rules: avoid collision with neighbors, match velocity with neighbors and stay close to the center of gravity of all neighbors. Each simulated bird is treated as a particle, attracted and repelled by others. On the one hand there is a desire to stay close to the flock, but on the other hand, there is a desire to avoid collisions. However, this model was limited only to the interaction of the agents, and did not allow for their individual goals (e.g., their own steering behavior).

Tu and Terzopoulos [33] simulated motion of artificial fish that addressed individual goals. Like Reynolds' "boids", the artificial fish are autonomous creatures which have simple behaviors and together are able to create a more complex, collective behavior. However, unlike Reynolds' boids, that selected their behavior based on the current state of their neighbors, each fish revealed habits and mental state (for example hunger, fear etc.) that also impact behavior selection. Indeed, Reynolds later expanded his work on collective movement in [26] but, this time allowing for a steering behavior for the autonomous agents. In the revised model, each agent has a set of simple steering behaviors such as seek, flee, pursuit, evade, etc. The combination of these simpler behaviors creates a complex steering behavior.

Similar ideas have been applied in swarm robotics. Matarić [22] sees collective (complex) behaviors as a combination of basic behaviors. Each robot has spatial behaviors (controllers) that are combined to create different kinds of group behavior: for example, flocking consisting of *safe-wandering* (moving around without bumping), *homing*, *dispersion* (moving away from other agents), and *aggregation* (moving towards other agents). The combined outputs of the basic behaviors provide a velocity vector which is used to control the robot.

Yamashita and Umemura [37] take a different approach in simulating panic behavior. While inspired by Reynolds' boid model, they propose a model where each simulated person moves by three instincts: escape instinct, group instinct and imitational instinct. According to Yamashita and Umemura, when a person is in panic, he or she acts based on their instincts which make their decision making process much simpler.

Henderson compared pedestrian movement to gaskinetic fluids. Based on experiments on real human crowds, he showed in [17] that crowd distribution is compatible with Maxwell-Boltzmann's distribution. Henderson [18] developed a pedestrian movement model based on the Maxwell-Boltzmann theory. Since each person has mass and velocity, the crowd may be transformed to liquid gas and under some assumption the Maxwell-Boltzmann theory may be applied. Based on Boltzmann-like equations, Helbing [13] developed a general behavior model for simulation of crowd dynamics. The proposed model takes into account social forces caused by interaction between the individuals and external or spontaneous forces which are caused by the physical environment.

Helbing et al. [15, 14] observed phenomena of self-organization in collective motion which can be caused by interaction among pedestrians. By self-organization, it means that there are some behavioral phenomena which were not planned: for example, creation of lane formation in pedestrian movement. These lanes are created as a result of pedestrians moving against the flow. When a pedestrian moves against the flow, he experiences an interaction which makes him move a little aside, in contrast to a pedestrian who moves with the flow and will not have an interaction. The number of lanes that are created cannot be planned. It depends on the width of the street and on pedestrian density.

Helbing and Vicsek [16] expanded their physical model by using game theory. The attraction force can be expanded to profitable force which may lead to optimal self-organization in pedestrian movement. Each entity calculates "expected success" per each possible action and the action with maximum success will be chosen. In pedestrian relations, actions are possible directions that an entity can move to and optimal self-organization is minimal interaction between entities.

Adriana Brown et al. [6] examined how individual characteristics impact crowd evacuation. They expanded Helbing's physical model by adding to each agent individual parameters, such as dependence level and altruism level. According to the model, there will be a creation of groups which are combined from altruism and dependent agents. By changing these attributes, they examined crowd evacuation by measuring the flow of people passing the door per second and population distribution in the flow.

Blue and Adler [4] proposed a different approach to module collective dynamics. They used Cellular Automata (CA) in order to simulate collective behaviors, in particular pedestrian movement. The focus is again on local interactions: Each simulated pedestrian is controlled by an automaton, which decides on its next action or behavior, based on its local neighborhoods. These rules are responsible for making a decision about lane changing and forward movement: If the way forward is free, then it is taken. If not, then the automaton seeks to go left or right. If both lanes are available, one is chosen arbitrarily. Blue and Adler showed that this simple rule results in the formation of lanes in movement, similarly to those formed in human pedestrian movement [35]. Toyama et al. [32] expanded the cellular automata model by adding different pedestrian characteristics, such as speed, gender, repulsion level, etc. The model was examined on bi-directional pedestrian movement behavior and on evacuation behavior. The experiment analysis shows that macroscopic behavior of homogeneous agents is different from heterogeneous agents.

Osaragi [24] proposed an agent-based model for simulating pedestrian flow by using the concept of pedestrian mental stress. Pedestrian mental stress increases as a result of other pedestrians (density) and whether the pedestrian is unable to move to her destination using the shortest pass. To decrease her mental stress, the pedestrian may dynamically change her direction or walking velocity. Because of these dynamic changes, the simulated pedestrians are heterogeneous. Unlike in other models, the model parameters were estimated using observed data.

Kretz [20] proposes the Floor field-and-Agent based Simulation Tool model (F.A.S.T) which is discrete in space and time model for pedestrian motion. The F.A.S.T model can be classified as an extension of probabilistic Cellular Automata (PCA). In this model there are three levels of decision making: 1. The choice of an exit. 2. The choice of a destination cell. and 3. The path between the current and destination cell. The F.A.S.T model has been validated against human data. In particular, the model simulation results of evacuation scenario was compared to results of evacuation exercise at a primary school.

In these previous works above, the behavior of crowds in every domain of study (pedestrian movement, flocking, evacuation, etc.) is computed using a different algorithm, yet the actions and perceptions remain largely invariant (e.g., distances to others, occupied spaces versus empty spaces, goal locations, etc.). Instead, the computation itself changes between modeled behaviors.

For instance, many models for crowd behavior utilize cellular-automata (CA), which differ between domains. One CA model for pedestrian movement [4] uses a set of 6 IF-THEN rules which work in parallel for all cells, to simulate the movement of pedestrians in cells. The rules utilize knowledge of the occupancy in adjacent (rules 1,3 in [4]) and farther cells (rule 2), as well as of the distance to oncoming pedestrians in the same lane (rules 4, 6). The rules set the forward velocity and position of the entities, by using a set of non-deterministic choices (sub-rules 5a,5b,5c), biased by distributions which differ depending on the environmental settings (e.g., choose from a uniform 50%/50% split distribution if two nearby cells are occupied, or from a 10%/80%/10% distribution when three cells are available). In contrast, a recent CA model for evacuation [31] uses knowledge of adjacent cells and distances to exits, and sets the position of the entities. Thus the actions and perceptions of each entity are similar to those used in the pedestrian model. But the algorithmic computation of the new position is done in two deterministic rules [31, pp. 17], which involve no arbitrary choices at all.

In contrast to these previous investigations, we seek a single cognitive mechanism that, when executed by individuals, would give rise to different crowd behaviors, depending on the perceptions and actions available to the agents. In other words, our goal is to unravel a *single computational mechanism*—a single algorithm—which would account for different crowd phenomena, by virtue of the actions and perceptions available to each individual.

Chapter 3

A Model of Social Comparison

We took Festinger's social comparison theory (SCT) [10] as inspiration for the social skills necessary for our agent in order to be able to exhibit crowd behavior. According to social comparison theory, people tend to compare their behavior with others that are most like them. To be more specific, when lacking objective means for appraisal of their opinions and capabilities, people compare their opinions and capabilities to those of others that are similar to them. They then attempt to correct any differences found. This section shows how SCT can be turned into a concrete algorithm, to be used for generating crowd behavior.

3.1 Festinger's Social Comparison Theory

Festinger [10] presents social comparison theory (SCT) as an explicit set of axioms. The following subset of axioms (re-worded) are particularly relevant:

- When lacking objective means for evaluation, agents compare their state to that of others;
- Comparison increases with similarity;
- Agents take steps to reduce differences to the objects of comparison.

Newell [23] classified each of Festinger's axioms with respect to the type of agent it assumes, and concluded that in fact, SCT may be used in principle to generate social behavior out of axioms that are largely non-social (in the sense that they do not cause agents to actively interact). However, Newell's discussion was essentially philosophical: No algorithm was suggested, nor any method for using SCT's axioms as the basis for a computational process.

To be usable by computerized models, SCT's axioms must be transformed into an algorithm that, when executed by an agent, will proscribe social comparison behavior. To do this, we re-examined Festinger's discussion and examples of how the axioms apply.

For instance, Festinger proposes that when lacking objective means for evaluation, people compare their opinions and capabilities to those of others. He then carefully notes that the comparison takes place at the level of the opinion or capability: "Thus, if a person evaluates his running ability, he will do so by comparing his time to run some distance with the times that other persons have taken." [10, p. 116].

Later on, in discussing how actions are selected to minimize differences, he again notes that the action is selected at the level at which the difference is found: "When pressures toward uniformity exist with respect to abilities, these pressures are manifested less in social process and more in action against the environment which restrains movement. Thus, a person who runs more slowly than others with whom he compares himself, and for whom the ability is important, many[1] spend considerable time practicing running. In a similar situation where the ability in question is intelligence, the person may study harder." [10, p. 126].

Based on these observations, we take another step towards the modeling of social comparison theory. We propose a concrete algorithmic framework for SCT that can be executed by an agent. Moreover, we propose the use of SCT algorithmic framework for modeling crowd behaviors. In social psychology there are several views on the mechanisms underlying individual that is a part of crowd behavior. However, to the best to our knowledge, social comparison theory has never been connected to crowd behavior phenomena. We believe that SCT algorithmic framework can provide social skills that are necessary for agents in order to exhibit crowd behavior phenomenons. The basis of our belief is that social comparison theory may account for Le Bon's [21] characteristics of crowd behavior:

[1]This is likely a typo in the original manuscript, to be replaced by "may".

Imitation. Using social comparison, people may adopt others' behaviors. Festinger notes [10]: "The drive for self evaluation is a force acting on persons to belong to groups, to associate with others. People, then, tend to move into groups which, in their judgment, hold opinions which agree with their own“.

Contagion. One implication of SCT is the formation of homogeneous groups. Festinger writes [10]: "The existence of a discrepancy in a group with respect to opinions or abilities will lead to action on the part of members of that group to reduce the discrepancy".

3.2 An SCT Algorithm

In order to build algorithmic framework for SCT, each observed agent is assumed to be modeled by a set of features and their associated values. For each such agent, we calculate a similarity value $s(x)$, which measures the similarity between the observed agent and the agent carrying out the comparison process. The agent with the highest such value is selected. If its similarity is between given maximum and minimum values, then this triggers actions by the comparing agent to reduce the discrepancy.

The process is described in the following algorithm, which is executed by comparing agent.

1. For each known agent x calculate similarity $s(x)$
2. $c \leftarrow \texttt{argmax}\ s(x)$, such that $S_{min} < s(c) < S_{max}$
3. $D \leftarrow$ differences between me and agent c
4. Apply actions to minimize differences in D.

In line 1, the comparing agent (*me*, for short) compares itself with other agents. We model each agent as an ordered set of features, where similarity can be calculated for each feature independently of the others. We use a weighted linear sum to compute the similarity measure $s(x)$:

$$s(x) = \sum_{i=0}^{k} w_i f_i$$

where k is the feature index, f_i similarity in feature i, $0 \leq f_i \leq 1$, and w_i the weight of the feature in overall similarity (non-negative).

For each calculated similarity value, we check in line 2 if it is bounded by S_{min} and S_{max}, and pick the agent that maximizes the similarity, but still falls within the bounds. S_{min} denotes values that are too dissimilar, and the associated agents are ignored. Festinger writes [10]: "When a discrepancy exists with respect to opinions or abilities there will be tendencies to cease comparing oneself with those in the group who are very different from oneself". Respectively, there is also an upper bound on similarity S_{max}, which prevents the agent from trying to minimize differences where they are not meaningful or helpful. For instance, without this upper bound, an agent that is stuck in a location may compare itself to others, and prefer those that are similarly stuck in place.

In line 3, we determine the list of features f_i that indicate a difference with the selected agent c. We order these features in an increasing order of weight w_i, such that the first feature to trigger corrective action is the one with the least weight. The reason for this ordering is intuitive, and we admittedly did not find evidence for it in the literature. However, no evidence was provided against this ordering, and it empirically worked better in the experiments (see below).

Finally, in step 4 of the algorithm, the comparing agent takes corrective action on the selected feature. Note that we assume here that every feature has associated corrective actions that minimize gaps in it, to a target agent, independently of other features. Festinger writes [10]: "The stronger the attraction to the group the stronger will be the pressure toward uniformity concerning abilities and opinions within that group". To model this, we use a gain function $g(o)$ for the action o, which translates into the amount of effort or power invested in the action. For instance, for movement, the gain function would translate into velocity; the greater the gain, the greater the velocity.

$$g(o) = \frac{S_{max} - S_{min}}{S_{max} - s(c)}$$

Chapter 4

Modeling Pedestrian Movement

The coordinated behavior of crowds has often been investigated in the context of pedestrian dynamics. Pedestrian motion (direction and velocity) is affected not only by physical elements (e.g., the sidewalk), but also by the motion of other pedestrians. Wolff [35] noted that pedestrians have a hight degree of cooperation and coordination which without it, walking on sidewalk would be impossible.

To learn more about microscopic and macroscopic pedestrians' behavior, Daamen and Hoogendoorn [9] performed empirical experiments on human crowds, in particular in terms of movement as pedestrians. In these experiments, participants were asked to walk through a monitored area, in both directions. Their movements were recorded. One conclusion was that "During capacity conditions, two trails or lanes are formed: pedestrians tend to walk diagonally behind each other, thereby reducing the head ways and thus maximizing the use of the infrastructure supply".

Since then, lane formations have been a hallmark of pedestrian movement models. Quicker lane formations typically leads to improved flow through the area, and the more agents organize into lanes, the less their need to spend efforts coordinating with others (change lanes). It is thus generally assumed that when measuring lane changes over time, improved models lead to a reduction in the number of lane changes.

We explore the use of SCT in generating pedestrian movements in different settings (individual, groups, with and without obstacles) and compare its performance to known models. Our goal is to explore if SCT model can account for common pedestrian behavior phenomenons like lane formations in bidirectional movement, and movement in groups, with and without obstacles.

To implement the model for pedestrians movement experiments, we used NetLogo [34]. We simulated a sidewalk where agents can move in a circular fashion from east to west, or in the opposite direction. Each agent has limited vision distance (beyond this distance it cannot see). It also has a cone-shaped field-of-view of 120 degrees. Each agent initially moves with a default walking velocity (in our case, 0.1). Agents are not allowed to move through other agents, and thus no two agents can occupy the same space.

Figure 4.1 shows the NetLogo sidewalk environment, in an initial state where simulated pedestrians are randomly placed about. Each small triangle is a simulated pedestrian, able to move left-to-right or right-to-left. Pedestrians exiting the sidewalk on any side appear on the other side, heading in the same direction. Figure 4.2 shows an end-result from one of the experiments (described below), where lanes have formed.

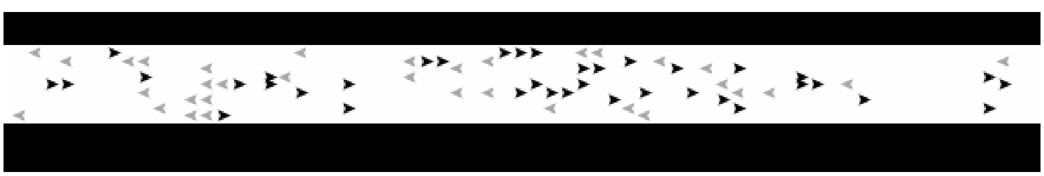

Figure 4.1: **Initial NetLogo sidewalk.**

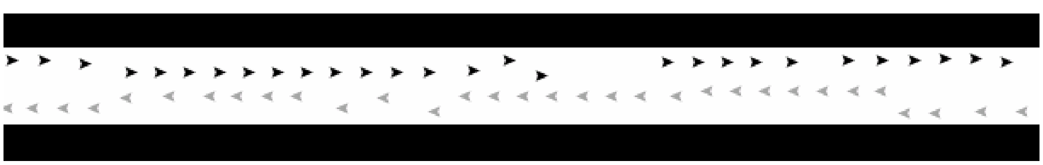

Figure 4.2: **Lane formations - Experiment end-results.**

Each agent has a set of features and its corresponding weight. For simulating pedestrian movement, we used the following features and weights:

Walking direction (weight: 2). Agents can move in two opposite directions, east and west.

Color (weight: 3). Each agent has a color (blue, pink, red, green, etc.)

Position (weight 1). Each agent has a position, given in terms of distance and angle. Distance - Represents the vicinity in position between me and other agent.

The similarities in different features (f_i) are calculated as follows. $f_{color} = 1$ if color is the same, 0 otherwise. $f_{direction} = 1$ if direction is the same, 0 otherwise. and finally, $f_{distance} = \frac{1}{dist}$, where $dist$ is the Euclidean distance between the positions of the agents.

The rationale for feature priorities, as represented in their weights, follows from our intuition and common experience as to how pedestrians act. Positional difference (distance) is the easiest difference to correct, and the least indicative of a similarity between pedestrians. Direction is more indicative of a similarity between agents, and color even more so. If an agent sees two agents, one in the same direction as it (and far away), and the other very close to it (but in the opposite direction), it will calculate greater similarity to the first agent, and try to minimize the distance to it (this may cause a lane change).

Each agent calculates $s(x)$ according to the model. If the chosen feature for closing the gap is distance, then the velocity for movement will be multiplied by the calculated gain $g(o)$. For other features (which are binary), the gain is ignored.

To evaluate the SCT model, we contrast it with a popular alternative model, often used in pedestrian crowd research [4, 15]. In this *individual choice* model, each agent chooses lanes arbitrarily if forward movement is blocked. This model was repeatedly shown to produce lane formations.

We compare these models as is commonly done in pedestrian movement experiments: We controlled for *crowd density*, calculated as the number of agents divided by the area. We follow the literature in measuring two principal characteristics of pedestrian movement: the total number of *lane changes* (lower numbers indicate improved lane formations), and the *flow* (average speed divided by the space-per-agent; higher flow is better).

In the following sections, we evaluate the social comparison model and its implementation in modeling pedestrian movement. The basic movement pattern that our simulated pedestrians follow, stemming from the social comparison model, is as follows: Each agent follows an initially set direction. It chooses moving in this direction, unless blocked. If forward movement is indeed blocked, the agent can change lanes to the left or right. It will choose the lane where there is an agent that is similar to it (if such an agent exists); otherwise, it chooses arbitrarily.

4.1 Experiment 1: Independent Pedestrians

Our first experiment contrasted the social comparison model with previous models. We begin by examining individual pedestrian movements, where each synthetic pedestrian is independent of others. Each agent had a unique color. Each agent's direction (east or west) and initial position was chosen randomly. We contrasted the social comparison model with that of individual choice, which was shown to produce lane formations [35, 14, 15] and is considered to be a base model for pedestrian models.

For the purpose of this experiment, we fixed the gain component at 1 (see below for experiments examining gain). S_{max} was set at 6 (which means any dissimilarity other than color triggers action, and S_{min} was set at 2 (which means that agents that differed only in distance (but not by color or direction) were not considered similar. Each trial was executed for 5000 cycles.

Figure 4.3 shows lane changes for the individual-choice and social comparison models. The X-axis measures density. The Y-axis measures the number of lane changes during the course of 5000 cycles. Each configuration was repeated 30 times. Figure 4.4 measures flow for the two models; the X-axis again measures density, the Y-axis measures the flow.

The figures shows that the number of lane changes is significantly lower than that of the individual-choice model, implying that lanes form faster and are maintained longer with the social-comparison models. However, as the flow results show, there are no meaningful differences in flow. In other words, the social comparison model performs better, but with essentially no cost to the flow.

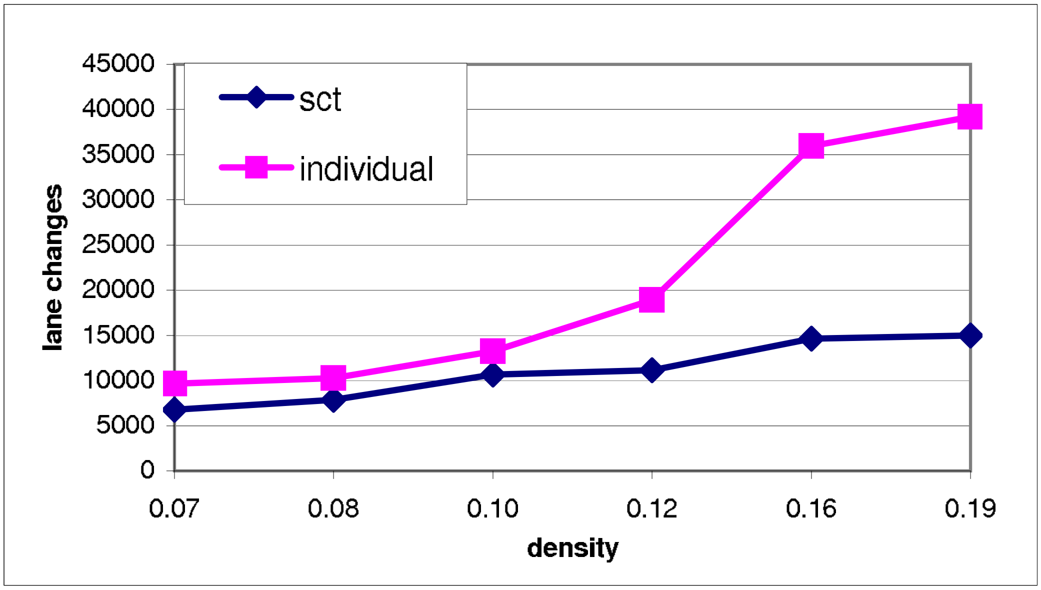

Figure 4.3: **Independent pedestrians' lane changes.**

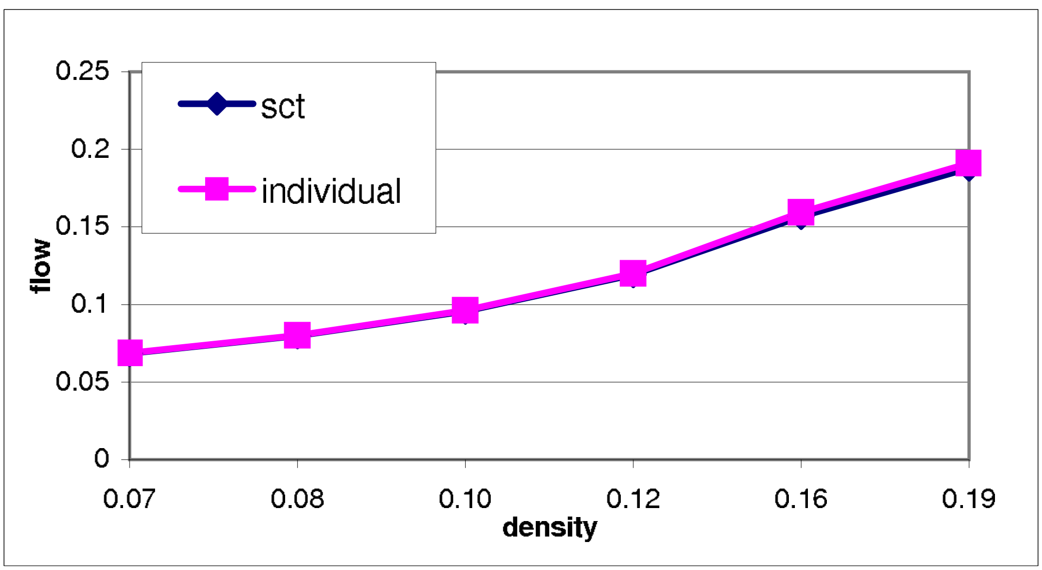

Figure 4.4: **Independent pedestrians' flow.**

4.2 Experiment 2: Independent Pedestrians with Varying Gain

The next set of experiments explored the performance of the model when the gain component was allowed to vary, per its definition in the social comparison model. We repeated the individual pedestrian experiments, though ignoring color: All agents moving east were colored red, and all agents moving west were colored blue. Because of this, agents really see only two kinds of agents: Those who have similarity of 1 (or less) , and those with similarity of 5 (or more). Thus the only way to vary the gain, is to vary the S_{min} and S_{max} values, as they set the enumerator in the gain calculation.

To evaluate the effect of the gain, we contrasted three variants of the social comparison model introduced earlier:

- $S_{max} = 5.5$, $S_{min} = 5$, i.e., $g(o) = 1$
- $S_{max} = 5.5$, $S_{min} = 4$, i.e., $g(o) = 3$
- $S_{max} = 5.5$, $S_{min} = 2$, i.e., $g(o) = 7$

Figure 4.5 shows the initial positions of the agents in one of the trials (4.5(a)), and typical results after 5000 cycles, with a gain of 1 (4.5(b)), gain of 3 (4.5(c)), and gain of 7 (4.5(d)). The figures show how the increased gain causes the agents to group more closely together.

Figures 4.6 and 4.7 show the lane-changes and flow in these experiments. The figures show that while again, there is no reduction in flow, there is significant improvement to the number of lane changes, as the gain increases.

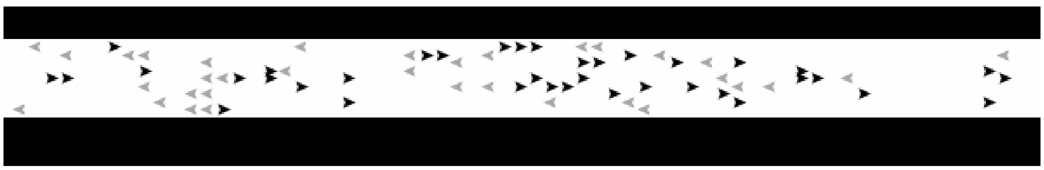

(a) Initial positions.

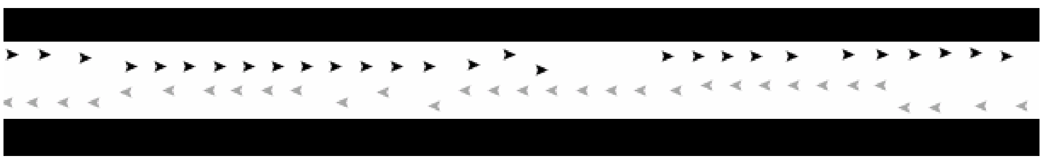

(b) $g(o) = 1$.

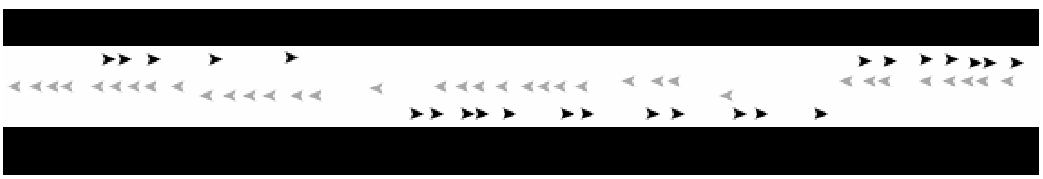

(c) $g(o) = 3$.

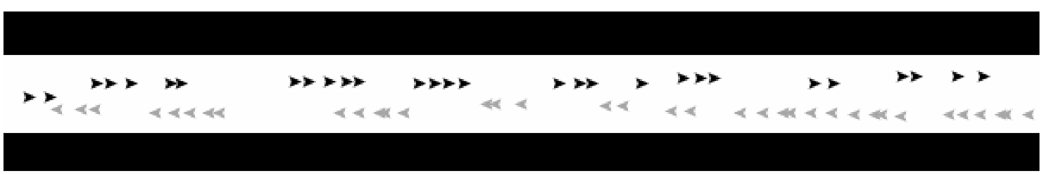

(d) $g(o) = 7$.

Figure 4.5: **Screen shots, Independent Pedestrians: Varied Gains.**

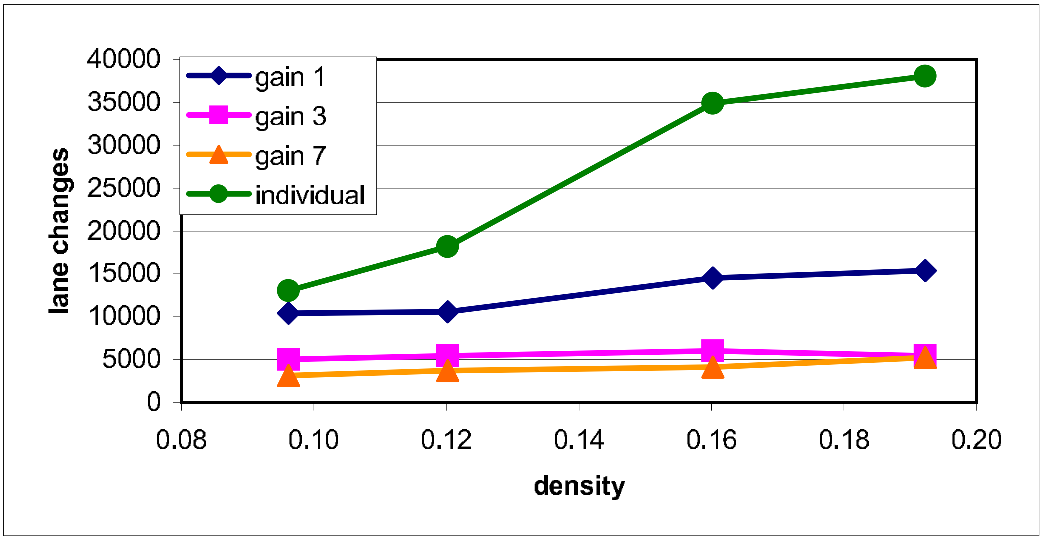

Figure 4.6: **Independent Pedestrians with Varying Gain: Lane Changes**

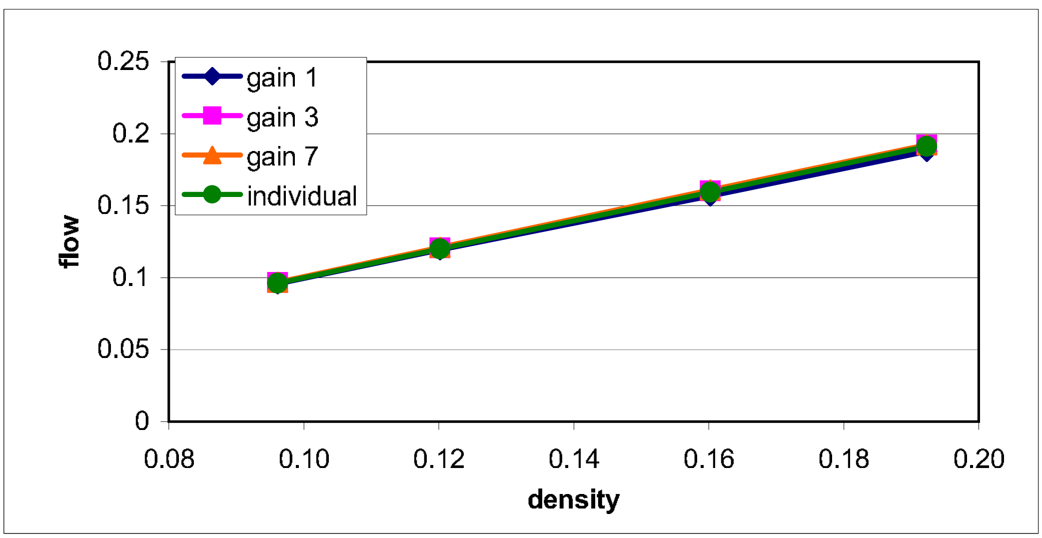

Figure 4.7: **Independent Pedestrians with Varying Gain: Flow**

4.3 Experiment 3: Pedestrians in Groups

We now move away from considering scenarios previously appearing in the literature, to exploring new types of movements. In particular, we experiment with pedestrian movement where the pedestrians belong to different groups internally. This type of situation arises, for instance, in pedestrians that are composed of families and/or friends. The individual-choice model does not account for such behavior, because it does not treat the group in any way. In contrast, we expect our social comparison model to treat groups (agents that belong to the same group would be more similar).

To examine this hypothesis, we carried out experiments in which color is meaningful: Agents belonging the same group have the same color. In these experiments, all agents move in the same direction, again, for 5000 cycles. Gain was allowed to vary per the model, as described above. The population contains 150 agents with a different number of colors (we experimented with 5, 10, and 20 color). Walking direction of all agents is West. S_{max} was set at 6.5, and S_{min} was set at 2.

To account for the western cultural intuition that friends (and family) walk side-by-side, rather than in columns, we added another feature: The similarity in position along the x-axis. The revised features and weights are as follows: **Direction**, with weight 2; **Distance**, with weight 0.5; **Color**, with weight 3; **X-Coordinate**, with weight 1.

The rationale behind these weights is that the agent will first close the distance gap with the agent selected as most similar, and only then try to locate itself on the same X-Coordinate.

There exists a significant challenge in being able to quantitatively measure the grouping results of the experiments. Normally, a simple clustering measure would do, as all agents of same color would group together. However, due to the initial random positions and the limited visual range of agents, agents of the same color may never group together, instead forming several groups that are far from each other.

Balch [3] has offered a clustering measure, *hierarchical social entropy*, that can address such cases. The key intuition behind this measure is to iteratively

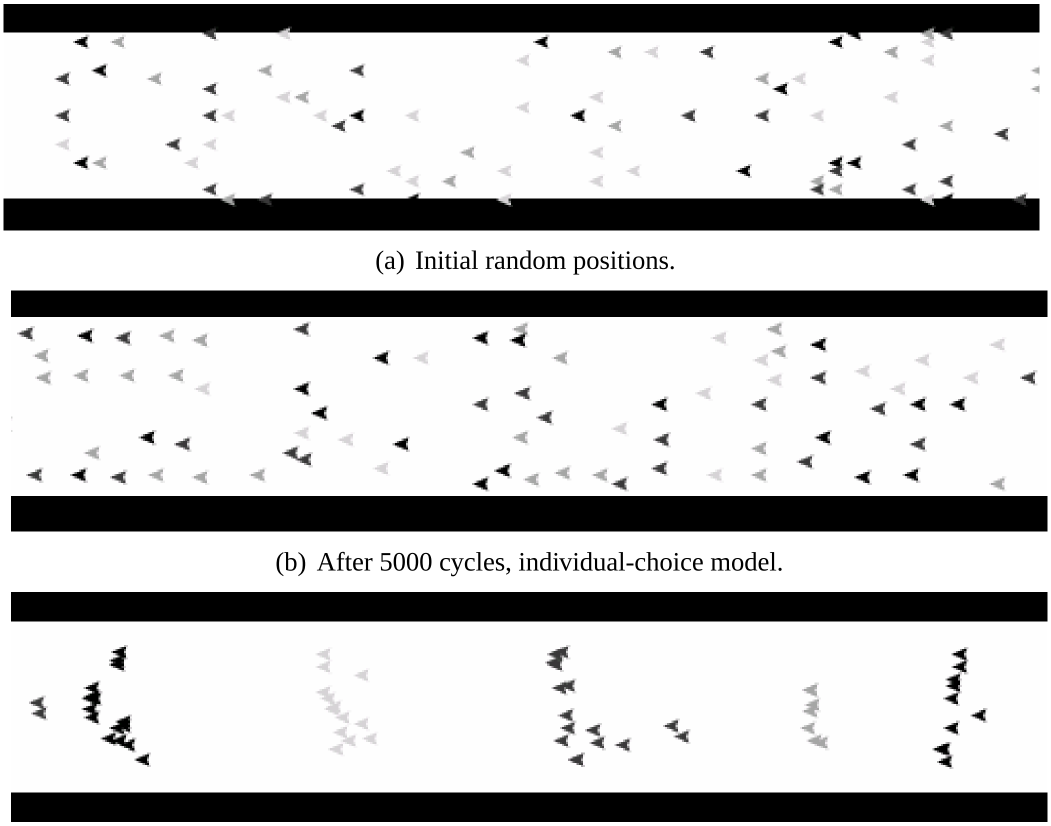

(a) Initial random positions.

(b) After 5000 cycles, individual-choice model.

(c) After 5000 cycles, social comparison model.

Figure 4.8: **Screen shots, Grouped Pedestrian Movement.**

sum entropy over increasing areas. The measure equals 0 when all agents are positioned in the exact same spot, and grows with their spreading around. Thus lower values indicate improved grouping. [3] provides the details.

Table 4.1 shows the hierarchical social entropy results for the individual-choice and social-comparison models. Each row corresponds to an experiment with a different number of colors. The table shows (third column) that the social comparison model provides for significantly improved grouping compared to the individual-choice model (one-tailed t-test, at $\alpha = 0.05$ significance level).

# Groups	Individual-Choice	Social Comparison
5	173.2	**87.4**
10	143.3	**85.8**
20	101.5	**60.1**

Table 4.1: **Grouping measurements of individual-choice and social comparison models. Lower values indicate improved grouping.**

4.4 Experiment 4: Groups and Obstacles

Our final set of pedestrian movement experiments addresses the response of groups within moving pedestrian crowds to obstacles. Intuitively, we recognize that such groups will choose to stick together in face of an obstacle (moving together to one side of it), while individual-choice pedestrians choose arbitrarily. We sought to examine whether the social comparison models would account for this behavior.

We created a sidewalk environment as described earlier, but this time with an elongated rectangular obstacle in the middle of it. When agents approach this obstacle, they must select to move to one of its sides. In the experiments, we allowed 100 agents of two colors (red and blue) to move west from their initial positions. Each agent has the following features: Direction, distance and color (weights: same as in the individual pedestrian experiments). Agents use comparison at all times, and not just when stuck. S_{max} was set at 6.5, S_{min} at 3.

Figure 4.9 shows the initial random positions of the agents (4.9(a)), their positions after going moving for a while using the individual-choice model (4.9(b), and their positions when moving using the social comparison model (4.9(c)). The figures show clearly that the the social comparison model causes similarly-colored agents to group together on one side of the obstacle, passing it together. In contrast, the individual-choice model has no such effect on the behavior of the agents.

Quantitative analysis again proved challenging, as here no clusters form. We needed, instead, to measure to what degree agents of the same color stay on one side of the obstacle. To do this, we defined virtual "gates" on either side of the obstacle, and monitored agents that move through them. Each trial allowed 100 agents to pass through the gates 10 times (i.e., 10 *waves*). At the end of each wave, we calculated (separately) the entropy of each color as its agents are divided between the two gates. A score of 0 indicates perfect grouping (all agents of same color pass through same gate). A score of 1 indicates perfect lack of grouping (the agents are evenly split between the two groups). The final result of each wave is the average entropy value across the two colors.

Figure 4.10 shows the average entropy value for each wave, for the ten waves. The results are averaged over 25 trials. The X-axis shows the wave number (1–10). The Y-axis measures the entropy. The figure shows that the entropy value of

(a) Initial random positions.

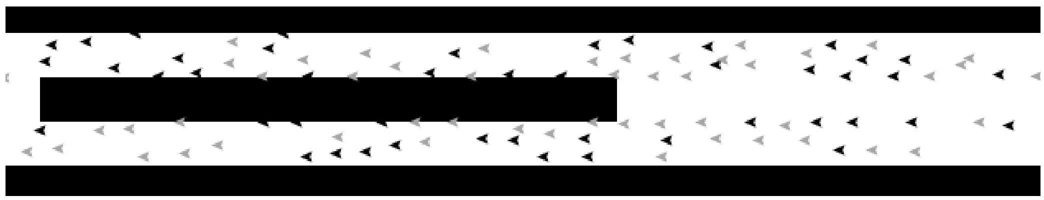

(b) Final positions, with individual-choice model.

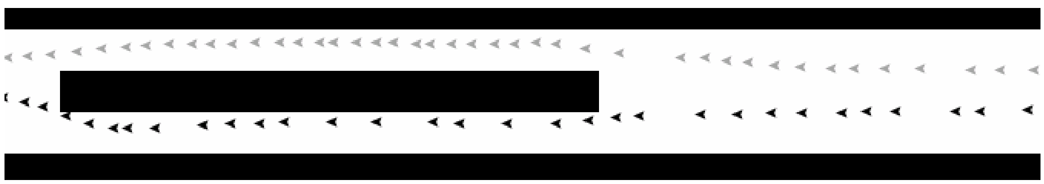

(c) Final positions, with social-comparison model.

Figure 4.9: **Screen shots, Grouped pedestrians' movement around the obstacle.**

the social-comparison model quickly goes down from 1 and approaches 0, while it remains around 1 for the individual-choice model. Indeed, after 10 waves, the average entropy value for the social comparison model is 0.131, while it is 0.992 for the individual-choice model.

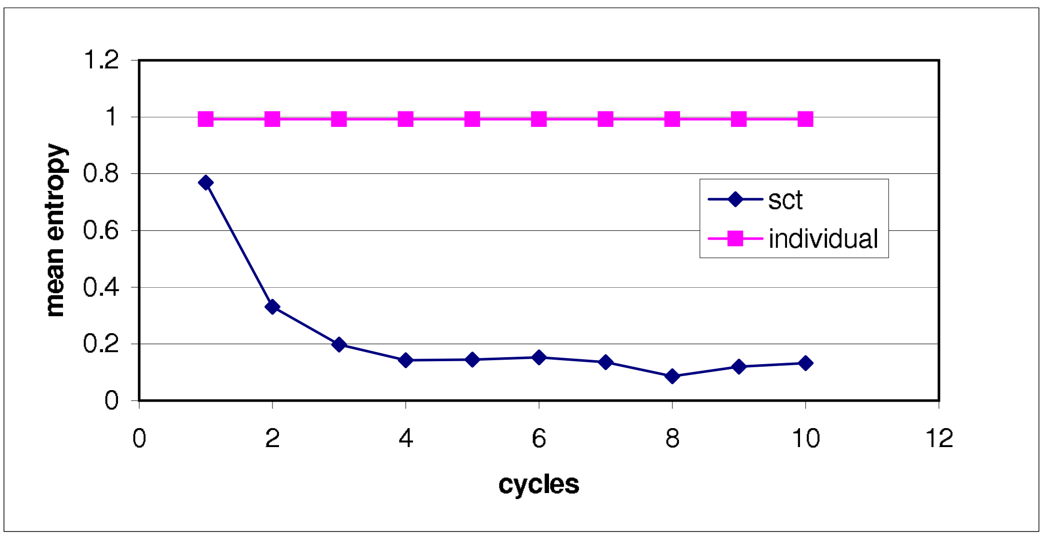

Figure 4.10: **Entropy of grouped pedestrians' movement around the obstacle.**

Chapter 5

Modeling Gathering Behavior

One of the common forms of collective interactions is gathering around a target location [36, 1, 8, 7]. Such gatherings are not necessarily organized in such a manner that all participants know the exact place and time. For example, in response to some significant event, there may be spontaneous demonstrations where most participants are uninformed of the exact location and may not even plan to become a part of such an event. According to Wright [36], such group forms occur with almost no verbal communication. Wright noted that "group forms are not only a product of collective interactions, but they also serve as a medium of nonverbal communication in those interactions."

The gathering problem, has been also extensively studied in autonomous multi robot systems. In particular, gathering at a point that is not fixed in advance[1, 8, 7]. The main motivation behind these studies is to identify the minimal capabilities that each robot should have in order to produce collective gathering regardless of specific target location. Thus, the research interest is on robots that are weak and simple in the sense that they have no common knowledge, no common coordinate system and no direct communication.

In this section, we focus on spontaneous gatherings around a target location, where all participants are uninformed about the location. Similar to autonomous robotic characteristics [1, 8, 7], we consider our agents to be without common knowledge and direct communication. However, unlike in [1, 8, 7] where the gathering should occur regardless of specific target location, in this research the target location is fixed in advance and recognizable to the agents if it is in their sensing range. We define one agent (leader), which is fixed in space, to define the target location. We explore whether the SCT model is able to account for such gathering behavior; in particular, whether all SCT model agents are able to reach the leader's location without previous information about it. We limited the agents' field of view, so only those closest to the leader are able to see it, thus enabling us to explore whether SCT model agents are able to find the leader.

In the simulation of gathering behavior, we used the same features set and their weights as in grouping pedestrian movements (for further information see section 4.3). Each agent has the following features and their corresponding weights:

- **Direction**, with weight 2
- **Distance**, with weight 0.5
- **Color**, with weight 3
- **X-Coordinate**, with weight 1

The features of similarity (f_i) and the gain value ($g(o)$) are calculated in the same way as in pedestrian movement model:

- $f_{direction} = 1$ if direction is the same, 0 otherwise.
- $f_{distance} = \frac{1}{dist}$, where $dist$ is the Euclidean distance between the positions of the agents.
- $f_{color} = 1$ if color is the same, 0 otherwise.
- $f_{x-coordinate} = 1$ if x coordinate is the same, 0 otherwise.

The gain value ($g(o)$) is relevant only for distance feature, for other features (which are binary), the gain is ignored.

To implement the model for gathering behavior experiments, we used NetLogo [34]. We define a terrain in a rectangular shape, with 64 patches in the width and 38 patches in the length. In this terrain, all agents are able to move freely to any direction within the boundaries. We placed a leader in the middle of the terrain and the agent's task was to gather around the leader. As in the pedestrian movement model, each agent has a limited field of view. In this simulation, each agent has a cone-shaped field of view of 120 degrees. The range of the field of view was limited to 5 patches, which is considered to be a low value compared with terrain size. Therefore, most of the agents are not able to see the leader in their initial positions.

Figure 5.1 shows the NetLogo terrain environment in an initial state where simulated agents are randomly placed about and have a random direction. Figure 5.2 shows an end result from one of the experiments (described below), where agents have finished gathering.

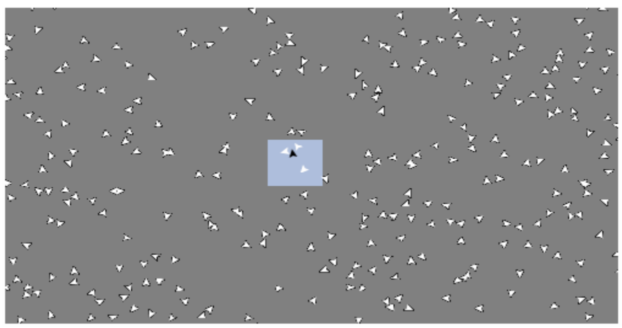

Figure 5.1: **Initial NetLogo sidewalk.**

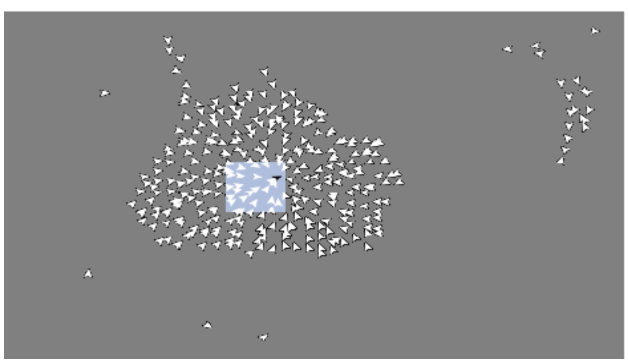

Figure 5.2: **Gathering around the leader - Experiment end-result.**

To gather around the leader, our current SCT model should be expanded to be able to simulate crowd behaviors under leadership presence and influence. There should be two main expansions. First, an agent needs to be able to calculate greater similarity value to a leader than to other agents. Another expansion is in choosing a corrective action. The agent should be able to take a different corrective action when compared with a leader than with other agents. For example, in gathering around the leader behavior, if the chosen feature for correction is direction, then the corrective action should be turning towards the same angle as the selected agent. However, if the selected agent is a leader, then the corrective action should be to face the leader (in this behavior). In general, the issue of leadership and its interactions with the SCT model is an open research issue, that we hope to investigate in more depth in future work. Here, we limit ourselves to this simple model of leadership as changing the preference of the selected corrective action.

We focus on the ability of agents to reach the leader without previous information about its location. Thus, in our simulation of gathering behavior we made a few assumptions. First, in order to be more visible to viewers, we presented the leader using a different color than for other agents. In the current simulation, color is the only feature that differentiates the leader from other agents. Therefore, with the current calculation formula of similarity value, agents calculate greater similarity to other agents with the same color than to the leader. Thus, by giving preference to leader's color, our agents were able to prefer the leader. Another assumption was in the ability of agents to face the leader. We defined a blue square shape around the leader (Figure 5.1). Each agent who reached this square box would turn towards the leader. These assumptions enabled the agents to gather around the leader.

We compared the SCT model with the individual knowledge model (i.e., each agent makes decision independent of its peers) and the full knowledge model (i.e., each agent knows the position of the leader). In the individual knowledge model, if an agent sees the leader, it moves toward it; otherwise, it searches for the leader by a random walk which means that it moves straight until it is blocked and then rotates at a random angle. In the full knowledge model, all agents know the location of the leader and move toward it. In the SCT model, agents move according to our algorithm. In these experiments, we controlled the crowd density and measured their clustering around the leader. Our hypothesis is that the SCT model will be ranked somewhere in-between the random and full knowledge model. However, we can improve the results of the SCT model by inserting full knowledge agents. In section 5.2 we will examine the impact of such agents on the SCT model results.

5.1 Gathering results

We first evaluated the SCT model in accounting for the gathering behavior. In particular, we wanted to evaluate whether the SCT model agents would be able to find the target location without previous information about it. We believed that those agents closest to the leader and able to see it would "pull" the others. Thus, eventually most of SCT model agents would find the gathering location, but not as quickly as it could be when knowing that location in advance. Therefore, our hypothesis was that the SCT model would be ranked somewhere in-between the individual knowledge model and the full knowledge model.

In this experiment all agents had the same color except for the leader. We placed the leader in the middle of the terrain and all the other agents were initialized with a random position and with random direction. We limited the agents' field of view in a way that only the agents closest to the leader were able to see it. The sample population comprised 350 agents. In this experiment, agents used comparison all the time, and not just when stuck. S_{max} was set at 7 (which means that there was no such thing as too similar), and S_{min} was set at 2 (which means that agents that differed only in distance were not considered similar). Each trial was executed for 2000 cycles.

Figure 5.3 shows the screen shots of gathering experiment results. Figures show the initial positions of the agents in one of the trials (5.3(a)), their positions after moving 2000 cycles using the individual knowledge model (5.3(b)), their positions after 2000 cycles using the SCT model (5.3(c)), and their positions after 2000 cycles using the full knowledge model (5.3(d)). The figures show that social comparison model agents are able to gather around a target location.

Figure 5.4 summarizes the measurement results of the models. The X-axis corresponds to the cycle number (1–2000) and the Y-axis corresponds to mean distance between all agents and the leader in that cycle. Each data-point is an average over 10 trials.

The results clearly demonstrate that the SCT model results lie in-between the

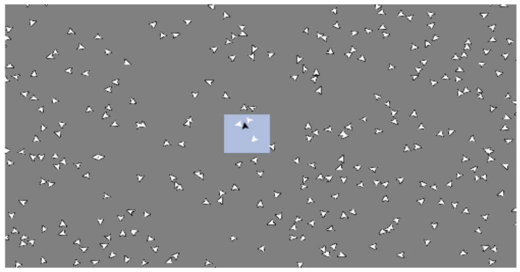

(a) Initial positions.

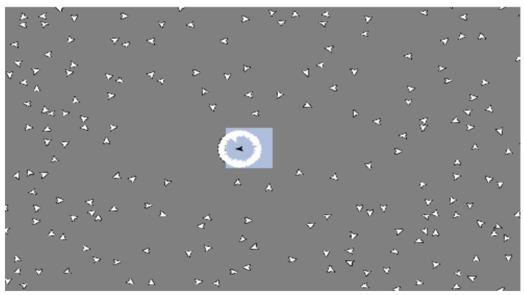

(b) Individual choice model

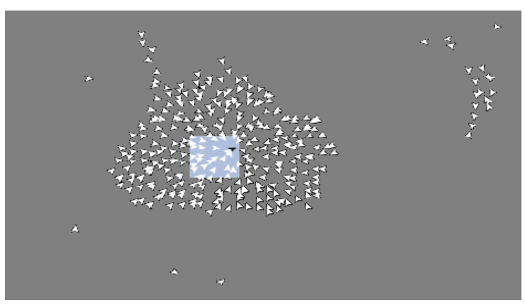

(c) SCT model

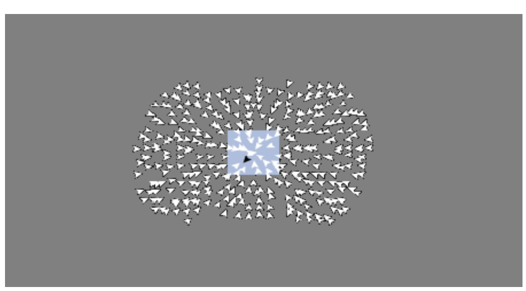

(d) Leader know model

Figure 5.3: **Screen shots, Gathering behavior.**

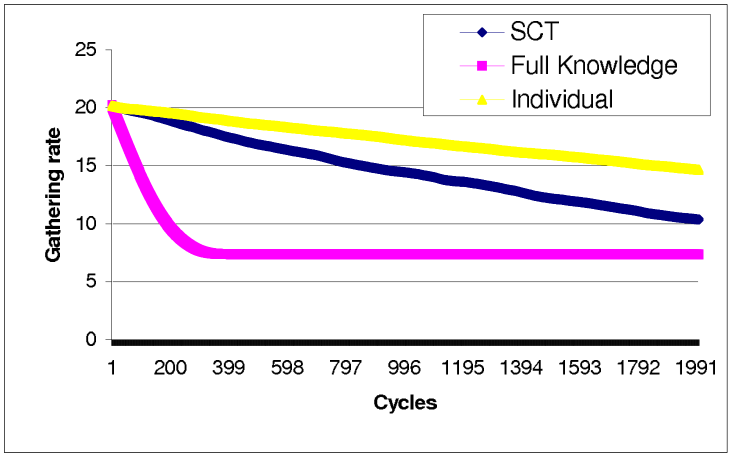

Figure 5.4: **Gathering results**

individual and the full knowledge model. In the full knowledge model, the gathering rate is faster, as agents quickly converge to the minimum distance, which remains stable. In the SCT model, the gathering rate is slower, as agents converge more gradually, but eventually get close to the minimum distance.

5.2 The effects of objective knowledge

The hypothesis underlying this experiment is that social comparison process can compensate for agent's lack of objective knowledge. According to Festinger: "To the extent that objective, non-social means are not available, people evaluate their opinions and abilities by comparison respectively with the opinions and abilities of others". Thus, as social psychology predicts, in some sense, the social comparison process may be considered a replacement for objective knowledge. In this section we show that this is true (at least in this domain) by showing that the gathering behavior of a population that is composed of relatively small percentage of agents that know the target location, and others that use SCT, is essentially equivalent to that of a population composed strictly of full-knowledge agents.

In this experiment, we evaluated whether we could improve the SCT model results. We wanted to explore whether we could increase the SCT model gathering convergence rate to the full knowledge model results. Thus, we added the full knowledge agents to the population of the SCT model agents. The full knowledge agents knew the position of the leader and moved towards it from the beginning of simulation, while our agents used the SCT model to find the leader. Our hypothesis was that by using social comparison in a population with agents who know the position of the leader, there would be improvement in the SCT model results.

To examine this hypothesis, we carried out experiments where a population of 350 agents contained different numbers of full knowledge agents. We experimented with 10, 20, 40, 60 and 80 numbers of such agents. All agents (except the leader) had the same color; therefore, the SCT model agents did not know which of the agents were full knowledge agents. The SCT model agents used social comparison all the time, and not just when stuck. S_{max} was set at 7 and S_{min} was set at 2. Each trial was executed for 2000 cycles.

Figure 5.5 shows the experiment results for the SCT model with a different number of full knowledge agents. The X-axis corresponds to the cycle number (1–2000) and the Y-axis corresponds to the mean distance between all agents and the leader in that cycle. Each data-point is an average over 10 trials.

The results clearly demonstrate that an increased number of full knowledge agents causes the gathering to more quickly converge to the minimum distance.

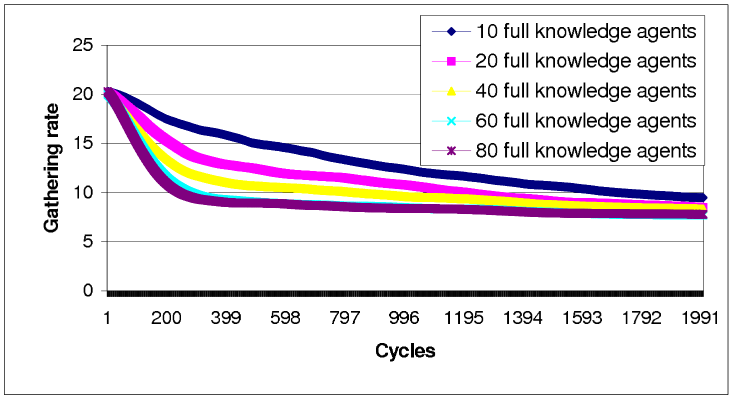

Figure 5.5: **SCT model improvement results**

Full knowledge model results	Correlation Coefficient
SCT with #10 full knowledge	0.689188815
SCT with #20 full knowledge	0.822569564
SCT with #40 full knowledge	0.918956756
SCT with #60 full knowledge	0.975423219
SCT with #80 full knowledge	0.984862296

Table 5.1: **Correlation Coefficient results between full knowledge model and Social Comparison model with different number of full knowledge agents.**

To explore whether and how strongly the gathering rate of the compared models are related, we calculated the correlation between them. Correlation coefficients were calculated between the full knowledge model and each of the SCT model results.

Table 5.1 summarized the correlation coefficient results. The results clearly demonstrate that with an increased number of full knowledge agents, the SCT model has a stronger relation to the full knowledge model. Moreover, there is a very strong correlation at 40 full knowledge agents number. Thus, if in population of 350 agents there will be only 40 full knowledge agents, the SCT model simulated behavior will be similar to the full knowledge model behavior.

Chapter 6

When is the SCT Process Triggered?

One of the open questions that we coped with in this thesis, is when should an agent trigger the social comparison process. Put differently, this is a question of implementing the SCT process at the agent architecture level.

In this section, we propose two approaches for such implementation. The first approach more directly follows Festinger's Social Comparison theory. According to Festinger, the SCT process is a response to uncertainty ("lack of objective means to evaluate opinions"). However, an alternative approach is to view the SCT process as an on-going process at the architecture level. Such a view requires that agent will be monitoring other agents all the time, and is thus more computationally expensive. It also seems to contrast with Festinger's theory. We will explore these approaches in the context of crowd behavior simulations. We conducted a set of experiments that will show the preference of one approach over the other. In particular, we show that the second view is correct.

It may appear easy to dismiss the implementation question as insignificant. However, the implementation choice carries significant implication: As SCT processes inherently rely on knowing about the behavior of others, the implementation question raises a more fundamental question about where modeling of others (e.g., using plan recognition) occurs in cognition: Is it a problem-solving activity, or is it carried out all the time, at an architectural level.

6.1 SCT as a Problem-Solving Activity

According to Festinger, people use social comparison when they have a lack of knowledge to make their decisions. Thus one way of implementing the SCT process in an architecture level is as a response to an uncertainty: When an agent is at an uncertain state, it may call on a comparison process that will be used to assess similarity and propose actions.

We thus may treat the social comparison theory as a new kind of uncertainty-resolution method. Unlike previous uncertainty-resolution (problem-solving) techniques, in which the agent focuses on using its own resources, here the agent uses knowledge of others as a basis for resolving the uncertainty.

We believe that treating of social comparison processes as generic uncertainty-resolution methods raise novel questions as to the role of social reasoning in cognitive architectures. Most cognitive architectures do not commit to social processes being a part of the architecture. Instead, most social reasoning is done by manipulating knowledge and beliefs as part of a task. This view is quite common in robotics and agent literature, which often treats reasoning about multiple agents as a process that is carried out at a higher, task-dependent, level of reasoning.

If, however, the continuous-monitoring view of social comparison is correct, then this implies that cognitive architectures must somehow specialize to cover rudimentary social reasoning at an architectural level. In particular, for social comparison processes to be possible, the architecture itself must distinguish between inputs that describe other agents from those that describe objects or features in the environment. Without such a distinction, any reasoning will necessarily be limited to where prior knowledge distinguishes the agents from other knowledge.

6.2 SCT as an On-Going process

Hakmiller [12] and Singer [28] expand Festinger's theory about when people use social comparison by adding another case. According to their theory people tend to confirm or reassure that their actions or beliefs are the correct ones. For this confirmation they use social comparison and such comparison is called constructed comparison. According to this approach people tend to use social comparison in parallel to their decision making process.

Thus, an alternative is to view the SCT as an on-going process, at an architectural level *in parallel* to any problem-solving activity. Whereas normally, actions are proposed (and selected) by cognitive architecture based on their suitability for a current goal (e.g., through means-end analysis), in our agent actions were also proposed based on their suitability for SCT. In other words, agent would consider actions that advance it towards its goal. In our implementation, it would also consider actions that seek to minimize perceived differences to other agents.

This may appear to contradict Festinger's theory that social comparison comes into play only when people are in an uncertain state. However, this is not the case. By preferring the SCT-proposed actions only when no task-oriented actions are available (i.e., in an uncertain state), one gets the behavior predicted by Festinger's theory. Further exploration of this issue is beyond the scope of this thesis.

6.3 Comparison of these approaches in regard to crowd behavior modeling.

In the following section, we argue that implementation of SCT as an on-going process is more suitable for modeling crowd behavior than SCT as a problem-solving activity process approach. We conducted a set of experiments to evaluate which of these two approaches is more applicable in the context of crowd behavior simulations. We examined these approaches in reference to pedestrian grouping behavior and to gathering behavior.

We used the same experiment framework for pedestrian grouping behavior as described in section 4.3 and for gathering behavior as described in section 5. In these experiments we compare between two types of SCT agents: the SCT-Problem-Solving agent and the SCT-On-Going agent. These agents have the same feature set and perform the SCT process in the same way as described previously in Chapters 4.3 and 5. However, the main difference between them is in *when* the SCT process is activated. While in a SCT-Problem-Solving agent the SCT process is activated when an agent is in an uncertain state, in a SCT-On-Going agent the SCT process is activated all the time.

SCT-On-Going approach	SCT-Problem-Solving approach
87.4	167.3

Table 6.1: **Grouping measurements of SCT-On-Going approach and SCT-Problem-Solving approach. Lower values indicate improved grouping.**

The question, what is an uncertain state for an agent is open and broad. We leave the discussion about different approaches in that field outside of this thesis. In our implementation of modeling different types of movements (pedestrian and gathering movement), the SCT-Problem-Solving agent initially follows a set direction which is straight. If forward movement is blocked, the agent will be considered to be in an uncertain state where it has to choose between alternative directions. In this state the SCT process will be activated and the agent will choose the lane based on the SCT algorithm.

Another open question that we also leave outside this thesis, is when an agent should ignore social comparison proposed actions and prefer the actions that promote it to its current goal. In the implementation of gathering behavior, the SCT-On-Going agent prefers the goal oriented actions (move toward the leader) when it sees the leader, in other cases the agent prefers the SCT actions. In the implementation of grouping pedestrian behavior, the SCT-On-Going agent always prefers the SCT actions.

Figure 6.1 shows pedestrian grouping behavior results. Figures show the initial positions of the agents in one of the trials 6.1(a), their positions after moving 5000 cycles using the SCT-On-Going approach 6.1(b) and their positions after 5000 cycles using the SCT-Problem-Solving approach 6.1(c). The figures show that the SCT-On-Going approach accounts for grouping behavior while the SCT-Problem-Solving approach provides behavior similar to the individual model (see section 4.3).

We use hierarchical social entropy for measuring the grouping behavior (see Chapter 4.3). Table 6.1 shows the measurement results for the SCT-On-Going approach and for the SCT-Problem-Solving approach. We experimented with population that compose from 5 different colors. The table shows that the SCT-On-Going approach provides improved grouping compared to the SCT-Problem-Solving approach.

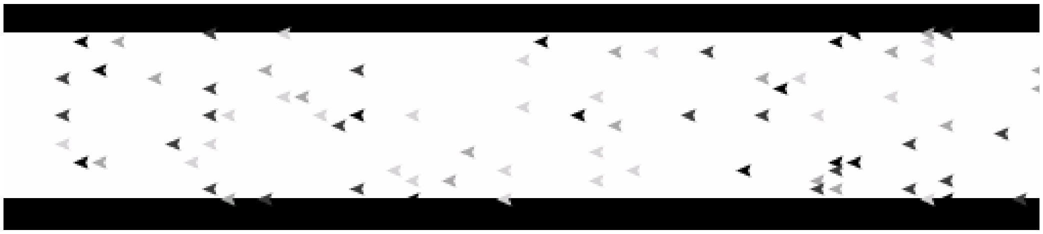

(a) Initial random positions.

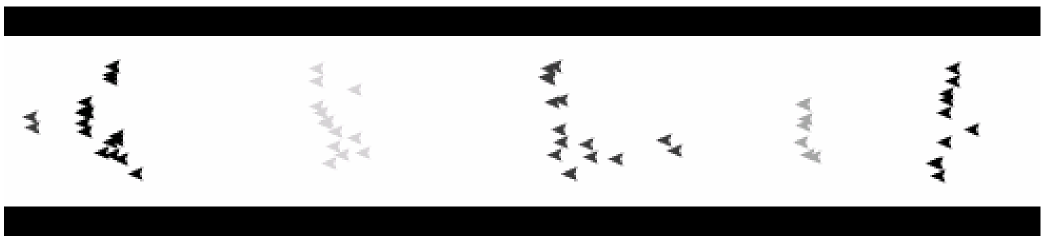

(b) After 5000 cycles, SCT-On-Going approach.

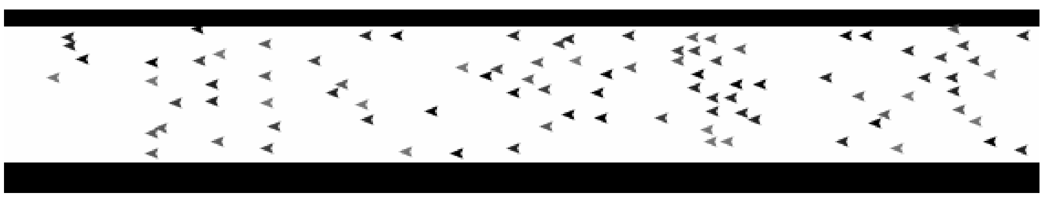

(c) After 5000 cycles, SCT-Problem-Solving approach.

Figure 6.1: **Screen shots, Comparison of Implementation approaches in regard to Grouped Pedestrian Movement.**

We now present the gathering behavior results. Figure 6.2 shows the screen shots of the gathering experiments. The figures show the initial positions of the agents in one of the trials 6.2(a), their positions after moving 5000 cycles using SCT-On-Going approach 6.2(b) and their positions after 5000 cycles using SCT-Problem-Solving approach 6.2(c). Again, the figures show that the SCT-On-Going approach is accounting for better gathering behavior than the SCT-Problem-Solving approach. However, as shown in the pictures, the SCT-Problem-Solving approach causes agents to divide into small separate clusters that move around the leader and not necessarily are pulled towards it. As opposed to grouping behavior where the SCT-Problem-Solving approach provide similar results to individual choice model, this provided behavior is not similar to individual choice model behavior where only the closest agents to the leader are gather around it, while the others are scattered in the terrain (section 5).

Figure 6.3 summarized the measurement results of the models. The category in the X-axis correspond to cycle number (1–2000) and the Y-axis correspond to mean distance between all agents to the leader in that cycle.

The results demonstrate that the gathering rate of the SCT-On-Going approach is better than the SCT-Problem-Solving approach. Moreover, in comparison to individual choice model (section 5) the SCT-Problem-Solving approach provide inferior results. The main reason for this distinction, is while the individual choice agents are scattered equable in the terrain, the SCT-Problem-Solving agents are gathered in small groups, distant from the leader.

We thus argue that implementation of SCT as an on-going process is more suitable for modeling crowd behavior than SCT as a problem-solving activity process approach. Moreover, it is more general approach, in the term that it also compatible with Festinger's theory. We thus advocate using this approach in modeling crowd behaviors.

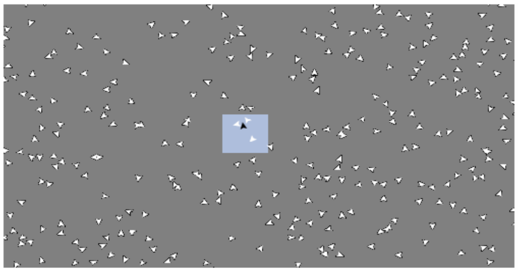

(a) Initial random positions.

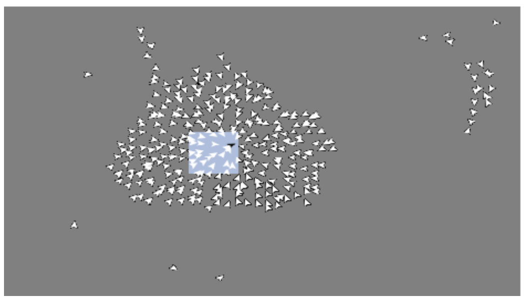

(b) After 5000 cycles, SCT-On-Going approach.

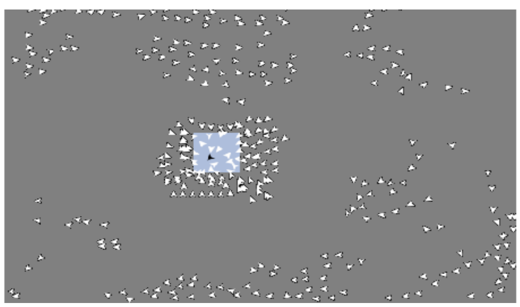

(c) After 5000 cycles, SCT-Problem-Solving approach

Figure 6.2: **Screen shots, Comparison of Implementation approaches in regard to Gathering behavior.**

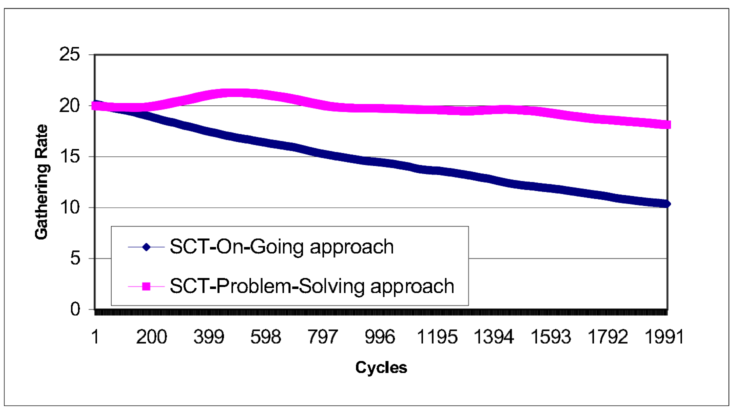

Figure 6.3: **Gathering results of SCT-On-Going approach and SCT-Problem-Solving approach**

Chapter 7

Implementations of SCT in Soar

We implemented SCT in the Soar cognitive architecture [23]. Soar was connected to the GameBots virtual environment [19]. Here, multiple agents, each controlled by a separate Soar process (each executing SCT) can interact with each other in a dynamic, complex, 3D virtual world (see Figure 7.1).

Figure 7.1: **Soar agents in the GameBots environment. Each agent has limited field of view and range, and may move about and turn.**

A detailed discussion of Soar's role as a cognitive architecture in implementing SCT is beyond the scope of this thesis. We provide a very brief overview here, and refer the interested reader to [23, 29] for additional details. Soar has two components: a graph-structured working memory, and a set of user-defined production rules that test and modify this memory. Efficient algorithms maintain the working memory by executing rules that match existing contents. All the agent's knowledge, sensor readings, and decisions are recorded in the working memory. Soar operates in a classic sense-think-act cycle, which includes a decision phase in which all relevant knowledge is brought to bear to propose, and then select, an *operator*, that will then carry out deliberate mental (and sometimes physical) actions. Once the operator finishes its actions, it is automatically de-selected (terminated), and the cycle repeats. Unlike simple production rules, whose effects on working memory are temporary, operator-induced actions on working memory (and in turn, on physical actions) are persistent, even after the operator has been de-selected. Overall, a Soar agent's behavior is the result of the sequential selection of operators, each performing an action on the environment and/or internal memory.

In section 6 we described two approaches for implementation of the SCT process in an architectural level. According to the first approach the SCT process is implemented as a response to an uncertainty. However, an alternative approach is to view the SCT process as an always-on process at the architecture level. Although we claimed that implementation of the SCT as an always-on process is more suitable for modeling crowd behavior, the SCT as a problem-solving activity process provide a novel approach for uncertainty-resolution techniques. Thus, in this section we present the implementation of both approaches in Soar.

7.1 Implementation of the SCT as Problem-Solving Activity in Soar

A key novelty in Soar is that it automatically recognizes situations in which the decision-phases is stumped, either because no operator is available for selection (*state no-change impasse*), or because conflicting alternatives are proposed (*operator tie impasse*). When impasses are detected, a subgoal is automatically created to resolve it. Thus, Social comparison theory as described by Festinger, seems to naturally fit Soar's impasse-driven operation. In particular, Festinger describes the trigger to using comparison as a situation in which people are unable to evaluate their opinions and capabilities, which seems to match an impasse situation.

We argued in 6.1 that SCT implementation as Problem-Solving Activity process may treat the social comparison theory as a new kind of uncertainty-resolution method. Thus, we can treat the SCT process as an alternative method for impasse-resolution (problem-solving) techniques. Our goal in this section is to determine a general way to describe social comparison processes in Soar, in such a way that they can be used for solving a wide variety of problems.

A snapshot from a log showing Soar using SCT implementation as Problem-Solving Activity process (here, to decide on movement) is shown below. Soar's decision cycles are denoted by numbers before colons. In the first and second decision cycles, operators called *init* and *explore-decision*, respectively, are selected by Soar. But then, more than 20 different instantiations of an operator called *elaborate-target* are proposed by the system; Soar is faced with the task of choosing one among them for execution. Since it cannot decide, an operator-tie impasse is declared; see the line marked

```
3:  ==> S: S3 (operator tie)
```

This triggers our social comparison process, which is carried out, in sequence, by the following operators: (i) *sct-init*, which sets up the new state, and copies relevant information. (ii) *sct-add-entities*, which copies information about other agents for use in ranking operators. *rank-item* then calculates a rank for all proposed operators, based on associated agents and their own choices. Finally *select-item* selects the highest-ranking operator and makes the decision. Indeed the last decision cycle (#8 in the log) shows a specific instance of the *elaborate-target* was chosen.

```
     1:    O: O2 (init)
root is active ->proposed child : explore-decision ->by : root
    2:    O: O4 (explore-decision)
->proposed child : elaborate-target ->by : explore-decision

->proposed child : elaborate-target ->by : explore-decision
```

[. . . . . 19 additional proposals for elaborate-target . . . .]

```
->proposed child : elaborate-target ->by : explore-decision

    3:    ==>S: S3 (operator tie)
    4:       O: O27 (sct-init)
    5:       O: O28 (sct-add-entities)
    6:       O: O51 (rank-item)

    7:       O: O68 (select-item)
SCT Done.   Chose  O21    Name:  elaborate-target
    8:    O: O21 (elaborate-target)
elaborate-target is active
```

[. . . . Soar continues . . . .]

We propose to use social comparison processes as generic impasse-resolution methods, that treat social reasoning as a problem-solving method.

7.2 Implementation of the SCT as an always-on process in Soar

A different way of treating SCT process as described in section 6.2 is as on-going process at the architecture level that activated in parallel to problem-solving activity. Thus, SCT was implemented as secondary parallel thread within Soar (Figure 7.2). At every cycle, operators are proposed (and selected) by Soar based on their suitability for a current goal (e.g., through means-end analysis), and also based on their suitability for SCT. Thus SCT-proposed operators compete with the task-oriented operators for control of the agent. By setting Soar's decision preferences to prefer the SCT-proposed operators, we get a very social agent. Conversely, by preferring the task-oriented operators, we get an individual choice agent which makes its decisions independently of its peers.

The SCT thread proposed operators by following the algorithm described previously in 3, though in a way that is adopted for Soar's decision cycle: At every

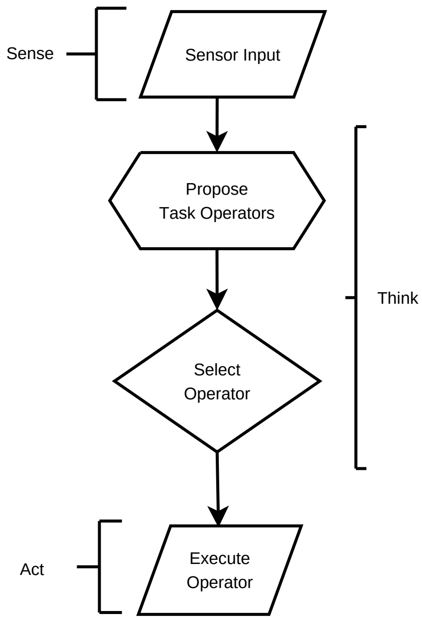

(a) The normal Soar sense-think-act decision cycle.

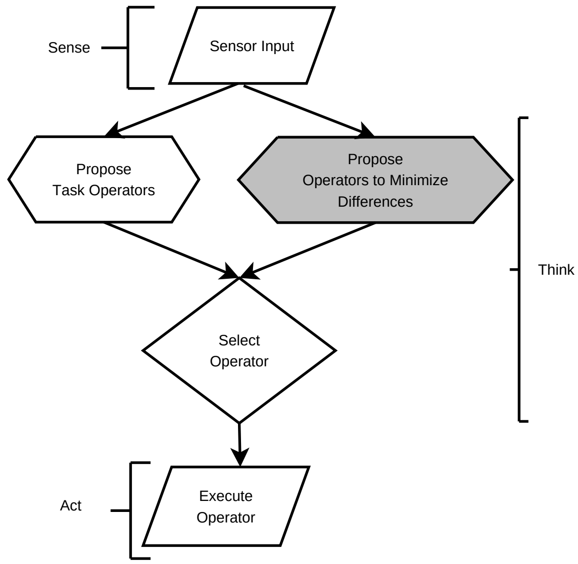

(b) Modified Soar decision-cycle. SCT process highlighted.

Figure 7.2: **The Soar sense-think-act decision cycle, SCT process highlighted.**

cycle, for each observed agent and for each difference, the SCT process would propose an operator that would minimize the difference. Then, a set of preference rules is triggered that ranks the proposals based on feature weight. Additional rules prefer the most similar agent (that is still not sufficiently similar). Thus at the end, only one SCT operator is supported.

Here an addition to the SCT model became necessary. Suppose an agent X decided to turn towards the same angle as an agent Y that is next to it. Due to the limited field-of-view of X, it would lose track of Y once it makes the turn. From that point on, it could no longer keep track of Y, to minimize additional differences. This would cause it to become overly reactive, turning about immediately to seek Y again, or to select a different operator altogether (now that Y could no longer be imitated).

We thus found it necessary to utilize two mechanisms: (i) a memory mechanism that keeps track of the whereabouts of agents, once seen; and (ii) an exploration mechanism that occasionally would turn towards remembered agents, to provide an update on their state (for the purpose of comparison). Both of these mechanisms (memory and exploration) are of course present in many cognitive architectures, and are not necessarily linked to SCT. We thus leave discussion of such mechanisms outside of this thesis.

We argued in section 6.3 that this implementation of SCT is better. It is compatible with Festinger's theory, and better accounts for pedestrian behavior in low-density areas. Thus, we will using this approach in Soar-based implementations of crowd modeling.

Chapter 8

Modeling Imitational Behavior

An attractive feature of social comparison is its hypothesized prevalence in human group behavior, i.e., its generality across different behaviors. Indeed, we believe that the SCT model we present in this thesis is sufficiently general to account for a wide variety of group behaviors. This is in contrast to many existing computational models, that typically focus on specific tasks.

This section provides additional evidence for such generality by describing the application of the SCT model to the problem of generating imitational behaviors in loosely-coupled groups. Unlike individual imitation, where one agent imitates a role model, crowd imitational behavior spreads across a group of individuals who dynamically select role models for imitation, from the level of observable actions to the level of unobservable internal mental attitudes (e.g., goals). Here, imitation occurs more loosely, as the role models do not necessarily intend to play their role, and indeed may not even know that they are being imitated. Also, the imitators potentially switch their role-model targets from one moment to the next. Psychology literature describes such imitational behavior as one of the keystones of crowd behaviors [21].

In order to simulate imitational behavior we used Soar cognitive architecture. The SCT process was implemented as an always-on process in Soar (for further information, see section 7.2). We used position and direction as the agents' features set. For each observed agent and for every difference found, the SCT process proposes a corrective operators to be performed in order to minimize the difference in selected feature. In this task the corrective operators were ’move-to’ (minimizing distance to the observed agent, correcting position differences) and ’turn-to’ (imitating angle of the observed agent).

In addition to the proposed SCT operators, Soar also proposes operators based on their suitability for the current goal, and based on an exploration mechanism which proposes operators seeking new information. In this task, goal operators were turn-to (a random angle); the exploration mechanism operators turned towards previously seen agents.

We used Soar preference rules to rank the features weights such that the position feature gets higher priority than direction. This means that a closest agent is consider to be more similar, however the chosen feature for correction is direction. The S_{max} value was unbounded, which means that there is no such thing as too similar. In our case Soar can propose corrective operator with value equal to zero if there is no correction to make with respect to the observed agent. We used additional Soar preference rules to give higher priority to exploration mechanism operators than to goal operators. Thus, each agent prefers the SCT operators (turn to) and in the case when there is no seen agents (i.e.. there is no proposed SCT turn-to operator) an agent will prefer the exploration mechanism operators, and only afterwards the goal operators. The resulting simulated behavior has the agents standing in their initial locations, turning to some direction or do nothing.

8.1 Evaluation of imitational behavior

We conducted experiments to evaluate whether SCT can indeed generalize to account for imitational behavior in groups. Unlike the pedestrian movement domain, where clear measures are available for objective measurement of a success of a model (e.g., flow, lane changes), imitational behavior does not have clear standards of evaluation.

We propose a method for evaluation of imitational behavior. We propose a questionnaire composed from general questions and specific tasks related questions. The general questions can be used as a common method for evaluation of all kinds of imitational behaviors. We rely on experiments with human subjects, which judged the human crowd behavior and the resulting SCT behavior in comparison to completely individual behavior (i.e., arbitrary decisions by each agent, independently of its peers), and to completely synchronized behavior (i.e., all agents act in complete unison).

The first hypothesis underlying the experiments was that groups controlled by SCT would generate behavior that would be ranked somewhere in-between the individual and perfect-coordination models, i.e., that SCT would generate behavior that would be perceived as coordinated, but not perfectly so. Another hypothesis is that human crowd behavior would also be ranked somewhere in-between the individual and perfect-coordinated behaviors.

To examine the first hypothesis, we created three screen-capture movies of 11 Soar agents in action. All movies were shot from the same point of view, and showed the agents in the same environment. In all screen-capture movies there is one blue agent that stands in a front and turn up to 90° left or right. All others are red agents that acted according to one of the models.

In one movie (*individual*), the red agents acted completely independently of each other, randomly choosing an angle and turning to it. In another (*unison*), the red agents acted in almost perfect coordination, turning towards the same angle as the blue agent almost instantaneously (small timing differences resulting from asynchronous responses of the simulated environment). Finally, in the *SCT* movie, the red agents acted according to our model as described above.

These experiments were carried out using 12 subjects (ages: 18–40, mean: 28; male: 6; additional 4 subjects dropped due to technical reasons). Each subject was given a brief description of the appearance of the environment and agents, sometimes aided by a snapshot from a movie (e.g., as in Figure 7.1). The subjects were told that the purpose of the experiment was to evaluate the use of perception models embedded in the agents; that there was a red dot—visible to the agents but not to the subject—that moves about on the walls surrounding the group. The agents' goal is to individually locate this dot, and then track it in place by turning around. The purpose of the cover story was to focus the attention of the subjects away from group behavior and imitation, so as to not bias the results. After the description, the movies were shown to the subject.

After each movie, the subjects were asked to fill a short questionnaire (described below) based on what they saw. Each movie was shown only once. The order of presentation of movies randomly selected for each subject, to control for learning and order effects. The questionnaire included the following questions:

1. If there is only one red dot in the room, to what degree did all agents see it? (1 - nobody saw the red dot; 6 - all agents saw it)

2. To what degree were the movements of the agents random? (1 - not random at all; 6 - very random)

3. To what degree was there cooperation between the agents? (1 - no cooperation at all; 6 - full cooperation)

4. To what degree was there agreement between the agents? (1 - no agreement at all; 6 - full agreement)

5. To what degree were the agents coordinated in terms of the direction of their movements? (1 -no coordination at all; 6 - fully coordinated)

6. How quickly did the agents find the red dot? (1 - dot not found at all; 6 - immediately found)

7. To what degree were the agents related to each other? (1 - no relation at all; 6 - tight relation)

8. Do you see any leaders? If so, how many? (1-11) (1- one leader; 11 - all agents are leaders, i.e., no leader).

In this experiment, the subjects were asked to grade the movies on an ordinal scale of 1–6, with 1 being a low score (typically associated with more individual behavior), and 6 being a high score (typically associated with perfect unison). In order to keep consistency in presentation of results, the scale of second question (Non-Random) was reversed. The results of the last question (Number of leaders) are presented separately due to inconsistency in scale with other questions.

8.1.1 Agents results

In general, the responses to the questions in this experiment have placed SCT between the individual and unison models. Results are summarized in Figure 8.1(a) and 8.1(b). The questions in Figure 8.1(a) are associated with agents performance on a given task. In presented questionnaire the number of questions are 1, 3, 4 and 6. Figure 8.1(b) refer to more general questions (i.e.. same questions that were used in human crowd movie). In questionnaire the relevant number of questions are 2, 5, and 7. The categories in the X-axis correspond to questions given to the subjects. The Y-axis measures the median result. Each bar correspond to compared model and as explained above we compare SCT model to Individual and Unison models.

The results clearly demonstrate that the SCT model lies in between the individual and perfect-unison model. While in some questions it appears to be somewhat closer to the individual model, it is significantly different from it at the $\alpha = 0.05$ significance level (t-test, one-tailed).

Figure 8.2 shows the results for the question on the number of leaders. The median result for the individual was 11 (i.e., every agent is a leader, or in other words, no leader). For the unison model, the median result was 1. For the SCT model, the median result was 3. In this question the SCT model result is very close to the Unison model. According to t-test (one-tailed) the SCT model significantly different than the Individual model $p = 0.02$. However, in comparison to Unison model there is no significance found ($p = 0.3$).

When we asked the subjects to qualitatively discuss their answer to these questions, many subjects reported on feeling that agents in the movie were organized in several subgroups, that were internally coherent, but not coordinated with the others.

8.1.2 Snapshots Experiment

The following section provides another examination of our first hypothesis that groups controlled by SCT would be ranked somewhere in-between the individual and perfect-coordination models. However, as opposed to screen-capture movies where the subject answered the questionnaire after seing each movie only once,

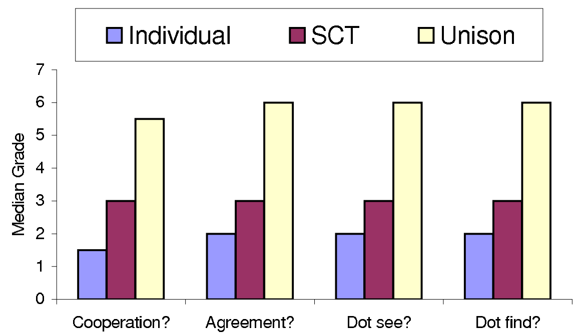

(a)

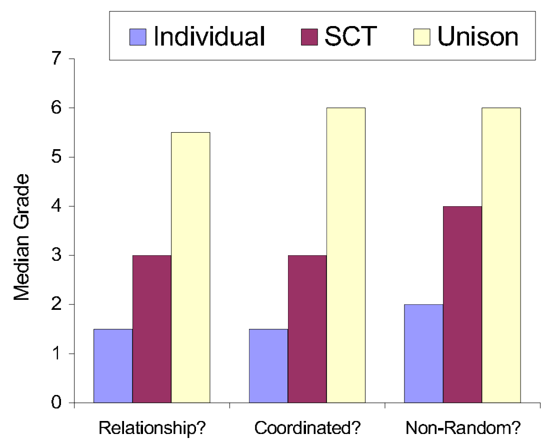

(b)

Figure 8.1: **Results of questionnaire on agents performance.**

in this experiment the subjects answered the questionnaire while watching the snapshot.

We took three snapshots from each movie. The snapshots were taken at same time slot from the beginning, middle and the end of the movies. After filling the screen-capture movies questionnaire, the snapshots from the movies were shown to the subjects. After each snapshot, the subjects were asked to fill a different questionnaire than in screen-capture movies. Like in screen-capture movies, the order of presentation was randomly selected for each subject, to control for learning and ordering effects.

The question presented to the subjects after watching each snapshot was: If

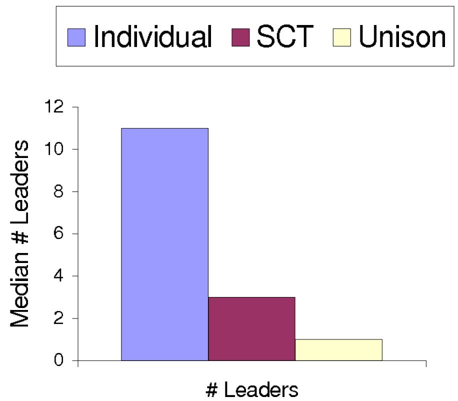

Figure 8.2: **Number of leaders in screen-capture movies**

there can be more than one red dot in the room and each agent look at it, how many red dots in the room? The subjects were asked to grade each snapshot on an ordinal scale of 1–11 with 1 being a low result (i.e., all agents look at the same red dot) and 11 being a high result (i.e., the number of red dots as the number of agents meaning all agents look at different dots).

The results of this experiment are summarized in Figure 8.3. Again the categories in the X-axis correspond to question given to the subjects. The Y-axis measures the average of median results that belong to same model.

Again the results demonstrate that the SCT model lies in between the individual and perfect-unison model and it significantly different from the individual model ($p = 0.011$, t-test, one-tailed) and from perfect-unison model ($p = 0.012$, t-test, one-tailed).

8.1.3 Human crowd experiment

Another hypothesis underlying the experiments is that human crowd behavior would also be ranked somewhere in-between the individual and unison models. To examine this, we search for a human crowd movie where individuals perform the same action as in simulated agents movies. We used a news clip movie where gathered people are standing and waiting for some event to occur, and the only

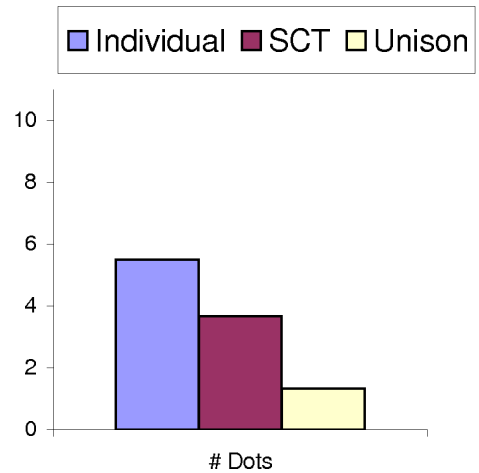

Figure 8.3: **Snapshots results.**

action they perform was turning to some direction.

This experiment was carried out using 12 subjects different than in the screen-capture movies experiments. Each subject, after viewing a human crowd movie (Figure 8.1.3) was asked to fill the same questionnaire as in previous experiments. However, since in human crowd movie there was no cover story about red dot, there were some irrelevant questions that were dropped out. The remaining questions are more general and not tied to specific task.

Figure 8.4: **Human crowd - clip movie.**

Results are summarized in Figure 8.5. As in previous results, the categories in the X-axis correspond to questions given to the subjects and the Y-axis measures the median result.

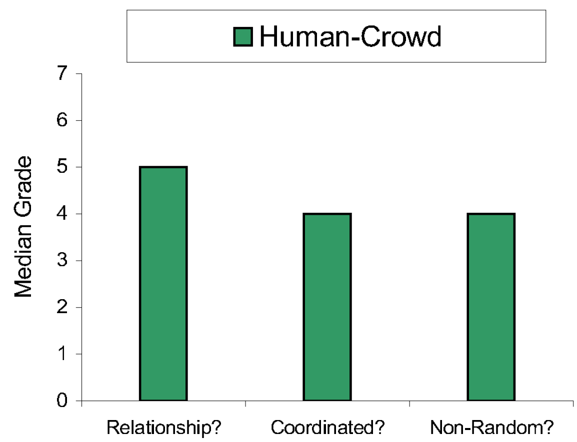

Figure 8.5: **Human crowd results - general questions.**

We compare the human crowd results to the individual and perfect-unison models results. It appears to be significantly different from individual model in all questions ($p = 0.000016$, $p = 0.000033$, and $p = 0.04$, respectively; t-test, one-tailed). However, in comparison to perfect-unison model, the results of the coordination and non-random questions are significantly different($p = 0.0034$, and $p = 0.0003$ correspondingly) significance level (t-test, one-tailed). But, in the results of the relationship question there is no significant found ($p = 0.44$). Therefore, the human crowd results also lies between the individual and perfect-unison model. However, in the relationship question results appears to be closer to the perfect-unison model.

In response to the question “Do you see any leaders? If so, how many?”, the median result in human crowd movie was 1.5. It appears to be significantly different from individual model ($p = 0.001$, t-test, one-tailed) but not in comparison to perfect-unison ($p = 0.374$). When the subjects were asked to qualitative discuss their answer to this question, many subjects reported on feeling that they don’t see any leader, however they sure that there is one since the crowd is waiting for some-thing or someone. However, when they asked to refer to only to people seen in the movie, the answer was that there were several subgroups in seen crowd. Therefore, we think that the question that should be asked in watching such movies: "Are there any subgroups? If so, how many?".

Chapter 9

Summary and Future Work

This thesis presented a model proscribing crowd behavior, inspired by Festinger's social comparison theory [10]. The model intuitively matches many of the characteristic observations made of human crowd behavior, and was shown to cover several distinct phenomena reported in the literature. Though there is lack of objective data against-which the model can evaluated, the results are promising and seem to match intuitions as to observed behavior.

We also presented implementation of the model for pedestrian movement experiments, gathering and imitational behavior. We describe two different ways of implementations the SCT process in an architectural level. Based on the results from experiments, we argue that implementation of SCT as an on-going process is more suitable for modeling crowd behavior than SCT as a problem-solving activity process approach.

In our future work we plan to extend the SCT model to include the repelling forces. Thus, each agent should not only be attracted to the similar but also should avoid the dissimilar. With our current SCT model the agent moves only through forces of attraction. Therefore, there are some collective behaviors that are difficult to simulate. For example, blended crowd behaviors. In order to evaluate this, we plan to generate a simulation of calm demonstration in which participants respond to threatening events (such as explosions, military/police presence, etc.). The crowd will exhibit different behaviors during different phases of simulation and with the ability to move independently and transparently from one behavior to another. By definition, this also includes behavior blending, where different parts of the crowd exhibit different behaviors simultaneously.

We also plan to explore the expression of leadership in social comparison theory, and expand our model to simulate collective behaviors with the influence of leaders. Our main goal that with SCT model we will be able to simulate crowd behaviors with and without leadership influence.

With current SCT model, agents behave very reactively since the only features for comparison is external and instantaneous one. For example, an agent is able to compare the location of other agent (instantaneous feature) and not able to compare the intended destination of the agent. Thus, we plan to integrate intention recognition model into a social comparison model in order to achieve more complex crowd behaviors.

SCT model require a high degree of computational complexity since each agent performs SCT with each seen neighbor, in every cycle of simulation. For example: suppose we have n agents, each agent sees k neighbors and simulation duration is c cycles. The simulation complexity is: $n^k c$. We plan to reduce this computational complexity by reducing the k value (number of compared agents) and the n value (number of agents that perform comparison each cycle).

Bibliography

[1] N. Agmon and D. Peleg. Fault-tolerant gathering algorithms for autonomous mobile robots. *SIAM Journal on Computing (SICOMP)*, 36(1):56–89, 2006.

[2] F. H. Allport. *Social Psychology*. Boston: Houghton Mifflin, 1924.

[3] T. Balch. *Behavioral Diversity in Learning Robot Teams*. PhD thesis, Georgia Institute of Technology, 1998.

[4] V. J. Blue and J. L. Adler. Cellular automata microsimulation of bidirectional pedestrian flows. *Transportation Research Record*, pages 135–141, 2000.

[5] A. Braun, S. R. Musse, L. P. L. de Oliveira, and B. E. J. Bodmann. Modeling individual behaviors in crowd simulation. In *Computer Animation and Social Agents*, pages 143–148, 2003.

[6] A. Braun, S. R. Musse, L. P. L. de Oliveira, and B. E. J. Bodmann. Modeling individual behaviors in crowd simulation. *casa*, 00:143, 2003.

[7] M. Cieliebak, P. Flocchini, G. Prencipe, and N. Santoro. Solving the robots gathering problem. In *30th Int. Colloq. on Automata, Languages and Programming*, pages 1181–1196, 2003.

[8] M. Cieliebak and G. Prencipe. Gathering autonomous mobile robots. In *9th Int. Colloq. on Structural Information and Communication Complexity*, pages 57–72, 2002.

[9] W. Daamen and S. P. Hoogendoorn. Experimental research of pedestrian walking behavior. *Transportation Research Record*, pages 20–30, 2003.

[10] L. Festinger. A theory of social comparison processes. *Human Relations*, pages 117–140, 1954.

[11] S. Freud. *Group Psychology and the Analysis of the Ego*. Liveright Publishing, 1951.

[12] K. L. Hakmiller. Threat as a determinant of downward comparison. *Journal of experimental social psychology*, 2:32–39, 1966.

[13] D. Helbing. Boltzmann-like and boltzmann-fokker-planck equations as a foundation of behavioral models. *Physica A*, 196:546–573, 1993.

[14] D. Helbing and P. Molnar. Self-organization phenomena in pedestrian crowds. In F. Schweitzer, editor, *Self-organization of Complex Structures: From Individual to Collective Dynamics*, pages 569–577. Gordon and Breach, London, 1997.

[15] D. Helbing, P. Molnar, I. J. Farkas, and K. Bolay. Self-organizing pedestrian movement. *Environment and Planning B*, 28:361–384, 2001.

[16] D. Helbing and T. Vicsek. Optimal self-organization . *New Journal of Physics*, 1:13–+, 1999.

[17] L. F. Henderson. The statistics of crowd fluids. *Nature*, 229:381–383, 1971.

[18] L. F. Henderson. On the fluid mechanics of human crowd motion. *Transportation research*, 8:505–515, 1974.

[19] G. A. Kaminka, M. M. Veloso, S. Schaffer, C. Sollitto, R. Adobbati, A. N. Marshall, A. Scholer, and S. Tejada. GameBots: A flexible test bed for multiagent team research. *Communications of the ACM*, 45(1):43–45, January 2002.

[20] T. Kretz. *Pedestrian Traffic: Simulation and Experiments*. PhD thesis, Universität Duisburg-Essen, 2007.

[21] G. Le Bon. *The crowd: A study of the popular mind.* Dunwoody, Ga., N.S. Berg, 1968.

[22] M. J. Matarić. Designing and understanding adaptive group behavior. *Adaptive Behavior*, 4(1):50–81, December 1995.

[23] A. Newell. *Unified Theories of Cognition*. Harvard University Press, Cambridge, Massachusetts, 1990.

[24] T. Osaragi. Modeling of pedestrian behavior and its applications to spatial evaluation. In *AAMAS '04: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 836–843, Washington, DC, USA, 2004. IEEE Computer Society.

[25] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques (SIGGRAPH-87)*, pages 25–34, New York, NY, USA, 1987. ACM Press.

[26] C. W. Reynolds. Steering behavior for autonomous character. In *Proceedings of the Game Developers Conference*, pages 763–782, 1999.

[27] S. J. Rymill and N. A. Dodgson. A psychologically-based simulation of human behaviour. In *Theory and Practice of Computer Graphics*, pages 35–42. 2005.

[28] J. E. Singer. Social comparison: progress and issues. *Journal of experimental social psychology*, 2:103–110, 1966.

[29] Soar. http://sitemaker.umich.edu/soar/home/, 2006.

[30] D. Thalmann. The foundations to build a virtual human society. In *Proceedings of Intelligent Virtual Actors (IVA-2001)*, pages 1–14. Springer-Verlag, 2001.

[31] P. C. Tissera, M. Printista, and M. L. Errecalde. Evacuation simulations using cellular automata. *Journal of Computer Science and Technology*, 7(1):14–20, April 2007.

[32] M. C. Toyama, A. L. C. Bazzan, and R. da Silva. An agent-based simulation of pedestrian dynamics: from lane formation to auditorium evacuation. In

AAMAS '06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems, pages 108–110, New York, NY, USA, 2006. ACM Press.

[33] X. Tu and D. Terzopoulos. Artificial fishes: physics, locomotion, perception, behavior. In *SIGGRAPH '94: Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 43–50, New York, NY, USA, 1994. ACM Press.

[34] U. Wilensky. NetLogo. Center for Connected Learning and Computer-Based Modeling—Northwestern University; http://ccl.northwestern.edu/netlogo/, 1999.

[35] M. Wolff. Notes on the behaviour of pedestrians. *In People in Places: The Sociology of the Familiar*, pages 35–48, 1973.

[36] S. Wright. *Crowds and riots : a study in social organization*. Beverly Hills, Calif. : Sage Publications, 1978.

[37] K. Yamashita and A. Umemura. Lattice gas simulation of crowd behavior. In *Proceedings of the International Symposium on Micromechatronics and Human Science*, pages 343–348, 2003.