Towards a Cognitive Model of Crowd Behavior Based on Social Comparison Theory

Natalie Fridman and **Gal A. Kaminka***
The MAVERICK Group
Computer Science Department
Bar Ilan University, Israel
{fridman,galk}cs.biu.ac.il

Abstract

Models of crowd behavior facilitate analysis and prediction of human group behavior, where people are affected by each other's presence. Unfortunately, existing models leave many open challenges. In particular, psychology models often offer only qualitative description, while computer science models are often simplistic, and are not reusable from one simulated phenomenon to the next. We propose a novel model of crowd behavior, based on Festinger's Social Comparison Theory (SCT). We propose a concrete algorithmic framework for SCT, and evaluate its implementation in several crowd behavior scenarios. Results from task measures and human judges evaluation shows that the SCT model produces improved results compared to base models from the literature.

Introduction

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 (Helbing *et al.* 2001; Le Bon 1968; Rymill & Dodgson 2005; Daamen & Hoogendoorn 2003). Existing models, in a variety of fields, leave many open challenges. In particular, 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 are therefore not reusable in simulating new phenomena.

We propose a novel model of crowd behavior, based on Social Comparison Theory (*SCT*) (Festinger 1954), a popular 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. We propose a concrete algorithmic framework for SCT.

We evaluate the use of SCT models in generation of pedestrian movement and imitational behavior, and show that SCT generates behavior more in-tune with human crowd behavior. Moreover, unlike many previous models, SCT generalizes across social phenomena. 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 the context of imitational behavior, the SCT model was evaluated in studies with human subjects. 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.

Background and Motivation

A phenomenon observed with crowds is its homogeneous nature. The participants act as if governed by a single mind, despite using little or no verbal communications (Le Bon 1968). Social psychology literature provides several views on the mechanisms underlying this phenomenon, but unfortunately, these are too abstract to be used algorithmically. For example, (Le Bon 1968) explains the behavior of crowd by two processes: (i) *imitation*, where people in a crowd imitate each other; and (ii) *contagion*, where people in a crowd behave very differently from how they typically would, individually.

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., (Rymill & Dodgson 2005). However, these models have rarely, if ever, been validated against human (or animal) data. Indeed, there is generally little quantitative data on the behavior of human crowds, at a resolution which permits accurate modeling. One exception is the formation of lanes (in opposing directions) in human pedestrian movements (Daamen & Hoogendoorn 2003), which has been extensively investigated, and for which specific performance measures are well-defined (reduced lane changes, flow, etc.).

Henderson modeled pedestrian movement on gas-kinetic fluid movement (Henderson 1971). Helbing et al. (Helbing *et al.* 2001) developed a model that takes into account interactions between the individuals and the physical environment. Blue and Adler (Blue & Adler 2000) used cellular automata. The focus of all of these is again on local interactions; each simulated pedestrian is a particle or an automaton, whose next action or behavior is determined by its local surroundings.

Indeed, in these previous works, the behavior of crowds

*This research was supported in part by IMOD and ISF.

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 (Blue & Adler 2000) 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 (Blue & Adler 2000)) and farther cells (rule 2), as well as to 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 (**?**) 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 (**?**, 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.

A Model of Social Comparison

We took Festinger's social comparison theory (SCT) (Festinger 1954) 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.

Festinger's Social Comparison Theory

Festinger (Festinger 1954) presents social comparison theory (SCT) as an explicit set of axioms. The following subset of axioms (re-worded) are particularly relevant: (i) When lacking objective means for evaluation, agents compare their state to that of others; (ii) comparison increases with similarity; (iii) agents take steps to reduce differences to the objects of comparison.

Newell (Newell 1990) 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." (Festinger 1954, 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." (Festinger 1954, 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 (Le Bon 1968) characteristics of crowd behavior:

Imitation. Using social comparison, people may adopt others' behaviors. Festinger notes (Festinger 1954): "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 (Festinger 1954): "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".

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

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 (Festinger 1954): "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 (Festinger 1954): "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.

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

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 (Wolff 1973) noted that pedestrians have a hight degree of cooperation and coordination which without it, walking on sidewalk would be impossible.

Indeed, human pedestrians, moving in opposing directions, quickly self-organize into uni-directional lanes, in which their movement is essentially undisturbed (Daamen & Hoogendoorn 2003). This phenomenon is a focal point for investigations: Quicker lane formations can lead 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). Thus crowd models for pedestrians are often evaluated based on their ability to reduce the number of lane changes (i.e., quicker lane formation).

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.

Each agent has a set of features and their corresponding weights. For simulating pedestrian movement, we used the following features and weights: *Walking direction* east or west (weight 2); *color* (weight 3); and *position* (weight 1), given global coordinates.

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. 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.

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 (which we use to denote sub-groups within the crowds) even more so. For instance, 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).

To implement the model for pedestrians movement experiments, we used NetLogo (Wilensky 1999). 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 range and field of view (120°). Each agent follows an initially set direction. If forward movement is blocked, the agent will choose the lane based on SCT algorithm above. We use a default uniform velocity for all agents. The velocity may change individually based on gain.

To evaluate the SCT model, we contrast it with a popular alternative model, often used in pedestrian crowd research (Blue & Adler 2000; Helbing *et al.* 2001). 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).

Experiment 1: Bidirectional Pedestrian Movement

We first evaluate SCT in accounting for lane formation in oppositely moving pedestrians; this is likely the most common experimental task for pedestrian crowd models.

In this first experiment all agents moving in the same direction had the same color (different than the other direction). To evaluate the effects of gain, we vary the S_{min} and S_{max} values, as they set the enumerator in the gain calculation. Setting S_{max} at 6 means that any dissimilarity other than color triggers action and setting S_{min} at 2 means that agents that differed only in distance (but not by color or direction) were not considered similar.

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., $g(o) = 1$ (no change).
- $S_{max} = 5.5$, $S_{min} = 4$, i.e., $g(o) = 3$
- $S_{max} = 5.5$, $S_{min} = 2$, i.e., $g(o) = 7$

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

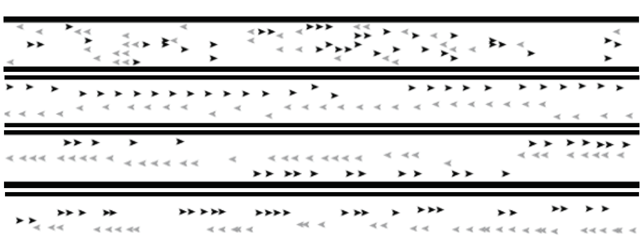

Figure 1: **Bidirectional Pedestrian Movement.**

Figure 2 shows lane formation and flow results for the individual and social comparison models. The X-axis in both subfigures measures density. The Y-axis in Figure 2(a) measures the number of lane changes; a higher number indicates more lanes are formed, since agents must still change lanes, having met opposing pedestrians. The Y-axis in Figure 2(b) measures flow. Each data-point is an average over 15 trials.

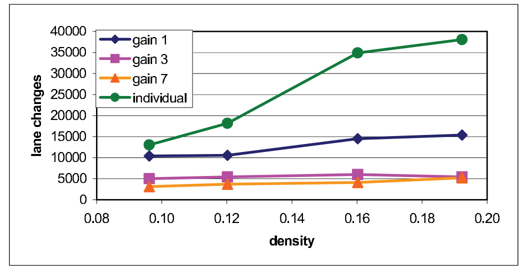

(a) Lane changes

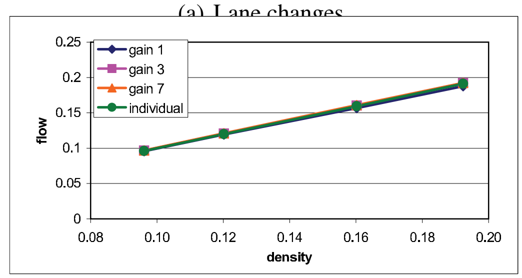

(b) Flow.

Figure 2: **Bidirectional pedestrian measurement results**

The figures show that the number of lane changes in SCT is significantly lower than that of the individual-choice model (one tailed t-test, 0.05 significance level). This is true even with a gain of 1, which effectively neutralizes the gain in comparison to the individual-choice model. Moreover, the difference with the individual-choice model increases with an increased gain. However, there are essentially no differences in flow. These results support the hypothesis that the use of SCT can lead to quicker lane formations, which would indicate an improved model of crowd behavior.

Experiment 2: Pedestrians in Groups

In this experiment, we evaluate the SCT model in settings where the pedestrians form subgroups. This type of situation arises, for instance, in pedestrians that are composed of families and/or friends, which then often move closer together even when undisturbed by opposing traffic. The individual choice model does not account for such behavior. In contrast, we expect SCT to address such settings, as subgroup members would be more similar.

To examine this hypothesis, we carried out experiments where agents belonging the same group have the same color. In these experiments, all agents move in the same direction, for 5000 cycles. The population contains 150 agents with a different number of colors (we experimented with 5, 10, and 20 and color). Agents use comparison at all times, and not just when stuck. 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*, weight 2; *distance*, weight 0.5; *color*, weight 3; and

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

Table 1: **Grouping measurements of individual-choice and social comparison models.**

X-coordinate, 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.

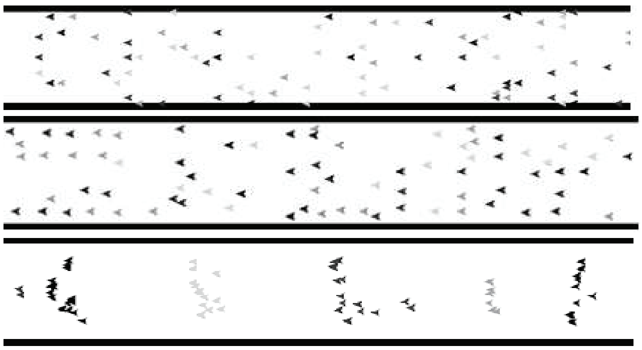

Figure 3: **Screen shots, grouped pedestrian movement.**

Figure 3 shows the initial random positions of the agents (top figure), their positions after moving for 5000 cycles using the individual model (second figure), and their positions after moving 5000 cycles using the social comparison model (third figure). The figures show that the social comparison model causes similar-colored agents to group together.

We used *hierarchical social entropy*—a measure of group heterogeneity (Balch 1998)—to measure the grouping results. The key intuition behind this measure is to iteratively sum entropy over increasing areas. The measure equals 0 when all agents are positioned in the exact same spot, and grows with their uniform spreading in space. Thus lower values indicate improved grouping.

Table 1 shows the measurement results for the individual-choice and social-comparison models. Each row corresponds to the average results over multiple trials, 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, 0.05 significance level). Again, these results support the hypothesis that SCT is an improvement over the basis individual-choice model.

Experiment 3: Groups and Obstacles

Our final set of pedestrian movement experiments addresses the response of groups within moving pedestrian crowds, to obstacles in their path. Intuitively, we recognize that such groups will choose to stick together in face of an obstacle (moving together to one side of it). We sought to examine whether the SCT model 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. The population contains 100 agents, with 2 colors (red or blue). Walking direction of all agents is West. Each agent has the following features: Direction, distance and color (weights: same as in the bidirectional 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: **Screen shots, movement around an obstacle.**

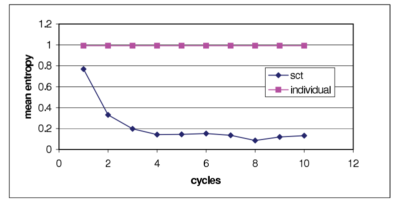

Figure 5: **Entropy of grouped pedestrians' movement around the obstacle.**

Figure 4 shows the initial random positions of the agents (top figure), their positions after 5000 cycles using the individual model (second figure), and their positions after 5000 using the SCT (third figure). 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 order to measure to what degree agents of the same color stay on one side of the obstacle, 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. The final result of each wave is the average entropy value across the two colors. Lower entropy values indicate increased grouping on one side.

Figure 5 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. Lower score indicates improved grouping. The figure shows that the entropy value of the social-comparison model quickly goes down from 1 and approaches 0. This indicates that the SCT model does indeed cause groups to stick together in face of obstacles.

SCT in Imitational Behavior

An attractive feature of social comparison is its hypothesized prevalence in human group behavior. Indeed, we believe that the SCT model is sufficiently general to account for a wide variety of group behaviors. This section provides initial 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. 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 (Le Bon 1968).

To explore imitational behavior, we utilized a 3D virtual world (see Figure 6), in which multiple agents can interact with each other. Each agent is controlled by a separate process, implemented SCT in the Soar cognitive architecture (Newell 1990; Soar 2006).

Figure 6: **Soar agents in the GameBots environment.**

A detailed discussion of Soar's role as a cognitive architecture is beyond the scope of this paper. We provide a very brief overview here, and refer the interested reader to (Newell 1990; Soar 2006) 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.

For our experiments, several basic task-oriented operators were implemented, to allow the agents to move about, turn towards each other, etc. Thus one thread of control, always running, is in control of the agent's actions towards whatever tasks it was given. Normally, Soar agents only utilize this single thread of control.

SCT was implemented as secondary parallel thread of control within Soar. Whereas normally, operators are proposed (and selected) by Soar based on their suitability for a current goal, in our agent, operators were *also* proposed based on their suitability for SCT. The SCT thread proposed operators by following our algorithm, though in a way that is adopted for Soar's decision cycle: At every 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.

In other words, at every cycle, a Soar agent would consider operators that advance it towards its goal (as it normally would), but these would compete for selection with SCT-proposed operators that seek to minimize perceived differences to other agents. This may appear to contradict Festinger's theorizing that social comparison comes into play only when people are at an impasse with respect to their goal (lack objective means to evaluate themselves). However, this is not the case. By setting Soar's decision preferences to prefer SCT-proposed operators only when no task-oriented operators are available, one gets the behavior predicted by Festinger's theory. Further exploration of this issue is beyond the scope of this paper.

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 paper.

Experiment 4: Imitational Behavior

We conducted experiments to evaluate whether SCT can indeed account for imitational behavior in groups, and whether it resembles human crowds. We rely on experiments with human subjects, which judged the human crowd behavior and the resulting SCT behavior in comparison to an individual-choice model (i.e., where each agent makes decisions independently of its peers), and to completely synchronized behavior (i.e., all agents act in complete unison).

We hypothesized that agents controlled by SCT would display 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. Furthermore, we hypothesized that human crowds would also be ranked in-between the individual and perfect-coordinated.

To examine these hypotheses, 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, the agents were fixed to their initial locations, and the only actions available to them were to do nothing, or turn at some angle. In all movies, one agent, colored blue, simply turned up to 90° left or right, arbitrarily. All others (red) acted according to one of the models. In one movie (*individual*), the red agents acted completely independently of each other, arbitrarily choosing an angle and turning to it. In another (*unison*), the red agents relied on communications to act in almost perfect synchronization, turning towards the same angle as the blue almost instantaneously (small timing differences resulting from asynchronous responses of the simulated environment). Finally, in the *SCT* movie, the red agents behaved according to the model described above.

The first hypothesis experiments were carried out using 12 subjects (ages: 18–40, 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 (as in Figure 6). After the description, the movies were shown to the subject. After each movie, the subjects were asked to fill a short questionnaire (described below). The order of presentation of movies was randomly selected for each subject, to control for learning and order effects.

To examine the second hypothesis, we used a TV news clip, which showed a group of people standing and waiting for some event to occur; the only action they performed was to occasionally turn. 12 new subjects were asked to fill the same questionnaire as in previous experiments, after seeing *only the news clip*.

For lack of space, we will focus here only on a subset of the results, summarized in Figures 7-a and 7-b. The categories in the X-axis in both figures correspond to questions given to the subjects: (i) Relationship: To what degree were the agents/humans related to each other; (ii) Coordinated: To what degree were the turning angles chosen by agents/humans coordinated; (iii) Non-random: To what degree were the observed movements non-random. In all of these questions, the subjects were asked to grade the movies on an ordinal scale of 1–6, with 1 being a low result (associated with more individual behavior, and 6 being a high result (associated with perfect unison). The three bars (Figure 7-a) for each question measure the median result for the individual, SCT, and unison models. In Figure 7-b, the results are shown for the news clip.

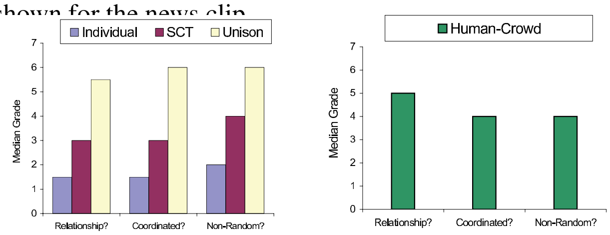

(a) Virtual environment. (b) News clip.

Figure 7: **Experiment 4 results (subset).**

The results clearly demonstrate that the SCT model lies in-between the individual and perfect-unison model. While it appears to be somewhat closer to the individual model, it is significantly different from it at the 0.05 significance level (t-test, one-tailed). We also compare the human crowd results to the simulated individual and perfect-unison models results. It appears to be significantly different from individual model in all questions at the 0.05 significance level (t-test, one-tailed). However, in comparison to perfect-unison model, it significantly different in coordination and non-random questions but not in relationship question. We thus conclude that in general, both the SCT and real human behavior lie between perfect coordinated and individual behavior, though some notion of relationship is observable in the news clip, but not in the virtual environment movies.

Summary and Future Work

This paper presented a model proscribing crowd behavior, inspired by Festinger's social comparison theory (Festinger 1954). The model intuitively matches many of the characteristic observations made of human crowd behavior, and was shown to cover several distinct crowd phenomena. Results from experiments in pedestrian movement and imitation domains are promising, and support the general applicability of the SCT model. Future work includes an improved memory model, and examination of leadership.

References

Balch, T. 1998. *Behavioral Diversity in Learning Robot Teams*. Ph.D. Dissertation, Georgia Institute of Technology.

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

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

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

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

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

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

Newell, A. 1990. *Unified Theories of Cognition*. Cambridge, Massachusetts: Harvard University Press.

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

Soar. 2006. http://sitemaker.umich.edu/soar/home/.

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

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