A Cognitive Model of Crowd Behavior Based on Social Comparison Theory *

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

Abstract

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 that 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, psychological 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. We show that our SCT model produces improved results compared to base models from the literature. We also discuss an implementation of SCT in the Soar cognitive architecture, and the question this implementation raises as to the role of social reasoning in cognitive architectures.

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 that are affected by each other's presence and actions. Accurate models of crowd behavior are sought in training simulations (Thalmann 2001), safety decision-support systems (Braun *et al.* 2003), traffic management (Helbing & Molnar 1997; Rymill & Dodgson 2005), 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., pedestrian movement on a sidewalk), 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*) (Festinger 1954), a popular social psychology theory that has been continuously evolving since the 1950s. The key idea 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 models compared to popular models in the literature, and show that SCT generates behavior more in-tune with human crowd behavior. In particular, the SCT models generate improved pedestrian movements and accounts for group formation in pedestrians that are inter-related, a phenomenon unaccounted for by previous models.

We describe the implementation of SCT in Soar, as an impasse-resolution method. We argue that SCT is a weak (read: general) problem-solving method, which is *social* in nature. This view of SCT raises questions as to the role of social reasoning in cognitive architectures and the mind. In particular, it may lead 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.

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 models (described below) tend to be simplistic, and ignore the psychological studies.

Social psychology. A phenomenon observed with crowds, and discovered early in crowd behavior research, is that people in the crowd act similar to one another, often acting in a

*This research was supported in part by IMOD.
Copyright © 2006, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

coordinated fashion, as if governed by a single mind (Le Bon 1968; Allport 1924). However, this coordination is achieved with little or no verbal communications.

Le Bon (Le Bon 1968) emphasized a view of crowd behaviors as "Collective Mind", and observed that an individual who becomes a part of the crowd is transformed into becoming identical with 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 the crowd behave very differently from how they usually behave, individually. Freud (Freud 1951) expanded on Le Bon by theorizing that individuals in the crowd identify with the leader and with each other, and that is they behave as one. The crowd behavior can be controlled by the leader, as the individuals imitate the leader.

Another phenomenon that occupies the researches is how a crowd is created from the first place or to be more specific is what causes an individual to be part of a crowd. According to Allport (Allport 1924) an individual becomes part of a crowd when he has "common stimulus" with people inside the crowd. For example, a common cause. Allport agrees with Le Bon (Le Bon 1968) about homogeneous behaviors of a crowd, but his explanation to this phenomena is that similar people act in similar ways, otherwise they are not 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 base our work on the Social Comparison Theory (Festinger 1954), which (to the best of our knowledge) has never been applied to modelling crowd behavior. Nevertheless, as we show in the next section, key elements of the theory are at the very least compatible with theories discussed above.

Computational models. Work on modelling crowd behavior has been carried out in other branches of science, in particular for modelling and simulation. Reynolds (Reynolds 1987) has simulated bird flocking using simple, individual-local rules, which interact 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.

Similar ideas have been applied in swarm robotics. Mataric (Matarić 1995) sees collective (complex) behaviors as a combination of basic behaviors. Each robot has spatial behaviors (controllers) which are combined to create different kinds of group behavior. For example, flocking combined of *safe-wandering* (move around without bumping), *homing*, *dispersion* (move away from other agents), and *aggregation* (move towards other agents). The combined outputs of the basis behaviors provide a velocity vector which is used to control the robot.

Much of the work we describe in this paper has been implemented and evaluated in the context of pedestrian movement. Several important previous investigations have examined this task.

Blue and Adler (Blue & Adler 2000) use 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 show that this simple rule results in the formation of lanes in movement, similarly to those formed in human pedestrian movement (Wolff 1973).

Helbing et al. (Helbing & Molnar 1997; Helbing *et al.* 2001) also focuses on simulating pedestrian movement. Each entity moves according to forces of attraction and repulsion. Pedestrians react both to obstacles and to other pedestrians. The study shows that this also results in lane formation.

Our work differs from those that have been described above in that we aim to develop a cognitive model of crowd behavior, one based on psychology, rather than particle physics. As a result, the model we present in this work covers phenomenons uncovered by the simplistic CA and particle models, such as that the choice of lane changing in pedestrians depends on subgroups within the crowd, and is not random.

A Model Based on Social Comparison Theory

The research question we address in this paper deals with the development of a computerized cognitive model which, when executed individually by many agents, will cause them to behave as humans do in crowds.

We took Festinger's Social Comparison Theory (Festinger 1954) as inspiration for the social skills necessary for our agent. According to the 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.

We believe that the social comparison theory may account for some characteristics of crowd behavior:

Common stimulus between crowd participants. One of the social comparison theory implications is group formation. Festinger notes (Festinger 1954): "To the extent that self evaluation can only be accomplished by means of comparison with other persons, 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 and whose abilities are near their own".

Imitational behavior. By social comparison, people may adopt others' behaviors. 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".

To be usable by computerized models, the social comparison theory must be transformed into a set of algorithms that, when executed by an agent, will proscribe social comparison behavior. A first step towards this goal has been take by Newell, which describes the social comparison theory as a set of axioms (Newell 1990):

1. If the agent can't evaluate its opinions and abilities objectively, then it compares them against the opinion and abilities of others.
2. Comparing against others decreases as the difference with others increases.
3. The more important a group for comparison of an opinion or ability, the more pressure there is toward uniformity on that opinion or ability.

We take another step towards the modelling of social comparison theory. Each observed agent is assumed to be modelled 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 the 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. We use a weighted linear sum to aggregate feature values into one similarity measure:

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

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 action 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.

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

Modeling Pedestrian Movement

To learn more about microscopic and macroscopic pedestrians' behavior, Daamen & Hoogendoorn 2003 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 lead to improved flow through the area, and the more agents organized into lanes, the less they 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.

Our goal is to explore the use of our social comparison model in accounting pedestrian movement phenomena like lane formations in bidirectional movement and movement in groups, with and without obstacles.

To implement the model for pedestrian movement experiments, we used NetLogo (netlogo). 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,

Figure 1: **Initial NetLogo sidewalk.**

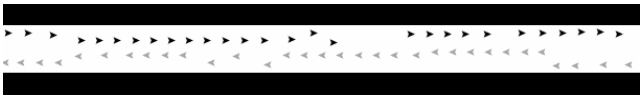

Figure 2: **Lane formations.**

0.1). Agents are not allowed to move through other agents, and thus no two agents can occupy the same space.

Figure 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 2 shows an end-result from one of the experiments (described below), where lanes have been formed.

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 the 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. 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 $c(x)$ 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.

Experiments in Pedestrian Movement

This section explores the use of the social comparison model and its implementation in modelling pedestrian movement. The basic movement pattern that our simulated pedestrians follow, stemming from the social comparison model, is as follows. The agents follow initially set directions. They choose moving in this direction, unless blocked. If forward movement is indeed blocked, the agents can choose between changing lanes to the left or right. It will choose the lane where there is an agent similar to it (if available). If there is no similar agent, it will arbitrarily choose any lane.

Individual Pedestrians Our first experiment contrasted the social comparison model with previous models. We began by examining individual pedestrian movements, where each synthetic pedestrian is independent of others. We contrasted the social comparison model with that of random choice, which was shown to produce lane formations (Helbing & Molnar 1997; Helbing *et al.* 2001) and is considered to be a base model for pedestrian models.

As is commonly done in pedestrian movement experiments, we controlled for *crowd density*, calculated as the number of agents divided by the area. To do this, we fixed the number of agents at 100, and changed the width of the sidewalk. Each agent had a unique color. Each agent's direction (east or west) and initial position was chosen randomly. We follow the literature in measuring two principal characteristics of pedestrian movement: the total number of *lane changes*, and the *flow* (average speed divided by the space-per-agent).

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. S_{min} was set at 2, which means that agents that differed only in distance were not considered similar. Each trial was executed for 5000 cycles.

Figure 3 shows lane changes for the random-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 multiple times. Figure 4 measures flow for the two models. The X-axis again measures density. The Y-axis measures the flow.

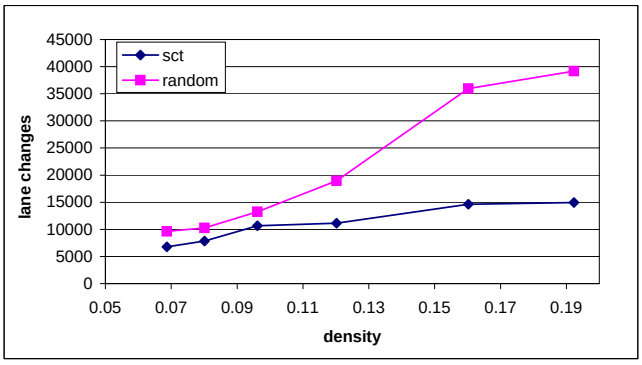

Figure 3: **Individual Pedestrians' lane changes.**

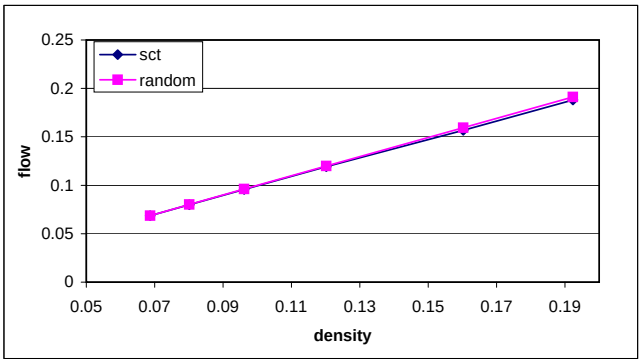

Figure 4: **Individual Pedestrians' flow.**

The figures shows that the number of lane changes is significantly lower than that of the random-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 significant differences in flow. In other words, the social comparison model performs better, but with no cost to the flow.

Individual Pedestrians, with Variable 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 denominator in 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., $g(o) = 1$ (ignoring the distance).
- $S_{max} = 5.5$, $S_{min} = 5$, i.e., $g(o) = 3$
- $S_{max} = 5.5$, $S_{min} = 2$, i.e., $g(o) = 7$

Figure 5 shows the initial positions of the agents in one of the trials (5(a)), and the typical results after 5000 cycles, with a gain of 1 (5(b)), gain of 3 (5(c)), and gain of 7 (5(d)). The figures show how the increased gain causes the agents to group more closely together.

Figures 6 and 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, with an increased gain.

Grouped Pedestrian Movements We now move away from considering scenarios that have previously appeared in the literature, and start 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

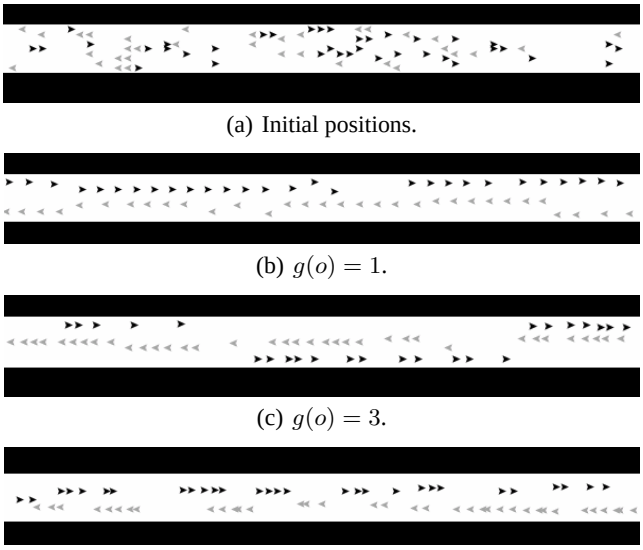

(a) Initial positions.

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

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

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

Figure 5: **Individual Pedestrians: Varied Gains.**

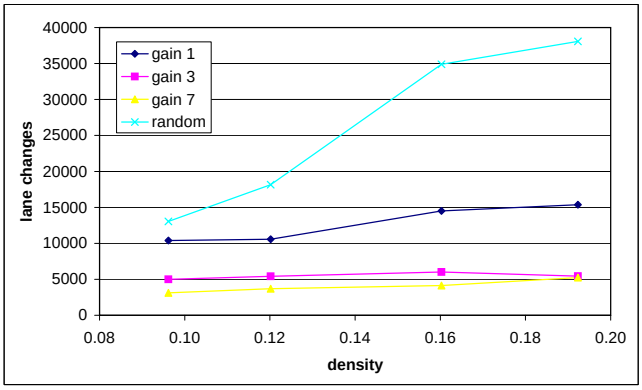

Figure 6: **lane changes**

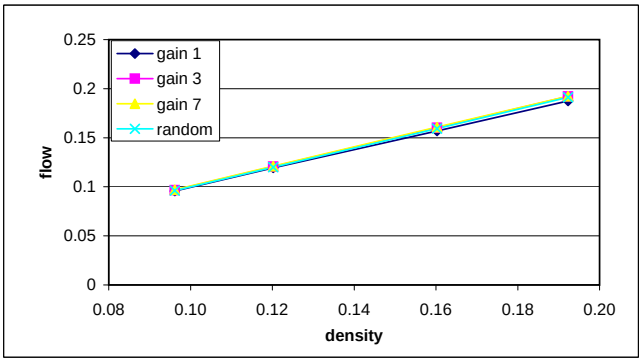

Figure 7: **flow**

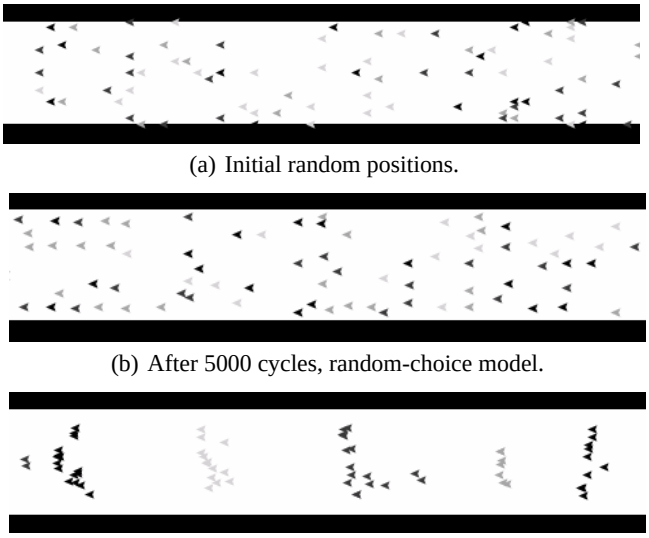

(a) Initial random positions.

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

(c) After 5000 cycles, social comparison model.

Figure 8: **Screen shots, Grouped Pedestrian Movement.**

random-choice model does not cover such phenomena, because it does not treat internal groups 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 to 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 and color). Agents use comparison at all times, and not just when stuck. Walking direction of all agents was West. S_{max} was set at 6.5, and S_{min} was set at 2.

To account for the 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.

Figure 8 shows the initial random positions of the agents (8(a)), their positions after moving for 5000 cycles using the random-choice model (8(b), and their positions after moving 5000 cycles using the social comparison model (8(c)). The figures show that the social comparison model causes similar-colored agents to group together. This group forming does not occur in the random-choice model.

# Colors	Random-Choice	Social Comparison
5	173.2	**87.4**
10	143.3	**85.8**
20	101.5	**60.1**

Table 1: **Grouping measurements of random-choice and social comparison models. Lower values indicate improved grouping.**

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 (Balch 1998) has offered a clustering measure, *hierarchical social entropy*, that can address such cases. While (Balch 1998) provides the details, 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 spreading around. Thus lower values indicate improved grouping.

Table 1 shows the hierarchical social entropy results for the random-choice and social-comparison models. Each row corresponds to an experiment with a different number of colors. The table shows (final column) that the social comparison model provides for significantly improved grouping compared to the random-choice model.

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 when face an obstacle (moving together to one side of it), while independent 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 approached the obstacle, they had 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 had the following features: Direction, distance and color (weights: same as in the individual pedestrian experiments). Agents used comparison at all times, and not just when stuck. S_{max} was set at 6.5, S_{min} at 3.

Figure 9 shows the initial random positions of the agents (9(a)), their positions after moving for a while using the random-choice model (9(b), and their positions when moving using the social comparison model (9(c)). The figures show clearly that the social comparison model causes similar-colored agents to group together on one side of the obstacle, passing it together. In contrast, the random-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 de-

(a) Initial random positions.

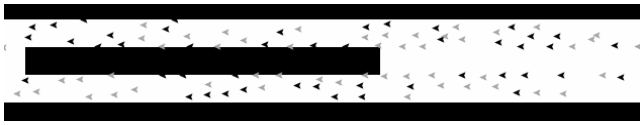

(b) Final positions, with random-choice model.

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

Figure 9: **Initial and final positions of agents in grouped pedestrian experiments.**

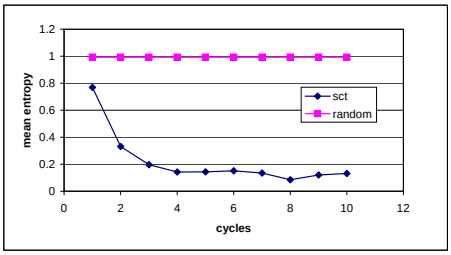

Figure 10: **Entropy of Grouped Pedestrian Movement around Obstacle.**

gree 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 10 shows the average entropy value for each wave, for the ten waves. The results are averaged over multiple trials. The X-axis shows the wave number (1–10). The Y-axis measures the entropy. The figure shows that the entropy value of the social-comparison model quickly goes down from 1 and approaches 0, while it remains around 1 for the random-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 random-choice model.

An Initial Implementation in Soar

We have developed an initial implementation of the social comparison theory model, as described above, in the Soar cognitive architecture (Newell 1990). Soar uses globally-accessible working memory, and production rules that test and modify this memory. Efficient algorithms maintain the working memory in face of changes to specific propositions. Soar operates in several phases, one of which is a decision phase in which all relevant knowledge is brought to bear to make a selection of an operator, that will then carry out deliberate mental (and sometimes physical) actions. A key novelty in Soar is that it automatically recognizes situations in which this 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. Results of this decision process can be chunked for future reference, through Soar's integrated learning capabilities. Over the years, the impasse-mechanism has shown to be very general, in that general problem-solving strategies could be brought to bear for resolving impasses.

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 thus chose to treat social comparison theory as a new kind of impasse-resolution method. Unlike previous impasse-resolution (problem-solving) techniques, in which the agent focus on using its own resources, here the agent uses knowledge of others as a keystone to resolving the impasse. Our goal is therefore 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 our current implementation (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* is 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 : elaborate-target
    8:    O: O21 (elaborate-target)
elaborate-target is active
```

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

We believe our treatment of social comparison processes as generic impasse-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, not as a problem-solving method. 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 our 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.

Summary and Future Work

This paper presented a preliminary algorithmic 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 phenomena reported on in the literature. Though there is lack of objective data against which the model can be tested, the results are promising and seem to match intuitions as to observed behavior. We also presented an initial implementation of the model in Soar, and show how it can be integrated very well with the impasse-resolution mode of this architecture. This view of social-comparison processes, as a social impasse-resolution method, is novel and raises important questions as to the role of agent modeling (from observations) in cognitive architectures.

References

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

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.

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

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.

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

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

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

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

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

NetLogo homepage. http://ccl.northwestern.edu/netlogo/.

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

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

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

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