People are agents too: A review of the use of agent based systems as descriptive models of human cognition and group dynamics

Eric Dimperio

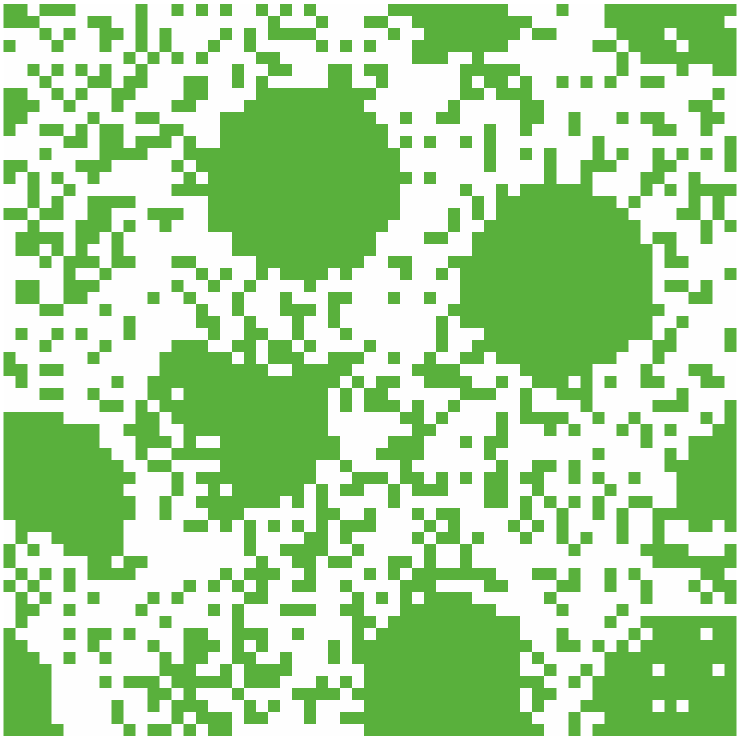

Qualifying examination essay 1 of 3.

Advisory committee:

- Professor Jerome Busemeyer, Ph.D.
- Professor Robert Goldstone, Ph.D.
- Professor John K. Kruschke, Ph.D.
- Professor Larry Yaeger

People are agents too: A review of the use of agent based systems as descriptive models of human cognition and group dynamics

Eric Dimperio

People and groups of people have been studied in many different branches of scientific inquiry (psychology, sociology, cognitive science, anthropology). Psychology and cognitive science have traditionally had the goal of investigating how an individual thinks by investigating behaviors of the individual. Anthropology and sociology have long studied collections of people by looking at patterns of behavior in those groups. Developments in the computational sciences have given rise to newer methods of inquiry for the study of both individuals and groups. Agent-based computational simulations allow scientists to study individual cognition by observing behaviors at the group level, as well as to study group dynamics by identifying behaviors at the level of the individual. This essay is intended to look at how traditions of studying individuals and groups are coming together to provide new insights into each other's respective domains. There will be an emphasis on the usage of agent-based modeling (ABM) to facilitate this convergence and how it relates to other work being done.

Cognitive scientists like to think of a person as a unit of cognition. Human behavior can be thought of as the result of some processing that occurs given some set of input stimuli (input stimuli → cognition → behavior). Research in the area of cognitive psychology generally consists of controlling the stimuli presented to subjects and measuring the responses. Cognitive modelers develop descriptions of the cognitive processes that occur within the subjects that would yield the observed responses given the stimuli presented. Many models have been presented that have provided much insight into psychological concepts of memory, learning, attention, language, and many others. Even though we have learned much from this style of investigation, it does have its limits. Humans are social animals. Many of our behaviors only make sense in the context of a social environment. Beyond that, many of the cognitive accomplishments achieved by mankind cannot be attributed to cognition by the individual. Certain cognitive processes seem to be the product of an individual and its environment, which may include one or more other cognitive agents (Hutchins, 1995). To get a true understanding of human cognition, one needs to understand something about group behavior.

Group behavior has traditionally been the domain of sociologists, economists, and anthropologists. Mathematical sociology tends to rely on equation-based modeling (Sawyer, 2003). In equation-based modeling (EBM), models are defined by a set of equations that are evaluated to provide descriptions/predictions of behavior. These types of models are often the most appropriate when the unit of decomposition is the macro-level observable variable. Unfortunately, some of the EBM techniques rely on strict assumptions that can lead to bad science (Moss, 2006; Sawyer, 2003). Often, EBM models of economic behavior assume that actions are taken by perfectly rational agents even though human often act irrationally. This may be useful in trying to determine what sort of decision should be made to achieve some result, but it is not useful when trying to understand human cognition.

A new breed of science

Science has usually relied on two methods of reasoning to reach conclusions about a subject of interest (Axelrod, 1997; Sun 2006). The deductive method begins with the specification of a formal model from which consequences are derived. The EDM methods mentioned from economics tend to use this approach. It is great when you understand the basic axioms of a situation and want to know the implications of those axioms. The second method is the inductive method. The inductive method generally involves the observation of empirical data from which generalizations can be formed. This method is widely used in sociology, where a large number of observations of group behavior can lead to qualitative descriptions of categories or to the identification of patterns of behavior.

Though both of these methods have their merits, there exists another form of study that has been gaining in popularity. Agent-based simulation modeling bridges the gap between deduction and induction. It begins with a formal model as in deduction, but uses the model to generate data from which patterns can be observed in similar fashion to the inductive method. In its most basic form agent-based modeling involves creating a model of an individual and placing many instances of that individual into a virtual environment and watching the behavior of the group unfold over time. Because the individuals are defined by mathematical and computational formalisms, the results are clear, qualitative, and objective (Goldstone & Janssen 2005). Although this might not yield results as concrete as closed-form mathematical equations, they allow for the investigation of complex non-linear models.

The power of these simulations is that we can see patterns *emerge* at the group level that were never programmed in at the individual level. One of the defining characteristics of a complex adaptive system is the process of *emergence*. Although the exact definition of emergence is often debated, we can see it when we see water flowing down stream and forming vortexes as it passes over obstacles. None of the water molecules in the flow are following a rule that tells them to join in on this circular dance; they only respond to local interactions. Nonetheless, when billions of the molecules interact in the stream, whirlpools *emerge* and describe the behavior of water, which is just a large grouping of water molecules. See Panzarasa & Jennings (2006) for more complete definition of emergence.

There are several characteristics that are often found in multiagent systems. These have been adapted from those given by Goldstone & Janssen (2005):

1) *Computational description at the level of agents* – The rules of the models are coded at the level of the individual
2) *Local interactions* – Agents can only directly communicate with others that are limited to a certain spatial distance, or a communication link in a network. All other communication takes place as stigmergic interactions. These interactions involve making changes in the local environment that can be observed by other agents who occupy a similar region of space in the future.
3) *Autonomy of agents* – Agents are not forced into action by other entities, but instead act in response to their environment. They operate with out intervention from external systems.
4) *Agents exist in some environment* – Agents have some environment that they can sense and/or act upon. This is usually a 2D or 3D spatial environment, but may also be a network, or a game/market environment.

Multiagent systems became a viable method of simulation in the 1990’s due to advances in computing technology. Agent-based models are being embraced in many fields of social science as well as computer science, complex adaptive systems, artificial life (Jennings, 2000; Holland, 1995; Epstein & Axtell, 1996). These fields focus more of the theoretical side of ABM. Many researchers are exploring these models to see what behaviors they produce rather than trying to yield specific observed phenomena. Most work that utilizes ABMs to explain group behaviors have been done so in reference to animal behavior. Schooling fish, swarming ants, and herding livestock have been studied in great detail (Kennedy & Eberhart, 2001). Even though it represents a small portion of the work done in ABM, there exists a significant body of work where researchers have modeled collections of individual humans to describe or predict behaviors in groups of people.

Modeling what we see

Agent-based modeling really began with seeking explanations for phenomena observed in groups we see in our everyday life. By creating these models, researchers can investigate the properties of the individuals that give rise to observed group dynamics. This applies to groups in specialized environments (Farkas & Vicsek, 2006, Bonabeau, 2002), groups in everyday activities (Helbing, Molnár, Farkas, & Bolay, 2001), and entire societies (Schelling, 1971). In fact, the use of ABM techniques in social sciences remains most popular in the investigation of such observed behavior.

Thomas Schelling (1971) was one of the first researchers to use a multiagent model to gain better insight into something he observed in society. Although primarily motivated by racial segregation, Schelling wanted to see how individual discriminatory behaviors (rather than organizational differences or imposing forces) influence segregation regardless of the nature of the separation. Beginning with a simple linear distribution of +’s and O’s, Shelling gave each one the same ability to shift positions based on the ration of same/different neighbors within a distance of 2 spaces. Schelling saw that a slight preference for living near others like you leads to large segregated groups. He was able to see how these groups changed as a function of the initial ratio of different types as well as the limitations on movement of agents. In the more realistic two-dimensional simulations, Schelling was able to see influences of varying degrees of ‘tolerance’ for other ‘races’ as well as the desire for integration. He was able to formulate very explicit relationships between these factors just from testing the simulations, an amazing feat given the lack of computing power in 1975. In fact Schelling asks his readers to verify his results using paper separated into grids and a couple rolls of different coins. Although Schelling was able to see the relationships in his simulation, it is very difficult to test the external validity of these findings. This ground breaking study represents how many sociologists use ABM today.

Researchers can still make useful conclusions in spite of this lack of external validity by sticking to the KISS (keep it simple, stupid) principle. Schelling’s model involved extremely simple agents with simple rules about how to move. We can be still reasonably confident that this represents a simple model of at least one influence of racial segregation in human societies if we believe people have preferences for being near others of the same race.

In his ground breaking research on cooperation, Robert Axelrod (1997) used agent based models of an iterated Prisoner’s Dilemma (IPD) task to explore strategies for promoting cooperation in international politics. The Prisoner’s Dilemma is a game in which 2 players are faced with a binary choice. According to the story of the problem, the two players are prisoners being questioned about a crime. They have the choice to cooperate with the other player and keep quiet or defect and implicate the other player. The payouts are arranged such that defecting is the dominant strategy (regardless of what the other player does in any particular turn, the dominant strategy yields a higher payoff). However, during multiple iterations of the game between two players, they can both gain larger rewards if they mutually abandon the dominant strategy and cooperate with each other. Table 1 demonstrates the payoffs in the Prisoner’s Dilemma game.

Table 1: Payoffs in a Prisoner's Dilemma task listed in the order A,B

	Player B	
Player A	Cooperate	Defect
Cooperate	3,3	0,5
Defect	5,0	1,1

Axelrod initially studied the success of different strategies by having leading scientists and computer hobbyists submit strategies to a regulated tournament (Axelrod, 1984). It was found that a simple strategy called TIT-FOR-TAT (TFT) was the most successful. The TFT strategy involved initially cooperating, but every successive move matched the opponent's previous move. The prisoner's dilemma task is a simple environment that captures the tensions of selfish behaviors that provide immediate gratification and cooperative behaviors that provide long term benefits (Axelrod, 1997). For this reason, the iterated prisoner's dilemma task is a great environment for studying how people deal with that tension.

Wedekind and Milinski (1996) have shown that people used one of two different cooperative studies. The first is a generous TFT strategy that is similar to TFT, but cooperates with some probability when the opponent defects. In a noisy situation where mistakes can be made, this generous strategy reduces the risk of a series of retaliations until another mistake is made. The second strategy used is basically a Win-stay, Lose-change strategy that has been nicknamed Pavlov. Using evolutionary algorithms, Nowak & Sigmund (1992) demonstrated the generous TFT strategy emerges as the dominant strategy when players of the Prisoner's dilemma make simultaneous moves. When a more natural game of alternating games is played, the Pavlov strategy emerges as dominant (1993). In the human experiments, the generous TFT strategy was observed in 30% of the subjects tested while the other 70% used a Pavlovian strategy. This distribution was observed regardless of whether a simultaneous or alternating game was being played. However, those that adopted a generous TFT strategy were more successful in the simultaneous game, whereas those using the Pavlov strategy had greater success in the alternating game.

The work on cooperative behaviors shows how simulations inspired the observation of humans. It is also possible for observations of social behaviors in the environment to inspire modeling efforts to explain observed group dynamics. One such phenomenon is walking. Pedestrian movement is easy to observe in cities throughout the world. Moreover, we can observe how specific changes in environment can influence pedestrian movement. Helbing et al. (2001) used an agent-based simulation to explain previously observed self-organization phenomena in pedestrians. These phenomena include the spontaneous formation of directional lanes, the oscillation of the direction of pedestrian flow at bottlenecks, and the formation of lanes when stationary pedestrian crowds need to be crossed. Pedestrians were modeled as fairly simple agents reacting to the environment in a fairly automatic way. They were treated as particles flowing in response to behavioral forces. These forces included a desired velocity, border avoidance, pedestrian avoidance, and social attractive forces. Parameters for the optimal behavior were learned using an evolutionary strategy. Although the details about the motion of any individual was meaningless in the absence of real destination and time constraints, the flow of large populations of pedestrians could be modeled. Without any assumptions of strategic considerations, communication, or imitative behavior in the simulations, the self-organization of agents was observed (see figure 1). Although all forces were programmed as symmetrical in respect to right-hand and left-hand, the non-symmetric formation of lanes was observed. Although there may be a cultural preference for one side over the other, no preference is needed to initiate the division. In simulated bottleneck situations where groups of simulated pedestrians need to go through a single opening from both directions, the direction of the flow of agents through the opening oscillated back and forth. When two openings are present, each opening was utilized by the simulated pedestrians walking in a single direction. The last significant observation of the spontaneous formation of roundabout patterns when pedestrians intersected at a four-way intersection. All of these observations of the simulation have analogous patterns that have been seen in groups of real pedestrians.

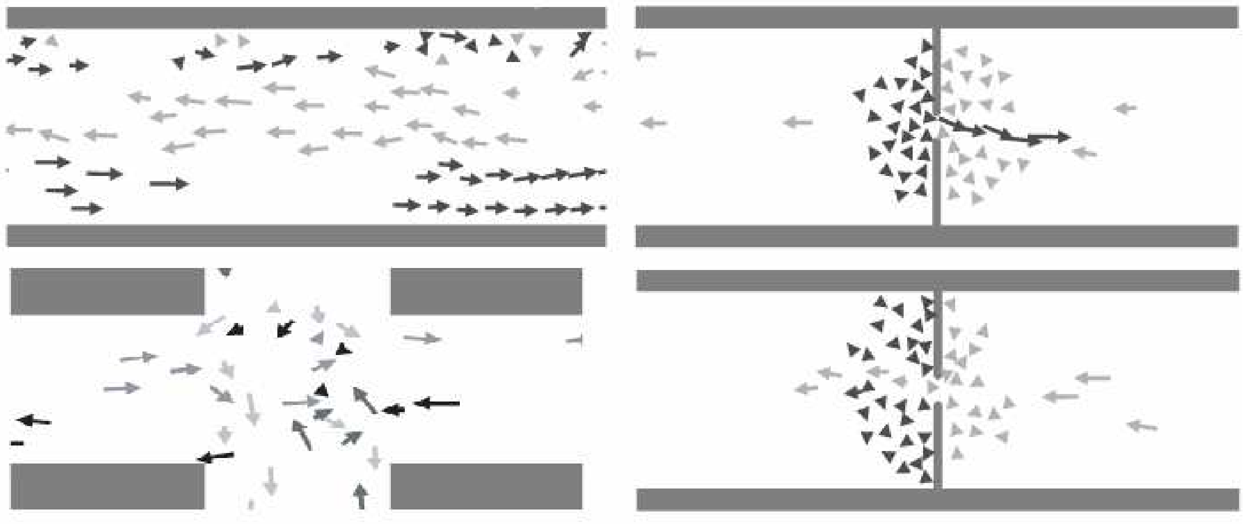

Figure 1: Simulations showing emergent patterns such as lane formation (upper left), oscillation of flow direction through a bottleneck (right), and rotational flow at 4-way intersection (bottom left). Adapted from Helbing, Molnár, Farkas, & Bolay, 2001.

Another readily observable example of self-organization is that of the Mexican wave. The Mexican wave refers to the systematic standing and sitting of spectators in a stadium so as to cause a wave of standing to travel around the arena. Although the event seems as simple as having many individuals merely stand up with their arms above their heads slightly after a neighboring fan, the Mexican wave is much more complex (Farkas & Vicsek, 2006). For instance, what does it take for a stable wave to start? Why does the wave always go in one direction, though not always the same direction? Farkas & Vicsek investigated the phenomenon and generated a model to try to answer such questions. They first began by observing video recordings of 15 different waves and collecting responses about the wave from an online survey. The videos showed waves that were one-directional where 7 propagated clockwise and 8 propagated counter-clockwise around the arena. The surveys suggested that the relative position of the initiators as well as obstacles near the starting point influenced the direction of the wave. Farkas & Vicsek concluded that participants in the wave relied on local cues (their neighbors) as well as distal cues (the wave as a whole) when deciding when to stand and cheer. Utilizing an excitable medium model (Farkas, Helbing & Vicsek, 2002) the authors concluded that it is the influence of global observations that support propagation in a single direction and stifle the wave's movement in the opposite direction. Although the model only required a small group of about 10 individual agents to initiate the wave, a real wave requires strong social cues that are not fully modeled.

How useful are these models?

It is completely fair to question the usefulness of a simulation that merely recreates patterns we observe in groups of people. Can we actually conclude anything about human cognition from these recreations? Later on I will point out how these simulations are being combined with laboratory experiments to make more definitive conclusions. However, even without experimental manipulation, these types of simulations are being used in many different ways. Civil engineers, event coordinators, the New York Stock Exchange, and several large corporations rely on similar simulations when making significant job related decisions (Helbing, 2001; Bonabeau, 2002). Placing many instances of a model in the environment and then observing simulated group dynamics that match observed dynamics in the real environment can boost confidence in the model of the individual. In addition, when the group is a complex, adaptive system, the nonlinearities inherent in the system make it difficult to predict without simulation. With confidence in your simulation, you can see how the people respond to changes in the environment that would be too costly to actually test.

Civil engineers utilize the predictions from traffic simulators to design roads and highways that minimize traffic congestion (Erol, Levy, &Wentworth, 1998). Poorly planned roadways can also be extremely expensive. Because of the nonlinear dynamics, sometimes changes to an environment lead to non-intuitive results. With an accurate simulator, engineers can see the effects of additional lanes, changes in speed limits, additional onramps, etc on a particular road system. They can also simulate different times of day with different densities of motorists. These observations lead to more confidant decisions and more productive tax dollars. A team at the Los Alamos National Laboratory has spent several years developing the very detailed TRansportation ANalysis SIMulation System (TRANSIMS) software package (Bonabeau, 2002). Using detailed census data, researchers were able to simulate a 25 square mile region of the Dallas/Fort Worth area. Results suggested that improving local arterials would better improve congestion than adding lanes to the highway. A larger simulation of a region of Portland, OR is now being conducted.

When panic overruns a large group of people and induces a stampede, individuals can easily get crushed or trampled to death while trying to escape. When faced with danger, people often moved faster than usual and tend to follow the behavior of others. This often leads to overlooked secondary exits and clogged doorways (Helbing, Farkas, and Vicsek, 2000). Helbing, Farkas, and Vicsek (2000) built a simulator that could account for many of these dangerous situations. They were able to make predictions that quantitative theories of group dynamics failed to make. The model is a variation of the generalized force model they used to describe pedestrian movement. For the escape panic model, physical forces like friction and body force were introduced. A detailed model allows for the exploration of ways to alleviate the dangers of escape panic. Real-life experiments of such a spontaneous and dangerous event are not practical. Limited evidence exists regarding human escape from panic situations in small groups of 4 to 7 people (Kelley, Condry, Dahlke, & Hill, 1965). Because of a lack of real-life data, the group behavior of the simulation can not be quantitatively verified. Nonetheless, as will be further explained later, specific flows and counter-flows of pedestrians in normal conditions and in time sensitive situations have been observed in controlled settings (Helbing, Buzna, Johansson, & Werner 2005).

It is now becoming popular for companies to use ABMs to study customer flows. Once a satisfactory simulation exists, businesses can seek out methods of optimizing sales by influencing when and where customers proceed. Axtell and Epstein have generated a simulation of theme parks for a major theme park resort company called ResortScape (Bonabeau, 2002). The ABM model allows a fast, relatively inexpensive method of watching what might occur if particular rides are turned off, how to distribute rides around the park, when hours should be extended, or even what type of wait times might be tolerated. Because the behavior of any particular customer is dependent on the actions and desires of many other people, the actual flow of customers and the flow of money are emergent properties within the park. The ABM gives a natural method to study these patterns.

Another software package called SimStore (www.simworld.co.uk) was created (Casti 1997) to model a British supermarket. The model includes optimization routines to try to balance the customers' desire for fast shopping and the store owner's desire to encourage impulse buying by having customers pass by as many products as possible. The department store Macy's hired PricewaterhouseCoopers to determine what the right number of salespeople on the selling floor would be (Bonabeau, 2002). An ABM was produced that not only answered that question, but allowed users to see the effects of moving around specific store displays. Also, certain patterns such as 'microbursts' of shopping where someone shops for an outfits then shops for matching accessories were observed. These types of behaviors are not visible from traditional spreadsheet analysis.

The sorts of simulations that have been described so far can be used to learn about emergent group behaviors. Unfortunately, they are not ideal for the scientific exploration of human cognition. They do not allow for experimental manipulation of variables in a controlled environment. Behaviors that rely on thousands of individuals in specialized environment are difficult to manipulate experimentally. Situations such as those dealing with escape panic or even traffic flow could be dangerous to experiment with. It can be difficult to even gather small groups of 30 to conduct group experiments. Nonetheless, several researchers have begun to enter this field using new experimental methods.

Observations of groups can influence models of individuals

Studying the strategies and decision making methods used in economic markets are much easier to study in the laboratory because the do not require the physical space for movement. Although they do not have the same spatial patterns that can be observed, fluctuations and trends in a market represent emergent behaviors. Such environments are traditionally studied using neoclassical economics. Neoclassical economics assumes that in a given environment, people always act rationally to try to maximize their rewards (McCain 2003). This assumption of rationality eases some of the predictions that can be made by limiting possible actions. It also provides an ability of measuring efficiency by comparing actual behaviors to optimal behaviors. However, many situations have been observed that question this assumption of rationality (Duffy 2005, Jager & Janssen 2003). Let's first look at some attempts to use ABMs to deduce necessary characteristics of agents to produce behaviors observed in experimental settings. The rational choice to maximize rewards does not seem to be a necessary assumption.

KISS – lessons from the double auction

One experimental environment that has been recently explored using ABMs is that of the double auction. In a double auction, sellers offer prices at which they are willing to sell a quantity of a good while buyers offer prices at which they will buy a quantity of a good. If no match is made, the players may update their offers. In the experiments, participants are told whether they are a buyer or a seller, and are told of their own valuations and costs, but not given any information about other participants (Duffy 2005). In most double auctions, only the best bid and ask prices are shown at any given time. When a match is made, goods are exchanged and those items are not available for the rest of the selling period. The entire history of the auction is available as public knowledge. The main result that has been consistently replicated in such an environment is the rapid convergence to the competitive equilibrium price. The competitive equilibrium price is where supply and demand are equal.

Much of the modeling work done in this environment has been performed by Gode & Sunder (1993, 1997, 2004). Gode and Sunder originally hypothesized (1993) that it might simply be the rules of the double auction experiment itself that led to the fast convergence to the competitive equilibrium price. To test such a hypothesis, they designed an agent-based simulation where the agents randomly chose bids and asks over some range. This simple agent was dubbed a zero-intelligence (ZI) agent. The simulation consisted of 6 simulated sellers and 6 simulated buyers. Two variations were run. In the first unconstrained variation (ZI-U), agents randomly selected bids/asks from a uniform distribution U[0,B] where be was a value exceeding the highest valuation among all buyer agents. The second constrained variation (ZI-C) simply prevented agents from making unprofitable trades by having buyers bid from the distribution U[0, minimum ask] and sellers ask from the distribution U[maximum bid, B].

The results of the simulation showed that the ZI-U agents led to extreme volatility in transaction prices and never converge upon the competitive equilibrium. However, The ZI-C agents did quickly converge to equilibrium. Also this is only a qualitative similarity, it is amazing that it comes from agents that have no memory of past events and choose valuations randomly. A more quantitative measure comes from looking at the market efficiency as calculated by summing the total profit earned over all trading periods and dividing by the maximum possible profit. The average efficiency of human subjects was calculated at 97.6%, the average efficiency of ZI-C agents was 98.7% and the average efficiency of ZI-U agents was 78.3%. It seems as if the restriction imposed by the ZI-C agents is the source of the main findings in human experiments. Gode and Sunder revisited this environment (2004) to show that the ZI-C agents also match human data when a non-binding price ceiling for bids and asks is set just above the equilibrium price.

Gode and Sunder specifically point out that they are not modeling human cognition (1993, 2004). Their work shows that the simple interactions that exist as part of the environment can lead to complex emergent behaviors without the need for intelligent agents. This is an important lesson for all users of ABMs that may try to improperly credit agent behaviors for producing group dynamics. When trying to generate useful multi-agent simulation, one has to balance the need for simplicity with the desire to follow established models of psychology.

Common Pool Resources

Janssen and Ostrom (2006) describe a particular market situation that has been studied in the laboratory over the last twenty years. The experiments were designed to look at how members of a society share common resources. At the beginning of an experiment, participants are allocated a set of tokens that they must allocate between two markets. The first market has a fixed return, while the second market has a value based on the collective investments of the group. Each person has and endowment of ω and chooses to place a portion x_i of the endowment into the second market (the rest $\omega - x_i$ goes into Market 1). The pay off for Market 2 follows a quadratic production function dependant on the sum of all participants' contributions to the second market. In the common-pool resource experiments without communication, it was observed that the average amount of tokens collected will approach a non-cooperative equilibrium where no one can change their strategy without receiving a direct penalty. This is referred to as a Nash equilibrium state and is the result when all players of a game are rational agents and assume the other players are too. When participants are allowed to communicate, they break from the Nash equilibrium and average harvests decrease to a cooperative level (Ostrom, Gardener, & Walker 1994). This behavior is not consistent with the use of a rational choice model of individual behavior.

Several researchers have used multiagent systems to explore what kinds of models of individual behavior explain the results seen in the common-pool experiments. Deadman (1999) modeled individuals based on strategies collected from exit interviews during the common-pool experiments. He was able to reproduce the fluctuations seen in Market 2 during real-life experiments. Communication was added to the simulation (Deadman, Schlager, & Gimblett 2000) so that all agents had a similar view of what strategies worked. This led to investment levels approaching the optimal level for cooperative behavior just as they did in the real-life experiments.

Jager and Janssen (2002) used a more general framework of psychological theories to determine what qualities were necessary to get the results seen in the experiments without communication. They were able to conclude that five elements of individual cognition were necessary for the common-pool resource environment.

1) A focused attention on the rewards the entire group receives
2) Preferences for particular distributions of outcomes for everyone
3) Satisfying behavior
4) Exploratory behavior when no change in payoffs is seen for a number of rounds
5) Heterogeneity of needs among the agents

Environments that depend on markets have been a focus of neoclassical economics. However, humans often make economic decisions in situations where the external circumstances include direct communication with others in a similar situation. Neoclassical economics treats people as if they are responding to a market as a value that changes over time. Game theory allows the same situation to be thought of as a game being played against other people playing the same game. The combination of our choices determines the rewards we receive. The games derived and studied under game theory are metaphors for common human interactions including arms races, market competition, and environmental pollution (McCain 2003). Studying such games gives an insight into human behavior that is much more realistic. Unfortunately, the analysis techniques that come from game theory rely on an assumption of rationality just like the neoclassical economics. Although the complexity of the games allow for the question of what constitutes rational behavior (maximizing your own reward, or the cumulative reward of a set of players), this assumption may not be necessary.

Observations of individuals can influence models of groups

We have just looked as some ways in which ABMs are used as tools to make some conclusions about agents based on observed group dynamics. However, multi-agent simulations have another use that has yet to be explored to its full potential. But utilizing established models of human cognition, we can use MASs to explore expected group behaviors given what we know about the individual. The applications mentioned such as ResortScape and SimStore are doing just this, but are not being verified in controlled environments.

Erev and Roth started with a model of reinforcement learning, and then tested how agents following the model behave in multiple environments (Erev & Roth, 1998). Reinforcement learning involves trying to maximize a delayed reward signal through a trial and error exploration of action (Sutton & Barto, 1998). The reinforcement learning model was tested in three different types of sequential games (a market game, a best shot game, and an ultimatum game). Each of the games has a perfect equilibrium where a single player gains most of the reward (Duffy, 2005). The agents in the simulation, however, displayed much more varied results. In two of the games, the Nash equilibrium points were gradually learned over time. However, in the ultimatum game, the Nash equilibrium point was found right away. This pattern of behavior was also observed in an experiment involving humans (Ochs, 1995).

Helbing et al. (2005) conducted experiments using students to verify and enhance aspects of their social force model of pedestrian dynamics. Tables were set up to provide paths with bottlenecks of different lengths to see how they influenced uni-directional and bi-directional traffic (see figure 2). They surprisingly found that the bi-directional traffic was much more efficient when examining the average number of people able to pass through the obstacle per time unit. They were also able to artificially create cross flows at specific angle to investigate how segregated bands of traffic pass through one another in a seemingly coordinated fashion. Even though the artificial environment is not a perfect representation of the walls and more diversified crowd seen in public streets, the data collected helped to refine parameters in the social force model originally created by observing pedestrians in public. Simulations initially made to explore the clogging of bottlenecks by a pushy crowd suggested that the placing of an obstacle just in front of the entrance can greatly increase flow. Although this seems counter-intuitive, Helbing et al. (2005) were able to verify that prediction in an experiment involving groups of 20 students placed equidistant from a room’s exit and asked to exit as quickly as possible. In one setup, there was no obstacle, while in the other an obstacle was placed in front of the exit. The knowledge gained from such simulations has helped Helbing et. al. (2005) make specific suggestions to improve designs for pedestrian facilities. To increase the safety in situations that are prone to panicking crowds, guidelines have been set up for exits of theaters, classrooms, & Lecture halls, as well as staircases (including stadium exits), and queues at entrances. Using a similar simulation, guidelines have been established for the safe evacuation of a bi-level ship or hotel (Werner & Helbing, 2003).

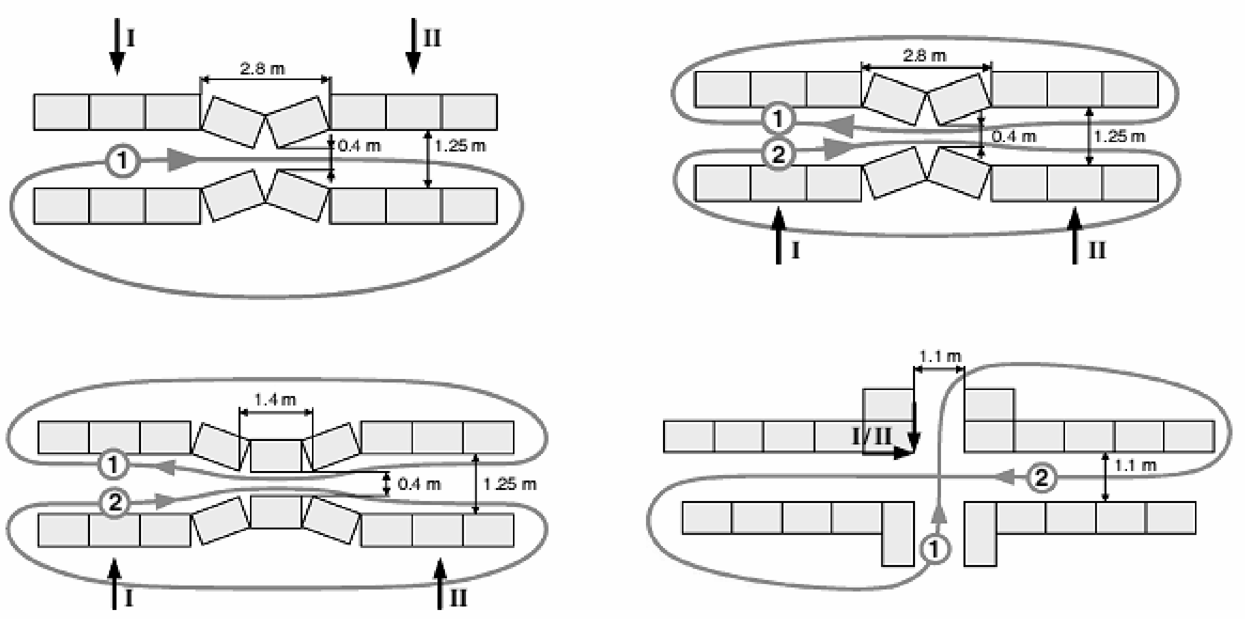

Figure 2: Human experiments testing uni-directional pedestrian flow through a bottleneck (upper left), bi-directional pedestrian flow through a bottleneck (upper right), bi-directional pedestrian flow through a longer bottleneck (bottom left), and perpendicular cross flow (bottom right). Adapted from Helbing, Buzna, Johansson, & Werner 2005.

Recently, there has been growing interest in utilizing complex models of human cognition to represent agents in multiagent simulations. West, Lebiere, & Bothell (2006) utilized the ACT-R architecture of cognition to develop an agent to play the paper-rock-scissors game. ACT-R is a productions system that maintains separate procedural and declarative memory stores using symbolic representations (Anderson & Lebiere, 1998). Sub-symbolic processes govern which procedure and what chunks of declarative knowledge will be utilized in an attempt to complete some goal. These sub-symbolic mechanisms utilize learning mechanisms to increase the chances of successfully reaching the desired goal in future trials. Although a fairly simple model of game play was used, results from repeated trials had similar qualitative results as human data such as the appearance of random behavior.

Work done by Naveh & Sun (2006) provides another example of the use of a complex cognitive architecture in a multiagent system. The task was and organizational decision making task where members of the organization had to classify an object as one of three different categories based on 9 features. Each member of the organization only received information regarding 3 of the features. Team structures were studies where all agents were autonomous and had an equal vote. In addition, hierarchical group structures were studied where agents had a direct chain of command and the agent at the top of the hierarchy made the final decision. The model was created within an architecture called CLARION. CLARION is a highly modularized system that makes a clear distinction between an explicit memory with a symbolic representation and an implicit memory with a neural network representation (Sun, 2006). Within such an architecture, complex models are described by large code sets and the specific model was not published. It was shown, however, that both teams and hierarchies made up of CLARION agents closely matched the success levels of humans in the classification task.

Conclusion

Despite the many uses for ABMs in studying human cognition, multiagent systems are not widely used in the field of psychology. Most of the examples we have looked at have come from economists, political scientists, & computer scientists, but new tools are changing the accessibility of ABMs for many new fields.

Experimental psychologists rely on heavily controlled environments to identify sources of variation in human behavior. Even simple experiments such as those performed by Helbing et al (2005) to investigate simple walking behaviors are quite complex. It is extremely difficult to control and experimental environment supporting the actions of many people. Nonetheless, the explosion of computer network accessibility allows large groups o people to interact through controlled virtual environments. A recent study has utilized a virtual environment to exam human behavior in simulated foraging and path formation experiments (Goldstone, Ashpole, & Roberts, 2005). Participants control the movement of their own agent in a two dimensional world while searching for food particles. Experimenters can control what participants are able to see (food eaten, all food, other players, etc) as well as when and where food will be located. The initial experiments were done to explore undermatching, where individuals distribute themselves such that a smaller proportion of individuals is found in more profitable regions than one would expect from an optimal distribution. The environment has since been used to explore human behavior in path formation experiments (Goldstone, Jones, & Roberts, 2006). Previous path formation data has only come from observations of trails and paths formed by humans in natural environments. Helbing, Schweitzer, Keltsch, & Molnár (1997) developed a model to simulate these observed patterns. Within the virtual environment, simulated agents exist in settings that are much more comparable to humans. Therefore, when Goldstone et al. found that the model makes predictions that are very near to human behaviors, they can be more confident in the model’s accuracy. Virtual environments such as these will help psychologists verify group behaviors in controlled settings. This can be done initially, to inductively identify models of individual behavior that can be tested using an agent based model of the task, or to test predictions of group dynamics based on models of individuals.

Agent based simulations also provide an under tapped resource for the scientific study of specific computational models. Agent-based models provide a novel form of model comparison.

Traditional model comparison compares how close model behaviors match observed human behaviors by means of qualitative measurements of performance. Sometimes these measures are not available, or not appropriate for group interactions. Comparisons based on maximum likelihoods can not be used since interacts do not keep observations as independent measures (Duffy 2005). Also, whereas two models may produce behavior that is indistinguishable at the level of the individual, emergence properties in group settings may provide clear ways of differentiating models. Naveh & Sun (2006) used a simulation of organizational decision making to compare the performance of a model developed in the CLARION architecture to one developed in ACT-R.

Multiagent simulations can get quite complex. Jager & Janssen (2003) push the need for a meta-theory of agents that is based on established psychological theories. The goal is to provide a common framework that allows researchers to develop simple agents to test the fundamental aspects of social interactions. Very much the way Allen Newell (1990) proposed cognitive architectures as a way to unify models of cognition, Jager & Janssen hope to unify multiagent systems. As agent-based modeling becomes more accepted as a means to study the social science, we may see others follow in their footsteps. Until then, researchers will continue to explore this exciting field in their own unique ways.

References

Anderson, J. R. & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Erlbaum.

Axelrod, R. (1984). *The Evolution of Cooperation*. New York: Basic Books.

Axelrod, R. (1997). *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton , N.J.: Princeton University Press.

Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences of the United States of America, 99* (Suppl. 3), 7280-7287.

Casti, J. (1997). *Would-Be Worlds: How Simulation is Changing the World of Science*. New York: Wiley.

Deadman, P. (1999). Modelling Individual Behaviour and Group Performance in an Intelligent Agent-Based Simulation of the Tragedy of the Commons. *Journal of Environmental Management, 56*, 159-172.

Deadman, P. J., Schlager, E., & Gimblett, R. H. (2000). Simulating Common Pool Resource Management Experiments with Adaptative Agents employing Alternate Communication Routines. *Journal of Artificial Societies and Social Simulation, 3*(2).

Duffy, J. (2005). Agent-based models and human subject experiments. in K. L. Judd & L. Tesfatsion (Eds.), *Handbook of Computational Economics*, Amsterdam: Elsevier.

Epstein, J. M. and Axtell, R. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. Brookings Institution/MIT Press.

Erev, I., & Roth, A. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review, 88*, 848–881.

Erol, K., Levy, R., & Wentworth, J. (1998). *Application of Agent Technology to Traffic Simulation*, http://www.tfhrc.gov/advanc/agent.htm.

Farkas, I., Helbing, D., & Vicsek, T. (2002) Mexican Waves in an excitable medium. *Nature, 419*, 131-132.

Farkas, I. J., & Vicsek, T. (2006). Initiating a Mexican wave: An instantaneous collective decision with both short- and long-range interactions**.** *Physica A: Statistical and Theoretical Physics, 369*(2), 830-840.

Gode, D. K. & Sunder, S. (1993), Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy, 101*, 119-137.

Gode, D. K. & Sunder, S. (1997). Lower bounds for efficiency of surplus extraction in double auctions. In D. Friedman and J. Rust (Eds.), *The Double Auction Market: Institutions Theories, and Evidence, Santa Fe Institute Studies in the Sciences of Complexity, Vol. 14*. (pp. 199-220). Reading, MA: Addison-Wesley.

Gode, D. K. & Sunder, S. (2004). Double auction dynamics: structural effects of non-binding price controls. *Journal of Economic Dynamics and Control, 28*, 1707-1731.

Goldstone, R. L., Ashpole, B. C., & Roberts, M. E., (2005). Knowledge of resources and competitors in human foraging. *Psychonomic Bulletin & Review, 12,* 81-87.

Goldstone, R. L., & Janssen, M. A. (2005). Computational models of collective behavior. *Trends in Cognitive Science, 9*, 424-430.

Goldstone, R. L., Jones, A., & Roberts, M. E. (2006). Group path formation. *IEEE Transactions on System, Man, and Cybernetics, Part A, 36*, 611-620.

Helbing , D. (2001). Traffic and related self-driven many-particle systems. *Reviews of Modern Physics, 73*, 1067-1141.

Helbing, D., Buzna, L., Johansson, A., & Werner, T. (2005). Self-Organized Pedestrian Crowd Dynamics: Experiments, Simulations, and Design Solutions. *Transportation Science, 39*, 1.

Helbing, D., Farkas, I., & Vicsek, T. (2000). Simulating dynamical features of escaping panic. *Nature, 407***,** 487–490.

Helbing, D., Molnár, P., Farkas, I. J., & Bolay, K. (2001). Self-organizing pedestrian movement, *Environment and Planning B*, *28*, 361-383.

Helbing, D., Schweitzer, F., Keltsch, J., & Molnár, P. (1997). Active walker model for the formation of human and animal trail systems. *Physical Review E*, *56(*3), 2527–2539.

Holland, J. (1995). *Hidden Order: How Adaptation Builds Complexity*. Perseus.

Hutchins, E. (1995). How a cockpit remembers its speeds. *Cognitive Science*, *19*, 265-288.

Jager, W., & Janssen, M. A. (2002). Using artificial agents to understand laboratory experiments of common pool resources with real agents. In M. A. Janssen (Ed.) *Complexity and Ecosystme Management. The Theory and Practice of Multi-Agent Systems* (pp. 75-102). Cheltenham, UK: Edward Elgar.

Jager, W., & Janssen, M. A. (2003). The Need for and Development of Behaviourally Realistic Agents. In: J.S. Sichman, F. Bousquet & P. Davidson (Eds.) *Multi-Agent Based Simulation II*. Berlin: Springer. P. 36-49.

Janssen, M. A., & Ostrom, E. (2006). Governing Social-Ecological Systems. In L. Tesfatsion & K. L. Judd (Eds.), *Handbook of Computational Economics, Volume 2.* (pp. 1466-1510). Amsterdam: Elsevier.

Jennings, N. R. (2000). On Agent-Based Software Engineering. *Artificial Intelligence, 117*(2), 277–296.

Kelley, H. H., Condry J. C. Jr., Dahlke A. E., & Hill A. H. (1965). Collective behavior in a simulated panic situation. *Journal of Experimental Social Psychology*, *1*, 20-54.

Kennedy, J. F., & Eberhart, R. C., Shi, Y., (2001). *Swarm intelligence*. San Francisco: Morgan Kaufmann.

McCain, R. (2003). *Game Theory: A Non-Technical Introduction to the Analysis of Strategy.* South-Western.

Moss, S. (2006). Cognitive Science and Good Social Science. In R. Sun (Ed.), *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation* (pp. 393-400). New York: Cambridge University.

Nevah, I., & Sun, R. (2006). Simulating a Simple Case of Organizational Decision Making. In R. Sun (Ed.), *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation* (pp. 124-150). New York: Cambridge University.

Newell, A. (1990). *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.

Nowak, M.A. & Sigmund, K. (1992). Tit for tat in heterogeneous populations. *Nature, 355*, 250-253.

Nowak, M. A. & Sigmund, K. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature*, 364, 56-58.

Ochs, J. (1995), Games with unique, mixed strategy equilibria: an experimental study. *Games and Economic Behavior, 10*, 202–217.

Ostrom, E., Gardener, R., & Walker, J. (1994). *Rules, Games, and Common-Pool Resources*. Ann Arbor: University of Michigan Press.

Panzarasa, P., & Jennings, N. R. (2006). Collective Cognition and Emergence in Multi-Agent Systems. In R. Sun (Ed.), *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation* (pp. 401-408). New York: Cambridge University.

Sawyer, R. K. (2003). Artificial Societies: Multiagent Systems and the Micro-Macro Link in Sociological Theory. *Sociological Methods and Research 31*, 325-363.

Schelling, T. (1971). Dynamic Models of Segregation. *Journal of Mathematical Sociology, 1*, 143-86.

Sun, R. (2006). Prolegomena to Integrating Cognitive Modeling and Social Simulation. In R. Sun (Ed.), *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation* (pp. 3-26). New York: Cambridge University.

Sutton, R. S., & Barto, A. G. (1998).*Reinforcement learning : an introduction.* Cambridge, MA: MIT Press.

Wedekind, C., & Milinski, M.(1996). Human cooperation in the simultaneous and the alternating prisoner's dilemma: Pavlov versus generous tit-for-tat. *Proceedings of the National Academy of Sciences of the United States of America, 93*, 2686-2689.

Werner, T., & Helbing, D. (2003). The social force pedestrian model applied to real life scenarios. In E. R. Galea (Ed.), *Pedestrian and Evacuation Dynamics*. (pp. 17–26). London: CMS Press.

West, R. L., Lebiere, C., & Bothell, D. J. (2006). Cognitive Architectures, Game Playing, and Human Evolution. In R. Sun (Ed.), *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation* (pp. 103-123). New York: Cambridge University.