

Collective Behavior

Robert L. Goldstone,^a Todd M. Gureckis^b

^a*Indiana University, Bloomington*

^b*New York University*

Received 13 January 2009; received in revised form 26 April 2009; accepted 28 April 2009

Abstract

The resurgence of interest in collective behavior is in large part due to tools recently made available for conducting laboratory experiments on groups, statistical methods for analyzing large data sets reflecting social interactions, the rapid growth of a diverse variety of online self-organized collectives, and computational modeling methods for understanding both universal and scenario-specific social patterns. We consider case studies of collective behavior along four attributes: the primary motivation of individuals within the group, kinds of interactions among individuals, typical dynamics that result from these interactions, and characteristic outcomes at the group level. With this framework, we compare the collective patterns of noninteracting decision makers, bee swarms, groups forming paths in physical and abstract spaces, sports teams, cooperation and competition for resource usage, and the spread and extension of innovations in an online community. Some critical issues surrounding collective behavior are then reviewed, including the questions of “Does group behavior always reduce to individual behavior?” “Is ‘group cognition’ possible?” and “What is the value of formal modeling for understanding group behavior?”

Keywords: Collective behavior; Group psychology; Computational models; Innovation diffusion

1. Introduction

Consider the arbitrarily selected concept of “Spam filter.” Like most of our concepts, it is very much the product of our culture. Despite its seemingly mundane nature, it is the culmination of a rich and complex series of conceptual bootstrappings. To understand this concept requires understanding computers, advertisement, money, attention, e-mail, value, the Internet, and the nature of canned meats. Each of these concepts, in turn, requires understanding many other concepts. No individual, no matter how smart or motivated,

Correspondence should be sent to Robert L. Goldstone, Psychological and Brain Sciences, Indiana University, Bloomington, IN 47405. E-mail: rgoldsto@indiana.edu

would be able to come up with this concept if born and raised in a substantially different culture, such as that of Easter Island circa 1850. An individual is still less likely to come up with this concept in complete isolation from others.

Yet most of the methods of cognitive science tacitly adopt the assumption of individual people as the unit of cognition. Researchers studying concept learning, perception, memory, attention, expertise, neuroscience, and consciousness typically isolate their subjects in cubicles and expose them to materials that they must categorize, recognize, organize, remember, or select. To be fair, there are certainly cognitive scientists studying language, group behavior, and community organization. However, there is still no general recognition of the ubiquitous influence of the collective on the individual's cognition. Structured cognitive behavior can be described at multiple levels, and our thoughts both depend upon and determine the social structures that contain us as elements.

The study of collective behavior is timely for several reasons. First, as many of the articles in this issue attest,¹ there have been recent and important developments in the formal modeling of collective behavior. These models have played a valuable role in sociology (Macy & Willer, 2002), economics (Kirman & Zimmermann, 2001), psychology (Kenrick et al., 2003; Latane & Bourgeois, 2000; Smith & Conrey, 2007), and anthropology (Kohler & Gumerman, 2002). Their utility is not restricted to only predicting individual and group behavior but also in organizing theories, highlighting idealized patterns, and determining what data should be collected next (Epstein, 2008).

A second reason for choosing to explore collective behavior now is that there has been exciting recent progress on empirical tools for measuring and manipulating the collective patterns that people produce. For the laboratory-based psychologist, there are tools that allow moderate-sized groups of people to be connected together via computers, calculators, cell phones, or clicker response systems. These technologies have made it relatively easy for experimenters to collect moment-by-moment data on the decisions of people as they are influenced by the decisions of their peers (Goldstone & Ashpole, 2004; Goldstone & Roberts, 2006; Kearns, Suri, & Montfort, 2006; Mason, Jones, & Goldstone, 2008). Virtual worlds have made it possible to construct rich scenarios in which people can interact in a shared synthetic world, and these interactions can be efficiently recorded and analyzed (Bainbridge, 2007; Friedman, Steed, & Slater, 2007). By creating separate virtual worlds for different groups, replicability is achievable, and the inevitability of group outcomes can be quantitatively assessed (Salganik, Dodds, & Watts, 2006; Salganik & Watts, 2009). Software environments such as Netlogo have greatly simplified the development of "participatory simulations" in which people play the role of agents in a system that can be inhabited by a combination of other real people or artificial agents (Wilensky & Stroup, 1999).

While all of these innovations have revolutionized the collection of experimental data in well-controlled and precisely manipulable laboratory settings, a parallel set of technological advances has radically improved the harvesting of data from real-world sources. Archival data available from online news groups, blogs, social network services, chat groups, and topical communities can effectively be used to explore naturally occurring coalition formation, idea spread, and group evolution (Berger & Heath, 2005). Gureckis and Goldstone (in press) show one application of modeling tools to the now-available large database of all

baby names in the United States over a 150-year period. Other data sets, such as the movement patterns of mobile phone users (González, Hidalgo, & Barabási, 2008) or paper currency (Brockmann, Hufnagel, & Geisel, 2006), have revealed general laws of cultural diffusion. Carley, Martin, and Hirshman (2009) provide a demonstration of how demographic data can be used to constrain a model of media influence. Although these real-world data sets are not as hygienically controlled as their laboratory counterparts, their sheer size often allows factors to be statistically pulled apart even when they cannot be manipulated.

A third and final reason for the recent resurgence of interest in collective behavior is humanity's increasing connectedness. For example, understanding how the consumption of fossil fuels contributes to climate change or how patterns of interconnectivity between groups facilitate the spread of disease highlights the collective challenges that cognitive science can, and should, be in a position to address. After all, many of these societal challenges are a reflection of individual decision making and behavior but manifest at the level of aggregates. Indeed, one of the key lessons from this issue is that patterns of individual behavior do not simply combine to determine the behavior of the group. As a result, a key challenge for cognitive science is to not only understand how one might effectively structure individual incentives consistent with societal goals but also the result this will have when aggregated in larger populations. These two questions are two sides of the same practical question and cannot be considered in isolation from one another, something that the traditional individual-centric perspective of cognitive science often overlooks. Where cognitive science may excel relative to other fields which focus on aggregate outcomes (such as economics or sociology) is that we start with a much richer and realistic model of the individual and can "build upward" leveraging what we, as a field, have learned about complex systems organized at multiple levels (such as the brain).

Another prominent example of this increasing connectivity is apparent on the World Wide Web. It has given us many compelling case studies of the nonlinear group dynamics by which ideas are exchanged and people are connected to one another. The recruitment of friends by friends on social networking sites such as Facebook makes it clear that one of the most valuable resources that a service can offer its users is other users. Masses-produced scholarly works such as Wikipedia make it clear that a large and decentralized collective can still produce highly structured and high-quality information. Sites such as Youtube that provide a repository for user-provided content make it clear that the popularity of a cultural artifact is not only determined by its own intrinsic value but also by its prior history of popularity. Videos that land on Youtube's "Most watched" page are almost guaranteed of becoming still more watched, often times dramatically so.

Although the technology underlying the Web is new, there is a deeper sense in which the Web is simply one of the most recent and socially significant manifestations of people's perpetual drive to become more connected. Through innovations like the printing press, far-reaching transportation systems, and telecommunications networks, our lives have become increasingly intermeshed. Whether this has been good for individuals is debatable, but this question is mostly beside the point for the human collective. In fact, the speed and momentum of the collective's push toward ever greater dependence suggests that it is a social force that is beyond the control of individuals to curtail. Richard Dawkins (1976) has argued that

the interest of humans may be quite distinct from the interest of our genes, resulting in people behaving in ways that are opposed to their own good, but are for the good of their genes. Examples include people sacrificing their lives to save their kin, and people having kin in the first place even though they believe it will reduce their happiness. Analogous to Dawkins' selfish genes below the level of the individual, there are also "selfish teams" above the level of the individual. Both of these levels can cause an individual to behave against his or her own self-centered interests. Historically, the influence of selfish teams seems to be precipitously increasing as societies become more organized and differentiated (Wright, 2001). Cognitive science risks marginalizing its relevance if it ignores the cognitive effects that teams ranging from families, friend networks, companies, communities, political parties, religious groups, and professional groups have on their members.

As mentioned earlier, one of the theses of this issue *Topics in Cognitive Science* is that groups do not simply affect their members' thoughts and behaviors. Indeed, one might go so far as to say that groups of people themselves can be interpreted as information, processing systems (Gureckis & Goldstone, 2006). Whereas individual humans have probably not increased substantially in the complexity of their internal structure, the groups that they are part of have. The Bureau of Labor Statistics currently lists 820 U.S. occupations, each representing a well-differentiated and stable professional role that a person can play in society. As an example of this collective complexity, when we install a virus checker on our computer, we use programs written by teams of people to prevent programs written by other teams of people from incapacitating programs written by still other teams of people. In 1984, there were only about 1,000 devices that could reach our global digital network. By 1992, about 1 million could. In 2008, over 1 billion can. From 1990 to 2003, mobile phone usage and global network usage each rose over 100-fold (Rawlins, in press). While social critics have argued that people are less deeply enmeshed in their local communities (Putnam, 2001), it is undeniable that people are becoming more broadly connected.

2. The dynamics of collective behavior

If we take seriously the premise that groups of people create emergent patterns that have an integrity of their own, then it should not come as a surprise that there are striking "individual differences" across different groups. Although there are several very general patterns that groups of people form, such as positive and negative feedback loops, it is also helpful to characterize critical dimensions of variation. Moussaid, Garnier, Theraulaz, and Helbing (2009) divide groups into those with members that pass information directly versus indirectly. Table 1 compares seven collective scenarios along a set of four alternative attributes.

2.1. Primary motivation of individuals

A traditional assumption in economic models of human behavior is that people behave exclusively in accord with their own self-interest. However, there is a growing appreciation that people are inherently social creatures, and that we often intrinsically care about the

Table 1
A comparison of some illustrative collective behavior scenarios

Scenario	Primary Motivation of Individuals	Interactions Among Individuals	Typical Dynamic of Individuals	Characteristic Outcome
Wise crowds	Selfish	None	None	Crowd average better than average crowd member
Bee swarms to find nest	Selfish and group	Cooperative recruitment	Copying good solutions	Fast and high-quality nest choices
Path formation	Selfish	Stigmergic	Facilitation leads to similarity	Growth of reinforced paths
Basketball teams	Group	Cooperative, high bandwidth	Differentiation	Coordinated play
Common pool resource	Selfish	Competitive	Monitors, sanctions, conflict, and coordination	Exhausted resources without organization
Foraging	Selfish	Competitive	Differentiation and copying	Approximate match to resource distribution
Open source software	Selfish and group	Cooperative	Borrowing and improving	Complex software

welfare not only of each other as individuals but also of the groups to which we belong (Haidt, Seder, & Kesebir, 2008). Our identities extend beyond our own skins, when we bask in pride for the hometown team or take personal umbrage when our colleagues are unfairly attacked. In these and many other cases, who we are is distributed across the many groups to which we belong, and our interests naturally extend beyond ourselves as individuals (Smith & Semin, 2007).

2.2. Interactions among individuals

Emergent group patterns depend crucially on how the agents within the group interact. Two dialectically related interactions are competition and cooperation, which are often synthesized together in the same system. The same motivations that lead to competition between companies, countries, or teams can lead to cooperation among the members of a single company, country, or team. Frequently there is a dynamic interplay between cooperation and competition. This dynamic is well illustrated by a spatial version of prisoner's dilemma in which agents organized in a lattice cooperate with or defect against their neighbors (Nowak & Sigmund, 2004). Pockets of cooperators do relatively well and expand their territory but then are ripe for invasion by opportunistic defectors. Often there is no stable final configuration but rather patterns of cooperation that move across space and time, followed by defectors.

Another form of interaction, called "stigmergy," is only indirectly between individuals (see Moussaid et al., 2009). Stigmergy is a form of indirect communication between agents that is achieved by agents modifying their environment and also responding to these modifications (Dorigo, Bonabeau, & Theraulaz, 2000). This effect has been well documented in ant swarms, in which ants lay down pheromones as they walk that attract subsequent ants

(Theraulaz & Bonabeau, 1995). However, stigmergy is a much more general form of interaction than this, appearing, for example, on the website Amazon when one customer's behavior in buying books X and Y can affect a subsequent customer who has bought X and is told by Amazon that other customers like them have also bought Y. Stigmergy is also evident on the video-sharing site Youtube. Popular videos appear conspicuously on the site's main page, assuring their further popularity (similar to Salganik & Watts, 2009).

2.3. *Typical dynamic of individuals*

The individuals within a group may engage in a number of behaviors that are contingent upon the behaviors of their peers. Both copying, and its converse behavior, differentiation, are common (Kennedy, 2009). Copying is abetted by the uniquely human ability to imitate. Far from being the last resort of dull or dim-witted individuals, imitation is a sophisticated skill requiring advanced cognitive capacities of motor perception, action planning, and analogical reasoning (Blakemore, 2000). Alan Bandura (1965) has referred to imitation as "no trial learning," even faster than the one-trial learning observed in animals who are predisposed to form associations between food tastes and stomach aches. Imitation plays an important role in spreading valuable innovations across a community. In fact, the nature of a culture largely consists of the ideas, artifacts, beliefs, concepts, and values that spread through a community by imitation and assimilation (Gureckis & Goldstone, in press).

There are other situations where people are motivated not to imitate, but rather to differentiate themselves from others. In many situations, the group's overall ability to solve problems is facilitated if the individuals divide the intellectual labor (Bettencourt, 2009), and one's individual lot is also improved if one occupies a relatively uninhabited, unique niche. Auditioning actors, foragers, pioneering scientists, and workers looking for employment all benefit from distancing themselves from other people who would otherwise be competing for the same resources. As with cooperation and competition, the choice between imitation and differentiation is often dynamic and simultaneously partakes of both processes. For example, successful scientists may choose to copy the general area and approach of a well-received piece of scholarship, but once in the general vicinity of this work, they may choose to explore a somewhat new issue with somewhat new methods (Gilbert, 1977). This kind of blend between imitation and differentiation could be well modeled by a reaction-diffusion process that diffuses influential innovations to neighboring regions in science but also reacts against exactly the same innovation being presented more than once.

2.4. *Characteristic outcome*

An important point about individuals' motivations is that they need not predict group-level outcomes. Individuals that are group oriented can create poorly functioning groups, as when two pedestrians end up colliding with each other because they both try to get out of the other person's way (Helbing, Molnar, Farkas, & Bolay, 2001; Moussaid et al., 2009). Conversely, as Adam Smith (1776) observed more than 200 years ago, well-functioning groups can arise from individuals motivated only by their own self-interests.

One of the most exciting prospects of current efforts toward modeling collective behavior is to systematize the development of more effective ways not only to predict but to control collective outcomes. For example, Moussaïd et al. (2009) describe an “Active Walker” model of pedestrian movements that posits that walkers’ movements are a compromise between going to their destinations and going where the travel is easiest. The success of the Active Walker model suggests a novel method of crowd control. The most common method of crowd control is through direct orders or laws. If we wish to direct pedestrian traffic, for example, we may institute rules or physical barriers that prohibit certain movements. The cost of such prohibitions is decreased pedestrian morale and the perception of excluded possibilities (Bonns & Carrus, 2004). An alternative method of crowd control is to change the structure of the environment such that certain navigational behaviors are facilitated while others are hindered. Even without instituting physical or abstract barriers, it may be possible to indirectly control collective behavior with substantial efficacy. Collective behavior is potentially more controllable than isolated individual behavior because of the strong influences among the individuals’ behavior. A small pressure can often be magnified by the positive feedback involved in individuals following other individuals (Dorigo et al., 2000; Salganik & Watts, 2009).

Under this new approach toward fostering effective collective organization, the aim would be to facilitate the development of self-organized patterns rather than dictate high-level structures via top-down control. This conceptualization of design planning as facilitating self-organization rather than dictating final form may have an important moral for social systems in general. Just as a school cafeteria can subtly nudge children toward good diets by putting the healthiest foods at the front of the food line (Thaler & Sunstein, 2008), if good enough models of collective behavior can be constructed, subtle nudges to collective dynamics can have ripples of influence as people are influenced by each other. It is exactly this reasoning that has led owners of some sports teams to hire a few “paid fans” to vigorously show their support for the local team, in the hope that their enthusiasm will prove infectious.

3. Illustrative case studies of collective behavior

With these four attributes in mind, we will now consider some case studies of collective behavior. The purpose is to demonstrate some of the universal collective dynamics that commonly recur when agents get together.

3.1. Wise crowds

James Surowiecki (2004) has made a spirited case for the ability of a group of people to perform better than the average, and in some cases the best, individual performer from the group. The canonical case for his argument is Galton’s (1907) report of fair-goers’ judgments of the weight of an ox. The average² of the 787 guesses was within 1% of the correct weight of the ox. According to Surowiecki, the conditions that allowed this excellent

group performance were the following: a diversity of opinions, independence of members' judgments, decentralization, and a good method for aggregating opinions.

From the current perspective, what is most distinctive about Surowiecki's examples is that they do not feature *any* interaction among the individuals within a group. Interactions among people are viewed as dangerous because they can amplify incorrect information and create speculative bubbles of people predicting other people's reactions rather than "true value." Early decision makers can have an undue influence on the group's behavior when subsequent decision makers are influenced by their own judgments as well as their predecessors' judgments (Bikhchandani, Hirshleifer, & Welch, 1992). Bettencourt (2009) formally models the importance of having sufficient independence among judges if the benefits of synergistic aggregation are to be achieved.

However, preventing all interactions among people may throw out the beneficial consensus-building baby with the speculative bubble bathwater (List, Elsholtz, & Seeley, 2008). The Delphi technique for group decision making attempts to avoid undue early entrant influences but still provide a mechanism for people to learn from each other's valid information. In this method, experts provide initially independent predictions with supporting rationales but afterwards are provided with an anonymous summary of these forecasts. These forecasts may then be used to improve individuals' predictions on subsequent rounds. In general, judges' opinions tend to converge over rounds of opinion exchange, and they tend to converge on predictions that are more accurate than their original predictions (Linstone & Turoff, 1975; Yaniv & Milyavsky, 2007). The empirical successes of information exchange provide justification for considering the benefits of group structures that relax the requirement of individual independence.

3.2. Bee swarms to find nests

Bee hives frequently need to relocate themselves to accommodate population growth within the hive or changes in resource distributions. In *Apis Mellifera*, this is achieved by individual bees communicating only locally with one another (Seeley, 2003). Scout bees explore their area, and upon returning to their hive express the quality of a discovered nest location through a "dance" that signals the distance and direction of the candidate site. The length and intensity of a dance tends to be proportional to the quality of the site and influences the number of other scouts that will investigate the site. This cooperative recruitment leads to a positive feedback that eventually leads to a consensus decision to move the swarm to a new site (Conradt & Roper, 2005; List et al., 2008). It is important that the decision to move nests be consensual because a split decision could weaken the two smaller groups' survival ability.

This underlying consensus-building dynamic has important applications that extend far beyond the bee hive (Couzin, Krause, Franks, & Levin, 2007). It features several attributes that frequently occur in human collectives. First, individual group members only communicate locally. Second, no individual is required to directly compare the relative quality of different options. Third, no individual needs to have an overview of the different options, but all of the information held by individuals contributes to the eventual decision. Even though

the decision is self-organized and determined in a decentralized fashion, it is nonetheless consensual and efficient. Informative signals are amplified and noisy responses tend to be deemphasized. Bettencourt's (2009) information theoretic approach provides a basis for predicting when information is beneficially amplified by group interactions versus lost because of extreme redundancies in possessed information. It is tempting to extend this logic to explain the historical advantages of democracies over dictatorships (no matter how benevolent the latter may be). Democracies allow a marketplace of opinions to determine the best ideas without requiring the bottleneck of a single decision maker who will compare the many attributes of multifarious options.

3.3. Path formation

The poet Antonio Machado reminds us that "Traveler, there is no path. Paths are made by walking." In collective systems ranging from ants to people, this is literally true. Path systems spontaneously form when agents are motivated to take advantage of the trails left by their predecessors. In the process of exploiting previously left trails, agents further reinforce these trails, potentially leading to a lock-in of originally tentative and faint paths. For example, early trail blazers through a jungle use machetes to make slow progress in building paths—progress that is capitalized on and extended by later trekkers, who may then widen the trail, then later put stones down, then gravel, and then asphalt.

The "Active Walker" computational model has done a good job of describing the paths that ants and people form (Helbing, Keltsch, & Molnár, 1997a; Helbing, Schweitzer, Keltsch, & Molnár, 1997b; Moussaid et al., 2009). Predictions of the models have also been confirmed by laboratory experiments with humans (Goldstone & Roberts, 2006). Furthermore, parametric variation within the Active Walker model can potentially be used to guide policy decisions. For example, two of the most important parameters of the model are the visibility of paths (the extent to which an agent is influenced by distant patches' travel ease) and a path's decay rate (how quickly the influence of a step on a patch's ease of travel dissipates). If path decay is set to a high level, then the paths that agents make disappear relatively rapidly. The Active Walker model suggests specific interventions depending upon a community's goals. For situations where conserving the total amount of pathway is desirable (e.g., when valuable vegetation must be cut down to create the paths), the model's advice is that planners should explore ways of increasing either path visibility or path decay (Goldstone, Jones, & Roberts, 2006). While increasing visibility intuitively makes a group better coordinate itself to create and exploit good paths, it may be less intuitive that making paths decay relatively quickly promotes better path systems. However, a moment's reflection reveals that without decay, a group's path system will become stuck on the early paths created by walkers, and these paths will be bee-line paths connecting destinations rather than path networks that effectively combine trails that are close. As shown in Fig. 1, the most efficient path network, shown in A and approximating a Minimal Steiner Tree,³ is achieved by a self-organized collective with highly visible, but short-lasting paths. The benefits of this organization has implications for building social structures with flexible path systems that reinforce, extend, and redirect preceding paths, be they physical or abstract




Fig. 1. The influence of parameters on path formation within the Active Walker model (Helbing et al., 1997a), as reported by Goldstone et al. (2006). The destinations are shown by colored circles. The darkness of a patch is positively related to its level of comfort for walking and indicates the eventual paths after 1,000 iterations. (A) Decay rate of a path = 0.1, impact of each step on comfort level = 1,000, path visibility = 100; (B) decay rate = 0.1, impact = 1,000, visibility = 1; (C) decay rate = 0.001, impact = 10, visibility = 10; (D) decay rate = 0.1, impact = 1,000, visibility = 10. The most efficient path system is found in A, in which walkers are strongly influenced by the comfort level of the patches (visibility of paths is high), but the influence of steps on travel comfort quickly dissipates (paths decay quickly).

paths. One moral for abstract paths (i.e., innovations) is that intentionally building in obsolescence into computer architectures, patents, and music may promote social structures that build upon earlier innovations but are not unduly locked into their specifics. More broadly, this model of group path formation provides a good example of how theory can provide advice on how to nudge social systems to achieve desirable ends.

3.4. Basketball teams

The purpose of considering a sports team example is to make explicit several important attributes of real teams. In contrast to group studies in the laboratory, naturally formed teams often have a long-standing purpose and membership, specific tasks to perform that are of vital importance to its members, a hierarchical structure that is nonetheless labile, and support structures that the group is responsible for creating and maintaining. In the case of a

basketball team, the messages passed between members have a far larger “bandwidth” than those typically encapsulated in formal models. Players respond almost instantaneously to each other and their opponents, taking into account not only their teammates’ positions but also their projected intentions, abilities, and roles. Coaches do provide strategic direction, but much more important are the in-the-moment negotiations and communications between players (Van Wormer, Besthorn, & Keefe, 2007).

Ed Hutchins (1995a, 1995b) and Hutchins and Johnson (2009) have methodically studied other teams that must show flexibility in dealing with environmental challenges, including cockpits and navigation teams on a naval vessel. Many of these teams use representations that are distributed across individuals, or involve a three-way interaction of individuals, the environment, and the tools they have built. As formal models of group behavior are further developed, it is worth bearing in mind the massively interactive nature of cockpit crews, jazz bands, and sport teams. As Hutchins and Johnson (2009) observe, many teams interact with a rich environmental context and in a highly multimodal fashion in which gestures, tones, and rhythms are as important as symbolic content. Given the recent successes in designing teams of robots to play soccer (Lakemeyer, Skylar, Sorrenti, & Takahashi, 2007), we are optimistic that principles for even highly interactive and effective team behavior can be formalized.

3.5. *Common pool resources*

One commonly recurring social pattern is for a collection of people to be sharing a resource in which it is costly or difficult to exclude potential beneficiaries from using the resource (Ostrom, Gardner, & Walker, 1994). In economics, these are known as common pool resource (CPR) situations, and they feature *subtractable* resources in which one person’s use of a resource detracts from the total amount of resource available to others. They are important in society because they occur in water management, pollution (clean air is a resource that is depleted by carbon emissions), pastures, forests, and fishing. The canonical example of a CPR is cows grazing in a grassy open commons. Individual farmers are motivated to increase the size of their herds and the caloric intake of each of their cows for economic gain, but there is a risk of overharvesting the commons to destruction if all farmers exploit the resource without restraint. An aquatic analog to this scenario is the lobster harvesters of Maine, who collectively risk annihilating their lobster “crop” if they overharvest (Acheson, 2003). As with many CPRs, the lobster harvesters have responded to resource conflicts by taking the law into their own hands, for example, by cutting the rope lines to traps of interlopers on their own traditionally held territories.

This example reveals an important and neglected group dynamic. Often times a contrast is drawn between the emergent patterns of self-organized groups and groups that are driven top-down by a leader, rule system, or hierarchical structure (Resnick, 1994). What this rhetorical antithesis misses is that some of the things that self-organized groups do are elect leaders, form rule systems, and institute hierarchies. Most groups that follow rules are typically self-organized, and the rule systems themselves are self-organized. The rules are the tangible products of courts, parliaments, congresses, and governments at city, county, state,

country, and world levels. In the absence of an existing governmental structure that effectively regulated lobster harvesting, the harvesters themselves created this structure. Rules and their less explicit cousin, norms, are complex systems in their own right, no less so than bee hives or traffic jams. They do not exist on their own, but rather depend upon supporting structures for their continuation. They require legal and governmental systems to be created, changed, and eliminated (Ostrom, Dietz, & Stern, 2003). They require monitor systems (e.g., police) to assure that they are being followed. They require sanctioning systems (e.g., jails) to assure that discovered rule violations are punished. Originally unorganized groups will propose, vote upon, and live under rule, monitoring, and sanction systems that they construct themselves (Janssen, Goldstone, Menczer, & Ostrom, 2008; Samuelson & Messick, 1995). In this manner, groups that face scarce resources are often importantly not *simple* decentralized systems, but rather decentralized systems that spontaneously create rule systems that are themselves decentralized.

3.6. Foraging

A problem faced by all mobile organisms is how to search their environment for resources. Animals forage their environment for food, Web-users surf the Internet for desired data (such as music files—Salganik & Watts, 2009), and businesses mine the land for valuable minerals. When an organism forages in an environment that consists, in part, of other organisms that are also foraging, unique complexities arise. The resources available to each individual are affected not just by their own behavior but also by the simultaneous actions of others.

Groups of animals often distribute themselves in a nearly optimal manner, with their distribution matching the distribution of resources. For example, Godin and Keenleyside (1984) distributed edible larvae to two ends of a tank filled with cichlid fish. The food was distributed in ratios of 1:1, 2:1, or 5:1. The cichlids quickly distributed themselves in rough accord with the relative rates of the food distribution before many of the fish had even acquired a single larva and before most fish had acquired larvae from both ends. Similarly, Harper (1982) observed that mallard ducks distributed themselves in accord with the rate or amount of food thrown at two pond locations. Similarly, humans distribute themselves appropriately (Goldstone & Ashpole, 2004), although all three species tend to exhibit *under-matching* such that the distribution of foragers is not as extreme as the distribution of resources.

Group foraging is a good example of a situation where group level properties emerge. Whether a group matches a resource distribution, how quickly the group achieves an efficient solution, and whether the group shows periodic waves of migration into and out of pools (Goldstone, Ashpole, & Roberts, 2005) are all properties of the group as a whole. It is no metaphor to talk of the group's problem-solving ability as a whole. The group's ability to adapt the distribution of its members to the distribution of resources is not simply reducible to its members' problem-solving abilities (Theiner, 2008). In fact, it makes no sense to talk of a single individual matching a resource distribution because it can only be in one place at a time. Matching is only a property of the group. In this case, it truly is that the group has a

mind of its own, or at least demonstrates simple problem-solving capacities. It may be essentially unknowable by us whether the groups that we take part in are conscious or not, just as the individual bee cannot fathom the decisions of the hive. However, if we define cognition as the property which allows systems to produce flexible and adaptive problem-solving behavior that is most felicitously interpreted as involving information processing, then it is not unrealistic to view this type of adaptive behavior as a kind of “group cognition” that can be evaluated distinctly from individual cognition.

3.7. Open source software

Although software is often expensive and strongly protected by copyrights and “digital rights management” systems that prevent its copying, there has also been a strong and growing movement to make software products, including the source code for the software, available to any interested party without restrictions. Many open source software projects have had more than 200 programmers contribute to them and are the product of over 50,000 collective programming hours. Why would a person volunteer her time to such a project? A first reason often cited by programmers is to contribute to the open source community because they believe in the collective value gained by making software freely accessible (Lerner, Tirole, & Pathak, 2006). Second and relatedly, programmers like to make programs that they develop for their own purposes available for others who have similar needs. By making a project open source, a programmer can benefit because other people extend the features of a project that they start. The operating system Unix is a striking example of this; the robustness and functionality of Unix has extended far beyond its original creator’s Linus Torvalds’ programming talents or time. Third, programmers may benefit personally by having their name associated with a prominent open source project.

Many of the dynamics in the open source community match those previously described. Like bee swarms establishing new nests, there is value to solidarity and consensus building (Conradt & List, 2008). Although a single programming project may split into separately developing projects any number of times, it is noteworthy how rarely this happens. Users benefit from a single robust software package rather than a Balkanized set of related packages, and so do programmers if one of their goals is to have their work used by as many people as possible.

Up until now, we have focused on collective behavior that emerges when individuals are interacting locally with one another, with no appreciation for the higher-order patterns that they are creating. The dynamics of the open source software movement do not fully fit this scenario, given that one self-professed motivation for individuals’ behavior is to promote the movement’s welfare. In this way, the human collective’s dynamic differs from those observed with ants and bees. When the individuals that comprise a collective are capable of developing concepts of the collective, then the collective’s identity and goal directedness are intensified. When individuals can entertain thoughts like “My efforts may make other people also volunteer their time to ‘The Cause’ too” and “We should develop technology that allows people to see, use, and extend what other people have contributed” then the groups formed by these individuals look increasingly like self-steering systems. Thus, there is not always a zero-sum

competition between levels of organization, such that the more “unit-like” one level is, the less unit-like higher and lower levels are. In part because of permeability across levels of organization, intelligent wholes are often associated with intelligent parts.

4. Issues concerning collective behavior

The above case studies are helpful in identifying a number of themes and controversies surrounding collective behavior. These issues are also highlighted by the current contributions to this issue of *Topics in Cognitive Science*.

4.1. Does group behavior reduce to individual behavior?

There are several senses of reduction that could apply to collective behavior. One sense is that “the behavior of the collective can be understood in terms of the behaviors of the individuals considered separately” can be easily dismissed. Understanding the interactions among individuals is critically important for understanding all of the examples of collective behavior in Table 1 except for the first. A second sense, that “the collective has no properties that are not also properties of individuals” can similarly be dismissed. Groups often have properties, like under- or overmatching resource distributions when foraging, that are not attributable to individuals.

A third sense of reduction is “Collective behavior does not require theoretical constructs above the level of the individual for its explanation.” By this account, any high-level descriptions of a collective’s behavior are unnecessary and could have been just as well described by referring to the properties of individuals instead. Contrary to this reductive claim, we view group-level constructs to be theoretically indispensable. In fact, one of the primary motivations for many agent-based models is to provide a theoretical bridge across different levels of description. Consider Schelling’s (1971) classic “simulation studies” of segregation. Schelling created agents belonging to two classes (represented by dimes and pennies) that are reasonably tolerant of diversity and only move when they find themselves in a clear minority within their neighborhood, following a rule like “If fewer than 30% of my neighbors belong to my class, then I will move.” Despite this overall tolerance, the agents still divide themselves into sharply segregated groups after a short time. What is surprising is that this occurs even though no individual in the system is motivated to live in such a highly segregated world. Although hardly a realistic model of migration, the model was influential in contrasting group-level results (i.e., widespread segregation) and individual goals. If group-level constructs like segregation, wealth disparity, monetary flow, social network topology (Kennedy, 2009), and intellectual climate are eliminated, then many of the most surprising and useful theoretical claims for how individual-level incentives affect these constructs would no longer be possible. Not only would we miss out on truly bridging theories that show how one kind of behavior creates behaviors at a completely different level, but we would also lose much of our ability to predict and control social structures at scales that are meaningful for society.

A fourth sense of reduction, that of “The patterns of collective behavior can be explained by only referring to individuals and their interactions” is more viable, and it is an assumption underlying many bottom-up agent-based models. Many agent-based models are framed only in terms of local interactions among agents and their environment. A possible violation of this kind of reduction, alluded to in the “open source software” section above, is that sometimes cognitively sophisticated agents develop an awareness of collective organizations and patterns to which they belong. In these cases, the collective’s behavior can be shaped by this awareness and can causally affect individuals’ behavior. However, an alternative interpretation of these scenarios is that individuals’ behavior is always influenced locally by information immediately available, even if these pieces of information include text that refers explicitly to groups, organizational hierarchies, or government regulations. Our impression is that this interpretation is not intellectually productive for two reasons: because it implies a hyperreductionist approach to these questions, and because it ignores that individuals are not inherently nondecomposable units either. Cognitive scientists have long wrestled with ideas concerning the appropriate levels of analysis (Marr, 1982). The general consensus of the field is that all behavior need not be explained in terms of the activity of individual neurons even while recognizing that such neurons ultimately do give rise to behavior. Likewise, a volvox can be optionally viewed as a single organism or a collection of single-celled organisms. A human body is composed of 10 times more cells that do not contain human DNA than cells that do, and many of these former cells are indispensable for human digestion and waste regulation (Frank et al., 2007). The focus on only the lowest, elemental level of analysis curtails our ability to uncover laws of organization that span levels from individuals to groups. As an example, one such principal may be that an organizational unit, once constructed, seeks to preserve its own existence. This applies no less to Microsoft, Israel, and the National Academy of Science than it does to individuals. Indeed, a large part of the purpose of many organizations is to perpetuate themselves, as indicated by an inspection of their mission statements, and their bylaws that define how the membership replaces itself and how the bylaws are permitted to change. By taking organizational constructs seriously, we open ourselves to the exciting possibility of a general science of unit construction that spans all the way from biological cells through individual people to collective organizations (see also Bettencourt, 2009).

There is a parallel between attitudes toward this fourth kind of reduction and programming frameworks for agent-based models. In some frameworks, such as Netlogo and Starlogo, the fundamental ontology consists of agents (called “turtles”) and environmental locations (“patches”). Models from chemistry, physics, biology, and the social sciences can all be implemented by specifying agent-to-agent, patch-to-patch, and agent-to-patch interactions. The resulting models certainly yield interesting global patterns, but there are no programming constructs that explicitly control and represent this global structure. The global-level patterns are supposed to emerge bottom-up from the lower-level interactions (Epstein, 2007; Resnick, 1994). By contrast, in frameworks such as Repast and Swarm, there are programming constructs that allow the user to explicitly refer to collectives as a group and to control the group’s behavior directly. The same impulses that drove Repast’s designers to allow group-level structures and hierarchical control underlie many

researchers' decisions to reject even this last form of reductionism. Carley et al. (2009) provide an example of this approach, in which media sources such as radio and direct mail are represented as agents even though they could also be subdivided into individual people.

4.2. Is "group cognition" possible, or a level confusion error?

As alluded to earlier, there are some researchers who argue that talk of "group cognition" is incorrect or cannot be taken literally. By this argument, the members of a group may be cognitive agents on their own, but it is a confusion to think of the group as a whole as a unified cognitive agent. One version of this argument is that groups do not, as far as we can tell, have mental states in the sense of consciously introspectable experiences (Harnad & Dror, 2006). A more general version of this argument is skeptical of any ascription of cognition to systems that include people as only one element (Adams & Aizawa, 2001; Rupert, 2004). This generalized argument was originally aimed at claims that minds can be distributed across people and the tools they use (Clark & Chalmers, 1998)—with other people's minds being just one example of a useful tool. If minds never extend outside of individual people's skulls, then a fortiori they do not extend to include multiple people.

On the other side of the controversy are researchers who argue that people often work together in such an integrated, interactive manner, that it is appropriate and useful to consider the whole group as an information processing system (Hutchins, 1995a,b; Theiner, 2008). One of the considerations in favor of this argument is that the group engages in representation building enterprises in which no individual has access to the complete representation. The group as a whole is needed to explain how the representations, often involving physical devices, are processed. Theiner (2008) also argues that Clark and Chalmers' parity arguments for distributed cognition apply to group minds: "If, as we confront some task, a group collectively functions as a process which, *were it done in the head*, we would have no hesitation in accepting as a cognitive process, then that group *is* (for that time) performing the cognitive process" (p. 313).

Consistent with Theiner's general perspective, we recommend identifying possible cases of collective cognition on the basis of information processing, rather than on the basis of whether the collective has conscious mental states or not (Gureckis & Goldstone, 2006). This is based simply on the pragmatic consideration that determining individual consciousness, let alone group consciousness, is a murky and presumptive enterprise at best. From an informational perspective, describing a group of people as a single functional unit is justified to the extent that the elements within the group (i.e., individual people) are highly connected with each other, and if there are relatively lower levels of connectivity between elements in the group and elements outside of the group. This is a general criterion for considering any set of elements to be part of a single unit. For example, a leading theory for the evolutionary origin of mitochondria and chloroplasts is that they were originally independent bacteria that became incorporated into the cytoplasm of cells, and once incorporated, conferred advantages for the cell because they allowed cellular respiration (mitochondria) and photosynthesis (chloroplasts) for energy production (Margulis, 1970). We are less likely to view mitochondria as the individual units they once were because of their strong dependencies

with other internal cell elements. We view it as an attractive feature of this characterization of “unithood” that it works at multiples levels, because we see nothing inherently unique about individual people as units. Unithood is graded, and legitimizing one level of unithood does not repudiate the legitimacy of other levels. Cells, individuals, and companies can all be real(ly useful) descriptions.

We believe that groups of people are often times cognitively interesting systems because they exist at the *cusp* of unithood. Before the bacteria has been incorporated into the cell at all, it is simply an independent environmental influence on the cell. Once the mitochondria loses its ability to make its own living in the world, it is no longer a unit by itself, but rather part of the eukaryotic cell unit. In between being a free-agent bacteria and a mitochondrial cog in the cellular wheel, the “bactondria” is both independent and dependent on the cell. This status, we argue, is particularly important when it comes to cognitive systems. Computational complexity, in terms of being able to transmit information, is at its greatest for systems made of partially dependent elements. Sporns, Chialvo, Kaiser, and Hilgetag (2004) have quantified the “information integration” of a system in terms of its total amount of mutual information. On the one hand, if a system’s elements are completely independent, then information cannot be transmitted from one part of the system to another. On the other hand, if a system’s elements are too tightly connected, then they all end up possessing the same information and communication is pointless. Human nervous systems have apparently evolved so as to maximize the usefulness of neural communication (Sporns, 2002). Similarly, we would argue that groups of people also adapt so as to create information-amplifying systems. Useful human collectives are those that promote robust information transmission across people yet avoid having everybody know the same things (see Mason et al., 2008 for empirical evidence, and Kennedy, 2009 for relevant modeling). Collectives that do this will maximize their computational capability.

4.3. *What is the value of formal models for understanding collective behavior?*

Many of this issue’s authors engage in formal computational and mathematical modeling to understand collective behavior. Models of the kind employed by Moussaid et al. (2009), Kennedy (2009), Gureckis and Goldstone (in press), Bettencourt (2009), and Carley et al. (2009) have a number of attractive features that supplement traditional methods for exploring group behavior. First, they are expressed with unambiguous mathematical and computational formalisms so that once they have been fully described, their predictions are clear, quantitative, and objective. Second, they provide true bridging explanations that link two distinct levels of analysis: the properties of individual agents (e.g., their attributes and interactions), and the emergent group-level behavior. When successful, agent-based models are particularly satisfying models because they show how coherent, group-level structures can spontaneously emerge without leaders ordering the organization, and sometimes despite leaders’ efforts. Third, because the models are typically either simple or informed by real-world data, they are appropriately constrained and cannot fit any conceivable pattern of data.

Many models of group behavior are conspicuously idealized and simplified, more so than models of individual cognition. For example, agents are often represented by a single value

or vector, the world is a two-dimensional grid, and interactions between agents simply involve exchanging these values. Many of these simplifications are due to practical considerations. If every agent in a model is as complicated as our state-of-the-science cognitive models of memory, attention, learning, and problem solving, then the collective that involves hundreds of these agents may well be extremely complicated and hard to understand. It will have too many degrees of freedom and could easily end up being insufficiently constrained.

Beyond this practical consideration, there are both costs and benefits of idealized models of collective behavior. Many researchers purposefully choose to create highly idealized models that boil down a collective phenomenon to its functional essence. Researchers pursuing idealized models are typically motivated to describe domain-general mechanisms with a wide sphere of application. Physicists have recently entered the arena of modeling social systems, and one of their attractions to the field is being able to apply the same kinds of models that have been successfully applied to the Brownian motion of particles, gasses under pressure, and interacting magnetic elements (Ball, 2004; Bettencourt, 2009; Helbing et al., 1997a, 2001). A good example of the effectiveness of these idealizations is the fertility of networks analyses (Kennedy, 2009; Salganik & Watts, 2009; Watts & Strogatz, 1998). Power-laws have been implicated in the distribution of connections within actor, neural, power grid, and telephone networks (Barabási & Albert, 1999). Preferential attachment has been posited to explain all of these networks, according to which the likelihood of a node in a network attracting still further connections is proportional to its current degree of connectivity. Certainly this mechanism is something of a caricature. Its simple mathematical formulation fits none of the real networks exactly. However, a compelling argument can be made that it captures a critical and powerful dynamic in each of them.

Other researchers have argued that most models of collective behavior are too simplified. Some have chosen to develop much more complex models that incorporate highly detailed, situation-specific parameter values. One example of this approach is Carley et al.'s (2009) model of the way in which public opinion shifts in the face of different kinds of media. They consider relatively rich interactions among agents who decide who they will communicate with, what they will communicate, how much they will communicate, and whether they are influenced by others' communications. The researchers incorporate real-world demographic information regarding race, income, and education, and the coverage zones of media such as advertisements, the Web, telephone calls, and radio in order to constrain these more complex models.

Choosing a similarly detailed approach to address the question, "Why did the Anasazi people of southwestern United States abandon their homeland around 1350 AD?" research teams have developed simulations that incorporate features grounded in historical records: maize production levels, ground water reserves, the 3-D geography of the Anasazi's Long House Valley homeland, populations established from archeological digs, and social trends regarding childbirth age, the average age of children leaving home, and food consumption needs, all based upon recent maize-growing societies of Pueblo Indians descended from the Anasazi (Axtell et al., 2002; Dean et al., 2000). While useful for answering specific historical or sociological questions, the disadvantages of highly specific models such as this are

that critical dynamics and parameters may remain hidden among other less critical model components, and it may be difficult to draw general implications for other future scenarios.

Another objection to the typical simplifications of collective behavior models is that the choices of real-world properties to exclude from models unfortunately have neglected truly essential aspects of human-to-human and human-to-environment interactions. Hutchins and Johnson (2009) make exactly this argument on the basis of highly simplified communication protocols typically used in models (including his own prior work on the self-organized emergence of language norms; Hutchins & Hazlehurst, 2002). When people, or bonobos, interact, they are not simply transmitting numeric values, or even digital symbols. They interact in a rich world through gesture, tone, shared physical representations, and bodily actions. Although these aspects could in principle be included in models of interaction, we concur with Hutchins that the strategy of most current models is to distill worlds to their simplest discrete structures, and interactions down to their simplest message-passing essence.

Our own opinion is that a pluralistic approach toward understanding collective behavior is in order, and that the field is sufficiently open that the relatively idealized modeling approaches of Moussaid, Kennedy, Gureckis, Bettencourt, Carley, and Salganik are as likely to produce fertile results as Hutchins' more richly contextualized anthropological approach. The large payoff that comes from creating general models that can apply across superficially dissimilar scenarios and hence unify them is too valuable to bypass. We are not arguing that an idealizing approach to collective behavior is superior to an in-depth, specialist's focus on a single domain's details. Both approaches are necessary for a complete science, and in fact it is only by understanding a system's details that we can determine the general principles by which it is governed. However, given the climate of progressive specialization in contemporary science, it is important to remember that many of the most noteworthy advances of science, from Einstein's unification of gravitational and electromagnetic acceleration to Darwin's unification of the principles by which snails and humans evolve, have involved finding deep principles shared by seemingly dissimilar phenomena.

4.4. What does cognitive science have to do with it?

A final question is more specific to the audience of this issue. As cognitive scientists, why should we care about collective behavior? Aren't issues of group behavior better addressed by economics, social psychology, sociology, and political science? The answer to this question is twofold. First, at its core, cognitive science has always been an interdisciplinary approach to complex, adaptive, intelligent systems. In the preceding sections, we have argued that collected units (of people, animals, and cells) also exhibit adaptive information processing. Thus, our belief is that studying such systems is a natural extension of the traditionally articulated goals of the field. Indeed, there is much that cognitive science can contribute and learn from studying such systems. Second, cognitive science is in a unique position to leverage theoretical tools for understanding individual behavior in order to understand collective outcomes. For example, traditional economic approaches to market behavior assume rational, utility-driven agents. In contrast, cognitive scientists can leverage

our understanding of the limited learning, memory, and decision-making capacities of individuals in order to understand aggregate outcomes. One example of this is presented in Gureckis and Goldstone (in press) where modeling psychological assumptions about novelty preferences and encoding of frequency information from the environment provides a deeper insight into the dynamics of baby names in the United States.

5. The future of collective behavior

There are several topics that will likely be major areas of development in collective behavior research. Some of our idiosyncratic suggestions for growth zones include the evolution of social networks, the evolution of human language, identifying factors that affect cooperation in public good and common pool resource problems, the dissemination of innovations in communities, consensus decision making, experiments and models of group selection, the spontaneous emergence of norms, traffic pattern analysis, classroom dynamics, organizing work teams without management, coalition formation, and long-term multiparticipant collaborations. Beyond these specific topics, we would also like to point to general directions for future research.

5.1. Methods

Much of the empirical and theoretical work to be done will be in bridging individual-level and group-level accounts of behavior. This will need to proceed by integrating individual experiments, group experiments, real-world social interactions, historical records, and model building. As suggested in the introductory justification for the timeliness of collective behavior, there are exciting developments along all of these fronts.

Experimentally, there are now tools that allow social scientists with little or no programming experience to connect groups of people and systematically record their interactions. The use of cell phones, connected calculators, computers, virtual worlds, and RF tagging will allow two traditions of psychological experimentation to be united. In one tradition, psychologists interested in social behavior have attempted to *control* the group and study individual decision-making characteristics while manipulating the group's apparent behavior. The most prominent example of this strategy is Ash's (1956) classic experiments on conformity. Participants judged unambiguous stimuli after hearing other opinions offering incorrect estimates. Sixty-nine percent of his participants conformed to the bogus majority. For our purposes, what is important about this method, and the multitude of experiments inspired by it, is that there is only one actual experimental subject per session. The other participants are accomplices of the experimenter, giving responses scripted ahead of time. The benefit of this approach is that it allows powerful experimental designs for unambiguously determining the influence of the group on the individual's behavior. The cost of this approach is that it eliminates the possibility of finding emergent group-level patterns, because all but one member of the group has a fixed behavior script. For this reason, we have emphasized experimental designs in which all members of a group are free to choose

their own behavior (see, e.g., the articles by Kennedy, 2009 and Salganik & Watts, 2009). In fact, both kinds of designs are needed. In the future, we expect much more interplay between group experiments that show important emergent patterns of behavior for an entire group, and individual experiments with programmed peers that pinpoint the individual decision strategies that produce these global consequences.

In terms of modeling methods, we expect to see continued progress in the development of both idealized and richly detailed simulations of social systems and experimental results. In general, there is likely to be a trend toward increasingly detailed models. However, it would be a mistake to jump directly to highly detailed models until many of the foundational collective patterns are better understood. In our opinion, the most obvious future direction for formal models is for a greater emphasis on validation. It is all too frequent to see agent-based models being proposed in the literature with little effort made toward showing the predictiveness of these models for actual data on social patterns. The articles by Carley, Gureckis, Moussaid, and Salganik all take validation very seriously, but in general, efforts to organize and predict large data sets through models is still in its infancy. One exciting method introduced by Salganik and Watts (2009) is to incorporate replicability into a naturally occurring scenario with strong group influences—downloading music files from the Web. By partitioning participants into independent groups, they were able to measure whether separate “re-runnings of history” would have produced the same most popular songs, or whether different songs would arise as most popular because of rich-get-richer dynamics operating on initially haphazard choices. Would The Beatles always have achieved rampant popularity because of the sheer quality of their music, or were they just lucky beneficiaries of cumulative advantage? Salganik and Watts’ surprising results, suggesting that The Beatles’ tremendous success was not predestined and that a significant helping of luck was involved, point to the power of combining laboratory-inspired replicability with observations of naturally behaving collectives.

One particularly ripe arena for future modeling will be the application of models to detailed data sets obtained from controlled laboratory settings. Often times, there is a disconcerting mismatch between the simplicity of formal models and the complexities of the real-world situations. One strategy for bridging the gap between computational models and group behavior phenomena is to create relatively simple laboratory situations involving groups of people interacting in idealized environments according to easily stated “game rules.” Some external validity is admittedly sacrificed in creating idealized experimental scenarios, but this loss is offset by the advantage of having the assumptions underlying the psychological experiments correspond almost exactly to the assumptions of the computational models, allowing the models to be aptly applied without sacrificing their concise explanatory value and genuine predictiveness.

Although we (somewhat predictably, as psychologists) are advocating greater appreciation and use of the tools of experimental psychology, we also view the analysis of real-world dynamics as a major research opportunity. Excellent data sets can either be developed or easily accessed that describe collective behavior patterns involving decisions that really matter to people. Examples from the present issue include music downloading (Salganik & Watts, 2009), decisions on what to name one’s baby (Gureckis & Goldstone, 2009), and

traffic patterns (Moussaïd et al., 2009). Other prominent data sets include scholarly citations, telephone calls, the movement of currency, disease spread, gossip spread, patterns of collaboration, patent uses and dependencies, jury decisions, and business transactions. These data sets are not typically as clean as those collected from experiments, involve nuisance factors and artifacts, and causality must be inferred from patterns of correlations rather than through more statistically powerful interventions. However, the sheer volume of data in many of these cases can compensate for a lack of experimental control, and for this reason they are likely to be an important tool for not only the sociologist and economist, but for the psychologist and computational modeler as well.

5.2. The collective cognition of collective cognition

A good case can be made for applying the science of collective behavior to the emergent science of collective behavior itself. Not unlike an ant colony in which an ant's role in the nest depends upon the roles assumed by others, a vigorous science of collective behavior depends upon differentiated roles for empiricists, statisticians, theorists, and modelers who interact to feed into and off of one another. Our understanding of collective behavior is expected to grow fastest when different researchers each have their expertise, but also know enough about each others' fields that they can hold a sophisticated discourse (Cowan & Jonard, 2001). Recent analyses of scientific collaborations, as revealed by scholarly citation networks and journal databases, indicate three important results. First, team sizes in science are increasing rapidly in terms of size and diversity (Börner, Maru, & Goldstone, 2004). Second, the fastest rate of increase is found for high-impact journals (Guimerà, Uzzi, Spiro, & Amaral, 2005). Third, the cliquishness of a team can be measured by seeing how dense the collaborative links are within the team relative to links that connect the team to other teams, and Guimerà et al. (2005) have shown that an intermediate level of cliquishness is ideal—not so cliquish so as to be inbred, but the team members should also not be so promiscuous with their connections so as to lose their ability to deeply communicate and connect with their team. In sum, along with the rest of science, research on collective behavior can benefit from increased communication among its participants. Ideally, a shared perspective will emerge that allows efficient transmission of information, but this perspective should not become so predominant that it stifles diversity. If researchers have exactly the same approach and perspective, then transmission of information is pointless.

There are two symmetric cases of perspective sharing that we feel are particularly valuable for future developments. One activity that will likely lead to useful applications and future progress is to increase the sophistication of biophysics models by incorporating richer psychological and sociological models. In many cases, agent-based models conceptualize agents as single values or scalars, or perhaps if they are more complicated, as vectors of numbers on a variety of attributes. These models often assume that communication between agents simply involves transmitting these numeric values from one agent to another. In fact, we view both of these simplifications as dangerously limiting. At the agent level, knowledge is organized in rich conceptual networks, not scalars. Human groups are networks of people,

each of whom is a network of concepts. Agents vary on important factors that change how they behave and think. At the community level, communication is often complicated. Agents lie, fail to communicate because their conceptual systems are too dissimilar, consume mass media and not just agent-to-agent communications (Carley et al., 2009), and are not perfectly rational. So, physicists, mathematicians, biologists, and computer scientists will need to talk to the social scientists.

Conversely, a second major area of progress will be for psychologists and sociologists to increase the sophistication of their models by borrowing formalisms from bio-physics. There are elegant treatments of diffusion, percolation, local interactions, and network dynamics that could go a long way toward systematizing social science. Psychologists ought to be at least striving for the kind of universal laws that are the physicists' bread and butter. Obviously people behave very differently, but a useful application of Occam's razor is to begin with the perspective that perhaps this diversity arises from a universal process of people adapting themselves to their local contexts. So social scientists will need to talk to the physicists, mathematicians, biologists, and computer scientists.

This call for cross-fertilization of traditional disciplines may be preaching to the cognitive science choir. Cognitive scientists have long appreciated the benefits of approaching minds and intelligent systems from multiple vantage points. Given the multitude of levels and approaches needed to understand collective behavior, it behooves us all to interact with each other to understand how people interact with each other.

Notes

1. In addition to the four articles on collective behavior in this physical issue of *Topics in Cognitive Science*, there will be three additional contributions to the topic in an upcoming issue.
2. The average was conceptualized as the median by Galton (1907) but mean by Surowiecki (2004).
3. A minimal steiner tree (MST) is the set of paths that connects a set of points (e.g., destinations) using the minimal amount of total path length. If we restrict ourselves to only creating direct connections between destination points, then the shortest total path network that connects a set of points is called the minimal spanning tree.

Acknowledgments

The authors wish to express thanks to Katy Börner, Georg Theiner, Peter Todd, and Thomas Wisdom for helpful suggestions on this work. This research was funded by Department of Education, Institute of Education Sciences grant R305H050116, and National Science REESE Foundation grant 0910218. More information about the laboratory and online versions of experiments related to the current work can be found at <http://cognitn.psych.indiana.edu>.

References

- Acherson, J. M. (2003). *Capturing the commons: Devising institutions to manage the Maine lobster industry*. Lebanon: University Press of New England.
- Adams, F., & Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology*, *14*, 43–64.
- Ash, S. E. (1956). Studies of independence and conformity: A minority of one against a unanimous majority. *Psychological Monographs*, *9*, (Whole No. 416).
- Axtell, R. L., Epstein, J. M., Dean, J. S., Gunerman, G. J., Swedlund, A. C., Harburger, J., Chakravarty, S., Hammond, R., Parker, J., and Parker, M. (2002). Population growth and collapse in a multiagent model of the Kayenta Anasazi in Long House Valley. *Proceedings of the National Academy of Sciences*, *99*, 7275–7279.
- Bainbridge, W. S. (2007). The scientific research potential of virtual worlds. *Science*, *317*, 472–476.
- Ball, P. (1999). *The self-made tapestry*. Oxford, England: Oxford University Press.
- Ball, P. (2004). *Critical mass*. New York: Farrar, Straus, and Giroux.
- Bandura, A. (1965). Behavioral modification through modeling procedures. In L. Krasner & L. P. Ulmann (Eds.), *Research in behavior modification: new development and implications* (pp. 310–340). New York: Rinehart and Winston.
- Barabási, A., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *286*, 509–512.
- Berger, J. A., & Heath, C. (2005). Idea habitats: How the prevalence of environmental cues influences the success of ideas. *Cognitive Science*, *29*, 195–222.
- Bettencourt, L. M. A. (in press). The rules of information aggregation and emergence of collective intelligent behavior. *Topics in Cognitive Science*.
- Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, *100*, 992–1026.
- Blakemore, S. (2000). *The meme machine*. Oxford: Oxford University Press.
- Bonns, M., & Carrus, G. (2004). Environmental psychology. In C. Spielberger (Ed.), *Encyclopedia of applied psychology*. (pp. 345–349) Amsterdam: Elsevier.
- Börner, K., Maru, J., & Goldstone, R. L. (2004). The simultaneous evolution of author and paper networks. *Proceedings of the National Academy of Science*, *101*, 5266–5273.
- Brockmann, D., Hufnagel, L., & Geisel, T. (2006). The scaling laws of human travel. *Nature*, *439*, 462–465.
- Carley, K. M., Martin, M. K., & Hirshman, B. R. (2009). The etiology of social change. *Topics in Cognitive Science*, DOI: 10.1111/j.1756-8765.2009.01037.x
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, *58*(1), 7–19.
- Conradt, L., & List, C. (2009). Group decisions in humans and animals: A survey. *Philosophical Transactions of the Royal Society B*, *364*, 719–742.
- Conradt, L., & Roper, T. J. (2005). Consensus decision making in animals. *Trends in Ecology and Evolution*, *20*, 449–456.
- Couzin, I. D., Krause, J., Franks, N. R., & Levin, S. A. (2007). Effective leadership and decision making in animal groups on the move. *Nature*, *433*, 513–516.
- Cowan, R., & Jonard, N. (2001). Knowledge creation, knowledge diffusion and network structure. In J. -B. Zimmermann (Ed.), *Economics with heterogeneous interacting agents* (Vol. Lecture Notes in Economics and Mathematical Systems 503, pp. 327–343). Berlin: Springer.
- Dawkins, R. (1976). *The selfish gene*. Oxford, England: Oxford University Press.
- Dean, J. S. et al. (2000). Understanding Anasazi cultural change through agent-based modeling. In T. Kohler & G. Gumerman (Eds.), *Dynamics in human and primate societies* (pp. 179–205). Oxford, England: Oxford University Press.
- Dorigo, M., Bonabeau, E., & Theraulaz, G. (2000). Ant algorithms and stigmergy. *Future Generation Computer Systems*, *16*, 851–871.
- Epstein, J. M. (2007). *Generative social science: Studies in agent-based computational modeling*. Princeton, NJ: Princeton University Press.

- Epstein, J. M. (2008). Why model? *Journal of Artificial Societies and Social Simulation*, 11, 12.
- Frank, D. N., St Amand, A. L., Feldman, R. A., Boedeker, E. C., Harpaz, N., & Pace, N. R. (2007). Molecular-phylogenetic characterization of microbial community imbalances in human inflammatory bowel diseases. *Proceedings of the National Academy of Science*, 104, 13780–13785.
- Friedman, D., Steed, A., & Slater, M. (2007). Spatial social behavior in second life. *Lecture Notes in Computer Science*, 4722, 252–263.
- Galton, F. (1907). Vox Populi. *Nature*, 75, 450–451.
- Gilbert, G. N. (1977). Competition, differentiation and careers in science. *Social Science Information*, 16, 103–123.
- Godin, M. J., & Keenleyside, M. H. A. (1984). Foraging on patchily distributed prey by a cichlid fish (Teleostii, Chichliidae): A test of the ideal free distribution theory. *Animal Behavior*, 32, 120.
- Goldstone, R. L., & Ashpole, B. C. (2004). Human foraging behavior in a virtual environment. *Psychonomic Bulletin & Review*, 11, 508–514.
- 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., Jones, A., & Roberts, M. E. (2006). Group path formation. *IEEE Transactions on System, Man, and Cybernetics, Part A*, 36, 611–620.
- Goldstone, R. L., & Roberts, M. E. (2006). Self-organized trail systems in groups of humans. *Complexity*, 11, 43–50.
- González, M. C., Hidalgo, C. A., & Barabási, A. (2008). Understanding individual human mobility patterns. *Nature*, 453, 779–782.
- Guimerà, R., Uzzi, B., Spiro, J., & Amaral, L. A. (2005). Team assembly mechanisms determine collaboration network structure and team performance. *Science*, 308, 697–702.
- Gureckis, T. M., & Goldstone, R. L. (2006). Thinking in groups. *Pragmatics and Cognition*, 14(2), 293–311.
- Gureckis, T. M., & Goldstone, R. L. (in press). How you named your child: Understanding the relationship between individual decision-making and collective outcomes. *Topics in Cognitive Science*.
- Haidt, J., Seder, P., & Kesebir, S. (2008). Hive psychology, happiness, and public policy. *Journal of Legal Studies*, 37, S133–S156.
- Harnad, S., & Dror, I. E. (2006). Distributed cognition: Cognizing, autonomy, and the Turing test. *Pragmatics & Cognition*, 14, 209–223.
- Harper, D. G. C. (1982). Competitive foraging in mallards: Ideal free ducks. *Animal Behavior*, 30, 575.
- Helbing, D., Keltsch, J., & Molnár, P. (1997a). Modeling the evolution of human trail systems. *Nature*, 388, 47–50.
- Helbing, D., Molnar, P., Farkas, I., & Bolay, K. (2001). Self-organizing pedestrian movement. *Environment and planning B*, 28, 361–383.
- Helbing, D., Schweitzer, F., Keltsch, J., & Molnár, P. (1997b). Active walker model for the formation of human and animal trail systems. *Physical Review E*, 56, 2527–2539.
- Hutchins, E. (1995a). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Hutchins, E. (1995b). How a cockpit remembers its speed. *Cognitive Science*, 19, 265–288.
- Hutchins, E., & Hazlehurst, B. (2002). Auto-organization and emergence of shared language structure. In A. Cangelosi & D. Parisi (Eds.), *Simulating the evolution of language* (pp. 279–305). London: Springer-Verlag.
- Hutchins, E., & Johnson, C. M. (2009). Modeling the emergence of language as an embodied collective cognitive activity. *Topics in Cognitive Science*, 1(3), 523–546.
- Janssen, M. A., Goldstone, R. L., Menczer, F., & Ostrom, E. (2008). Effect of rule choice in dynamic interactive spatial commons. *International Journal of the Commons*, 2, 288–312.
- Kearns, M., Suri, S., & Montfort, N. (2006). An experimental study of the coloring problem on human subject networks. *Science*, 313, 824–827.
- Kennedy, J. (2009). Social optimization in the presence of cognitive local optima: Effects of social network topology. *Topics in Cognitive Science*, 1(3), 498–522.

- Kenrick, D. T. et al. (2003). Dynamical evolutionary psychology: Individual decision rules and emergent social norms. *Psychological Review*, *110*, 3–28.
- Kirman, A., & Zimmermann, J. (2001). *Economics with heterogeneous interacting agents*. Berlin: Springer-Verlag.
- Kohler, T., & Gumerman, G. (2002). *Dynamics in human and primate societies*. Oxford, England: Oxford University Press.
- Lakemeyer, G., Skylar, E., Sorrenti, D. G., & Takahashi, T. (2007). *RoboCup 2006: Robot Soccer World Cup X*. Berlin: Springer.
- Latane, B., & Bourgeois, M. J. (2000). Dynamic social impact and the consolidation, clustering, correlation, and continuing diversity of culture. In R. S. Tindale & M. Hogg (Eds.), *Handbook of social psychology: Vol. 4. Group processes* (pp. 235–258), New York: Blackwell.
- Lerner, J., Tirole, J., & Pathak, P. (2006). The dynamics of open source contributors. *American Economics Review Papers and Proceedings*, *96*, 114–118.
- Linstone, H., & Turoff, M. (1975). *The Delphi method: Techniques and applications*. Ontario, Canada: Addison-Wesley.
- List, C., Elsholtz, C., & Seeley, T. D. (2008). Independence and interdependence in collective decision making: an agent-based model of nest-site choice by honeybee swarms. *Philosophical Transactions of the Royal Society B*, *364*, 755–762.
- Macy, M. W., & Willer, R. (2002). From factors to actors: Computational sociology and agent-based modeling. *Annual Review of Sociology*, *28*, 143–166.
- Margulis, L. (1970). *Origin of eukaryotic cells*. New Haven, CT: Yale University Press.
- Marr, D. (1982). *Vision*. San Francisco: Freeman.
- Mason, W. A., Jones, A., & Goldstone, R. L. (2008). Propagation of innovations in networked groups. *Journal of Experimental Psychology: General*, *137*, 422–433.
- Moussaid, M., Garnier, S., Theraulaz, G., & Helbing, D. (2009). Collective information processing in swarms, flocks and crowds. *Topics in Cognitive Science*, *1*(3), 469–497.
- Nowak, M. A., & Sigmund, K. (2004). Evolutionary dynamics of biological games. *Science*, *303*, 793–799.
- Ostrom, E., Dietz, T., & Stern, P. (2003). The struggle to govern the common. *Science*, *302*, 1907–1912.
- Ostrom, E., Gardner, R., & Walker, J. (1994). *Rules, games, and common-pool resources*. Ann Arbor: University of Michigan Press.
- Putnam, R. D. (2001). *Bowling alone: The collapse and revival of American community*. New York: Simon & Schuster.
- Rawlins, G. J. E. (in press). *Standing on the shoulders of giants*. Cambridge, MA: MIT Press.
- Resnick, M. (1994). *Turtles, termites, and traffic jams*. Cambridge, MA: MIT Press.
- Rupert, R. (2004). Challenges to the hypothesis of extended cognition. *The Journal of Philosophy*, *101*, 389–428.
- Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, *311*, 854–856.
- Salganik, M. J., & Watts, D. J. (2009). Web-Based Experiments for the Study of Collective Social Dynamics in Cultural Markets. *Topics in Cognitive Science*, *1*(3), 439–468.
- Samuelson, C. D., & Messick, D. M. (1995). When do people want to change the rules for allocating shared resources? In D. A. Schroeder (Ed.), *Social dilemmas: Perspectives on individuals and groups* (pp. 143–162). Westport, CT: Praeger.
- Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology*, *1*, 143–186.
- Seeley, T. D. (2003). Consensus building during nest-site selection in honey bee swarms: The expiration of dissent. *Behavioral Ecology and Sociobiology*, *53*, 417–424.
- Smith, A. (1776 [1977]). *An inquiry into the nature and causes of the wealth of nations*. Chicago: University of Chicago Press.
- Smith, E. R., & Conrey, F. R. (2007). Agent-based modeling: A new approach for theory-building in social psychology. *Personality and Social Psychology Review*, *11*, 87–104.

- Smith, E. R., & Semin, G. R. (2007). Situated social cognition. *Current Directions in Psychological Science*, 16, 132–135.
- Sporns, O. (2002). Network analysis, complexity and brain function. *Complexity*, 8, 56–60.
- Sporns, O., Chialvo, D., Kaiser, M., & Hilgetag, C. C. (2004). Organization, development and function of complex brain networks. *Trends in Cognitive Sciences*, 8, 418–425.
- Surowiecki, J. S. (2004). *The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business, economies, societies and nations*. New York: Doubleday.
- Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. New Haven, NJ: Yale University Press.
- Theiner, G. (2008). *From extended minds to group minds: Rethinking the boundaries of the mental*. Doctoral dissertation, Department of Philosophy, Program in Cognitive Science, Indiana University.
- Theraulaz, G., & Bonabeau, E. (1995). Modeling the collective building of complex architectures in social insects with lattice swarms. *Journal of Theoretical Biology*, 177, 381–387.
- Van Wormer, K., Besthorn, F. H., & Keefe, T. (2007). *Human behavior and the social environment*. New York: Oxford University Press.
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of “small-world” networks. *Nature*, 393, 440–442.
- Wilensky, U., & Stroup, W. (1999). Learning through participatory simulations: Network-based design for systems learning in classrooms. In *Proceedings of the Computer Supported Collaborative Learning Conference*. Article no. 80. Stanford, CA: International Society of the Learning Sciences.
- Wright, R. (2001). *Nonzero*. New York: Vintage.
- Yaniv, I., & Milyavsky, M. (2007). Using advice from multiple sources to revise and improve judgments. *Organizational Behavior and Human Decision Processes*, 103, 104–120.