Review TRENDS in Cognitive Sciences Vol.9 No.9 September 2005

Computational models of collective behavior

Robert L. Goldstone and Marco A. Janssen

Psychology Department, Indiana University, Bloomington, IN 47405, USA

Computational models of human collective behavior offer promise in providing quantitative and empirically verifiable accounts of how individual decisions lead to the emergence of group-level organizations. Agent-based models (ABMs) describe interactions among individual agents and their environment, and provide a process-oriented alternative to descriptive mathematical models. Recent ABMs provide compelling accounts of group pattern formation, contagion and cooperation, and can be used to predict, manipulate and improve upon collective behavior. ABMs overcome an assumption that underlies much of cognitive science – that the individual is the crucial unit of cognition. The alternative advocated here is that individuals participate in collective organizations that they might not understand or even perceive, and that these organizations affect and are affected by individual behavior.

Introduction

Cognitive scientists tend to focus on the behavior of single individuals thinking and perceiving on their own. However, interacting groups of people also create emergent organizations at a higher level than the individual. Interacting ants create colony architectures that no single ant intends. Populations of neurons create structured thought, permanent memories and adaptive responses that no neuron can comprehend by itself. Similarly, people create group-level behaviors that are beyond the ken of any single person. The emergence of higher-level organizations from the interactions of lower-level units is surprising in the case of group behavior because we are the lower-level units, and the higher-level organizations typically emerge spontaneously, without our knowledge. Social phenomena such as rumors, the emergence of a standard currency, transport systems, the World Wide Web, resource harvesting, crowds, and scientific establishments arise because of individuals' beliefs and goals, but the eventual form that these phenomena take is rarely dictated by any individual.

There is a growing realization across the social sciences that one of the best ways to build useful theories of group phenomena is to create working computational models of social units (e.g. individuals, households, firms or nations) and their interactions, and to observe the global structures that these interactions produce. In the past few years, the use of computational models of collective behavior has grown tremendously in sociology [1], economics [2], psychology [3,4], and anthropology [5]. This approach is relevant to cognitive science because it integrates computational modeling and understanding of human behavior. This relevance is timely because these models provide balance to cognitive science's bias to view cognition as a property of an individual mind rather than as resulting from interactions among people and their environments [6].

We will focus on computational models called Agent-Based Models (ABMs), which build social structures from the 'bottom-up', by simulating individuals by virtual agents, and creating emergent organizations out of the operation of rules that govern interactions among agents [7,8]. ABMs have several 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' effort. 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. The self-organization process itself exerts strong constraints on the kinds of patterns likely to be observed [9]. In this review of ABMs, we will characterize the approach; describe crucial decisions that a modeler must make; present case studies of ABMs from literatures on organization, contagion and cooperation, and assess the future opportunities and challenges for ABMs.

Characteristics of agent-based models

ABMs tend to possess four characteristics:

Computational description at the level of agents

ABMs consists of a large number of interacting agents, operating within an environment. Each agent's behavior is governed by rules triggered by their local condition rather than global information [10]. High-level summary descriptions emerge from the unfolding agent interactions,

Corresponding author: Goldstone, R.L. (rgoldsto@indiana.edu).
Available online 8 August 2005

 1364-6613/$ - see front matter doi:10.1016/j.tics.2005.07.009

but they are not explicitly programmed (see also Supplementary data online).

Stigmergic interactions

Agents act on and are influenced by their local environment. Stigmergy is a form of indirect communication between agents that is achieved by agents modifying their environment and also responding to these modifications [11], for example, ants following pheremone trails left by other ants [12]. Analogous stigmeric effects are achieved by 'swarms' of humans that make a terrain more attractive to others by creating paths with their own steps [13,14], book recommendations on Amazon.com based upon similar readers' buying habits, or robots that create large-scale architectural structures even though they cannot directly communicate, by reacting to and building upon structures left by others [15].

Autonomy of agents

Each agent is capable of autonomous behavior, and possesses individual, albeit frequently simplified, representations of beliefs, goals and strategies [16]. Agents typically do not calculate optimal or rational courses of action, but rather use heuristics [17,18], reinforcement learning [19], opportunistic adaptation [20], or cross-generational evolution [21] to change their strategies.

Spatially distributed populations of agents

ABMs often consist only of agents and a 2D or 3D landscape of environmental 'patches,' both of which may have several attributes [10]. The visuo-spatial and animated nature of the resulting simulations makes the most of people's natural aptitude for visual pattern recognition.

Idealized and detailed models

The most fundamental decision an ABM researcher makes is how detailed their model will be. 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. A good example of this strategy is Robert Axelrod's Culture Model [22]. The goal of this model is to explain how beliefs or attitudes converge or diverge in a population over time. Agents are placed at fixed locations within a 2D grid, and initially have random trait values on each of several features. The likelihood of neighboring agents interacting with one another is proportional to their similarity across all features. When two agents interact, one of the trait values of one of the agents is copied to the other agent, in a process that simulates cultural imitation or social influence. Over time, spatial clusters of like-minded agents develop, although some diversity of opinion is often maintained as dissimilar agents are unlikely to interact even if they are neighbors (see Figure 1). This simple model can explain (i) the spatial clustering of opinions, (ii) bandwagon effects, and (iii) the spontaneous division of a culture into sub-cultures.

The 'similarity begets even more similarity' dynamic has applications to the spread of smoking in teenagers [23] and geospatial political patterns [24], to take just two examples. The original Culture Model has given rise to follow-up simulations showing how the Culture Model can be extended to disseminating solutions to objective problems rather than opinions [25], incorporating global media rather than only neighbor-to-neighbor interactions [26], and also analyzed to show rapid phase transitions from disordered to ordered configurations of opinions [27]. The Culture Model may be too simplified to fully explain any specific real-world pattern of opinion spread, but it also formally captures an essential commonality of many situations in which homogeneity within cliques co-exists with striking heterogeneity across cliques.

Other ABMs are intimately tied to a specific domain because they include a considerable amount of detail derived from real world datasets and their goal is answering a specific real-world question. One such question is, 'Why did the Anasazi people of southwestern United States abandon their homeland around 1350 AD?' To find the answer, research teams [28,29] have developed ABMs that incorporate features grounded in historical records: maize production levels, ground water reserves, the 3D 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. Specific runs from the eventual model [29] (see Figure 2) capture aspects of the rise and fall of the Anasazi population for more than a millennium period from AD 200–1300, although the modeling would have been even more impressive if it had been systematically compared to alternative models.

The juxtaposition of these models allows us to critically assess the costs and benefits of idealized and detailed models. When successful in isolating a universal pattern, idealized models have widespread application to many real-world domains, and generate comprehensible explanatory accounts by focusing on only a few crucial causal elements. Revealing idealized models have been formulated for the diffusion of innovations [30], collective action [31], the transmission of cultural elements over generations [32], the development of social conventions [33], and language change [34,35]. The downsides of these idealized models are that without extensive tailoring, they might not map onto any actual case study, and they can oversimplify to the point of leaving out crucial details [36]. By contrast, detailed models hold the promise of making faithful predictions by being grounded in a case's particular data. There are ABMs that effectively incorporate considerable detail about university tenure systems [37], electricity markets in England [38], and hunting behavior in eastern Cameroon [39]. The downside of detailed models is that they may be able to predict too many possible outcomes if they have many parameters that are insufficiently constrained. If the models become too detailed, they may become as complex as the modeled phenomenon itself, and hence serve as poor explanatory aids. Given these considerations, it is best to choose the level of model detail based on the (1) importance of

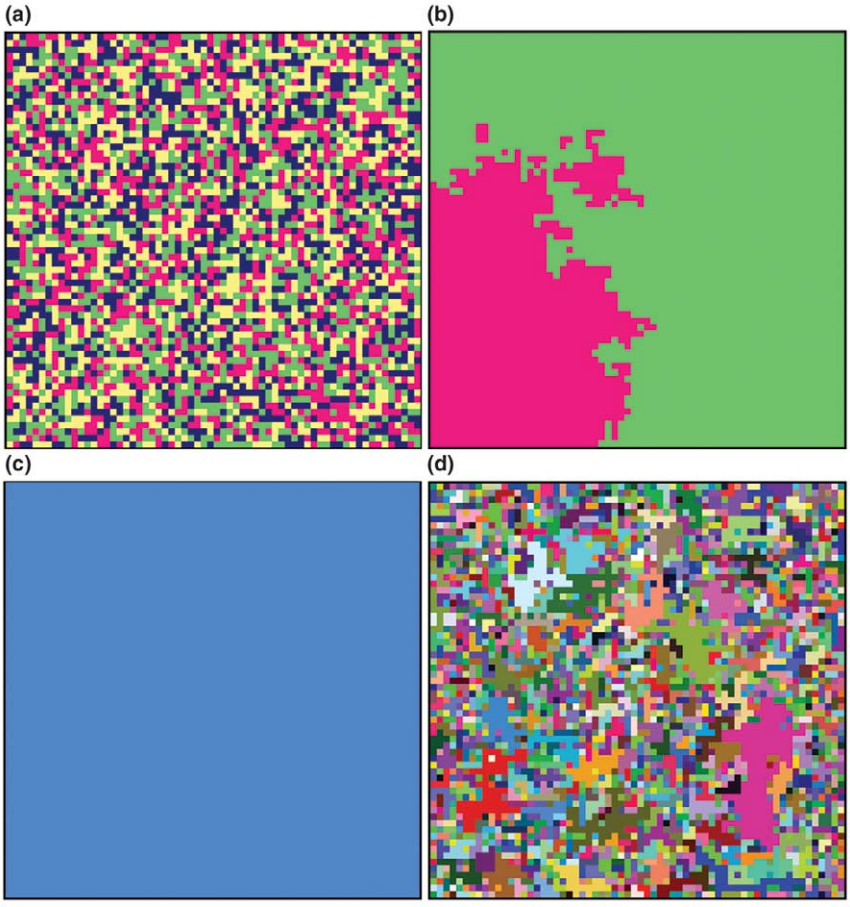

Figure 1. Four simulations of Axelrod's Culture Model [22]. Each agent is represented by a color that represents their entire set of features (e.g. hobby) and trait values (e.g. chess, badminton, violin). Agents with identical colors have identical traits. **(a)** an initial, randomly generated population of agents, each possessing two traits (T=2) on each of two features (F=2). **(b)** The same population after 4000 generations of interactions in which neighboring agents copy each others' traits, with the probability of interaction proportional to the agents' similarity. The population is frozen because the red and green agents have no traits in common and hence will never interact. **(c)** After 4000 generations, a simulation starting with agents characterized by two trait values (T=2) along 15 features (F=15). Increasing the number of features increases the probability of a homogeneous, like-minded population emerging. **(d)** When F=3 and T=15, the population quickly becomes frozen into small cliques that have no cross-group interactions. Increasing the number of traits per feature decreases the size of cliques. (A web-based simulation of the Culture Model can be found at http://www-personal.umich.edu/~axe/)

predicting future behavior in a particular case study (advocating detail), (2) importance of generalizing the model's behavior to new scenarios (advocating idealization), (3) difficulty in assessing what elements of a case study are crucial versus inconsequential in determining the global behavior (advocating detail), and (4) desire for a concise and comprehensible explanatory account (advocating idealization).

ABMs and aggregate models

The ABMs described above all work by synthetically constructing virtual versions of social phenomena from low-level descriptions of the individual agents. This approach is in contrast to descriptive models, which use equations to describe aggregate level phenomena. To create tractable descriptive equations it is often necessary to make misleading assumptions that fail to capture essential aspects of natural phenomena. One example is the Mean Field Approximation, according to which all individuals in a group are assumed to be in the same location and experience the same local environment. ABMs that incorporate space and local variability frequently produce much more realistic models. For example, giving agents unique rather than aggregate positions has proven invaluable in modeling the continued stability of host-pathogen populations [40] the genetic diversity in a population [41], and preserved pockets of cooperation surrounded by defectors [42,43]. More generally, ABMs often provide more satisfying accounts than purely descriptive approaches because they posit mechanisms by which aggregate qualities emerge.

Three core themes for agent-based models

Three prevalent themes for computational models of collective behavior have been spatial and temporal patterns, social contagion and cooperation. The human agents in these models are represented by a wide range in complexity, from particles to simple rule-following devices to rich cognitive architectures.

Patterns and organization

Political economist Thomas Schelling is one of the founders of computational models of collective behavior in the social sciences, although his original experiments on segregation were done by hand with dimes and pennies [44]. Schelling created agents belonging to two classes

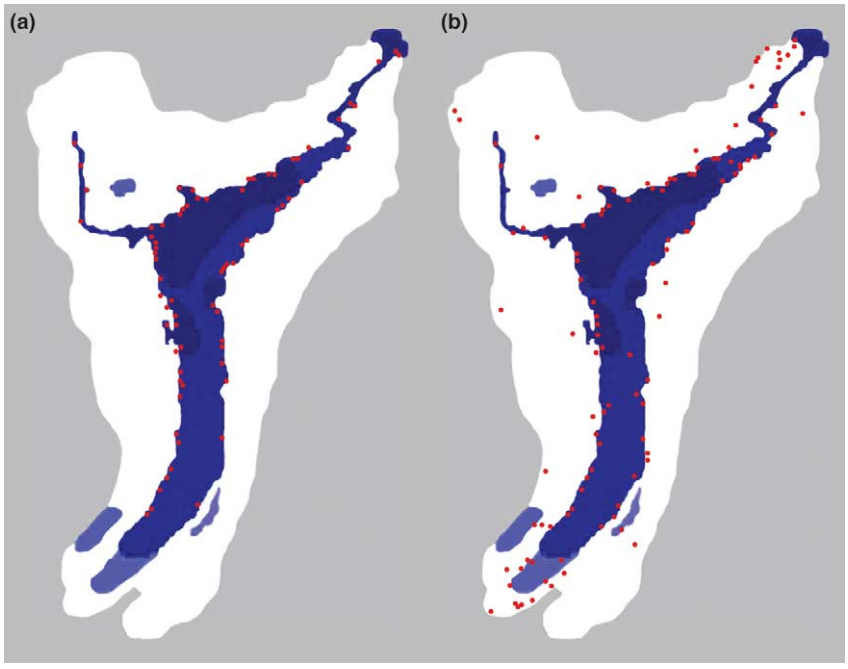

Figure 2. Simulated **(a)** and historical **(b)** patterns of settlement for the Anasazi in the Long House Valley around AD 1125. Red circles indicate settlements, and the shade of blue depicts the annual groundwater level for a location. Figure adapted by Rowan Johnston from [29] with permission from the author and publisher.

(his 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*'. The agents still divide themselves into sharply segregated groups after a short time, even though no individual is motivated to live in such a highly segregated world. The work of Schelling stimulated the development of other models of sorting where micro-motives lead to surprising macrobehavior, especially within political economic processes [43,45].

Recently, several physicists have used simulation models to study patterns that can emerge when many humans interact, including human trails [13], traffic jams [46], Mexican waves [47] and panic behavior of pedestrians [48]. In these models humans are represented as particles with variation in speeds or position, and without any requirement of cognition for the agents. In some applications of financial markets, agents are explicitly called 'zero intelligence agents' to show that full rationality is not required to explain observed patterns in economic statistics [49]. These simple reactive agent-based simulation models have often provided surprisingly apt accounts of empirically observed behavior.

Contagion

Social contagion is the spread of an entity or influence between individuals in a population via interactions between agents. Examples are the spreading of fads, rumors and riots. Computational approaches to simulate social contagion are based on thresholds models [50]. Each agent has a threshold that, when exceeded, leads the agent to adopt an activity. This threshold represents the number of other agents in the population or local neighborhood following that particular activity. Threshold models can be either deterministic [51] or stochastic [51]. Recent work in this area assumes that thresholds are applied to the adoption rate within a local neighborhood, rather than the whole population [52,53]. This has led to the study of the impact of different social network configurations on contagion [54,55].

Cooperation

A social dilemma is a situation where sub-optimal group outcomes are achieved if all agents do the action that is optimal for themselves. If self-centered rational agents do not cooperate in social dilemmas why do we often find cooperation in actual case studies? From an ABM perspective, are the roots of cooperation in the model of the individual, environmental conditions, or information exchange between agents? Robert Axelrod pioneered the use of computational models by showing that strategies which lead to conditionally cooperative behavior are effective in a tournament of repeated prisoner dilemma games, giving better overall performance than uncooperative strategies even though in head-to-head competition with non-cooperative strategies, the non-cooperative strategies prevail [56]. This work led to a large literature on extensions of the original models to include the addition of space [57], indirect reciprocity [58], and more complex strategies [59]. Most of this work uses simple reactive agents. Some recent studies [60,61] focus on more cognitively sophisticated agents with designs informed by psychological theories, such as social comparison and bounded rationality.

The future of computational models of collective behavior

These early explorations have given us enough data to make some prescriptions for the development of the next

generation of ABMs. Broadly speaking, there are several limitations with the current crop of ABM models from a cognitive science perspective. We have four recommendations for realizing the promise of ABMs:

Genuine predictiveness

If the first generation of ABMs can be characterized as generally 'post-dicting' existing data, the next generation should aspire to genuinely predicting the outcome of future patterns of collective behavior.

Create a computational lingua franca

Early work with ABMs has been somewhat unsystematic, with different researchers developing idiosyncratic systems. One result has been the lack of replicability of some ABM results. For example, early demonstrations of the persistence of cooperation in a spatial Prisoner's Dilemma game [57] have been shown to hinge crucially on the unstated assumption of agents that update synchronously (e.g. all at the same time) rather than asynchronously [62]. If ABM modelers all use the same simulation environments, then interesting differences between models can be missed [63], but establishing common modeling methods will promote comparison between models, improve replicability of results, and facilitate researchers' efforts to build upon prior work.

Greater synthesis between experiments and models

When ABMs have been compared to empirical data, they have often been applied to case studies. A problem with case studies is that they are not genuinely replicable events, although role-playing games can sometimes capture their major elements. Instead, we would recommend validating ABMs against data obtained from experiments. Ideally, laboratory experiments, field observations and computational models will not only be integrated, but they will inform and improve one another over several iterative cycles. Relatively simple laboratory situations can be constructed to involve groups of people interacting in idealized environments according to easily stated game rules. Experiments can bridge the oft-noted [36] gap between computational models and group behavior because the assumptions underlying the experiments can be tailored to correspond almost exactly to the assumptions of the computational models, and so the models can be aptly applied without sacrificing the concision of their explanatory accounts.

Greater sophistication of internal representations of agents

The majority of current ABMs incorporate rather impoverished representations of agents. Often, each individual is represented by a single, albeit time-varying, number (e.g. 'Probability of cooperation = 0.8') [3]. In more complicated models, agents are represented by a vector of independent values across a set of dimensions [22]. Actual knowledge has much richer structures than either of these representations. For example, an evolutionary theorist has concepts about natural selection, sexual reproduction and genetic variability within a population, but these concepts are not independent elements, but rather support and contextualize one another. Concepts gain their meaning by their relations to other concepts [64]. Finding a good balance between incorporating these influences and achieving constrained and elegant models is an excellent challenge for ABMs. If each person is to be modeled as a conceptual network, then a social group is to be modeled as network of networks. From a modeling perspective, the intellectual interest is in the study of how these two levels of networks interact [65]. Communicating is not simply transmitting individual concepts. Communication involves aligning the conceptual systems of agents [33]. One implication of this alignment process is that as concepts migrate across people, they will be systematically altered to fit their owners' conceptual network.

Opportunities

Despite these challenges to current ABMs, the future looks bright for computational modeling of collective behavior. We highlight three opportunities for future modeling efforts.

Modeling large-scale collectives

Recently, there has been a phenomenal increase in archival data on groups. Archival data available from on-line 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 [66].

Computational models as test-beds

Computational models of collective behavior can explore in advance the possible consequences of public policy changes. ABMs can be used to address '*what if*' scenarios like 'What is the consequence of the spread of HIV if policies are implemented to affect stigmatization of HIV infected persons?' and 'What would be the impact on world demographics if parents can choose the sex of their child?' As computational simulations become increasingly realistic, they will serve as increasingly useful test-beds for exploring potential consequences of public policies that have complex, non-linear dynamics.

Group control through indirect manipulation rather than explicit rules

Perhaps 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 [67]. ABMs suggest an alternative method of crowd control by changing the structure of the environment such that certain behaviors are facilitated and others are indirectly hindered without instituting physical or abstract barriers. Small changes in environments can often have a major change on the flow because of the positive feedback involved in individuals following other individuals. The need for direct force is reduced by this approach. As we gain confidence in our computational models, they will provide useful advice not only for predicting but also for controlling collective behavior (see also Box 1).

Box 1. Questions for future research

- What are the mechanisms by which individuals within groups learn to cooperate, compete, spontaneously specialize and divide labor, form coalitions, distribute resources, propagate innovations, create social networks and coordinate complex activities?
- How is the gulf between the real-world complexities of social interaction and the relative simplicity of ABMs best bridged: by creating simplified experimental environments that constrain human interaction, by incorporating as much real-world data into ABMs as possible, or by identifying crucial real-world elements and selectively incorporating only these into ABMs?
- When can models inspired from physics be effectively used to explain collective behavior, and when must the 'humans-as-particles' idealization be enriched to incorporate people's beliefs, memories, plans, strategies and creativity?
- Can ABMs be successfully used to advise organizational design or public policy, for example, by predicting the implications of new voting or auction rules before they are publicly implemented?

Conclusion

There are exciting opportunities for recent efforts in ABMs to both benefit from, and inform, cognitive science. ABMs can benefit from the advanced statistical tools and empirical methods that cognitive scientists have developed for assessing the quality of the fit between computational models and the world. The emphases of cognitive science on neural and cognitive constraints, replicability, comparison between models, laboratory-controlled validation and realistic cognitive processes of individual decision makers are much needed insights for the ABM community. Conversely, ABM methods advance cognitive science by providing a generative, proof-by-construction approach to understanding social behavior. Cognitive scientists often act as though individuals are the sole loci of organized thought, but ABMs remind us that organized behavior can be described at multiple levels, and that our thoughts both depend upon and determine the social structures that contain us as elements within those structures.

Acknowledgements

We are grateful to Robert Axelrod, Robert Axtell and Elinor Ostrom for helpful comments on earlier drafts of this article. More information about the laboratory can be found at http://cognitrn.psych.indiana.edu. This research was funded by NSF grant 0432894 and Department of Education, Institute of Education Sciences grant R305H050116.

Supplementary data

There is a Box listing some ABMs and useful programs in the supplementary data to this article, available online at doi:10.1016/j.tics.2005.07.009

References

1 Macy, M.W. and Willer, R. (2002) From factors to actors: computational sociology and agent-based modeling. *Annu. Rev. Sociol.* 28, 143–166
2 Kirman, A. and Zimmermann, J. (2001) *Economics with Heterogeneous Interacting Agents*, Springer-Verlag
3 Latane, B. and Bourgeois, M.J. (2000) Dynamic social impact and the consolidation, clustering, correlation, and continuing diversity of culture. In *Handbook of Social Psychology* (Vol. 4) (Tindale, R.S. and Hogg, M., eds), pp. 235–258, Blackwell
4 Kenrick, D.T. *et al*. (2003) Dynamical evolutionary psychology: Individual decision rules and emergent social norms. *Psychol. Rev.* 110, 3–28
5 Kohler, T. and Gumerman, G. (2002) *Dynamics in Human and Primate Societies*, Oxford University Press
6 Hutchins, E. (1995) *Cognition in the Wild*, MIT Press
7 Epstein, J.M. and Axtell, R. (1996) *Growing Artificial Societies: Social Science from the Bottom Up*, Brookings Institution/MIT Press
8 Goldspink, C. (2002) Methodological implications of complex systems approaches to sociality: Simulation as a foundation for knowledge. *J. Artif. Soc. Soc. Simul.* 5, 1–19
9 Goodwin, B. (1994) *How the Leopard Changed its Spots: The Evolution of Complexity*, Weidenfield & Nicolson
10 Resnick, M.R. (1994) *Turtles, Termites, and Traffic Jams*, MIT Press
11 Dorigo, M. *et al*. (2000) Ant algorithms and stigmery. *Fut. Gen. Comput. Syst.* 16, 851–871
12 Bonabeau, E. *et al*. (2000) Inspiration for optimization from social insect behavior. *Nature* 406, 39–42
13 Helbing, D. *et al*. (1997) Modeling the evolution of human trail systems. *Nature* 388, 47–50
14 Goldstone, R.L. *et al.* Group path formation. *IEEE Transactions on System, Man, and Cybernetics. Part A* (in press)
15 Howsman, T.G. *et al*. (2004) A stigmeric cooperative multi-robot control architecture. In *Artificial Life IX* (Pollack, J. *et al*., eds), pp. 88–93, MIT Press
16 Kaufman, S. (1996) *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity*, Oxford University Press
17 Simon, H. (1998) *The Sciences of the Artificial*, MIT Press
18 Gigerenzer, G. *et al*. (1999) *Simple Heuristics that Make Us Smart*, Oxford University Press
19 Sutton, R.S. and Barto, A.G. (1998) *Reinforcement Learning: An Introduction*, MIT Press
20 Holland, J. (1995) *Hidden Order: How Adaptation Builds Complexity*, Perseus
21 Holland, J.H. (1975) *Adaptation in Natural and Artificial Systems*, MIT Press
22 Axelrod, R. (1997) The dissemination of culture: A model with local convergence and global polarization. *J. Confict Resolut.* 41, 203–226
23 Kobus, K. (2003) Peers and adolescent smoking. *Addiction* 98, 37–55
24 Feldman, M.P. and Massard, N. (2002) *Institutions and Systems in the Geography of Innovation*, Kluwer
25 Kennedy, J. (1998) Thinking is social: Experiments with the adaptive culture model. *J. Conflict Resolut.* 42, 56–76
26 Shibanai, Y. *et al*. (2001) Effects of global information feedback on diversity. *J. Conflict Resolut.* 45, 80–96
27 Castellano, C. *et al*. (2000) Nonequilibrium phase transition in a model for social influence. *Phys. Rev. Lett.* 85, 3536–3539
28 Dean, J.S. *et al*. (2000) Understanding Anasazi cultural change through agent-based modeling. In *Dynamics in Human and Primate Societies* (Kohler, T. and Gumerman, G., eds), pp. 179–205, Oxford University Press
29 Axtell, R.L. *et al*. (2002) Population growth and collapse in a multiagent model of the Kayenta Anasazi in Long House Valley. *Proc. Natl. Acad. Sci. U. S. A.* 99, 7275–7279
30 Rosenkopf, L. and Abrahamson, E. (1999) Modeling reputational and information influences in threshold models of bandwagon innovation diffusion. *J. Comput. Math. Org. Theory* 5, 361–384
31 Chwe, M.S-Y. (1999) Structure and strategy in collective action. *Am. J. Sociol.* 105, 128–156
32 Boyd, R. and Richerson, P.J. (2005) *The Origin and Evolution of Cultures*, Oxford University Press
33 Barr, D.J. (2004) Establishing conventional communication systems: Is common knowledge necessary. *Cogn. Sci.* 28, 937–962
34 Hutchins, E. and Hazelhurst, B. (1995) How to invent a lexicon: The development of shared symbols in interaction. In *Artificial Societies: The Computer Simulation of Social Life* (Gilbert, N. and Conte, R., eds), pp. 157–189, UCL Press
35 Steels, L. (2002) Grounding symbols through evolutionary language games. In *Simulating the Evolution of Language* (Cangelosi, A. and Parisi, D., eds), pp. 211–226, Springer
36 Doran, J.E. (2001) Can agent-based modelling really be useful? In *Cooperative Agents: Applications in the Social Sciences* (Saam, N.J. and Schmidt, B., eds), pp. 57–81, Kluwer
37 Phelan, S.E. (2004) Using agent-based simulation to examine the robustness of up-or-out promotion systems in universities. *Nonlinear Dyn. Psychol. Life Sci.* 8, 177–203
38 Bunn, D.W. and Oliveira, F.S. (2001) Agent-based simulation: an application to the new electricity trading arrangements of england and wales. *IEEE Trans. Evol. Comput.* 5, 493–503

39 Bousquet, F. *et al*. (2001) Multi-agent simulations of hunting wild meat in a village in eastern Cameroon. *Ecol Model.* 138, 331–346
40 de Aguiar, M.A.M. *et al*. (2004) Invasion and extinction in the mean field approximation for a spatial host-pathogen model. *J. Stat. Phys.* 114, 1417–1451
41 Rauch, E.M. and Bar-Yam, Y. (2004) Theory predicts the unevent distribution of genetic diversity within species. *Nature* 431, 449–452
42 Fehr, E. and Fischbacher, U. (2004) Social norms and human cooperation. *Trends Cogn. Sci.* 8, 185–189
43 Nowak, M.A. and May, R.M. (1992) Evolutionary chaos and spatial games. *Nature* 359, 826–829
44 Schelling, T.C. (1971) Dynamic models of segregation. *J. Math. Sociol.* 1, 143–186
45 Kollman, K. *et al*. (1997) Political institutions and sorting in a Tiebout model. *Am. Econ. Rev.* 87, 977–992
46 Helbing, D. (2001) Traffic and related self-driven many-particle systems. *Rev. Mod. Phys.* 73, 1067–1141
47 Farkas, I. *et al*. (2002) Social behaviour: Mexican waves in an excitable medium – the stimulation of this concerted motion among expectant spectators is explained. *Nature* 419, 131–132
48 Helbing, D. *et al*. (2000) Simulating dynamical features of escape panic. *Nature* 407, 487–490
49 Farmer, J.D. *et al*. (2005) The predictive power of zero intelligence in financial markets. *Proc. Natl. Acad. Sci. U. S. A.* 102, 2254–2259
50 Granovetter, M. (1978) Threshold models of collective behavior. *Am. J. Sociol.* 83, 1420–1443
51 Strang, D. and Macy, M.W. (2001) In search of excellence: fads, success stories, and adaptive emulation. *Am. J. Sociol.* 107, 147–182
52 Valente, T.W. (1996) Social network thresholds in the diffusion of innovations. *Soc. Networks* 18, 69–89
53 Solomon, S. *et al*. (2000) Social percolation models. *Physica A (Amsterdam)* 277, 239–247
54 Abrahamson, E. and Rosenkopf, L. (1997) Social network effects on the extent of innovation diffusion: A computer simulation. *Org. Sci.* 8, 289–309
55 Janssen, M.A. and Jager, W. (2003) Simulating market dynamics: Interactions between consumer psychology and social networks. *Artif. Life* 9, 343–356
56 Axelrod, R. (1984) *Evolution of Cooperation*, Basic Books
57 Nowak, M.A. and May, R.M. (1992) Evolutionary games and spatial chaos. *Nature* 359, 826–829
58 Nowak, M.A. and Sigmund, K. (1998) Evolution of indirect reciprocity by image scoring. *Nature* 393, 573–577
59 Gotts, N.M. *et al*. (2003) Agent-based simulation in the study of social dilemmas. *Artif. Intell. Rev.* 19, 3–92
60 Jager, W. *et al*. (2002) How uncertainty stimulates over-harvesting in a resource dilemma: Three process explanations. *J. Environ. Psychol.* 22, 247–263
61 Mosler, H.J. and Brucks, W.M. (2003) Integrating commons dilemma findings in a general dynamic model of cooperative behavior in resource crises. *Eur. J. Soc. Psychol.* 33, 119–133
62 Huberman, B.A. and Glance, N.S. (1993) Evolutionary games and computer simulations. *Proc. Natl. Acad. Sci. U. S. A.* 90, 7716–7718
63 Axtell, R. *et al*. (1996) Aligning simulation models: a case study and results. *Comput. Math. Organ. Theory* 1, 123–141
64 Goldstone, R.L. and Rogosky, B.J. (2002) Using relations within conceptual systems to translate across conceptual systems. *Cognition* 84, 295–320
65 Börner, K. *et al*. (2004) The simultaneous evolution of article and author networks in Proc. Natl. Acad. Sci. U. S. A.. *Proc. Natl. Acad. Sci. U. S. A.* 101(Suppl. 1), 5266–5273
66 Berger, J.A. and Heath, C. (2005) Idea habitats: how the prevalence of environmental cues influences the success of ideas. *Cogn. Sci.* 29, 195–222
67 Bonnes, M. and Carrus, G. (2004) Environmental psychology. In *Encyclopedia of Applied Psychology* (Spielberger, C., ed.), pp. 801–814, Elsevier