Cognitive Simulation of Academic Science

Ron Sun and Isaac Naveh

***Abstract*— This work describes a cognitively realistic approach to social simulation. It begins with a model created by Gilbert [4] for capturing the growth of academic science. Gilbert's model, which was equation-based, is replaced here by an agent-based (neural network) model, with the (neural network based) cognitive architecture CLARION providing greater cognitive realism. Using this agent model, results comparable to previous simulations and to human data are obtained. It is found that while different cognitive settings may affect the aggregate number of scientific articles produced by the model, they do not generally lead to different *distributions* of number of articles per author. It is argued that using more cognitively realistic models in simulations may lead to novel insights.**

I. Introduction

An important development in social sciences has been *agent-based* social simulation (ABSS). This approach consists of constructing models of societies of artificial agents. Agents are autonomous entities with well-defined rules of behavior. Running such a model amounts to instantiating a population of agents, allowing the agents to run, and observing the interactions between them. ABSS thus differs markedly from traditional (equation-based) approaches to simulation, where relationships among conceptual entities (e.g., social groups and hierarchies, or markets and taxation systems) are expressed through a set of mathematical equations. Agent-based modeling has a number of advantages over equation-based modeling, including, notably, the ability to represent a heterogeneous population and to realistically model social networks.

Interestingly, the evolution of simulation as a means for computational study of societies has been mirrored by developments in computational modeling at the individual level. Whereas earlier models of cognition tended to emphasize one of the aspects of cognition (for instance, memory or learning), some recent approaches have been more integrative, with a focus on putting the pieces together. The products of this integrative approach are known as *cognitive architectures*, and are essentially models that capture different aspects of cognition and their interaction. Such models tend to be generic and task-independent. Cognitive architectures have greatly grown in expressive power in recent years, and now capture a variety of cognitive phenomena, including various types of memory/representation, modes of learning, and sensory-motor capabilities (see, e.g., Anderson and Lebiere [1]; Sun [13]). At the same time, there has been a growing interest in studying multi-agent interactions, particularly issues of coordination and cooperation among cognitive agents.

So far, however, the two fields of social simulation and cognitive architectures have developed in near-isolation from each other (with some exceptions; e.g., Carley and Newell [2]). Thus, most of the work in social simulation assumes very rudimentary cognition on the part of the agents. At the same time, while the mechanisms of individual cognition have been the subject of intensive investigation in cognitive architectures, the relationships between sociocultural forces and individual cognition remain largely unexplored (again with some exceptions).

We believe, however, that the two fields of social simulation and cognitive architectures can be profitably integrated. As has been argued before (Sun and Naveh [16]; Moss [8]; Castelfranchi [3]), social processes ultimately rest on the choices and decisions of individuals, and thus understanding the mechanisms of individual cognition can lead to better theories describing the behavior of aggregates of individuals. So far, most agent models in social simulation have been extremely simple (in the form of very simple automata with a few ad-hoc assumptions) or entirely absent (in the case of equation-based modeling). However, we believe that a more realistic cognitive agent model, incorporating realistic tendencies, inclinations and capabilities of individual cognitive agents can serve as a more realistic basis for understanding the interaction of individuals (Moss [8]). Although some cognitive details may ultimately prove to be irrelevant, this cannot be determined *a priori*, and thus simulations are useful in determining which aspects of cognition can be safely abstracted away.

At the same time, by integrating social simulation and cognitive modeling, we can arrive at a better understanding of individual cognition. By studying cognitive agents in a social context, we can learn more about the sociocultural processes that influence individual cognition.

In this work, we describe a cognitively realistic approach to social simulation. We begin with a model created by Gilbert [4] for capturing the growth of academic science. Gilbert's model is then replaced by an agent-based (neural network) model, with the (neural network based) cognitive architecture CLARION providing greater cognitive realism. Using this agent model, results comparable to previous simulations and to human data are obtained. It is found that while different cognitive settings may affect the aggregate number of scientific articles produced by the model, they do not generally lead to different *distributions* of number of articles per author. We argue that using more cognitively realistic models in simulations may lead to novel insights.

Ron Sun and Isaac Naveh are with Department of Cognitive Sciences, Rensselaer Polytechnic Institute, Troy, New York (email: rsun@rpi.edu).

This work has been supported in part by the ONR grant N00014-08-1-0068 (to Ron Sun) and the ARI contract W74V8H-05-K-0002 (to Ron Sun and Bob Mathews).

II. Gilbert's Model of Academic Science

Science develops in certain ways. In particular, it has been observed that the number of authors contributing a certain number of articles to a scientific journal follows a highly skewed distribution, corresponding to an inverse power law. This distribution, which is known as a Zipf distribution, is common to a number of other phenomena in information science. In the case of scientific publication, the tendency of authorship to follow a Zipf distribution was observed by Lotka [6] and is known as Lotka's law.

Simon [12] developed a simple stochastic process for approximating Lotka's law. One of the assumptions underlying this process is that the probability that a paper will be published by an author who has published i articles is equal to a/i^k, where a is a constant of proportionality.

Using Simon's work as a starting point, Gilbert [4] attempts to model the growth of science, including Lotka's law. He obtains his simulation data based on the assumption that the system selects a focal paper randomly first, which can be represented as a point in a multi-dimensional (two-dimensional) space of ideas, and then it randomly selects a number of other papers, each of which occupies a different point in the nearby region and pulls the original point in its direction slightly. The resulting paper can be located on that multi-dimensional space based on the above factors. Papers are randomly assigned authors, based on a stochastic process that takes the ratio of papers to published authors into account. To capture the constraint that academic papers must be original, a newly published paper must lie at least m coordinate units away from any other existing paper, where m is a constant.

Another assumption is that the number of papers produced in a given time period is determined by the number of papers in existence during the previous time period, by specifying a small probability of each existing paper acting as the seed for a new paper (and then by selecting an author for that paper). Thus, it is papers that spawn more papers, with authors playing only an ancillary role in the process.

Using this model, Gilbert obtained an idea space divided into clusters, which he identified as corresponding to different scientific specialties. Each cluster originates in a few seminal papers and accumulates additional papers at an increasing rate over time. This model yielded a set of publication trends that accord with human data, including the power curve described above (as can be seen in Tables 1-2). A highly uneven distribution of number of publications per author was observed, with the majority of authors publishing but one paper. A similarly skewed outcome was obtained for the number of citations received per author, with most authors receiving a modest number of citations, and a minority of authors receiving a large number of citations.

However, to a significant extent, Gilbert's model is not cognitively realistic. The model assumes that authors are non-cognitive and interchangeable; it therefore neglects a host of cognitive phenomena that characterize scientific inquiry (e.g., learning, creativity, the evolution of field expertise, etc.). Using a more cognitively realistic model, we can envision addressing some of these omissions, as well as exploring other emergent properties of the model and their correspondence to real-world phenomena. The challenge, then, is to develop a model that explains macro-level phenomena in terms of micro-level cognitive processes.

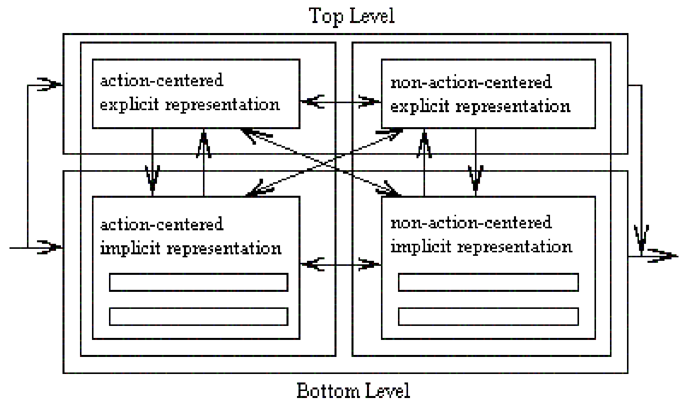

Fig. 1. The CLARION architecture

III. Our Cognitive Architecture

A. Explicit vs. Implicit Learning

In an attempt to understand the processes underlying human learning, various categories of knowledge have been proposed. Among them, one enduring distinction is that between explicit and implicit—or conscious and unconscious—learning (e.g., Reber [10]). While both implicit and explicit learning have been actively investigated, the complex interaction between these two modes of learning has largely been downplayed or discounted (with a few exceptions; e.g., Mathews et al [7]; Sun et al [15]).

However, despite the lack of study of such interaction, recent evidence suggests that it is difficult to find a situation in which only one type of learning is employed (Reber [10]; Seger [11]). Our review of empirical data suggests that while one can manipulate conditions so that one type of learning is favored over the other, in nearly every case, both types are involved, with varying degrees of contributions from each (see, e.g., Sun et al [15]; Mathews et al [7]).

In the next subsection, we describe a cognitive architecture, CLARION, which seeks to capture the interaction between explicit and implicit learning (Sun et al [15]). CLARION learns in bottom-up fashion, by extracting explicit rules from implicit knowledge, in accordance with what has been observed in humans (e.g., Karmiloff-Smith [5]).

B. A Summary of CLARION

CLARION is a general cognitive architecture with a dual representational structure (Sun [13]). It consists of two levels: the top level encodes explicit knowledge, and the bottom level encodes implicit knowledge. See Figure 1 for a sketch of the model.

At the bottom level, the inaccessible nature of implicit knowledge is captured by a subsymbolic distributed representation provided by a backpropagation neural network. This is because representational units in a distributed environment are capable of performing tasks but are subsymbolic and generally not individually meaningful (see Sun [13]). Thus, they are relatively inaccessible. Learning at the bottom level proceeds in trial-and-error fashion, with the neural networks being guided by reinforcement learning (i.e., Q-learning; see Watkins [17]).

At the top level, explicit knowledge is captured by a symbolic or localist representation, in which each unit is easily interpretable and has a clear meaning. This characteristic captures the property of explicit knowledge being more accessible and manipulable (Sun [13]). Learning at the top level involves, first, constructing a rule that corresponds to a "good" decision made by the bottom level. This rule is subsequently refined, either by generalizing or specializing it. Here, the learning process is guided by an "information gain" measure that compares the success ratio of various modifications of the current rule.

The overall algorithm of CLARION's action decision making is the following:

1) Observe the current state x.
2) Compute in the bottom level the value of each of the possible actions (a_i's) associated with the state x: $Q(x, a_1)$, $Q(x, a_2)$, ..., $Q(x, a_n)$.
3) Find out all the possible actions $(b_1, b_2, ..., b_m)$ at the top level, based on the state x and the rules in place at the top level.
4) Compare the values of a_i's with those of b_j's (which are sent down from the top level), and choose an appropriate action a.
5) Perform the action a, and observe the next state y and (possibly) the reinforcement r.
6) Update the bottom level in accordance with the *Q-Learning-Backpropagation* algorithm, based on the feedback information (as will be explained later).
7) Update the top level using the *Rule-Extraction-Refinement* algorithm (explained below).
8) Go back to Step 1.

At the bottom level, a Q-value measures the "quality" of an action in a given state; that is, $Q(x, a)$ indicates how desirable action a is in state x. Actions are selected based on Q-values. To acquire the Q-values, Q-learning, a reinforcement learning algorithm (Watkins [17]), is used. Q-learning is implemented in backpropagation networks. See Sun [13] for full details. Q-values are then used to decide probabilistically on an action to be performed (using a distribution of Q-values).

At the top level, explicit knowledge is captured by simple prepositional rules. We devised an algorithm for extracting rules using information culled from the bottom level (the *Rule-Extraction-Refinement*, or RER, algorithm; see Sun et al. [15]). The basic idea is as follows: whenever an action decided by the bottom level is successful, a rule (with conditions corresponding to the current input state and an action corresponding to the one selected by the bottom level) is created and added to the top level. Then, in subsequent interactions with the world, an agent may refine a rule by considering its outcome: if successful, an agent may try to generalize a rule by relaxing its conditions to make it more universal. If a rule is unsuccessful, an agent may try to specialize a rule by imposing further constraints on the rule and making them exclusive of the current state. This is an on-line version of hypothesis testing processes studied in other contexts.

To integrate the outcomes from the two levels, a number of methods are possible. Here, levels are selected stochastically, with a base probability of selecting each level. Other selection methods are possible as well (see Sun et al [15]).

IV. Simulation with Cognitive Agents

In our simulation of academic science, we move to an agent-based model. Different from Gilbert's assumptions in his simulation, we treat authors as cognitive agents. Thus, authors are not merely passive placeholders, but cognitively capable individuals whose success or failure depends on their ability to learn in the scientific world. Successful authors (that is, agents who manage to identify promising research leads early on) will go on to publish many papers in their area, whereas unsuccessful authors will be removed from the system and replaced.

Similar to Gilbert's approach, our model characterizes the scientific world as consisting of "papers," each of which proposes a new piece of knowledge, and of authors who combine these papers to form new papers. This coincides with the constructivist view of scientific inquiry, which sees scientific knowledge first and foremost as the product of a *constructive* process. On this view, the fruits of scientific inquiry are produced in a predetermined technological, linguistic, and social context. This is reflected in our model, in which papers are constructed from previous papers, and themselves serve as a basis for future constructions.

To publish a paper, an agent adopts a focal idea (as represented by an existing paper), non-randomly, in accordance with some cognitive processes. The cognitive processes may be implicit or explicit. The agent then uses other ideas (published papers) that pull the original idea in different directions, also non-randomly, based on similar cognitive processes.

In addition, apart from utilizing existing ideas, an agent also performs local search to "optimize" the resulting idea. This reflects the fact that authors do not merely cobble together ideas from existing sources, but also try to integrate the different ideas and to refine the final product.

Because our simulation involves learning agents, there is the possibility of failure; this is important, because humans can produce papers that prove to be unpublishable. This is in contrast to Gilbert's approach, in which ideas are undifferentiated in their quality. Instead, in our simulation, each agent has a set of evaluation functions that determine the quality of ideas in the multi-dimensional idea space.

These functions specify the most important considerations in terms of evaluating a scientific idea (e.g., clarity, insightfulness, empirical evidence, theoretical results, and application potential). Agents are aware of these functions. However, just as researchers in the real world cannot predict precisely when the result of their research will meet with approval and interest, so the agents' individual valuations of these functions may differ from the community-wide valuation. This is reflected in a set of *individual* evaluation functions for each agent, consisting of a varied version of the global, or *communal* evaluation functions.

The author population consists of CLARION-based agents. The feedback to agents is based on paper acceptance or failure. In addition, agents are awarded *partial* feedback at each step of the paper generation process, amounting to a fraction of the unfinished paper's evaluation (as determined by the agent's own evaluation functions). This reflects the fact that agents do not stumble blindly through the publication process, but rather are guided to a certain extent by their experience and intuition.

An agent uses the bottom level of CLARION to select a focal idea and a number of pull ideas. These two tasks are carried out by one network. The network "learns" using the *Q-learning-Backpropagation* algorithm, which corresponds to a simple form of reinforcement learning and naturally captures sequences of actions (i.e., selecting the focal idea, then the first and second pull ideas).

On the other hand, an agent uses RER rule learning in CLARION to extract rules that determine: (1) how to choose focal ideas, and (2) how to choose pull ideas. These rules are used in conjunction with other rules already existing at the top level concerning local search, which represent *a priori* knowledge.

Agents are pre-trained for a certain number of cycles before entering the system. A cycle corresponds to a single attempt by an agent to publish a paper, whether successfully or not. Reflecting a "publish or perish" academic environment, agents are evaluated every few cycles based on their publication record (success rate). If the latter falls below a minimum expected standard, the agent is removed from the academic world. If the agent passes all the evaluations, it retires upon reaching the maximum allowable age. Whenever an agent retires (or is removed), a new agent takes its place.

The results of our simulation are shown in Tables 1-2, along with results (reported by Simon) for *Chemical Abstracts* and *Econometrica*, and estimates obtained from previous simulations by Simon [12] and Gilbert [4]. The figures in the tables indicate number of authors contributing to each journal, by number of papers each has published.

The CLARION simulation data for the two journals can be fit to the power curve curve $f(i) = a/i^k$, resulting in an excellent match.

The results of the curve fit are shown in Table 3, along with correlation and error measures. Aggregate paper counts for authors with 11 or more papers were not included in the fit.

TABLE I
NUMBER OF AUTHORS CONTRIBUTING TO *Chemical Abstracts*.

# of Papers	Actual	Simon's estimate	Gilbert's simulation	CLARION simulation
1	3991	4050	4066	3803
2	1059	1160	1175	1228
3	493	522	526	637
4	287	288	302	436
5	184	179	176	245
6	131	120	122	200
7	113	86	93	154
8	85	64	63	163
9	64	49	50	55
10	65	38	45	18
11 or more	419	335	273	145

TABLE II
NUMBER OF AUTHORS CONTRIBUTING TO *Econometrica*.

# of Papers	Actual	Simon's estimate	Gilbert's simulation	CLARION simulation
1	436	453	458	418
2	107	119	120	135
3	61	51	51	70
4	40	27	27	48
5	14	16	17	27
6	23	11	9	22
7	6	7	7	17
8	11	5	6	18
9	1	4	4	6
10	0	3	2	2
11 or more	22	25	18	16

TABLE III
RESULTS OF FITTING CLARION DATA TO POWER CURVES. CA STANDS FOR CHEMICAL ABSTRACTS AND E STANDS FOR ECONOMETRICA.

Journal	a	k	Pearson R	R-square	RMSE
CA	3806	1.63	0.999	0.998	37.62
E	418	1.64	0.999	0.999	4.15

In our simulation, the number of papers per author reflected the cognitive ability and the cognitive suitability of an author, as opposed to being based on auxiliary assumptions such as those made by Gilbert [4]. This explains, in part, the slightly greater divergence of our results from the human data: whereas Gilbert's simulation consists of equations selected to match the human data, our approach relies on much more detailed and lower-level mechanisms—namely, a cognitive agent model that is generic rather than task-specific. The result of the CLARION-based simulation is therefore emergent, and not a result of specific and direct attempts to match the human data. We put more distance between mechanisms and outcomes, which makes it harder to obtain a match with the human data. Thus, the fact that we were able to match the human data shows the power of our cognitive agent-based approach compared to traditional methods of simulation.

V. Varying the Cognitive Parameters

Because CLARION captures a wide variety of cognitive phenomena, we can vary parameters that correspond to specific cognitive factors, and observe the effect on performance as function of *cognition*. This confers an important advantage over other, more task-specific models, where differences in performance tend to be artifacts of the particular model used and may be of little independent interest. With CLARION, the parameters being altered are the fundamental building blocks of cognition, and thus observed differences in performance are far more likely to stem from testable differences in individual cognition.

The effect of learning rate on performance is shown in Figure 2. An agent's learning rate essentially determines its responsiveness to success or failure. As can be seen, the best performance is obtained under a moderately high learning rate (0.1-0.3). If the learning rate is too high, an agent's recent experiences tend to disproportionately impact the learning process. This leads the agent to submit articles that are too similar to recently successful articles, which leads to more non-unique articles that are in turn rejected; equally, an agent is too swift to abandon a previously promising line of research as the result of a single rejection. If, on the other hand, the learning rate is too low, an agent will be slow to capitalize on recent successes and failures.

The fact that a balance between implicit learning and explicit rule learning is desirable can be seen in Figure 3. While increasing the reliance on implicit learning can lead to modest gains, over-emphasizing it at the expense of explicit rule learning slows down the learning process dramatically. This is especially true during the initial stages of learning, when neural networks are still imperfectly trained, and rules, as crisp guidelines that are based on past success, are useful for speeding up learning.

Crucial in this connection is the existence of a high-quality rule base. This can be ensured, among other things, by the proper selection of a "density" parameter, as shown in Figure 4. If the parameter is too low, then rules persist even when they are no longer needed (for instance, when an agent has already exhausted a particular line of research encoded as a rule, and has moved on to other fields, represented by different combinations of ideas). On the other hand, when the parameter is too high, even successful rules are often deleted before they can be fully utilized. In both cases, performance suffers.

An agent's exploration of the idea space is modulated, in considerable part, by the built-in randomness (i.e., "temperature") of its search process. As can be seen in Figure 5, agents are at their most prolific under a moderately high temperature setting—that is, when they show a willingness to experiment (i.e., to pursue new leads, which would not occur under a low temperature setting) while still being guided by their experience in the majority of cases. This observation accords with what we know of the role of serendipity in scientific discovery. In many areas of science—for instance, medicine—most of the major discoveries have been serendipitous, the result of seemingly irrelevant investigations. While such discoveries were formerly attributed to good fortune, it has since been argued (e.g., in Oliver [9]) that serendipity is, in part, a cognitive faculty that can be nurtured and developed. Our model captures this characteristic by adopting a modest degree of randomness in the decision-making process.

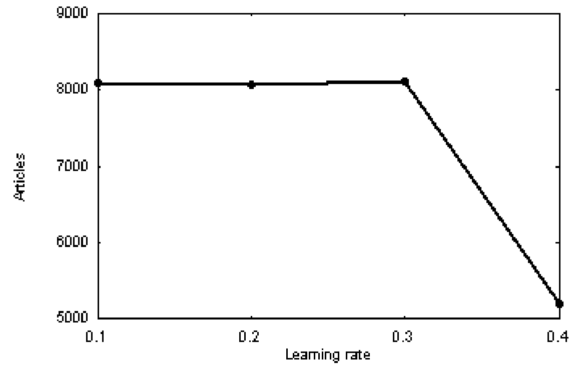

Fig. 2. The effect of learning rate on collective number of articles published.

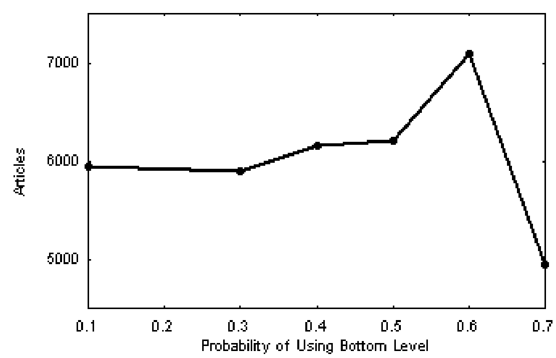

Fig. 3. The effect of explicit vs. implicit learning on collective number of articles published.

As with other parameters considered so far, an agent's generalization threshold must be carefully selected (see Figure 6). If it is set too low, even less successful rules will be generalized, leading to a lower-quality rule base. Too high, and it will prevent the generalization even of successful rules. In the latter situation, an agent will rigidly apply successful ideas only in the precise context in which they initially appeared (for instance, as the second "pull" idea in generating a paper) without recognizing their more general applicability.

As the preceding discussion shows, the cognitive parameters of individual agents are crucial in determining the rate of scientific progress in a society. By varying these parameters, we can arrive at communities that produce lesser or greater numbers of papers. Apart from this aggregate measure of scientific productivity, however, it is also interesting to see if the patterns of individual contribution observed in earlier simulations will be preserved. In particular, we want to see if the power curve obtained earlier will be obtained under

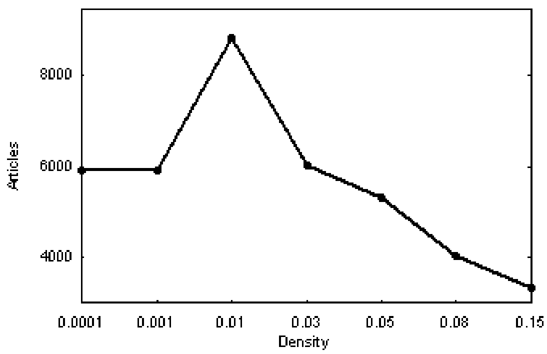

Fig. 4. The effect of density on collective number of articles published.

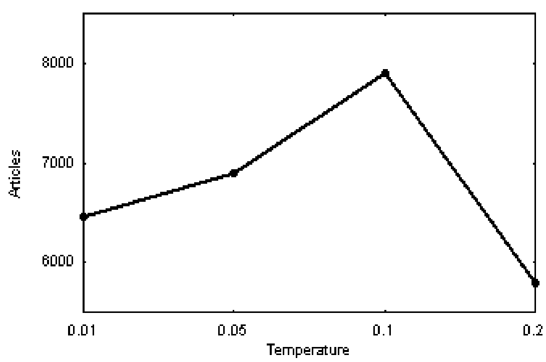

Fig. 5. The effect of temperature on collective number of articles published.

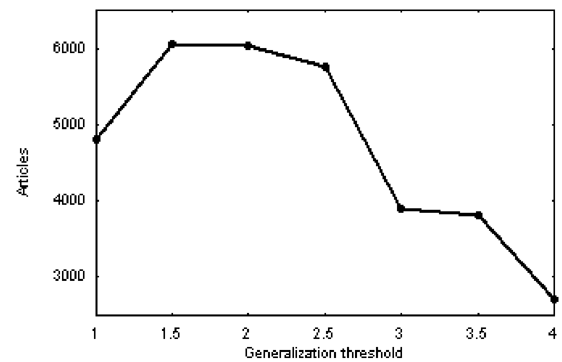

Fig. 6. The effect of generalization threshold on collective number of articles published.

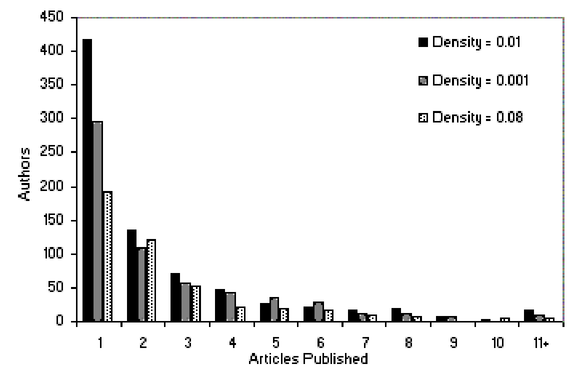

Fig. 7. Authors contributing to final paper count, by number of articles that each has published. CLARION simulation results for different settings of the density parameter.

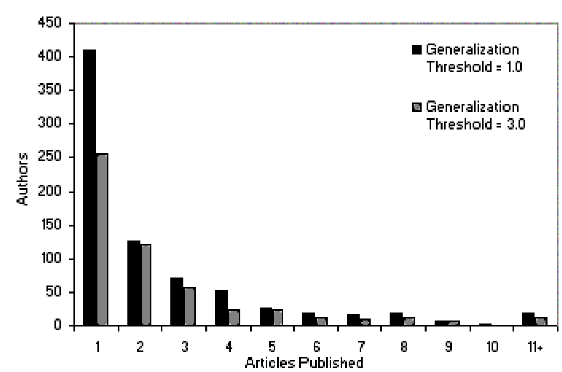

Fig. 8. Authors contributing to final paper count, by number of articles that each has published. CLARION results for different settings of the generalization threshold parameter.

different cognitive settings. Our results show that it is. As can be seen in Figures 7-8, different settings of the density and generalization threshold parameters lead to larger or smaller numbers of papers in aggregate, but they do not fundamentally change the authorship curve, which follows an inverse power distribution in all cases. Similar results were obtained for some other, though not all, ranges of cognitive parameters.

This result, which may be termed social-cognitive invariance, is an important one, since it shows that some regularities that characterize societies are to some extent invariant with respect to agent cognition (within a reasonable parameter range). While some societies may prove more successful than others in terms of absolute scientific productivity, the same large-scale patterns may be observed regardless of cognitive differences. This reduces the likelihood that the patterns observed are a byproduct of a particular set of cognitive parameters. In contrast, in Sun and Naveh [16], we have shown that some other patterns are indeed directly related to the settings of cognitive parameters.

VI. Discussions

One important aim of this study has been to determine whether the results of a previous model of academic science (of Gilbert [4]) can be reproduced without resorting to the broad simplifications of an equation-based model. The results of our simulation suggest that the observed growth of academic science can indeed be captured even if one migrates to an agent-based model. Such a migration offers several important benefits. First, it allows us to leave behind certain artificial assumptions (for instance, that papers automatically spawn more papers). Second, it affords us the opportunity of studying the macro-level repercussions of behavior at the micro level. Third, it allows us to study patterns of interaction between individual agents. Although the latter interactions occur only indirectly in our model (either through the collision of too-similar papers generated by different authors, or through the exploitation of others' ideas in generating new papers) they nonetheless result in a model that is more socially realistic than Gilbert's equation-based model.

What makes our approach unusual, however, is not that it represents actors as agents, but that it takes agent cognition seriously. So far, most agent models in simulations have been rather simple, with little attention being paid to the mechanisms of individual cognition. This study shows that a more cognitively realistic simulation, with CLARION, can replicate the results of earlier simulations. It thus provides a dual corroboration of these models, by showing them to be independent both of whether or not an agent-based model is used, and of whether or not cognitive representations are involved. Therefore, while some cognitive details clearly cannot be abstracted away, others can, and along the way, we discover important cognitive-social invariances.

Apart from validation, however, cognitive realism in simulations may lead us to better representations of the target phenomenon. For instance, we identified earlier a possible way of representing the role of serendipity in science: namely, as a researcher's willingness to explore apparently suboptimal combinations of ideas, rather than adhering to "tried-and-true" sequences. The ability to represent such aspects of observed phenomena in terms organic to the agent model, rather than through auxiliary assumptions (for instance, by adding a "randomizing" function to the idea selection process in Gilbert's simulation) is an advantage of cognitively realistic simulations.

Another advantage of cognitive realism is that it allows us to theorize about the relative role that individual cognitive factors play in the emergence of large-scale social phenomena. Thus, we were able to vary parameters of CLARION that correspond to aspects of cognition and test their effect on outcomes. Our investigation showed, for instance, that the tendency to engage in inductive reasoning (that is, an agent's generalization threshold) can dramatically influence the number of papers generated by the community. It moreover suggested that this phenomenon could be described by a u-shaped curve. Such results suggest how patterns of communal thinking may change as a consequence of shifts at the individual cognitive level.

Finally, a complete cognitive model of academic publishing would also need to incorporate a model of what motivates a scientist to carry out research and publish, as well as a model of the socio-political forces at work in science. The motivational component does exist in CLARION, so the current model can be extended to deal with these issues.

VII. Conclusions

While interest in simulation as a way of describing social phenomena continues to grow apace, the issue of choosing a realistic cognitive representation has largely been ignored. Against this trend we propose using more complex cognitive models, known as cognitive architectures, to capture human behavior. Our model of scientific publication assumed that, in constructing new ideas from previous ones, authors are guided chiefly by cognitive processes. The simulations were done under different cognitive settings and suggest that some of the patterns obtained (for instance, Lotka's Law) are to some extent independent of the cognitive parameters selected.

By paying more attention to the details of individual cognition, we can arrive at more accurate representations of target social phenomena. We can also learn which cognitive mechanisms are significant in shaping social interactions, and which are not. Finally, we can study the emergence of large-scale social behavior from micro-level cognitive processes.

References

[1] J. R. Anderson and C. Lebiere, *The Atomic Components of Thought*. Lawrence Erlbaum Associates, Mahwah, NJ. 1998.

[2] K.M. Carley and A. Newell, "The Nature of the Social Agent," *Journal of Mathematical Sociology*, 19 (4), 221-262. 1994.

[3] C. Castelfranchi, "The Theory of Social Functions: Challenges for Computational Social Science and Multi-Agent Learning," *Cognitive Systems Research*, special issue on the multi-disciplinary studies of multi-agent learning, 2 (1), 5-38. 2001.

[4] N. Gilbert, "A simulation of the structure of academic science," *Sociological Research Online*, 2 (2). 1997. Available online at http://www.socresonline.org.uk/socresonline/2/2/3.html.

[5] A. Karmiloff-Smith, "From Meta-Processes to Conscious Access: Evidence from Children's Metalinguistic and Repair Data," *Cognition*, 23, 95-147. 1986.

[6] A. J. Lotka, "The Frequency Distribution of Scientific Productivity," *Journal of the Washington Academy of Sciences*, 16, 317-323. 1926.

[7] R. Mathews, R. Buss, W. Stanley, F. Blanchard-Fields, J. Cho and B. Druhan, "Role of Implicit and Explicit Processes in Learning From Examples: a Synergistic Effect," *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15, 1083-1100. 1989.

[8] S. Moss, "Relevance, Realism and Rigour: A Third Way for Social and Economic Research," CPM Report No. 99-56. Center for Policy Analysis, Manchester Metropolitan University, Manchester, UK. 1999.

[9] J. E. Oliver, *The Incomplete Guide to the Art of Discovery*. Columbia University Press, NY. 1991.

[10] A. Reber, "Implicit Learning and Tacit Knowledge," *Journal of Experimental Psychology: General*, 118 (3), 219-235. 1989.

[11] C. Seger, "Implicit Learning," *Psychological Bulletin*, 115 (2), 163-196. 1994.

[12] H. A. Simon, *Models of Man, Social and Rational*. Wiley, NY. 1957.

[13] R. Sun, *Duality of the Mind*. Lawrence Erlbaum Associates, Mahwah, NJ. 2002.

[14] R. Sun, *A Tutorial on CLARION 5.0*. 2003. Available online at `http://www.cogsci.rpi.edu/~rsun/sun.tutorial.pdf`

[15] R. Sun, E. Merrill and T. Peterson, "From Implicit Skills to Explicit Knowledge: A Bottom-Up Model of Skill Learning," *Cognitive Science*, 25 (2), 203-244. 2001.

[16] R. Sun and I. Naveh, "Simulating Organizational Decision-Making Using a Cognitively Realistic Agent Model," *Journal of Artificial Societies and Social Simulation*, 7 (3). 2004.

[17] C. Watkins, *Learning with Delayed Rewards*. Ph.D. Thesis, Cambridge University, Cambridge, UK. 1989.