

The retriever-connector model: Matching Classroom Data and Agent-Based Computer Models to Simulate Students' Use of Multiple Epistemological Resources

Paulo Blikstein (paulob@stanford.edu)

TLT Lab, School of Education and (by courtesy) Computer Science Department, Stanford University
520 Galvez Mall, CERAS 232, Stanford, CA 94305 USA

Abstract

Utilizing data from a classroom intervention with 8th graders, I employ agent-based computer modeling to simulate the cognitive processes at play during the intervention, in which students transition between using multiple epistemological resources. The model substantiates the hypothesis of manifold epistemological resources, which can be activated with simple prompts and have a non-linear impact on learning.

Keywords: Cognitive modeling; agent-based modeling; classroom research; epistemological resources.

Introduction

Agent-based modeling (ABM) has been increasingly used by scientists to study a wide range of phenomena such as the interactions of species in an ecosystem, the collisions of molecules in a chemical reaction, and the food-gathering behavior of ants (Bonabeau, Dorigo, & Theraulaz, 1999; Wilensky & Reisman, 2006). Such phenomena, in which the *agents* in a system (molecules, or ants) follow simple rules and interaction patterns, but exhibit complex emergent macroscopic behaviors, are studied in a young interdisciplinary field called complex systems or complexity studies (Holland, 1995). Although complex-systems perspectives initially arose from the natural sciences, complexity, emergence, and multi-level descriptions of phenomena are all highly relevant to social science research. In fact, recent decades have observed a surge in social-science studies employing ABM (Epstein & Axtell, 1996; Axelrod, 1997). Recently, ABM has also been used to illustrate aspects of cognitive development (Abrahamson & Wilensky, 2005; Blikstein, Abrahamson & Wilensky, 2006; Smith & Conrey, 2006), and collaboration and group work in classrooms (Abrahamson, Blikstein, & Wilensky, 2007).

ABM has the potential to advance theory in multiple ways, which I illustrate in this paper: (a) explicitizing—ABM demands a high level of specificity in expressing a theoretical model, and it provides the tools and standard practices to express those models; (b) dynamics—ABM enables researchers to mobilize an otherwise static list of conjectured behaviors and observe the macroscopic patterns that may unfold; (c) emergence—ABM can examine cognition and social behaviors as a collection of decentralized, simple rules; and (d) interdisciplinary collaboration—the *lingua franca* of ABM enables researchers from different fields to understand, critique and challenge each other's theories by modifying and extending the computational algorithms that underlie their theoretical models.

Relevance to cognitive research

Agent-based modeling in cognitive research could address the limitations of current methodologies. First, because experiments with human subjects cannot be indefinitely conducted, replicating findings or exploring a wide parameter space is costly and oftentimes impossible. In the case of research in schools, once the classroom data are collected, the researchers can revisit the videotapes and transcriptions; however, they can never relive the situations. Second, as the field moves toward theories that conceptualize learning as a dynamic and adaptive phenomenon, the traditional medium of scientific discourse—static linear text—becomes limited in its capacity to express these theories. Both of these flaws could be addressed with a set of dynamic, adaptive computer models of learning. Third, tools such as brain imaging cannot yet offer the speed and resolution required to evaluate complex learning processes at a neuronal level, so such models are still far from being applicable to real classrooms. Lastly, ethnographic or micro-genetic methods still cannot offer a 'runnable,' systemic, task-independent account of human learning.

The ultimate goal of using agent-based simulation to explore human learning is to enable researchers to generalize and play "what-if" scenarios using in-depth interviews and ethnographic data and to help them investigate internal cognitive structures by observing external behaviors.

This work builds on previous seminal contributions to the field, in which theoretical models of cognition were implemented by using computer programs to attempt to predict human reasoning (Newell & Simon, 1972) in tasks such as shape classification (Hummel & Biederman, 1992), language acquisition (Goldman & Varma, 1995), memory (Anderson, Bothell, Lebiere, & Matessa, 1998); and other more general-purpose models (Anderson, 1983; Anderson & Bellezza, 1993; Anderson & Lebiere, 1998). My design, however, differs from extant approaches in two ways: (1) *Grain Size: Selecting a unit of analysis toward bridging the micro and macro perspective on learning.* Theories which slice human learning into diminutive pieces, when reintegrated into the larger context of classroom learning, could not account for any meaningful macro-cognitive phenomena, and (2) *Accessibility: Democratizing modeling-based research.* Most computational theories of the mind are so mathematically complex that only specialized researchers can examine and critique them; the intricacies

and jargon of these theoretical models render them incomprehensible for teachers, educators, and policymakers. Conversely, the computer language that I have used for modeling, NetLogo (Wilensky, 1999), has been developed for non-programmers so that users could not only run models but also modify their rules and compare scenarios.

My theoretical inspiration comes from the work of Minsky (1986), and Collins (1978). My computer-based models of human learning postulate non-intelligent cognitive entities with simple rules from which intelligent behavior emerges, or simple individual classroom behaviors that result in complex group-level patterns. To generate and validate such models, ABM tools enable researchers to initially feed a computer model with data from real-world experiments, such as classroom observations or clinical interviews and to subsequently simulate hypothesized scenarios in a safe virtual environment. Researchers from diverse disciplines (and with little, if any, programming background) can embody and articulate their theoretical models in a shared medium with shared nomenclature and shareable/replicable data, thus facilitating interdisciplinary discourse and critique.

However, the work described in this paper is not attempting to *reproduce reality*, which is oftentimes understood to be the goal of a computer model. My objective is to instantiate possible theories of learning in the agent-based form and use the data to qualitatively validate the models, with the goal of advancing theory. However, unlike classical cognitive models, this category of ABM models needs to be much more stylized and simple, as this paper will describe.

Personal epistemologies & resources

Traditional research on personal epistemologies (Hofer & Pintrich, 2002) has considered them as stable beliefs. However, evidence of variability in student epistemologies suggests the need for more complex models (Hammer & Elby, 2002). The activation of the students' different epistemological resources might depend on context, as shown by Rosenberg, Hammer, & Phelan (2006). In other words, students might instantiate different epistemologies as they perceive contextual cues about the most efficient approach in a given situation. In the Rosenberg et al. case study, a brief epistemological intervention by an 8th-grade science teacher led to the students abruptly shifting from one epistemological mode to another. The narrative tells the story of a group of students who were given the task of explaining the rock cycle. For the first few minutes, before the teacher's intervention, they fail to engage in any productive work or to construct a coherent explanation of the rock cycle. Students employ a 'brute force' approach by quickly trying out several short explanations without evaluating if the elements of their explanations make sense together. They generate fragmented descriptions, which do not survive simple logical inference. Rosenberg et al. state that the reason is epistemological and that "They are treating knowledge as comprised of isolated, simple pieces of

information expressed with specific vocabulary and provided by authority." (Rosenberg, et al., 2006, pp. 270)

The authors provide three pieces of evidence for this hypothesis: (1) the students organize their efforts around retrieving information from worksheets, (2) they focus on terminology, and (3) they combine information and construct sentences to present a formal ordering rather than a causal sequence. The narrative goes on to describe how the teacher, realizing the ongoing failure, stops the activity and tells the students: "So, I want to start with what you know, not with what the paper says."

Abruptly, the students change their approach toward engaging in the activity. They immediately start to focus on the elements of the rock cycle that they understand, and they rebuild the story from there. Within minutes, one of the students comes up with a rather complete explanation:

"OK, the volcano erupts, and lava comes out. Lava cools and makes igneous rock. Rain and wind cause small pieces of rock to break off. Sediments form, and rain and wind carry it away, and rain and wind slow down and deposit sediments and this happens over and over again to form layers." (Rosenberg, et al., 2006, pp. 274)

It is impressive how the students, focusing on a single element of the story ("Lava comes out"), correctly connect all of the other pieces of the explanation. Although the "lava comes out" piece was the first to be mentioned, they realized that for lava to come out, the volcano has to erupt; similarly, if the lava comes out and is hot, it has to cool down. For the students to generate a coherent explanation, it was crucial for them to concatenate information while making sense of the connection rules, and they resorted to worksheets fewer times than in the previous activity.

In this paper, my goal is to employ ABM to help model what occurred during those 15 minutes and to answer two research questions concerning the abrupt epistemological shift observed: (1) what caused the two modes to generate very diverse student performance? and (2) how could a brief intervention cause such dramatic change? I built a model that simulates the construction of declarative knowledge in terms of two basic cognitive operations: retrieving information from external/internal sources and applying concatenation rules to join *content pieces*. I expect to answer the research questions by exploring the parameter space of the model for number, type, and efficiency of retrievers and connectors; this might result in emergent behaviors similar to those observed by Rosenberg et al. I warn the reader that the goal is to match an overall reference pattern based on simple theoretical assumptions about learning. The nature of ABM is such that this simplicity is required to generate a manageable parameter space.

The Agent-Based Model

In the model, the world outside of the mind is composed of various disconnected *content pieces*, represented as green agents. A piece could be a simple statement, such as "Lava comes out of volcanoes," "Lava shoots up," or "Water erodes rocks." These pieces are retrieved by special agents,

called *retrievers* – represented as red agents – and accommodated into the simulated mind. Here, they interact with pre-existing structures until they connect to one of them, making use of a third type of agent, the *connectors*. These pre-existing structures could form an emergent, dynamic network with “hub ideas” (highly connected ideas) and peripheral ideas.

Therefore, the model consists of three independent elements: connectors, retrievers, and content pieces (line-shaped white agents, red agents, and green agents, respectively). Content pieces could be gathered by retrievers to simulate external information being exposed to the mind. Furthermore, the retrievers’ movement simulates the exposure and the searching for different content pieces over time. The recently retrieved information could be connected to preexisting mental structures to form a new knowledge “assembly.” This knowledge-construction process is represented by content pieces being connected to each other via connector agents. Accordingly, the students’ explanations are the ad hoc result of *pieces* of content collected by *retrievers* outside of the mind and assembled by *connectors* inside the mind. For the sake of the model’s simplicity, the internal and external processes occur in the same environment and have not been distinguished visually.

In the simulated world, the content pieces can have a different *stickiness* to the retrievers, thus the model can differently evaluate content from various sources. Content given from authority can have a different effect than previous knowledge in the virtual mind. In addition, retriever agents have been defined hypothesizing that content cannot simply enter the mind as raw information. Information needs to be retrieved and subsequently connected by internal knowledge structures (Piaget, 1952). Furthermore, retrieved content cannot be accessed until it is evaluated and internalized by connectors (Figure 1).

In the real classroom situation, students would assemble a textual explanation, such as: [*the volcano erupts + lava comes out + lava cools + lava makes igneous rock*]. In this model, textual explanation pieces are replaced with numbers, and a correct explanation is simply an ascending sequence of numbers (1, 2, 3, etc.), which reflects the nature of the task that students had at hand (sequencing information). I am aware of the limitations of the chosen approach, but this design principle was appropriate given the research questions and the target results.

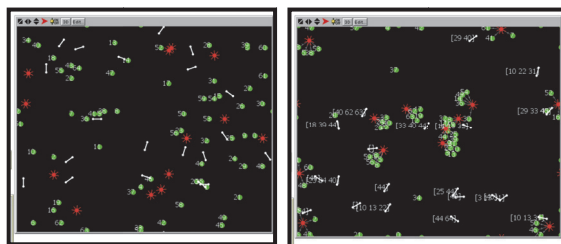


Figure 1. The model’s initial state (left) and after some steps (right), with ‘clusters’ of content form around retrievers.

Experiments and Data

Contrary to most cognitive modeling software, this model does not attempt to simulate human thinking in its immense range of complexities and detail. I selected, conversely, the particular features of learning processes that will possibly enable me to pair the model’s data and the observations of Rosenberg et al. Because I am only modeling the agent’s ability to construct correct connections between pieces, I am ultimately investigating the computational cost and accuracy in building probabilistic cognitive structures. “Success” in the model is defined by the correct assemblage of a *sentence* without errors (i.e., with all numbers correctly places in an ascending order). I measure the *time to completion* of the sentences as well as the *error rate* in building them. The final performance measure is the ratio between the *average time to completion* and the *average error rate*, which I call the *cost of accuracy*.

First experiment: Effect of Retrieval Skills

The first round of model runs compares retrievers with different performances or *stickinesses*. When retrievers encounter content pieces, they stick to them and carry them to the connectors. Low-performing retrievers, however, might collide with a piece but fail at sticking to it. The net effect of a low-performing retriever is to bring fewer pieces to the connectors per unit of time. The performance property of retrievers is loosely analogous to the students’ short-term memory skills or the amount of sheer information they can gather within the environment. One conclusion from this simple experiment is that retrievers have a small impact on overall task performance. Dropping retriever success rates from 90% to 20% results in a modest 16% increase in time for the completion of the task; therefore, within the initial parameters and assumptions of the model, retrievers appear not to be the *controlling phase* of the process. This is a key *qualitative* result of the model: good information retrieval skill does not cause abrupt gains in learning. It is important here to restate that my goal is not to present a calibrated computer model that would emulate precise response times of the human brain. Rather, I am advancing the understanding of a cognitive task by using computer modeling to explicitize certain assumptions and refine our understanding of the problem. In this case, for example, the simple experiment drew my attention to the fact that retrieving information should be relatively faster than connecting information to existing knowledge structures, and therefore, an improvement in the quantity of available information or the speed of retrieval would have a relatively low impact on performance. The data from Rosenberg et al. qualitatively corroborate this hypothesis: during the first narrative, with books and worksheets readily accessible but with weak *connecting skills*, students were unable to weave a coherent explanation. From the narrative, it is clear that if students were given more time or more informational resources to complete the task, the impact in task performance would *not* have been significant. In other words, my model replicates one of the classroom

observations of Rosenberg et al.: the controlling phase of the students' cognitive work was not *information retrieval*, and the cause of students' failure in explaining the rock cycle was not due to a lack of information, a lack of time to retrieve the correct information, or weak memorizing skills.

Second experiment: Effect of Connecting Skills

The goal of the second experiment was to investigate the influence of the connectors' performance on overall task completion time and accuracy. Connectors, in the model, represent more elaborate cognitive agents, which can evaluate different pieces of information and link them based on a simple rule (build ascending sequences of numbers.) Connectors can make "mistakes," for example, wrongly appending the number 41 to the otherwise correct ascending sequence [3, 45, 67]. The probability of such mistakes is controlled by an internal variable within each connector agent (*connector-strength*). The following plots show the impact on time to completion, and accuracy, for different values of *connector strength* (from 10% to 95% of probability of a wrong connection).

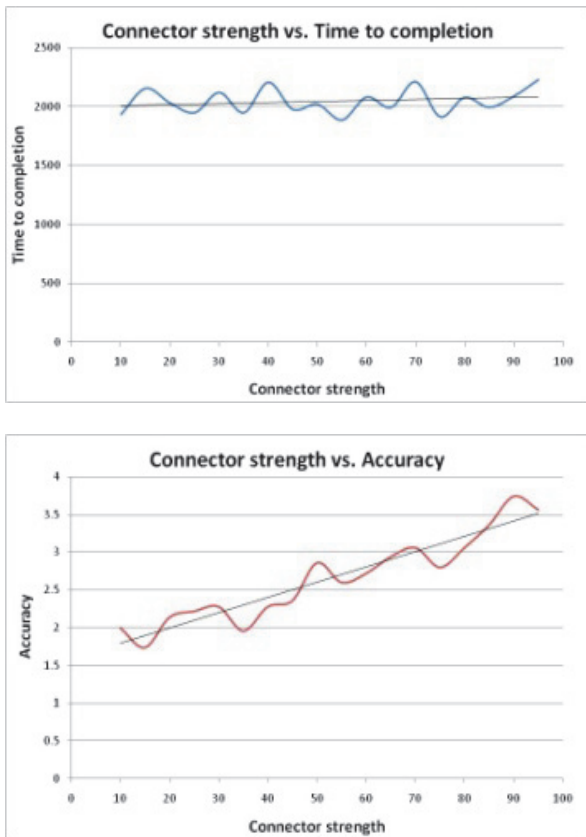


Figure 2. Connector strength, completion, and accuracy¹.

¹ Note that each data point is an average of 50 model runs. Given the qualitative interpretation of the results and the limited space, I considered that error bars and more detailed statistics would not be informative to the research questions and would add unnecessary information.

At first sight, the *Connector strength vs. Time to Completion* plot (Figure 2, top) suggests that "Connector strength" has no impact on overall performance. However, even though the time to complete the task remains roughly the same, accuracy increases significantly (Figure 2, bottom). Combining the two plots (not shown) suggests a reasonable linear fit between computational cost of accuracy and connector strength, which suggests that increasing the *skill* of the connectors has a much greater impact on overall task performance than increasing the retrievers' skill (see previous experiment). This result confirms a second expected behavior, which is also qualitatively in agreement with the data from Rosenberg et al. When the students were told to "start from what they already knew" and evaluate the connections among the different phases of the rock cycle using previous knowledge (i.e., 'if lava is hot, it must cool down'), their performance increased significantly.

This second experiment hints that *connecting skills* are more significant for task performance than *retrieving skills*. However, the cost of training skilled connectors is still unknown; hence, comparing "unskilled but fast" and "skilled but slow" is crucial, which I attempt to illuminate in the next section.

Third experiment: Explanation Complexity

The third experiment was aimed at discovering the impact of explanation complexity on performance. In this model, the complexity of the explanations is represented by the 'sentence-size,' which is the target number of *knowledge pieces* that the connectors need to put together (e.g., sentence size 3 would be "volcano erupts" + "lava comes out" + "lava cools"). The following plot shows a comparison between sentence sizes 2 and 3, for different values of connector strength.

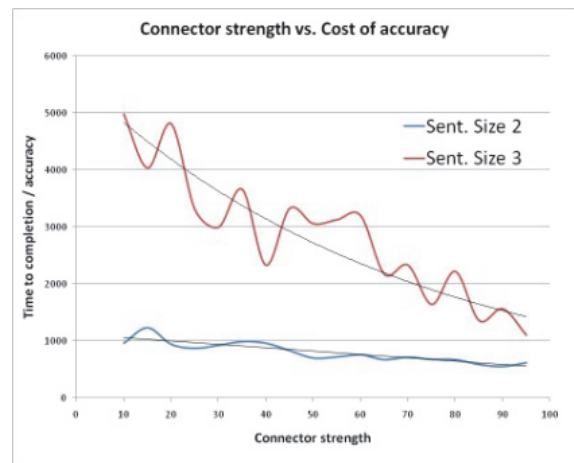


Figure 3. Time to completion divided by the correctness (y axis) and the connector accuracy (x axis).

One result of this experiment is that while the impact of increasing values of connector strength is linear for sentence size 2, it is roughly exponential for sentence size 3. This

finding suggests that for assembling ‘simple’ content the gain that students obtain from improved connecting skills is much lower than when they are struggling with complex knowledge.

Again, this finding seems to fit with Rosenberg et al.’s narrative. Even in the first moment of the narrative, when students are trying assemble explanations based on worksheets using a brute-force approach (quickly trying many different pairs), they were able to assemble a number of “sentence-size 2” explanations, such as [igneous rock forms] + [weathering occurs]. However, in that first part of the narrative, the students were never able to form “sentence size 3” explanations, which would require extra steps: connecting that initial pair of pieces to a third piece and evaluating all possible pieces for their fit. In the second part of the narrative, after only a few minutes, by trying to expand their explanation *making sense* of the connections between pieces (and not using the brute force approach), students formed a sentence size 4 explanation, and a few minutes later they formed a sentence size 10 explanation:

“Bethany: Listen up! OK, the volcano erupts [1] and lava comes out [2]. Lava cools [3] and makes igneous rock [4]. Rain and wind cause small pieces of rock to break off [5]. Sediments form [6], and rain and wind carry it away [7], and rain and wind slow down and deposit sediments [8] and this happens over and over again to form layers [9]. OK, so water is added to this [10]...” Rosenberg, Hammer, & Phelan (2006), pp. 274

To further investigate the role of the increase in sentence sizes to the overall cost of accuracy, I ran the model for sentence size 4 as well. The results, comparing sizes 2, 3 and 4, are in the following two plots:




Figure 4. Time to completion divided by the correctness (y axis) and the connector accuracy (x axis).

To understand Figure 4, it is important to comprehend the intuition behind the results. Essentially, I am comparing a “brute force” versus a “smart” approach for assembling sequences of different sizes. For sentence size 4 (SS4), with low values of connector strength (CS), it is virtually

impossible to assemble a correct explanation (see the very high values of the top curve). For CS 10%, increasing SS from 2 to 4, the accuracy drops by a factor of 100. Increasing SS from 2 to 3, the accuracy drops 5 times. Figure 4, therefore, shows that increasing sentence sizes has a dramatic impact on performance. The important finding here is that this differently impacts “long” and “short” explanations. For SS 2, brute force assemblage is not so costly and works relatively well, so there would be no benefit for developing connecting skills. However, for SS 3 and 4, this ‘brute force’ (low CS) assemblage breaks down.

The events in Rosenberg et al. narrative tell a similar story. In the first half of the class, when students were using brute force methods instead of their own connecting skills, they could not go much further than assembling simple, “SS 2,” explanations. When they activated their ‘connectors,’ prompted by the teacher’s intervention, they switched from a brute force to a “sense-making” mode, in which most of their energy was spent connecting pieces instead of retrieving and randomly connecting them. That shift enabled them to assemble seamlessly explanations of SS = 10.

Conclusion, limitations, and implications

Throughout this paper, I tried to pair computer model data with real classroom data. In the three experiments, I searched for instances that would resemble what Rosenberg et al. described in their classroom observations. The model seems to validate key elements of those observations:

1) The students’ failure in the first half of the narrative was epistemological (i.e., resulting from a particular approach toward learning) and not due to a lack of memorization or information retrieving skills (the first experiment).

2) The fundamental mathematical basis of the model, from which all other behaviors emerge, is that *brute-force methods are efficient for short sequences, but for long sequences, as the combinatorial space greatly increases, their performance drops accordingly*. In the high connector strength mode, the size of the sentence has a much lesser impact because of the evaluative rule of the connector: any connection will take the exact same computational time for any sentence size. This seems to be the case in the classroom, where the students could assemble long explanations quickly once they were in a ‘high connector strength’ mode.

3) In this simulated environment, I was able to verify that for learning intricate content (here, I equate that to assembling long explanations), there is a significant non-linear payoff from investing in sense-making skills (connector strength) as opposed to memorizing skills (retrieving speed). For simple content (involving the connection of two content pieces), however, *sheer memorization can even outperform sense-making skills*. The data show that the payoff of improved connector strength only manifests itself after CS 80% (Figure 2, 3, 4).

4) Abrupt, non-linear shifts in student understanding are indeed possible, even within very short periods of time, by

activating different cognitive resources and different epistemological modes.

Limitations and implication for design

The classroom data used in this paper was chosen because it described a relatively uniform macroscopic behavior that clearly derived from a change in simple, local rules. I acknowledge that many other typical classroom interventions might not exhibit such a uniform behavior. The goal of this model and paper, however, was not to match a computer model to a precise mechanism in the brain. Rather, my goal was to produce the “simplest possible” model that would exhibit the observed behaviors and generate further insight into the research questions. In that sense, this was a theoretical exercise made possible by formalizing the problem as agent rules. Therefore, given the assumptions of the model, I suggest that some possibly overlooked elements in classroom implementation might be more important than one would suspect: (1) the radically different payoffs for improving the *speed of retrieval* versus *sensemaking*, and the determination of which is the controlling phase in the learning process in different scenarios, (2) the non-linear impact of *sentence-sizes* on performance and accuracy, (3) the unexpected success of “brute force” methods for small sentences.

Given the limited space, it is impossible to go into detail about all possible implications, but one implication is very significant. *In earlier grades, exposed to simpler content, students might learn that brute-force methods ‘work.’ In later grades, they might insist on using this method, which would break down because the content is more complex.*

The computational task is of course an approximation of a real classroom task, and as with any model, it can only capture a portion of the real-world complexity. However, my goal here was to demonstrate the potential of agent-based models as a powerful and useful formalism for cognitive theory. This work could potentially have implications for the practice of curricular designers, teachers, and policy makers by offering researchers accessible, transparent tools to simulate, model and test hypotheses about human cognition in social contexts and to pair model data with real classroom data.

Acknowledgements

Thanks to Uri Wilensky, Dor Abrahamson and Shima Salehi for their input on previous versions of this work, and David Hammer and colleagues for the classroom data.

References

- Abrahamson, D., & Wilensky, U. (2005). *Piaget? Vygotsky? I'm Game: Agent-Based Modeling for Psychology Research*. Paper presented at the Annual Meeting of the Jean Piaget Society, Vancouver, Canada.
- Abrahamson, D., Blikstein, P., & Wilensky, U. (2007). Classroom model, model classroom: Computer-supported methodology for investigating collaborative-learning pedagogy. *Proceedings of the Computer Supported Collaborative Learning (CSCL) Conference*, pp. 46 - 55.
- Anderson, J. R. (1983). *The Architecture of Cognition*. Cambridge, MA: Harvard University Press.
- Anderson, J. R., & Bellezza, F. S. (1993). *Rules of the mind*. Hillsdale, N.J.: L. Erlbaum Associates.
- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, N.J.: Lawrence Erlbaum Associates.
- Anderson, J. R., Bothell, D., Lebiere, C., & Matessa, M. (1998). An integrated theory of list memory. *Journal of Memory and Language*(38), 341-380.
- Axelrod, R. M. (1997). *The complexity of cooperation: Agent-based models of competition and collaboration*. Princeton, NJ: Princeton University Press.
- Blikstein, P., Abrahamson, D., & Wilensky, U. (2006). *Minsky, Mind, and Models: Juxtaposing Agent-Based Computer Simulations and Clinical-Interview Data as a Methodology for Investigating Cognitive-Developmental Theory*. Paper presented at the Jean Piaget Society Annual Meeting, Baltimore, USA.
- Blikstein, P., Abrahamson, D., & Wilensky, U. (2009). *Toward a Framework for Cognitive Research Using Agent-Based Modeling and Complexity Sciences*. Paper presented at the American Educational Research Association Annual Meeting, San Diego, CA, USA.
- Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm Intelligence: from natural to artificial systems*. London: Oxford University Press.
- Collins, A. M. (1978). *Fragments of a theory of human plausible reasoning*. Urbana-Champaign, Illinois: ACL.
- Epstein, J., & Axtell, R. (1996). *Growing artificial societies: Social science from the bottom up*. Washington: BIP.
- Goldman, S. R., & Varma, S. (1995). CAPing the construction-integration model of discourse comprehension. In *Discourse Comprehension: Essays in Honor of Walter Kintsch*. Erlbaum.
- Hammer, D., & Elby, A. (2002). On the form of a personal epistemology. In B.K. Hofer & P. R. Pintrich (Eds.), *Personal epistemology: The psychology of beliefs about knowledge and knowing* (pp. 169–190). Mahwah, NJ: Lawrence Erlbaum.
- Hofer, B. K., & Pintrich, P. R. (Eds.). (2002). *Personal epistemology: the psychology of beliefs about knowledge and knowing*. Mahwah, NJ: Lawrence Erlbaum.
- Holland, J. (1995). *Hidden order: How adaptation builds complexity*. Reading, MA: Helix Books/ Addison-Wesley.
- Hummel, J. E., & Biederman, I. (1992, Jul 1992). Dynamic binding in a neural network for shape recognition. *Psychological Review*, 3(99), 480-517.
- Minsky, M. L. (1986). *The society of mind*. New York, N.Y.
- Newell, A., & Simon, H. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Piaget, J. (1952). *The Origins of Intelligence in Children*. New York, NY: International University Press.
- Rosenberg, S., Hammer, D., & Phelan, J. (2006) Multiple Epistemological Coherences in an Eighth-Grade Discussion of the Rock Cycle. *J. of the Learning Sciences*, 15(2), pp. 261-292.
- 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(1), 87-104.
- Wilensky, U. (1999). NetLogo. Evanston, IL: Center for Connected Learning and Computer-Based Modeling. <http://ccl.northwestern.edu/netlogo>.
- Wilensky, U., & Reisman, K. (2006). Thinking like a Wolf, a Sheep or a Firefly: Learning Biology through Constructing and Testing Computational Theories. *Cognition & Instruction*, 24(2), pp. 171-209.