Towards a Framework for Cognitive Research using Agent-Based Modeling and Complexity Sciences: two case studies.

Paulo Blikstein
Stanford University
paulob@stanford.edu

Dor Abrahamson
Berkeley University
dor@berkeley.edu

Uri Wilensky
Northwestern University
uri@northwestern.edu

INTRODUCTION

Complexity sciences and agent-based modeling 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, or the food-gathering behavior of insects (Bonabeau, 1999; Wilensky & Reisman, 2006). Such phenomena, in which the elements within the system (molecules, or ants) have multiple behaviors and a large number of interaction patterns, have been termed *complex* and are collectively studied in a relatively young interdisciplinary field called *complex systems* or *complexity studies* (Holland, 1995). Typical of complex phenomena is that the cumulative ('aggregate') patterns or behaviors at the macro level are not premeditated or directly actuated by any of the "lower-level" micro elements. For example, flocking birds do not intend to construct an arrow-shaped structure (Figure 1), or molecules in a gas are not aware of the Maxwell-Boltzmann distribution. Rather, each element ("agent") follows its local rules, and the overall pattern arises as epiphenomenal to these multiple local behaviors—the overall pattern *emerges*. In the mid-nineties, researchers started to realize that agent-based modeling could have a significant impact in education (Resnick & Wilensky, 1993; Wilensky & Resnick, 1995). For instance, to study the behavior of a chemical reaction, the student would observe and articulate *only* at the behavior of individual molecules — the chemical reaction is construed as emerging from the myriad interactions of these molecular "agents." Once the modeler assigns agents their local, "micro" rules, the model can be set into motion and the modeler can watch the overall patterns that emerge.

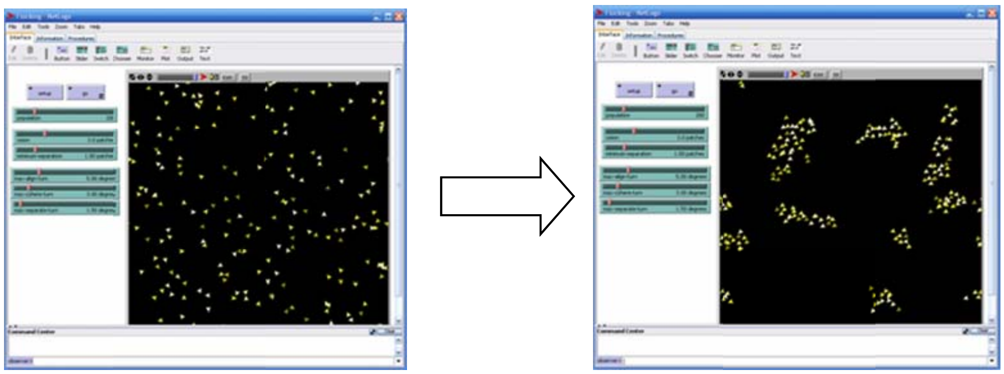

Figure 1: An agent-based model of the flocking behavior of birds.

Whereas initially complex-systems methods and perspectives arose from the natural sciences, complexity, emergence, and micro and macro levels of description of phenomena are all highly relevant to research in the social sciences. Indeed, the recent decades have seen a surge in social-science studies employing ABM (Epstein & Axtell, 1996; Diermeier, 2000; Axelrod, 1997). Behaviors of social groups or entire economies are modeled by extracting micro-level behavioral and interactional rules from more or less simplified human agents. Since the models are trying to uncover and explain just a few aspects of human behavior, these models avoid unmanageable levels of complexity and still render useful results – either as 'hard' quantitative data, or as computational thought experiments which can feed further cycles of theoretical development. Our main theoretical inspiration comes from the work of Minsky, Papert and Collins (Collins, 1978; Minsky, 1986), in which our computer-based simulations of human learning postulate non-intelligent cognitive entities with simple rules from whence emerges intelligent behavior. These software tools enable researchers to initially feed a computer model with data from real-world experiments, such as classroom observations or clinical interviews, and subsequently simulate hypothesized scenarios in the 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.

We have argued (Abrahamson & Wilensky, 2005, Blikstein, Abrahamson & Wilensky, 2006) that ABM has potential to contribute to the advancement of the learning sciences in multiple ways, as we have illustrated in previous work and in this paper:

(a) **explicitizing**—ABM computational environments demand an exacting level of clarity and specificity in expressing a theoretical model and provide the tools, structures, and standard practices to achieve this high level;

(b) **dynamics**—the computational power of ABM enables the researcher to mobilize an otherwise static list of conjectured behaviors and witness any group-level patterns that may enfold through multiple interactions between the agents who implement these conjectured behaviors;

(c) **emergence**—investigate intelligence as a collection of emergent, decentralized behaviors;

(d) **intra/inter-disciplinary collaboration**—the lingua franca of ABM enables researchers who otherwise use different frameworks, terminology, and methodologies to understand and critique each others' theory and even challenge or improve the theory by modifying and/or extending the computational procedures that underlie the model.

(e) **making models accessible—**various authors established the importance of practitioners' mental models of the learning process itself as determinant for their classroom action (Strauss, 1993; Strauss & Shilony, 1994). Therefore, using computer models to conduct research in education and make those models approachable and accessible to teachers could influence and transform their everyday work.

Additionally, ABM could address long-standing limitations of current methodological paradigms. First, experiments with human subjects cannot be indefinitely re-run, so replicating findings or exploring a wide parameter space are costly and oftentimes impossible tasks. Once the classroom data is collected, at most researchers can revisit the videotapes and transcriptions, but never re-live the situations. Second, as we move towards theories that conceptualize learning as a dynamic and adaptive phenomenon, the traditional media of scientific discourse—static linear text—becomes limited in its capacity to express these theories (Abrahamson & Wilensky, 2005; Blikstein, Abrahamson, & Wilensky, 2006). Thirdly, tools such as fMRIs cannot yet offer the speed and resolution needed to evaluate any complex learning process close to what we would find in a classroom. Lastly, ethnographic or micro-genetic methods oftentimes cannot offer a solid, "runnable", generalizable, task-independent account on how humans learn.

The ultimate goal of using agent-based simulation to explore human learning is to enable researchers to generalize and play "what-if" scenarios departing from in-depth interviews and ethnographic data, as well as investigate internal cognitive structures departing from external, observed behaviors, in other words, ABM could bridge quantitative and qualitative methods in a unique way. The two experimental obstacles mentioned above (the limitations of tools such as fMRI and qualitative methods), as we will explain throughout this paper, could be overcome by employing a variable 'grain size' for delimitating the cognitive tasks, together with simple interaction rules, within a coherent theoretical and empirical framework. Over the last years, indeed, we have jumpstarted this research agenda in several topics and fields:

(a) Theories of cognitive development, namely, the piagetian vs. vygotskyan perspectives (Abrahamson & Wilensky, 2005).

(b) An agent-based explanatory model for a classical piagetian task (the conservation experiment), based on Minsky and Papert's model (Minsky, 1986), and paired with data (bifocal modeling, Blikstein & Wilensky, 2006b) from interviews (Blikstein, Abrahamson & Wilensky, 2006).

(c) Collaboration and group work in classrooms, juxtaposing our simulation with real classroom data (Abrahamson, Blikstein & Wilensky, 2007).

(d) The emergence and feasibility of multiple epistemological resources (Blikstein, & Wilensky, 2006, 2007).

(e) A general theory and ontology for models of cognition based on agent-based modeling and network theory (Blikstein, & Wilensky, 2006, 2007).

Our work builds on previous seminal contributions to field, in which theoretical models of cognition were implemented in the form of computer programs in attempt to predict human reasoning (Newell & Simon, 1972; Rose & Fischer, 1999), in tasks such as shape classifications (Hummel & Biederman, 1992), language acquisition (Goldman & Varma, 1995), and memory (Anderson, Bothell, Lebiere, & Matessa, 1998), and other more general-purpose models (Anderson, 1983; Anderson & Bellezza, 1993; Anderson & Lebiere, 1998; Just & Carpenter, 1992; Polk & Rosenbloom, 1994). Our design, however, differs from extant approaches in two fundamental ways:

1. **Grain Size: Selecting a unit of analysis toward bridging the micro and macro perspective on learning, and quantitative and qualitative perspectives—**Those theories, slicing human learning into diminutive pieces, when reintegrated into the larger context of classroom learning, could not account for any meaningful macro-cognitive phenomena.
2. **Accessibility: Democratizing modeling-based research—**Most computational theories of mind were so mathematically complex that only specialized researchers could discuss them. The intricacy and language of these theoretical models rendered them incomprehensible for teachers, educators, and policymakers. Conversely, the computer language with which we have developed the models, NetLogo (Wilensky, 1999), was built from the ground up for non-programmers, so that users can not only run simulations, but modify their internal

rules and compare scenarios. Our models, too, were carefully conceived to follow established models for learning.

In what follows, we will present and discuss two examples of this work

Case study 1: The “Society of More” Model

Conservation of volume is probably the best known Piagetian experiment. It has been extensively studied and reproduced over the past decades (Piaget, Gruber, & Vonèche, 1977). Minsky & Papert (1986) proposed a computational algorithm to account for children’s responses during this experiment. It is based on their construct of the intelligent mind as an emergent phenomenon, which grows out of the interaction of non-intelligent cognitive agents. Minsky’s theory has been particularly influential for overcoming the ‘homunculus’ paradox: if intelligent behavior is controlled by more primitive intelligent behaviors, we get enmeshed in a recursive explanation which cannot ultimately account for a reasonable theory of the mind. Minsky, therefore, insists on using agents that are essentially non-intelligent and obey simple rules—intelligence, therefore, emerges from these interactions.

The simplicity of Minsky’s model is, actually, its main strength – and a perfect fit for the agent-based modeling paradigm. The first important principle in his model is that agents might conflict. For example, at a given time, a child might have Eat, Play and Sleep as predominant agents. Play could have subagents, such as Play-with-blocks and Play-with-animals. If both of these subagents are equally aroused (in other words, the child is equally attracted to both activities), the upper agent, Play, is paralyzed. Then a second important principle comes into play: non-compromise. The longer an agent stays in conflict, undecided, the weaker it gets compared to its competitors. If the conflict within Play is sustained long enough, its competitors will take control (in this case, Eat or Sleep).

Minsky’s fundamental rule is, thus: “whenever in conflict, a mental entity cannot (or takes longer to) decide”. Although relatively simple, this model, as we will see, is surprisingly powerful and opens up many interesting possibilities for investigation, some of which will be described in the paper. Minsky’s and Papert’s model of Piagetian experiments stresses the importance of structure to cognitive evolution, especially its reorganization (the ‘Papert Principle’). Within the context of the conservation task, younger children would have ‘one-level’ priority-based structures: one aspect would always be more dominant (tall would always take priority over thin and over confined - see Figure 1) and compensation, which requires a two-level structure, is thus inexistent. Minsky suggests that, as some perceptual aspects would be more present in the child's life at a particular age, they would be more prevalent. For example, being more or less "tall" than parents or other children would be a common fact for children since a very early age. On the other hand, being more fat or thin would not be as prevalent.

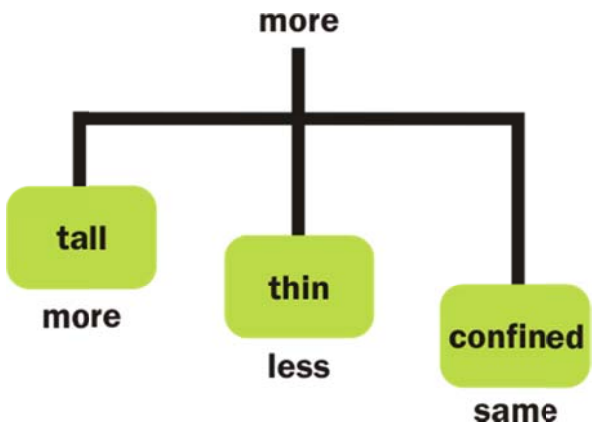

Figure 2: A one-level model for evaluating "who has more"

Later, states Minsky, the child develops a new "administrative" layer that allows for more complex decisions: in Figure 2, for example, if tall and thin are in conflict (i.e., both agents were activated by the child's cognitive apparatus), the "appearance" administrator cannot decide and shuts off, then the history administrator will take over the decision, as it has one one activated agent below it.

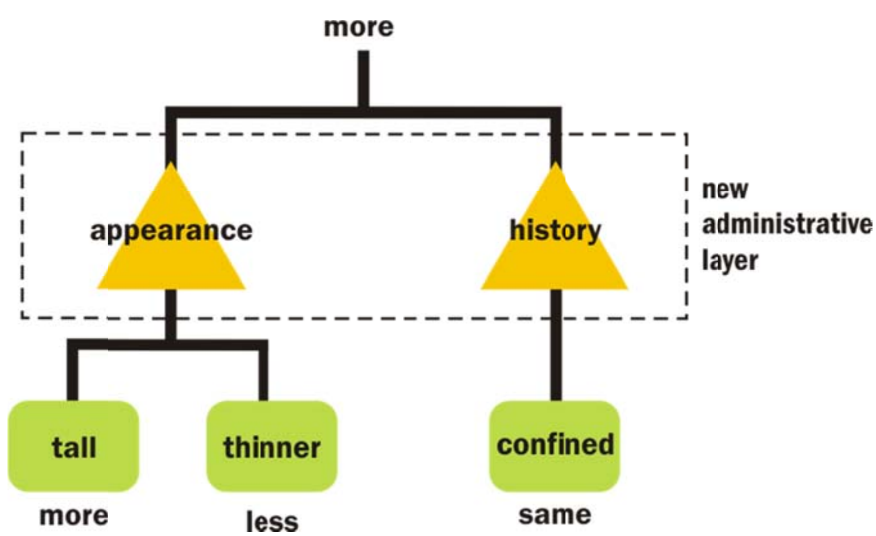

Figure 3: New administrative layer

Our interviews were based on the conventional format of the conservation of volume Piagetian experiment. Two elongated blocks of clay of same shape but different color are laid before the child. One is "the child's," and the other is "the experimenter's." After the child agrees that both are the same size, the experimenter cuts one block in two, lengthwise, and joins the two parts so as to form a block twice as long, then cuts the other block in two, widthwise, to form a block twice as thick as before. The child is asked whether the blocks are still "the same" or whether either person has more than the other. According to the child's response, the interaction then becomes semi-clinical, with the experimenter pursuing the child's reasoning and challenging him/her with further questions.

The approximate time of each interview was 20 minutes. All interviews were videotaped and transcribed, and the data were coded in terms of parameters of the computer simulation. The model (see Figure 4 and Figure 5) was then fed these coded data. We were able to perform different kinds of experiments:

- Playback the interview and the computer model side-by-side, trying to identify behavior patterns and couch them in terms of the simulated model;
- Model validation: investigate whether the child's decision-making process can be predicted by the model. We set the model with the child's initial responses, "run" it through to completion, and try to identify whether the simulated cognitive development matches the processes observed.
- Emergence of structures: investigate if some "society of mind" structures are more prone to emerge than others. For example, would a large number of agents organized into a one-level 'society' be more efficient than a less numerous population of agents organized in two levels?

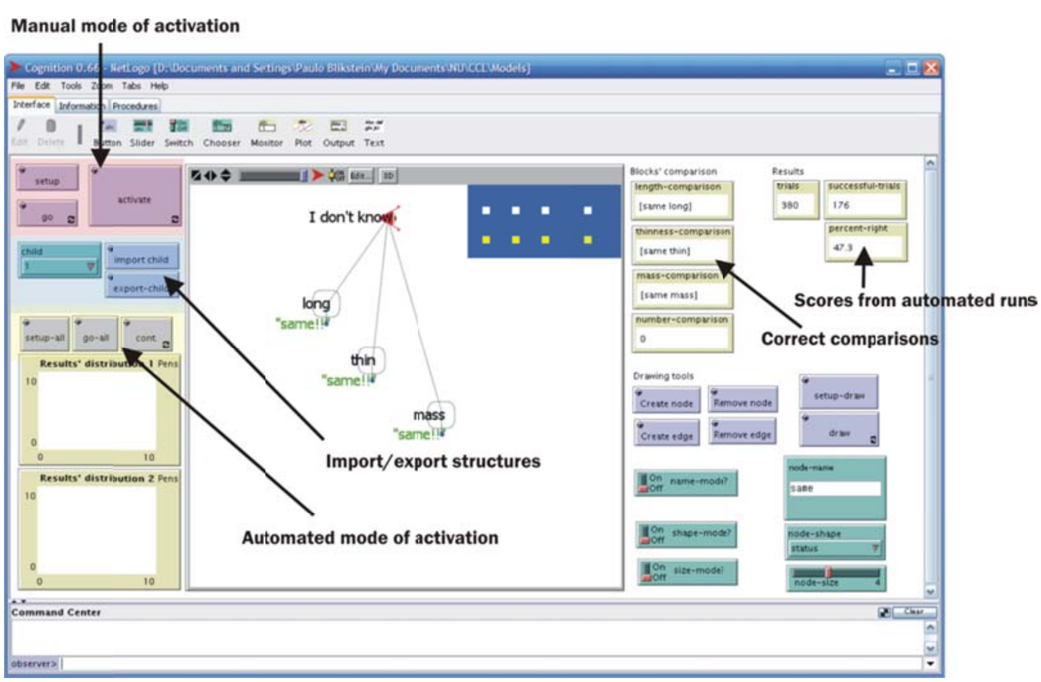

Figure 4: A screenshot of the computer model and its main components.

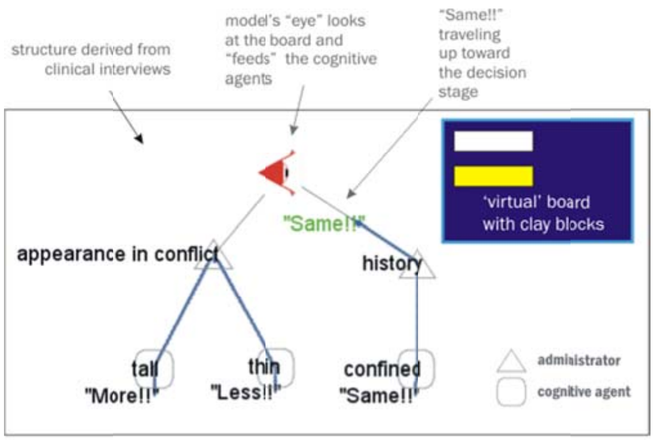

Figure 5: A screenshot of the model's main window. Similar to the child, the computer 'sees' blocks of clay and tries to determine which block is 'more.'

From the first set of studies, we collected and paired data for the side-by-side playback of interviews and the computer models, trying to match outcomes, as seen in the next table:

Table 1: Paired data from the computer model and the clinical interviews

Computer model (screen captures)	**Transcriptions/pictures**
From Child1's **(6yo)** interview, we inferred the simple model below. Cognitive agents presumed to be active are marked with a green outline. Dominance is represented in the model by the vertical distance to top. For this child, whenever **Number** – the cardinal dimension of the stimulus – is contextually salient, it dominates the decision-making process. Also **Tall** appears to dominate **Thin**.	"Because you cut in half, so there **is two pieces**, but... It's **not as fat as that**. This is kind of fat, but this is taller. I have more".

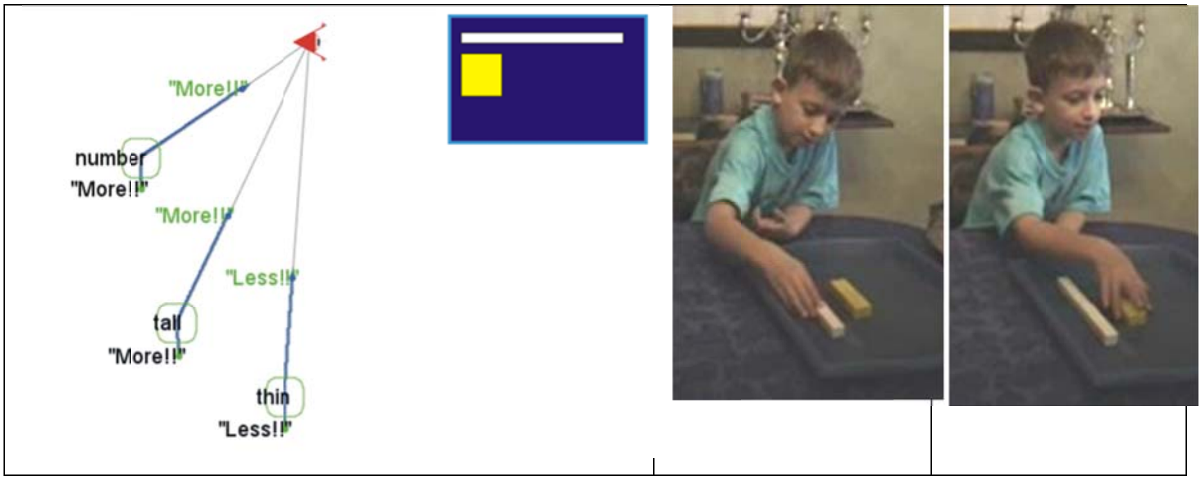

To study the emergence of structures, we ran several possible cognitive structures with varying levels of relaxation (adherence to the initial structure), to evaluate which structures were more sensitive to randomness, closer to the performance of the 'real' child, as seen in Table 2:

Table 2: Different levels of relaxation for different structures

Structure 1 As relaxation increases, the model gets more accurate. The relaxation compensates for the inadequacy of the model to evaluate situations in which 'long' is not determinant of 'more'. With relaxation up to 30%, the model scores as well as the child. After 30%, it gets better, but the gain tapers off – at some point, it is just as good as chance.	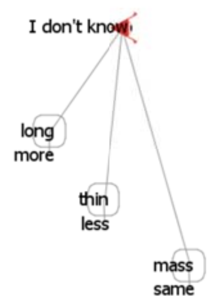	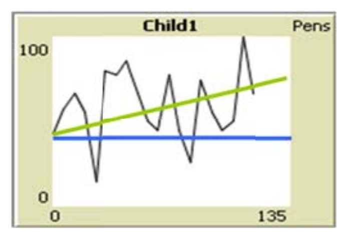
Structure 2 In this model with administrators, accuracy starts at 100%, and decreases with increase randomness. The average score of this model, even with high randomness levels, is far better than the child's.		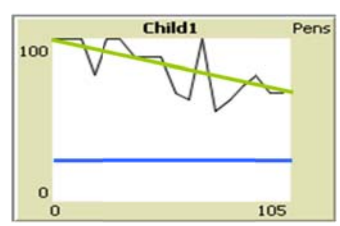

Structure 3

This model is identical to **Model 2**, except that **mass** and **thin** were switched. **Mass** is under **appearance**, and **thin** under **history**. The overall score drops dramatically from 100% to 30%, approaching the child's. This is corroborates Minsky's hypotheses about the importance of having the 'right' agents under each administrator. If we were to 'evolve' a structure using standard GA algorithms, probably **Model 2** would be rapidly selected over **Model 3**.

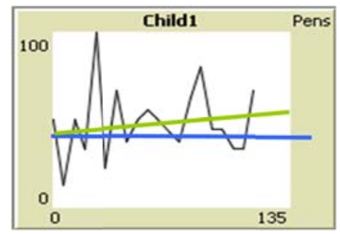

The different effect of relaxation in the models' performance was an important result of this experiment. Simple, one-level models increase their performance with increased relaxation (relaxation can be understood as a controlled introduction of error in the cognitive structure). Complex, specialized, high-accuracy models lower their performance with high levels of relaxation. This result might seem trivial: deforming an accurate structure causes it to perform badly, and introducing high error rate and randomness to a weak structure will benefit it. However, the usefulness of this result is that it can be used as a criterion to evolve cognitive structures. Another consequence of this finding is that it suggests that learning might benefit from relaxation of constraints in different ways, depending on the developmental level, knowledge domain, and age. We could hypothesize that, when children are first learning principles of a knowledge domain, the learning environment should promote "random" connection, wrong moves, unlikely choices. The primitive structure would benefit from those to evolve administrators. Once administrators are in place, perhaps, a more structured environment is beneficial.

These computer models were an useful vehicle both to illustrate the Piagetian theoretical model and to simulate it departing from interview data. Through the lens of agent-based models, new properties of Minsky's model are revealed—namely, the mature, hierarchical structure of the cognitive model is stochastically determined, in the sense that across combinatorial initial conditions, and over sufficient interactions, the same meta-structures ultimately emerge. Collecting and analyzing data from actual (not simulated) interviews is an essential phase in the improvement of the computer simulation of a theoretical model, such as Minsky's model: The data sensitize us to the crucial components and dimensions of the interactions and to the nature of the transformations.

Conventional "paper and pencil" representations of Piagetian structures might miss some of the dynamic factors in play. For example, we were able to identify in several children some 'embryonic' agents, which were present in just part of the interaction. Child 2, for instance, would oscillate between a "re-joinable" and a "conserved-mass" explanation in many interactions. Without a probabilistic approach, we would be obliged to just assume that those children were in a transitional stage. With the computer representation, we could actually calculate the number of times that different embryonic agents are aroused, and estimate the developmental stage of the child. This data could then be fed into the computer model for further confirmation – we could even envision, for future work, simulations which could predict the appearance and evolution of embryonic explanations.

A natural and promising path for this work is to evolve structures automatically. We suggested earlier that the dynamics of this simulation is such that favorable outcomes would be reinforced. As we observed in the experiments, random reconnections of agents do not render random results—structure matters. The mechanism which we demonstrated shows that there is a higher probability for related agents ("long" and "thin") to group together under one particular agent – this is the configuration that delivers the best performance. One can imagine that, along many years of cognitive development in the world, the child will group some sensorial and cognitive experiences into certain categories: i.e., "thin and long belong to appearance", "taken-away and spilled relate to history of the transformation". What Minsky states, and we verified, is that the actual content of such agents less irrelevant than it's placement within the structure, if they are under a closely related agent. Thus, the categorization process itself emergently generates intelligent behavior, without any interference from an external "intelligent" entity. This appears to be an indication that the 'Society of Mind' framework could be used with predictive power in developmental psychology, especially when coupled with clinical interview data.

Case study 2: Manifold epistemological resources

Traditional research on personal epistemologies (Hofer & Pintrich, 2002) has conceptualized them as stable, constant beliefs. However, evidence of variability in student epistemologies suggests the need for more complex models (diSessa, 1993; Hammer & Elby, 2002). The activation of students' different epistemological resources could depend on context, as shown by

Rosenberg, Hammer, & Phelan (2006). In their case study, a brief epistemological intervention by an 8th-grade science teacher led to students' abrupt shift from one epistemological 'mode' to another. Rosenberg *et al.* 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. Their explanations are fragmented, use the wrong vocabulary, and do not survive even simple logical inference. Rosenberg *el at.* 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, Hammer, & Phelan (2006), pp. 270.*

The authors provide three pieces of evidence for this hypothesis: (i) students organize their efforts around retrieving information from worksheets; (ii) they focus on terminology, and (iii) students combine information and construct sentences to present a formal ordering rather than a causal sequence. But the narrative goes on. Realizing the ongoing failure, the teacher stops the activity, and tells students:

> *"So, I want to start with what you know, not with what the paper says."*

Abruptly, students change their ways of engaging in the activity. They immediately start to focus on elements of the Rock Cycle that they understand and rebuild the story from there – in few minutes, one of the students was able to come up with a reasonable 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, Hammer, & Phelan (2006), pp. 274*

Particularly impressing is how students, departing from a single element of the story ("Lava comes out"), could correctly connect all the other pieces of the explanation. Even though 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. Concatenating pieces of information making sense of the connection rules was crucial for students to generate a coherent explanation, resorting even less times to their worksheets than in the previous half of the narrative.

We set out to employ ABM to model what took place during those 15 minutes, answering two research questions concerning the abrupt epistemological shift observed:

1) What caused the two ‘modes’ to generate very diverse student performance?

2) How could a brief intervention effect such dramatic change?

We built a model that simulates the construction of declarative knowledge in terms of two basic cognitive operations: retrieving information from external/internal sources, then applying concatenation rules to join information “pieces” (the *retriever/connector model*, Blikstein & Wilensky, 2006). We expected to answer the two research question aforementioned by exploring a significant part of the combinatorial space of initial conditions of the model, with different values for number, type, and efficiency of retrievers and connectors, which might result in emergent behaviors similar to those observed by Rosenberg *et al.*

In our model (see Figure 6), the world outside the mind is represented as an ‘ocean’ of disconnected content pieces of various kinds. These pieces are retrieved by special agents (“retrievers”) and accommodated into the simulated mind, where they interact with pre-existing structures until they connect to one of them, making use of another special type of cerebral agent (“connectors”). These pre-existing structures form an emergent, dynamic network with “hub ideas” (highly connected ideas) and peripheral ideas. Students’ explanations are the ad hoc result of pieces of content and ideas that were connected on the fly. For example, the patterns of change in the network can account for different kinds of learning, from short-term changes to deep conceptual change. These networks can assume different topologies, from uniform to scale-free networks, depending on the parameters of the simulation.

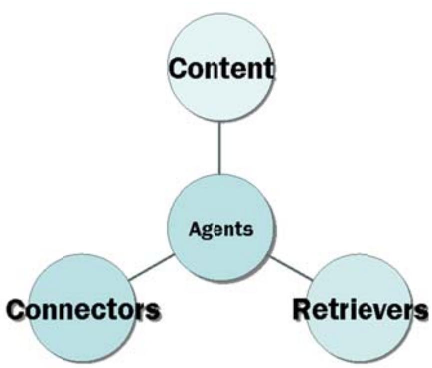

Figure 6: The three types of agents of the model: content pieces, retrievers and connectors

In Figure 7 and Figure 8 general a schema of the model and its components, as well as two screenshots of one particular model run.

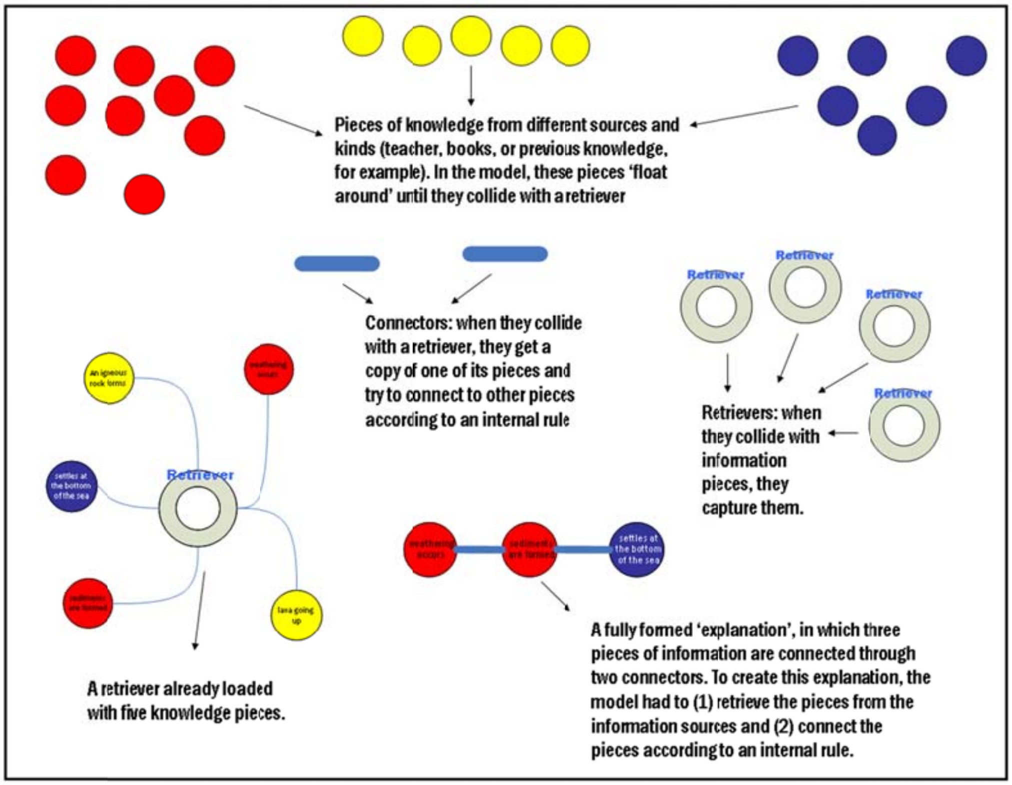

Figure 7. Illustrated explanation of the pieces/retriever/connector computer model.

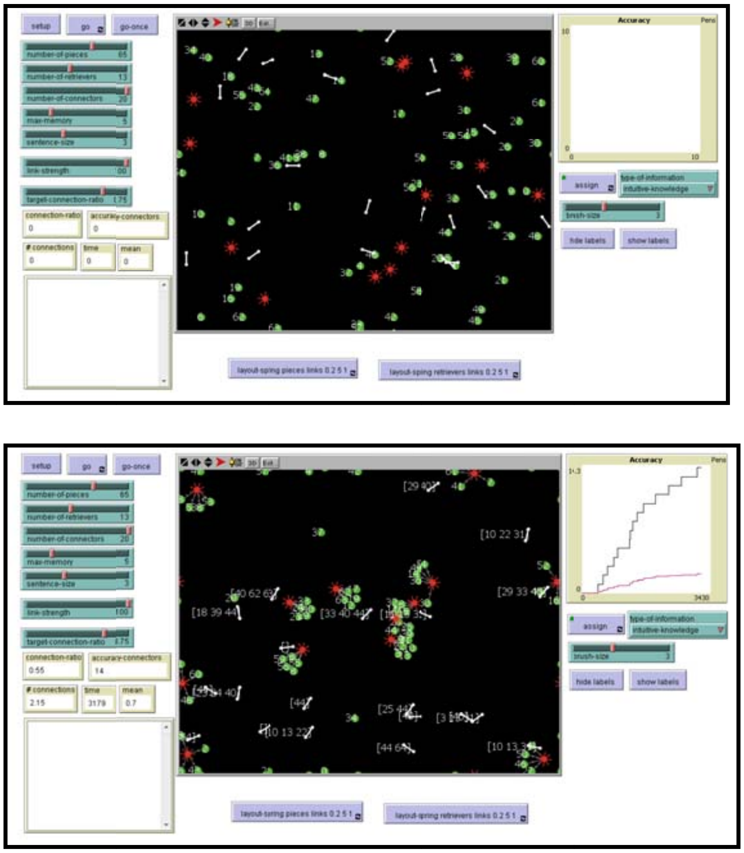

Figure 8. The model's interface at its initial state (top), and after some steps of the simulation (bottom), in which 'clumps' of content pieces around retrievers (represented as hubs) are noticeable.

Experiments

We conducted several experiments with the model (described in detail in Blikstein & Wilensky, 2006). One of them was aimed at finding out the impact in performance of the complexity of the desired explanation. In the model, the complexity of the explanations is represented by the 'sentence-size', which is the target number of knowledge pieces which connectors need to put together. The following examples show explanations with sentence size two, three and four:

Sentence size = 2

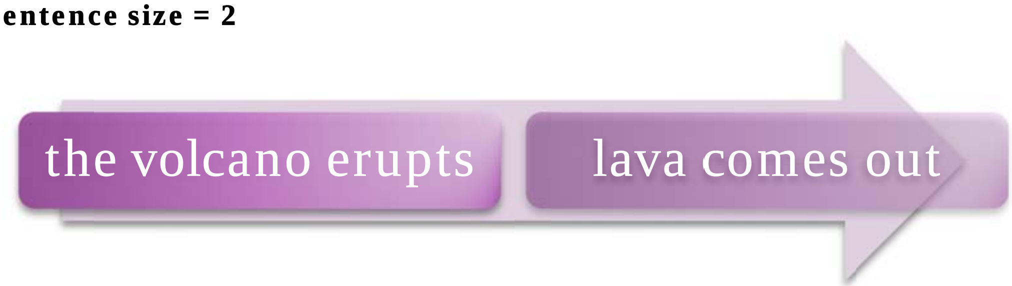

Sentence size = 3

Sentence size = 4

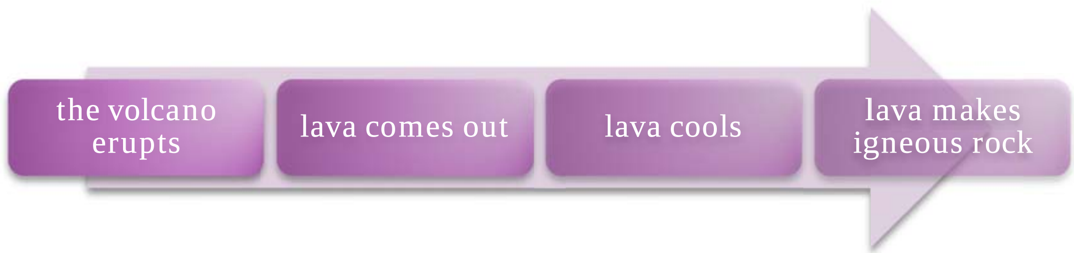

The following plot shows a comparison between sentence sizes 2 and 3, for different values of connector strength (connector strength is the measure of how strict are connectors when evaluating content pieces. High connect strength means that the connectors do now allow wrong pieces to be connected).

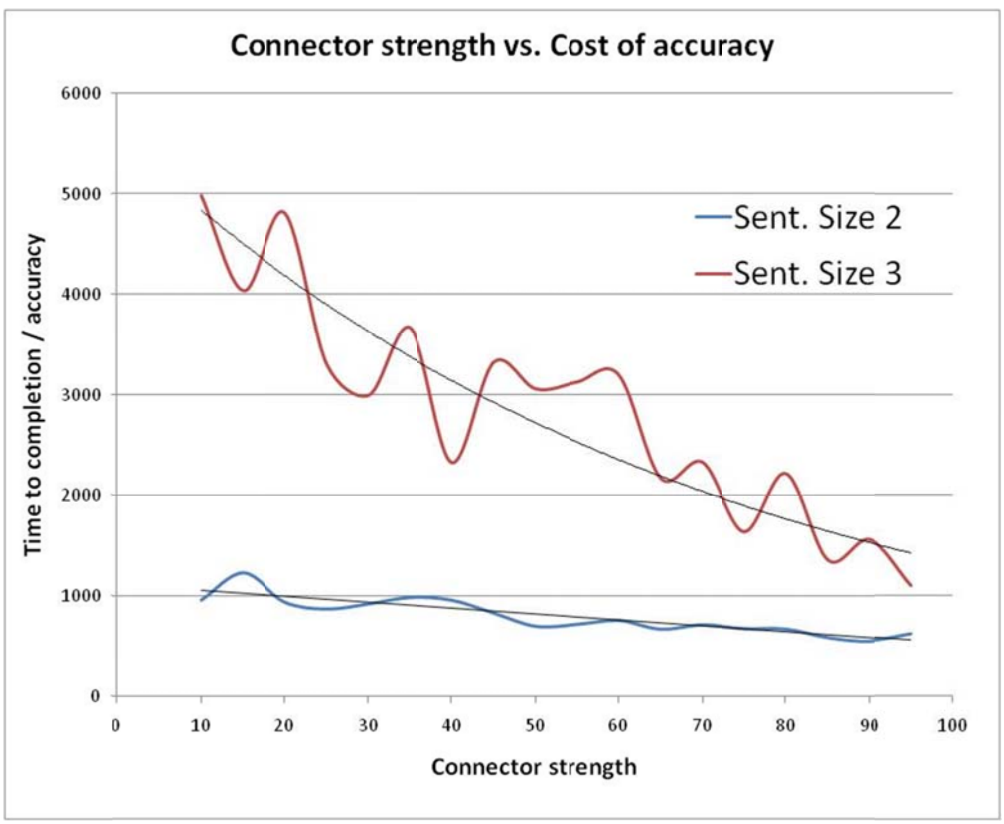

Figure 9. The graph represents the time to completion of the task (i.e., an explanation construed) divided by the accuracy of the explanation, on the Y axis, and the connector strength (how well trained the connectors are to

identify viable connection between two content pieces) on the X axis. Explanation comprised of few content pieces are relatively insensitive to the connectors' training (sentence size 2, blue line), whereas the drop is more dramatic when explanations are longer (sentence size 3, red line).

A striking result 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 (the best fit for the curve was exponential, but even a linear fit would have an much higher angular coefficient). This suggests that, for assembling 'simple' content, *the gain that students get from improved connecting skills is much lower than when there are struggling with complex knowledge.*

Again, this finding seems fitting with *Rosenberg et al.* narrative. Even in the first moment of the narrative, when students are trying to assemble explanations based on worksheets and other authority-based sources, with more consideration for formal ordering and a quasi-random approach, they were able to assemble a number of "sentence-size 2" explanations. The following four examples were extracted from the transcriptions of students' dialogues:

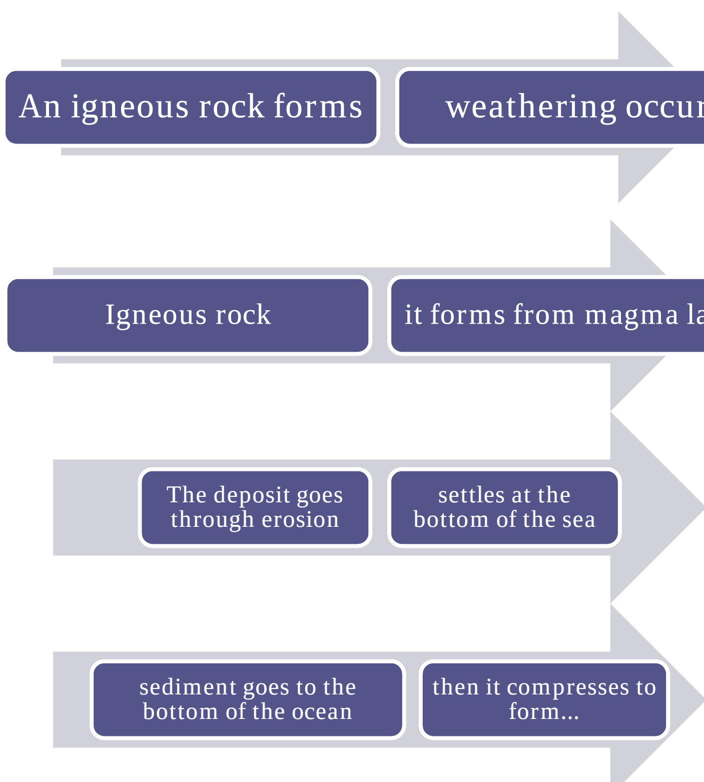

However, in that first part of the narrative, students were never able to form "sentence size 3" explanations, which would require an extra step: connecting a relatively simple pair of pieces to a third piece, evaluating all possible pieces for their fit. In the second part of the narrative, after just some minutes, by trying to 'enlarge' their explanation making sense of the connection between pieces, students formed a sentence size 4 explanation, and just some minutes later a sentence size 10 explanation.

Error! Reference source not found. shows that increasing sentence sizes has a dramatic impact on performance and on the important of 'connecting skills'. For sentence size (SS) 3 and 4, 'brute force' (low connector strength, or CS) assemblage breaks down. For SS 2, brute force assemblage is not so costly, and the benefit of developing connecting skills is not so pronounced.

The events in Mrs. Phellan's classroom tell a similar story. In the first half of the class, when students were using brute force methods and not investing on their own connecting skills, they couldn't 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 energy was spent on connecting pieces, and not retrieving them. That shift enabled them to assemble seamlessly explanations of SS as high as 10.

In a second set of experiments, we ran a more generic version of the model over a very large parameter space. We were able to identify a consistent non-linear behavior of cognitive structures – the model suggests that understanding actually takes place in "quantum leaps", and not in a purely linear progression (see, in Figure 10, a typical output of the model). Also, in this simulated environment, we were able to verify that indeed, for learning intricate content (i.e., assembling long explanations), there is a significant, non-linear, payoff to invest in "thinking skills" as opposed to "memorizing skills." For simple content (involving the connection of a few content pieces), however, sheer memorizing can outperform thinking skills.

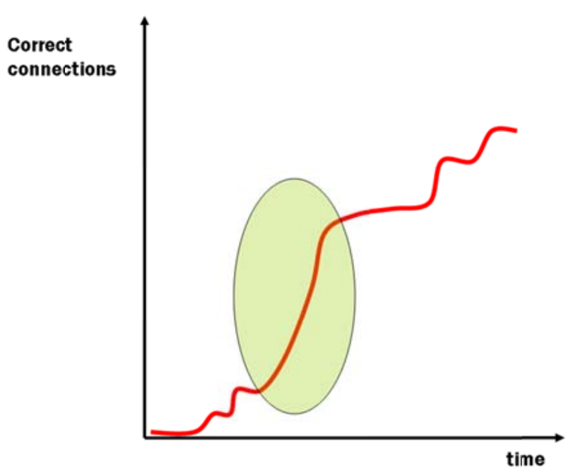

Figure 10: A stylized output of the model – typically, after a certain number of time steps, the number of correct connections (i.e., students' explanations correctly constructed) exhibits a sudden jump, instead of a linear progression.

Conclusions

Along the examples of this paper, we tried to pair our model data with interview or classroom data (*Bifocal Modeling*, Blikstein & Wilensky, 2006). The results from the "Society of More" study showed that conventional "paper and pencil" representations of Piagetian structures might miss some of the dynamic factors in play, such as 'embryonic' cognitive changes. Without a probabilistic approach, we would be obliged to just assume that those children were in a transitional stage, but with the computer representation, we could actually calculate the number of times that different embryonic agents are aroused, and estimate the developmental stage of the child. Another promising path is to evolve structures automatically, since favorable outcomes would be reinforced in the model. As we observed in the experiments, random reconnections of agents do not render random results—structure matters. The mechanism which we demonstrated shows that there is a higher probability for related agents ("long" and "thin") to group together under one particular agent – this is the configuration that delivers the best performance. Therefore, this confirms what Minsky states: the categorization process itself emergently generates intelligent behavior, without any interference from an external "intelligent" entity.

In our second case study, we 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) Students' failure in the first half of the narrative was epistemological, and not due to lacking memorizing or information retrieving skills.

2) The fundamental mathematical basis of the model, from which all other behaviors emerge, is that brute-force methods are fast for short sequences, but for long sequences, as the combinatorial space increases exponentially, their performance drops accordingly. In the high connector strength mode, however, once the connector is trained, the size of the sentence has a much lesser impact, since the evaluative rule of the connector filters out the combinatorial space, and one single successful connection (given an unlimited supply of pieces), will take the exact same computational time for any sentence size. This seems to be the case in the classroom, where students could assemble long explanations quickly, once they were in a 'high connector strength' mode.

3) In this simulated environment, we were able to verify that for learning intricate content (i.e., assembling long explanations), there is a significant, non-linear, payoff to invest in "sense-

making skills" (connector strength) as opposed to "memorizing skills" (retrieving speed). For simple content (involving the connection of 2 content pieces), however, sheer memorizing can even outperform "sense-making skills". The data shows that the payoff of improved connector strength only manifests itself after CS 80%.

4) Abrupt, non-linear shifts in student understanding are indeed possible even within very short periods of time, by activating different cognitive resources. If we consider "previous knowledge" as a strong connector, it follows that its activation following the teacher intervention could cause a sudden change in student performance.

Both cases studies show that, this work, although still in its infancy, could potentially have broad implications for the practice of curricular designers, teachers, and policy makers – by offering researchers "glass box," accessible tools to simulate, model and test hypothesis about human cognition in social contexts, as well as to pair model data with real classroom data. In particular, many of the ontologies and key ideas that are currently used on the Complexity Sciences to investigate abroad set of phenomena – from scale-free networks to non-linear behavior – were show to be useful analysis tools for better understanding human cognition, either as quantitative data confirming existing models, or as thought experiments and hypothesis-generators to advance theory building. The fact that agent-based model can be built from simple behaviors derived from qualitative analysis, but can generate quantitative data, is also a promising methodological tool for cognitive scientists.

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.

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.

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 annual meeting of the Jean Piaget Society, Baltimore, MD.

Bonabeau, E., Dorigo, M., Théraulaz, 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. Paper presented at the Proceedings of the 1978 workshop on Theoretical issues in natural language processing.

Epstein, J., & Axtell, R. (1996). Growing artificial societies: Social science from the bottom up. Washington: Brookings Institution Press.

Goldman, S. R., & Varma, S. (1995). CAPing the construction-integration model of discourse comprehension. In C. A. Weaver, S. Mannes & C. R. Fletcher (Eds.), Discourse Comprehension: Essays in Honor of Walter Kintsch. Hillsdale, NJ: Erlbaum.

Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. Psychological Review, 3(99), 480-517.

Hurford, A. (2004). A dynamical systems model of student learning and an example from a HubNet Participation Simulation. Paper presented at the annual meeting of the American Educational Research Association, San Diego, CA.

Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: individual differences in working memory. Psychological Review(99), 122-149.

Minsky, M. L. (1986). The society of mind. New York, N.Y.: Simon and Schuster.

Polk, T. A., & Rosenbloom, P. S. (1994). Task-independent constrains on a unified theory of cognition. In F. Boller & J. Grafman (Eds.), Handbook of Neuropsychology (Vol. 9). Amsterdam, Netherlands: Elsevier.

Rosenberg, S., Hammer, D., & Phelan, J. (2006). Multiple Epistemological Coherences in an Eighth-Grade Discussion of the Rock Cycle. The Journal of the Learning Sciences, 15(2), 261-292.

Strauss, S. (1993). Teachers' Pedagogical Content: Knowledge About Children's Minds and Learning: Implications for Teacher Education. Educational Psychologist, 3(28), 279-290.

Strauss, S., & Shilony, T. (1994). Teachers' models of children's minds and learning. In L. A. Hirschfeld & S. A. Gelman (Eds.), Mapping the mind : domain specificity in cognition and culture (pp. 455-473). Cambridge ; New York: Cambridge University Press.

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—an embodied modeling approach. Cognition & Instruction, 24(2), 171-209.