Multi-Agent Simulation as a Tool for Investigating Cognitive–Developmental Theory

Paulo Blikstein
Northwestern University
paulo@northwestern.edu

Dor Abrahamson
UC Berkeley
dor@berkeley.edu

Uri Wilensky
Northwestern University
uri@northwestern.edu

Abstract

We discuss an innovative application of computer-based simulations in the study of cognitive development. Our work builds on previous contributions to the field, in which theoretical models of cognition were implemented in the form of computer programs in attempt to predict human reasoning (Newell & Simon, 1972; Fischer & Rose, 1999). Our computer serves two distinct functions: (1) illustrate the Piagetian theoretical model and (2) simulate it departing from clinical interview data. We focused on the Piagetian conservation experiment, and collected and analyzed data from actual (not simulated) interviews with children from 4 to 10 years old.

The interviews were videotaped, transcribed, and coded in terms of parameters of the computer simulation. The simulation was then fed these coded data. We were able to perform different kinds of experiments:

1) Playback the interview and the computer model side-by-side, trying to identify behavior patterns;
2) Model validation: investigate whether the child's decision-making process can be predicted by the model.
3) Evolving cognitive structures departing from purely simulated data.

We conclude with some remarks about the potential for agent-based simulation as a methodology for making sense of the emergence of self-organized hierarchical organization in human cognition.

Introduction

We discuss an innovative application of computer-based modeling in the study of cognitive development. Our work builds on previous seminal contributions to the field, in which theoretical models of cognition were implemented in the form of computer programs in an attempt to predict human reasoning (Newell & Simon, 1972; Fischer & Rose, 1999). One particular type of computer modeling offers powerful methods for exploring the emergence of self-organized hierarchical organization in human cognition: *agent-based modeling* (ABM; e.g., 'NetLogo,' Wilensky, 1999; 'Swarm,' Langton & Burkhardt, 1997; 'Repast,' Collier & Sallach, 2001) enables theoreticians to assign rules of behavior to computer "agents," whereupon these entities act independently *but* with awareness of local contingencies, such as the behaviors of other agents. Typical of agent-based models 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 i.e., 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 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 put into motion and the modeler can watch the overall patterns that emerge.

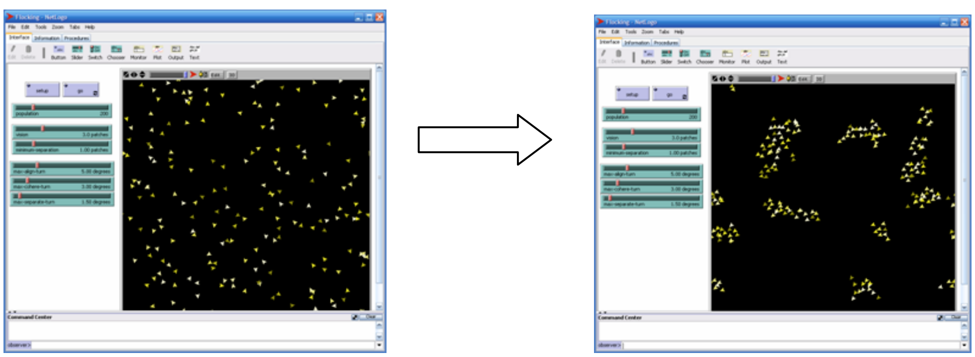

Figure 1: An agent-based model of the flocking behavior of birds.

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).

We argue that ABM has potential to contribute to the advancement of theory in multiple ways that we illustrate 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 and (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.

In this paper we focus on the potential of ABM as a research tool for formulating and critiquing cognitive development theory. ABM has been used to illustrate aspects of cognitive development (see Abrahamson & Wilensky, 2005, Blikstein, Abrahamson & Wilensky, 2006) and collaboration and group work in classrooms (Abrahamson, Blikstein & Wilensky, 2007). We, too, propose to use ABM to simulate human reasoning, yet we move forward by *juxtaposing our simulation with real data* using the Bifocal Modeling framework (Blikstein & Wilensky, 2006).

Previous research on cognitive modeling has generated many frameworks to model different tasks, such as shape classifications (Hummel & Biederman, 1992), language acquisition (Goldman & Varma, 1995), memory (Anderson, Bothell, Lebiere, & Matessa, 1998), as well as more general-purpose models (Anderson, 1983; Anderson & Bellezza, 1993; Anderson & Lebiere, 1998; Just & Carpenter, 1992; Polk & Rosenbloom, 1994). But it was in Minsky's "*Society of Mind"* theory (1986), elaborated in collaboration with Seymour Papert, that we found an adequate foundation of our agent-based models of cognition, due to its dynamical, hierarchical, and emergent properties, enabling the use of simple, programmable agent rules. We chose the classical Piagetian conservation task to model, because Minsky and Papert modeled this task with his theory; and we worked with children in both transitional and stable phases so as to elicit richer data. We will provide examples of step-by-step *bifocal* narratives – computer simulation vs. videography – of children's performance on a conservation task. In the remainder of this paper, we will introduce Minsky's and Papert's theory, explain our experiment (a variation on the classical conservation-of-volume task, Piaget, 1952), and present case studies where simulation and real data are juxtaposed.

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.

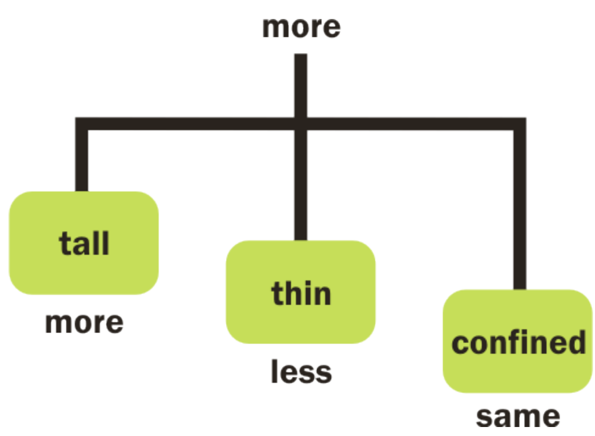

Figure 1 – A one-level model for evaluating "who has more"

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.

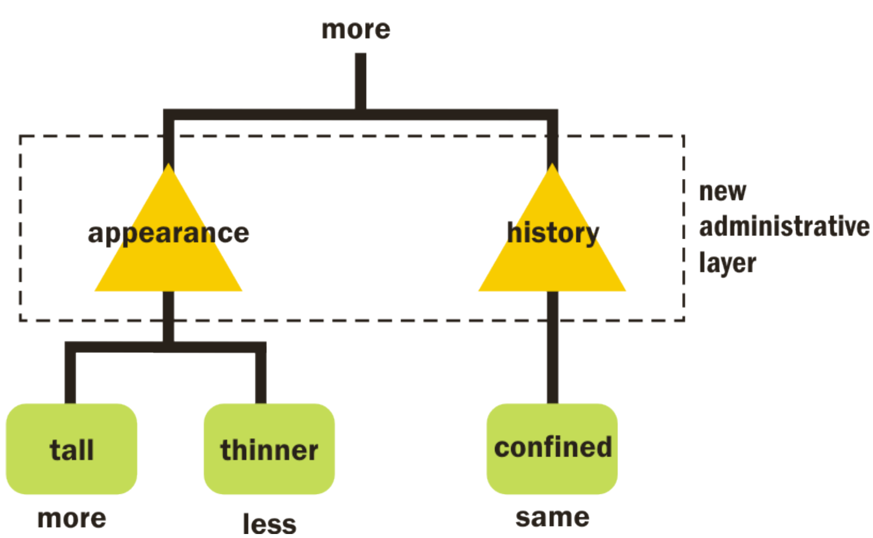

Figure 2 – New administrative layer

Experiments/Methods

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 (see Table 1). The simulation 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?

The computer model

The model tries to reproduce the clinical interview situation. We first define the "society of mind" (SOM) structure of a 'virtual child'. Then this virtual child is presented with random pairs of virtual blocks, and evaluates if one of the two is 'more', 'less', or 'same'. The model is able to automatically run multiple times, presenting the virtual child with different blocks, and also changing the rigidness of the structure (in other words, introducing random variations in each branch of the structure). In Figure 3, we have a screenshot of the model. Figure 4 shows the details of the main window.

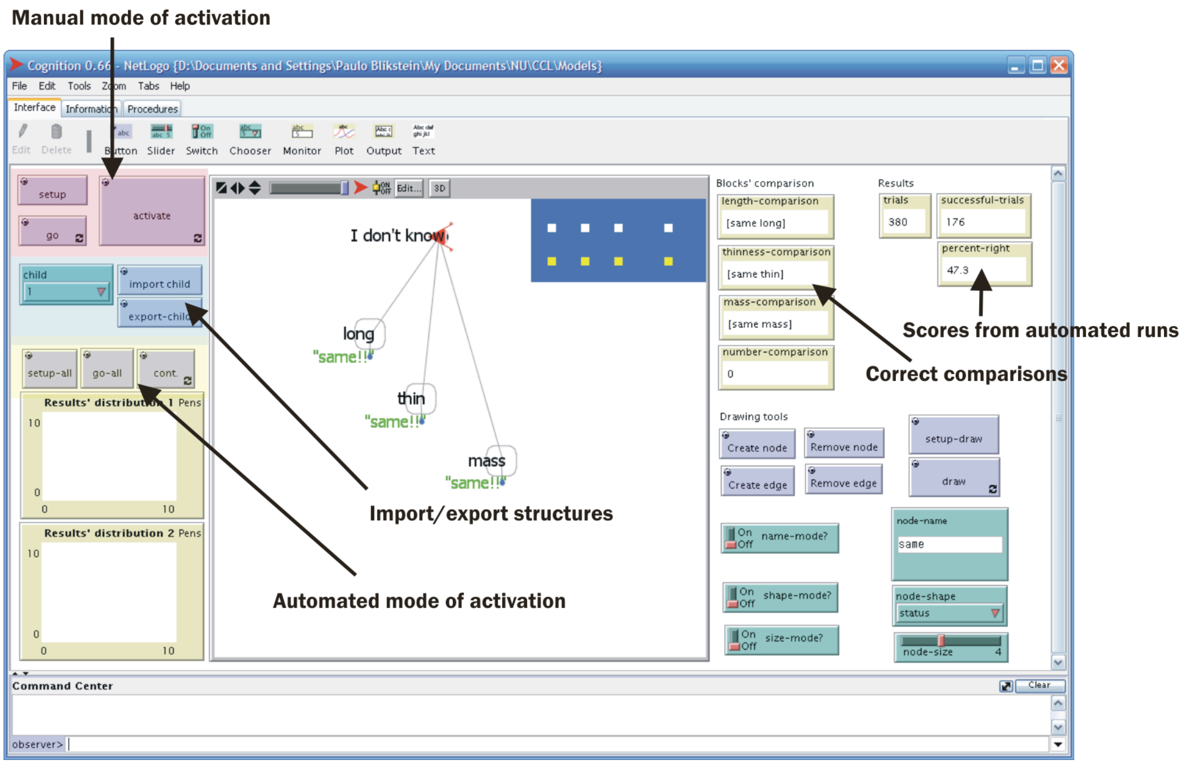

Figure 3 – A screenshot of the computer model and its main components.

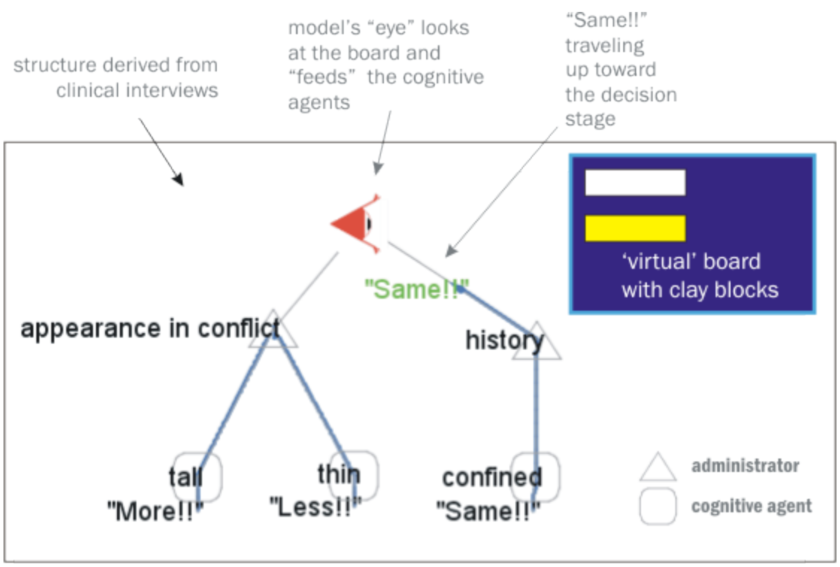

Figure 4. A screenshot of the model's main window. Similar to the child, the computer 'sees' blocks of clay and tries to determine which block is 'more.'

To use the model, the first step is to draw a structure in the central area of the screen (a 'society-of-more'). The drawing tools on the bottom right enable users to add nodes and edges, as well as change their labels, shapes, and sizes.

There are four possible types of nodes, each with a different shape and role:

- **RESULT (eye icon):** the final destination of the ‘turtles’, normally placed at the top part of the structure. This node will show the result of the computation, i.e., the final response of the virtual child. The default label for a result is “I don’t know”, which might change to “more!!”, “less!!”, or “same!!”. They can have **agents** or **managers** attached to them.
- **MANAGER or administrator (triangles):** these nodes might have **cognitive agents** attached below them, and **a result** node attached above.
- **Cognitive AGENTS (rounded squares):** these agents represent some perceptual element, such as “tall”, “thin” or “number”.
- **Cognitive agents’ STATUS (small dot and a word):** the status of an agent, which can be “more!!”, “less!!”, or “same!!”.

Once the structure is built, the second step is to “activate” the correct agents. This can be done manually of automatically:

- **Manual mode of activation:** the user assigns the correct word to the **agent status** (‘more!!’, ‘less!!’, or ‘same!!’, one by one, using the drawing tools), clicks on “Activate”, and clicks on the agents that should be active. Upon activation, a new “messenger” will be created under the **agent**, with a green label. For example, in Figure 3, all of the three agents are activated (note the three green words), as if the child did evaluate length, thinness and mass at the same time. Those green words are messengers that will travel upwards along the connecting lines when the model runs.
- **Automated mode of activation:** No user intervention in necessary. In this mode, pairs of blocks are randomly picked from a preprogrammed ‘block repository’ and displayed in the blue area inside the window (see Figure 4). The model automatically ‘sees’ the blocks and activate the correspondent agents.

Finally, for the computer to 'see' and evaluate each pair of blocks, each configuration of blocks has an associated list of 5 parameters, which are automatically compared by the model. They are: *[length of each piece, width of each piece, length of the whole arrangement, width of the whole arrangement, number of pieces]* (see Table 1). By comparing the parameters of each block, the model is able to determine which block is 'more' in total length, width, number of pieces, and mass.

Table 1. Parametrization of the blocks

Parameters	Description	Appearance of block
[8 1 8 1 1]	Each block is 8 units long, 1 unit thick, the full arrangement is also 8 x 1, and there is just one block	
[1 1 15 1 8]	Each block is 1 unit long, 1 unit thick, the arrangement occupies the total area of is 15 x 1, and there are 8 of them.	
[2 2 5 2 2]	Each block is 2 units long, 2 units thick, the total area they occupy is 5 x 2, and there are 2 of them,	
[4 1 4 1 1]	Each block is 4 unit long, 1 unit thick, the total area occupied is 4 x 1m and there is just 1 unit,	

First study: qualitative bifocal validation of the model

The goal of the first experiment is to validate the model qualitatively, i.e., evaluate if the model can minimally account for the different stages of cognitive development seen in the interviews. Below, we show the results, using 'bifocal' data (computer simulation alongside human behavior). We will show how the different models (programmed with data from the interviews with three children) yield a surprisingly similar probabilistic cluster of responses as the interviews themselves.

Computer model (screen captures)	**Transcriptions/pictures**
Child 1	
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”.
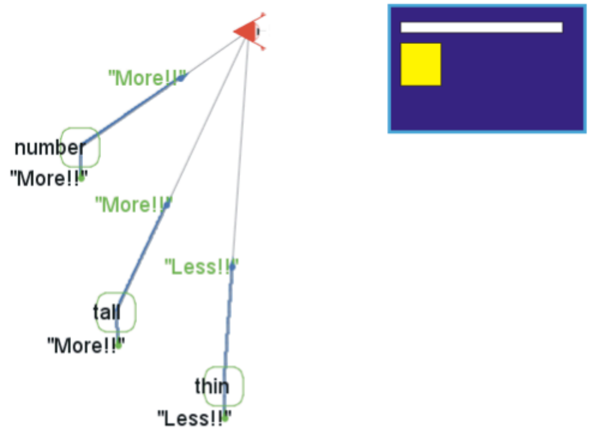	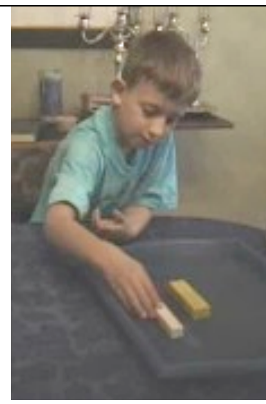 
Number is absent from this second interaction. Even when two other measurements conflict, one is always dominant. In this case, **tall** is more salient.	*Researcher: Who has more?* *Child1: It's hard to tell now. [tries to measure the fat one with his fingers, then compares his fingers with the thin and tall one]. This one [the taller].*
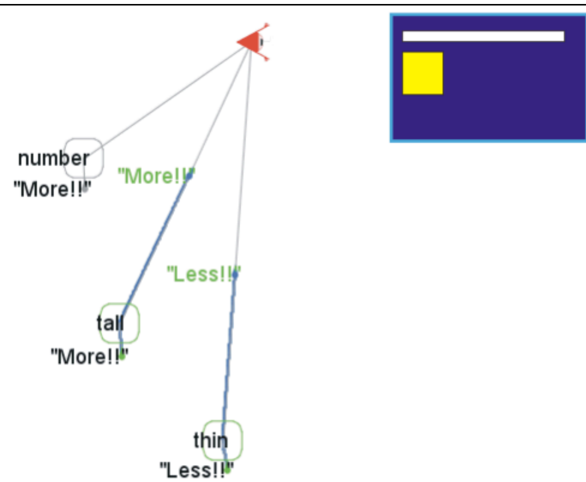	 

<table>
<tr><td colspan="2"></td></tr>
<tr><td>In the third interaction, the experimenter reintroduces Number by cutting his piece in four: as predicted by the model, number takes priority again over tall and thin. When number is present, the child does not even try to measure the two sets of blocks.</td><td>“You have more, because you have four quarters, I have only two halves.”</td></tr>
<tr><td>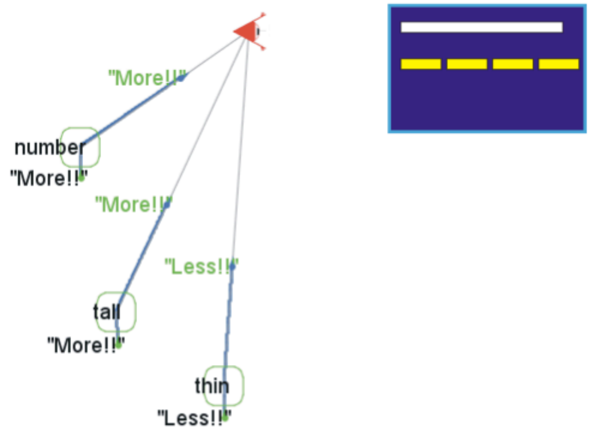
</td><td></td></tr>
<tr><td colspan="2">Interpretation: The ‘priority’ model can account for the responses of Child1: he cannot coordinate two or more measures. In the computer model, also, two measures cannot be coordinated. Given the same inputs, the computer model and the interview data yield comparable results.</td></tr>
<tr><td colspan="2"></td></tr>
</table>

<table>
<tr><td colspan="2">Child 2</td></tr>
<tr><td>Child 2 (8yo) has a model with Minsky's "administrators" (appearance and history of the transformations). With one in conflict, the other takes control. If the Tall agent reports 'more' and the Thin agent reports 'less', then the Appearance administrator will say nothing - it is in conflict and cannot decide. However, this child provided different answer to similar block configurations. He would alternate between a mass-conservation explanation (no material was taken away or added) and a 'joinable' one (two previously cut pieces can be joined together to form the original object). It appears that, even having a more developed SOM structure, this child is also in a transitional phase, in which 'mass' and 'joinable' take turns dominating.</td><td>"If you put them back together, you'll have the same"</td></tr>
<tr><td>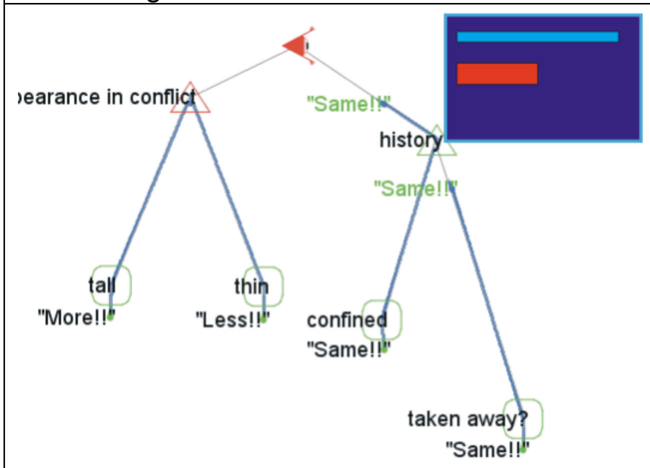
</td><td></td></tr>
<tr><td colspan="2">Child 2 has a level of administrators, which enables him to background the appearance and focus on the history of the objects. The blue is 're-joinable', so both blocks are the same. During the interview, Child 2 occasionally said that nothing was added or taken away - a static, rigid model is insufficient to account for those oscillations, as we will later discuss. The model, again, correctly determines the combinatorial space and predicts response frequency distribution.</td></tr>
</table>

Child 3	
For Child 3 **(10yo)**, **material taken away?** was far more dominant that **joinable?** or **appearance**.	"It's the same, because you still have the same amount, even if you cut in half in different ways, because it's still in half."
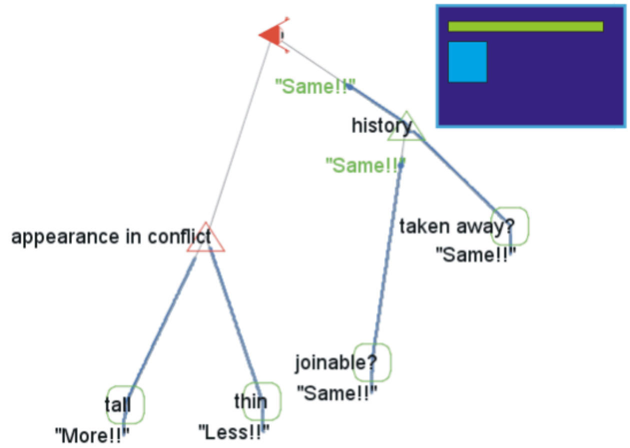	
Child 3 backgrounds **appearance** from the start (see, in the model, that these agents are lower than others) and focuses on **confinement** (nothing was *taken away* or added), and thus concludes that the blocks are still the same.	

In this part of the study, we were able to describe the cognitive development of child 1, 2 and 3 solely in terms of the variables of the computer model: the number of layers in the structure and the relative prominence of certain agents. Child 1's responses could be fit into a one-level structure, Child 2's responses fit into a two-level structure but without a clear 'leveling' of the agent – which we will only see in Child 3. Moreover, in both Child 1 and 2 we observed elements of a transitional phase, which the model can also account for, by promoting slight random variations in its own structure, until a stable configuration is reached (see next section).

Second study: understanding the interviews with non-rigid computer models

To investigate in more depth the relationship between the computer model and the performance of Child 1, we coded the whole interview in terms of computer-understandable parameters. We used the same parameterization employed in the model to describe each pair of blocks laid before the child. Thus, if the child was presented with a 4 x 2 and a 1 x 8 blocks, the coding would be [4 2 4 2 1] and [1 8 1 8 1], following the convention already mentioned in this paper. Then we rated the child's answer as right or wrong, and also noted which block was considered to be 'more'. The result for Child 1 is in the Table 2.(Under the "Blocks" column, the child's choice for the 'more' block has a light green background.)

Table 2 - Responses of Child 1

Blocks	Correct?	Transcription
[4 2 4 2 1] [4 2 4 2 1]	Yes	C1: "It's the same."
[4 1 4 3 2] [4 2 4 2 1]	No	I: "Who has more?" C1: "I have more. This is two, but it's smaller."
[8 1 8 1 1] [4 2 4 2 1]	No	C1: "I have more. Look at this. We cut it in half and it got taller. This is kind of fat but this is taller."
[8 1 8 1 1] [2 4 2 4 1]	No	C1: "It's hard to tell. It's hard to tell [measures with his fingers the thickness of both] This one [the longer and thinner]."
[8 1 8 1 1] [1 2 7 2 4]	No	C1: "You have four quarters, I only have two halves. If I would do that [join], we would have the same amount."
[8 1 8 1 1] [8 1 8 1 1]	No	C1: "You have more. It's a little bit taller."
[4 2 4 2 1] [4 1 4 1 1]	Yes	C1: "I have more. I have two halves, you only have one. I could break it apart."
[4 1 4 3 2] [1 1 15 1 8]	No	C1: "You have more. Because that's 1, 2, 3, 4, 5, 6, 7, 8 [counting]. And the other is 1, 2. But if we cut that there, that there, that there, it will be the same amount (cutting the 2 pieces into 8].
[1 1 15 1 8] [1 1 15 1 8]	Yes	C1: "[Counting] 1, 2, 3, 4, 5, 6, 7, 8, and 1, 2, 3, 4, 5, 6, 7, 8. It's even."
Child 1: 30% (3/9) correct		

We fed the computer model with the same 9 pairs of blocks, and rated the performance of the computer model in comparing the blocks. Four different structures were tested (see Table 2):

1) Agents: "long", "thin" and "mass", no administrative layer
2) Agents: "number", "long", "thin" and "mass", no administrative layer
3) Agents: "long", "thin" under appearance, "mass" under history.
4) Agents: "long", "mass" under appearance, "thin" under history. (note the displaced position of "mass" and "thin"

We further tested each structure with different relaxations – a total of 21 simulations with the level of 'relaxation' increasing from 0 to 100% in steps of 5%. In the model, the 'relaxation' corresponds to randomly deforming the distances which the agents have to travel upwards. In practical terms, the more 'relaxed' the structure (closer to 100%) the less determinant the structure is.

In the following plots, the black line represents the accuracy (in percent) of the model. The blue line represents the score of the actual child, and the green line is the trend line of the model's performance curve.

Table 3 - Different levels of relaxation for different structures

<table>
<tr>
<td>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.</td>
<td>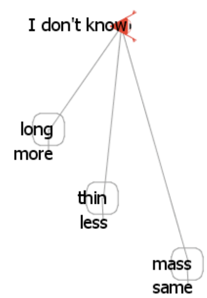
</td>
<td>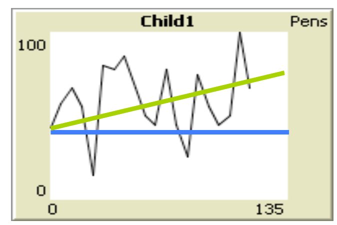
</td>
</tr>
<tr>
<td>Structure 2
Adding 'number' to the structure, the overall accuracy increases from approximately 35% to 55%. Interestingly, in this more complete model, increased randomness causes accuracy to decrease. A likely explanation is that "number" is an efficient "specialized" evaluator of 'more', at least for the population of blocks laid before the child.</td>
<td>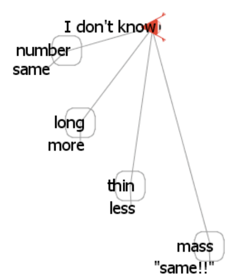
</td>
<td>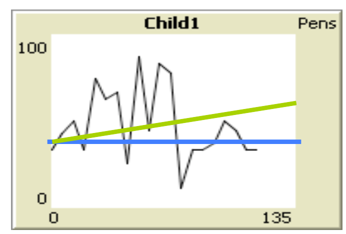
</td>
</tr>
<tr>
<td>Structure 3
In this model with administrators, accuracy starts at 100%, and decreases with increase randomness. The average score of this model, even with high randomness levels, is far better than the child's.</td>
<td>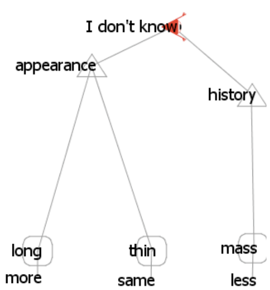
</td>
<td>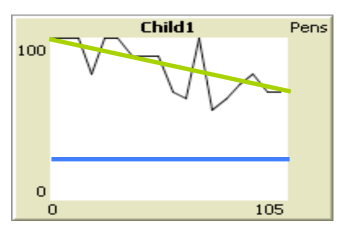
</td>
</tr>
<tr>
<td>Structure 4
This model is identical to Model 3, except that mass and thin were switched. Mass is under appearance, and thin under history. The overall score drops dramatically from 100% to 30%, approaching the child's. This is corroborates Minsky's hypotheses about the importance of having the 'right' agents under each administrator. If we were to 'evolve' a structure using standard GA algorithms, probably Model 3 would be rapidly selected over Model 4.</td>
<td>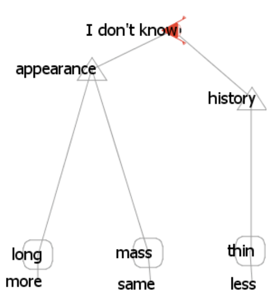
</td>
<td>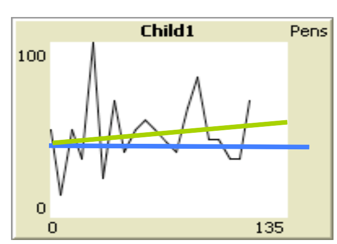
</td>
</tr>
</table>

Interpretation: The different effect of relaxation in the models' performance is an important result of this experiment. Simple, one-level models increase their performance with increased relaxation. 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 deforming a weak structure benefits from random correct hits. 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 hypothise 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. To further investigate this issue, we conducted a third study with more comprehensive runs.

Third study: effect of relaxation on different structures

To conduct this study, we generated a repository of 16 different 2-block configurations. Those blocks were randomly selected and evaluated by the computer using different SOM structures. Relaxation ranged from 0% to 300%, with 5 runs for each data point.

Model 1, as expected, had its accuracy increased with increased randomness. Accuracy tapered off around 50-60%. This is probably around chance, but the number is different from the expected 33% (three random outcomes) because the distribution of long and thin blocks was not uniform. Consequently, Model 1 is very context-specific – if there are more long objects around the child, it would work more than chance, but when that is not the case, it's worse than chance.

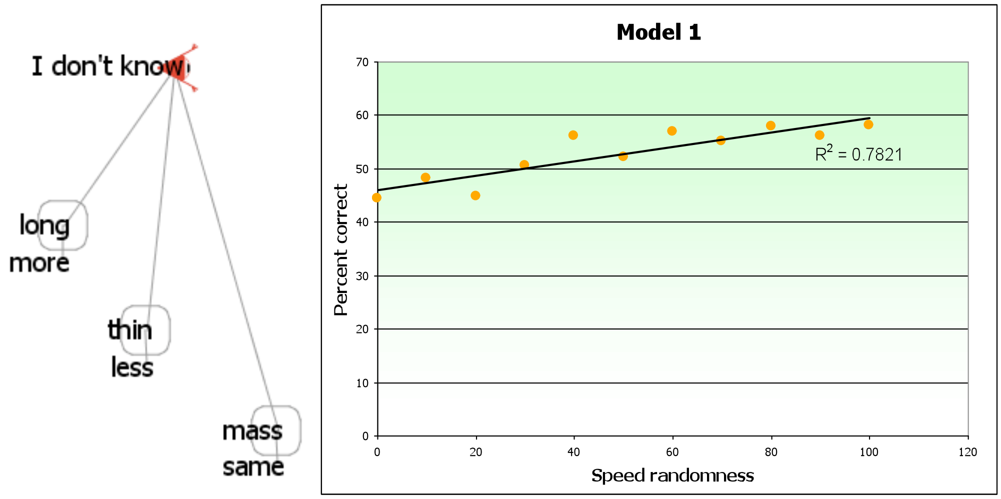

Figure 5. Speed-randomness vs. accuracy for Model 1

Model 3, however, presents a different picture. It begins with 100% of accuracy, which decreases dramatically with relaxation. The following plot show results from 0 – 140% of relaxation: we can observe that accuracy declines and tapers off around 60%.

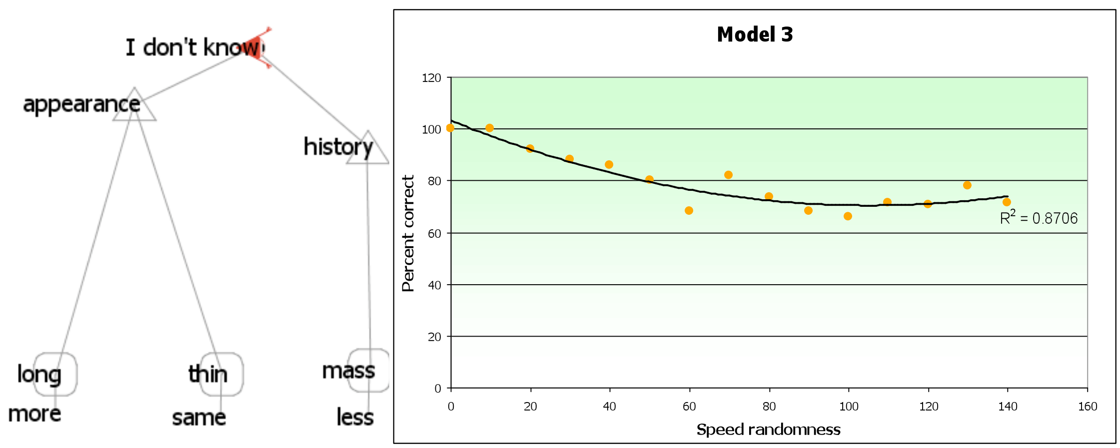

Figure 6. Relaxation vs. accuracy for Model 3

Model 4 confirmed the results from the previous study: placing the agents in wrong places in the structure has dramatic effects. In this experiment, accuracy dropped from 100% to around 35%, and increases with randomness.

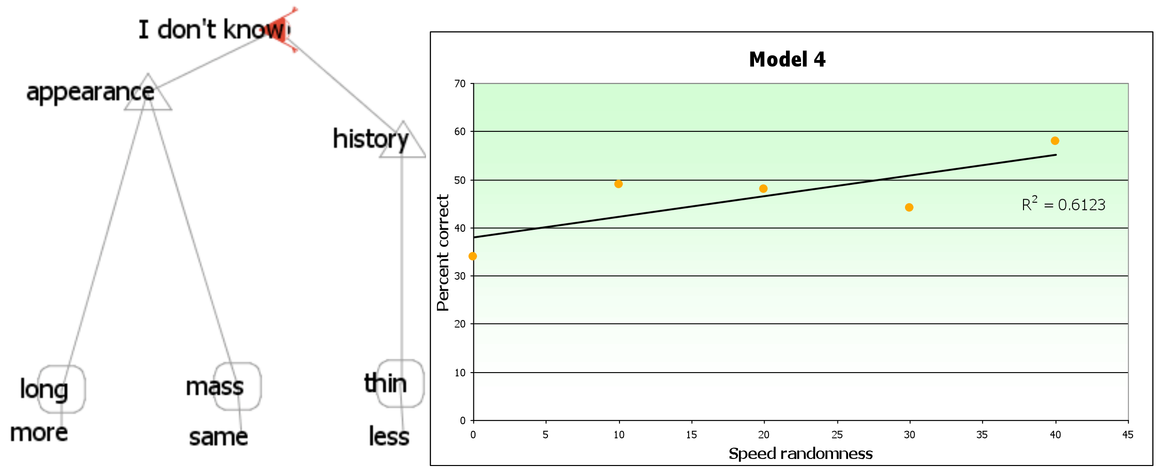

Figure 7. Speed-randomness vs. accuracy for Model 4

Conclusions

The computer model can be a 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 ongoing 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. We are currently exploring the entire combinatorial space of all hypothetical children's initial mental states and activating the simulation per each of these states. From that perspective, our data of real participants become cases out of the combinatorial space.

The following are some conclusions from the three experiments described in this paper:

1) Relaxation has different effects on structure with and without administrators. This suggests that relaxation, trial-and-error, and changes in the environment could be factors leading to a natural selection of structures of Minskyan non-intelligent agents.

2) 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 Child 1 and Child 2 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 specific.

3) A natural and promising path for this work is to evolve SOM structures automatically. We suggested earlier that the dynamics of this simulation (see Conclusion 1) is such that favorable outcomes would be reinforced. As we observed in experiments 2 and 3, 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.

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.
Axelrod, R. M. (1997). *The complexity of cooperation: Agent-based models of competition and collaboration*. Princeton, NJ: Princeton University Press.
Blikstein, P., & Wilensky, U. (2006). *Learning about learning: Using multi-agent computer simulation to investigate human cognition*. Paper presented at the International Conference of Complex Science, Boston.
Blikstein, P., & Wilensky, U. (2006). The Missing Link: A Case Study of Sensing-and-Modeling Toolkits for Constructionist Scientific Investigation. In *Proceedings of the 6th IEEE International Conference on Advanced Learning Technologies (ICALT 2006)*, Kerkrade, The Netherlands, 980-982.
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.
Collier, N., & Sallach, D. (2001). Repast. University of Chicago (2001). http://repast.sourceforge.net .
Diermeier, D. w. A. M. (2000). Government turnover in parliamentary democracies. *Journal of Economic Theory, 94*, 46-79.
Epstein, J., & Axtell, R. (1996). *Growing artificial societies: Social science from the bottom up*. Washington: Brookings Institution Press.
Fischer, K. W., & Rose, S. P. (1999). Rulers, models, and nonlinear dynamics: Measurement and method in developmental research. In G. Savelsbergh, H. van der Maas, & p. van Geert, (Eds.), *Nonlinear developmental processes* (pp. 197 - 212). Amsterdam: Royal Netherlands Academy of Arts and Sciences.
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.
Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: individual differences in working memory. *Psychological Review*(99), 122-149.
Langton, C. & Burkhardt, G. (1997). Swarm. Santa Fe Institute, Santa Fe. http://www.swarm.org/release.html
Minsky, M. (1985). The society of mind. London, England: Hienemann.
Minsky, M. L. (1986). *The society of mind*. New York, N.Y.: Simon and Schuster.
Newell, A., & Simon, H. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
Piaget, J. (1952). *The child's conception of number*. London: Routledge & Kegan Paul.
Piaget, J. (1952). *The child's conception of number*. London: Routledge & Kegan Paul.
Piaget, J., Gruber, H. E., & Vonèche, J. J. (1977). *The essential Piaget*. New York: Basic Books.
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.
Resnick, M., & Wilensky, U. (1993). *Beyond the Deterministic, Centralized Mindsets: A New Thinking for New Science*. Paper presented at the American Educational Research Association conference, Atlanta, GA.

Wilensky, U. (1999a). NetLogo. The Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL. http://ccl.northwestern.edu/netlogo/

Wilensky, U., & Resnick, M. (1995). New Thinking for New Sciences: Constructionist Approaches for Exploring Complexity. *Presented at the annual conference of the American Educational Research Association*.