

# A Dialogue on the Role of Computational Modeling in Developmental Science

Vanessa R. Simmering,<sup>1</sup> Jochen Triesch,<sup>2</sup> Gedeon O. Deák,<sup>3</sup> and John P. Spencer<sup>4</sup>

<sup>1</sup>University of Wisconsin–Madison, <sup>2</sup>Goethe University Frankfurt, <sup>3</sup>University of California at San Diego, and <sup>4</sup>University of Iowa Delta Center

**ABSTRACT**—All sciences use models of some variety to understand complex phenomena. In developmental science, however, modeling is mostly limited to linear, algebraic descriptions of behavioral data. Some researchers have suggested that complex mathematical models of developmental phenomena are a viable (even necessary) tool that provide fertile ground for developing and testing theory as well as for generating new hypotheses and predictions. This article explores the concerns, attitudes, and historical trends that underlie the tension between two cultures: one in which computational simulations of behavior are an important complement to observation and experimentation and another that emphasizes evidence from behavioral experiments and linear models enhanced by verbal descriptions. This tension is explored as a dialogue among three characters: Ed (Experimental Developmentalist), Mira (Modeling Inclusive Research Advocate), and Phil (Philosopher of Science).

**KEYWORDS**—computational modeling; philosophy of science; experimental methodology; quantitative methods; development

Mira, Ed, and Phil are in Santa Fe, New Mexico, for a conference. After attending some of the colloquia, they meet up at the Georgia O’Keeffe museum. Mira walks up to Ed and Phil who are standing outside the museum.

Mira: Hey, I just saw a great set of talks—did you go to the symposium called “To Model or Not To Model”? I thought they hit on some really important points.

Ed: I saw that in the schedule, but I didn’t go. I’m not really interested in modeling.

Mira: That’s too bad, you should have gone; this symposium was designed for people like you.

Ed: What do you mean, “people like me”?

Mira: I mean people who don’t do modeling. The point of the symposium was to highlight how modeling and empirical approaches can support one another if they stay connected.

Ed: I think empirical approaches are doing just fine, thank you very much.

Mira: I wasn’t trying to criticize nonmodeling approaches. My point—rather, the point of the symposium—was to discuss why developmental scientists from all perspectives should care about models. It’s not supposed to be an either–or argument; it’s about bringing the perspectives together.

Phil: Sounds like you two need to move closer together! Let’s move this inside. Although I’m sure the talks were fascinating, I’m eager to see some paintings. Hey, check out that quote:

*They look to where Phil is pointing. The wall reads:*

Nothing is less real than realism . . . Details are confusing. It is only by selection, by elimination, by emphasis, that we get at the real meaning of things.

Georgia O’Keeffe

We thank Peter Vishton, Chris Moore, and Rob Goldstone for participating in the 2007 Cognitive Development Society symposium that was the basis of this invited article. We also thank Martha Alibali, Sammy Perone, and Jenny Saffran for feedback on earlier versions of the work. John P. Spencer was supported by the National Science Foundation (HSD 0527698) and the National Institutes of Health (R01-MH062480). Jochen Triesch was supported by the European Union through project MEXT-CT-2006-042484 (“PLICON”) and by the German Federal Ministry of Education and Research (BMBF) within the “Bernstein Focus: Neurotechnology” through research Grant 01GQ0840. Gedeon Deák was supported by the National Science Foundation (HSD-0827040), and the Temporal Dynamics Learning Center.

Correspondence concerning this article should be addressed to Vanessa R. Simmering, Department of Psychology, University of Wisconsin–Madison, 1202 W. Johnson St., Madison, WI 53706; e-mail: simmering@wisc.edu.

© 2010, Copyright the Author(s)  
Journal Compilation © 2010, Society for Research in Child Development

*Once inside, the three walk around looking at paintings and continue their discussion.*

Phil: Why is there controversy surrounding the use of models in psychology? I thought models were part of all sciences.

Mira: It depends on what you mean by models. Most psychologists have embraced models of some kind, but they are often specified at the conceptual level, like box-and-arrow diagrams. I'm talking about computational models, models that use precise, mathematical implementations to try to explain the processes that underlie behavior.

Phil: I'm not sure I know what you mean—can you give an example of this type of model?

Mira: Sure. Let's say I wanted to explain how long you will look at this particular painting. That could be useful information for the owners who need to figure out how people will move through the museum.

*Mira opens her program and writes:*

$$L_{t+1} = L_t + \sum E(f, t) - \sum WM(f, t)$$

Mira: This specifies the relationship between looking ( $L$ ), encoding ( $E$ ), and working memory ( $WM$ ) formation as you look at the painting. It says that all of these factors now at time " $t$ " contribute to looking in the future ( $t + 1$ ), for instance, in the next second. This equation tells us how you will look from second to second, and can predict when you will look away from this painting and move on to the next one.

Ed: This is exactly what I don't like about modeling. That equation makes absolutely no sense to me!

*Mira sighs.*

Mira: Let me try this again.

*Mira writes:*

Looking at the next second is a function of: Looking now + (Encoding of the features, objects, etc., in the painting) – (Forming a working memory of the features, objects, etc., in the painting)

Mira: Here it is, restated in words. The model I wrote before says that you will keep looking at a painting as long as there are still interesting new things to perceive and learn about—new features, new details in the painting, new emotions to experience. As you look at the painting, you start to perceive and encode what's there. This process leads to some active representation of the painting as a whole. Once you have a good representation of the painting as a whole, you'll stop detecting new things and eventually, you'll stop looking.

Phil: That makes sense—you're using some combination of current looking, encoding, and working memory to predict how looking will change over time.

Ed: That definitely makes more sense now, which makes me wonder why you don't just write equations in words in the first place—that would help a lot!

Mira: Although it's possible to write down the basic parts of this model in words, the verbal description doesn't go far enough.

For instance, for the equation I wrote, I will need to specify a process that does the encoding and working memory part. So the primary equation has to be expanded to be complete. This highlights an important benefit of modeling: Models force us to specify our assumptions and exactly how we think things work.

Phil: Right, in your example, it's not enough to just say that how long you look depends on what you have left to encode—you need to spell out how that encoding happens, and how working memory formation helps determine when it's done. Otherwise, it's not really specific enough to test if this explanation is right.

Mira: Exactly! In psychology, when trying to describe a complex system verbally, we often stop with the verbal description and don't recognize that we have to push things to a greater level of detail to explain and predict behaviors.

Ed: Now wait, you make it sound like all conceptual theories, but no models, are poorly specified—that's certainly not the case. There are plenty of good, specific conceptual theories that have led to novel insights. For instance, Darwin didn't need complicated mathematics to come up with his theory of evolution by natural selection. Meticulous observation and description were the foundation of his theory.

Phil: Yes, and don't forget the incredible leap of creativity it took to go from the data to such an elegant yet encompassing theory.

Ed: Right: It's hard to imagine how any computational model could have brought together such a wide array of evidence within a single theory.

Mira: That's a great point. I would never argue that models should replace sound theoretical reasoning.

Phil: I think Ed is also right that, even with mathematical specificity, some models have vague or ambiguous connections to the data.

Mira: Sure, there is a range of specificity in both approaches, especially when it comes to connecting to data. But developmental science *in general* needs greater specificity, because a lot of time is spent arguing about vague constructs and vague theories. In my view, you can't get any more specific than detailing variables and their relationships using mathematical equations: If it's not in the equation, it can't affect the outcome. Equations can be demonstrated to be right or wrong.

Ed: Doesn't that just move the ambiguity to the next level? You still have to figure out what all those parts of the equation mean, and how they connect up with the data that you are collecting. Like in your example, you can't measure working memory formation directly, so isn't that just as vague as if I have a verbal theory saying the same thing?

Phil: That's true, models can be too abstract. But the same can be true of verbal accounts. Either way, your explanation really needs to be grounded in theory, to stay connected with the behavior you're studying.

Mira: That brings up one of the main points of the symposium today: Models, experiments, and theories are all interconnected

components of our science. Without theories to tie everything together, neither models nor experiments can make much progress on their own.

Ed: Why, then, does it seem as though some people try to elevate models above the rest?

Phil: It's funny to hear you two argue about this because in other sciences there isn't such conflict between different approaches. For example, in early astronomy, the positions of planets and stars were mapped from observation—they needed a foundation of information before models could get started. But they soon turned to precise theoretical tools. A great example is Le Verrier's work in the mid-1800s: He was a mathematician who noticed—through observation—some irregularities in the orbit of Uranus.<sup>1</sup> Then, by modeling the orbit mathematically, he predicted that an eighth planet must exist, with an orbit outside of Uranus'. Based on his calculations, astronomers knew just where to look in the sky and discovered Neptune!

Mira: That's a great example of the second contribution of models: to generate novel predictions. If you have a model that is well connected to behavior—in this case, the planets' "behavior"—then you can use the model in new ways, and see what it predicts in new situations, or to deduce that you need different variables in your model to really account for behavior.

Ed: Of course prediction is important, that's a basic part of the science. But I generate specific predictions all the time, without computational models. It's hard to see the parallel between the orbit of Uranus and the study of human development, which has many more influencing factors. Why represent complex human development with complex mathematical equations? Those simple verbal concepts you criticize do a great job simplifying and synthesizing real data sets. Isn't simplification one of our goals? Like O'Keeffe said, "through selection, elimination, and emphasis, we can get to the real meaning of things."

Phil: Yes, the role of *all* theories, with or without mathematical models, is to organize and synthesize data in a coherent explanation. The question is whether you can achieve those goals without a well-specified model. As phenomena get more complex, it's easy to see how purely verbal descriptions fall short.

Ed: That's why we use statistics, to give us standards for evaluating data. Instead of simply explaining the pattern I see in my data, I report descriptive statistics, and compare different conditions in an ANOVA, to interpret those data and infer what they mean, in a conventionalized mathematical framework. So I have mathematical specificity without modeling.

Phil: That's an interesting point—in a way, all researchers use mathematical models when they apply statistical methods.

Mira: Hm, I suppose that's true. I wouldn't normally include statistics in my definition of computational models. I'm thinking about models that emphasize the *processes* underlying behavior and development.

Phil: Just because statistical models have a different focus than your model doesn't mean they're not informative.

Mira: Sure, I get your point. I guess this also avoids the problem Ed brought up earlier—with statistical models you have an obvious connection to your data. But statistics don't necessarily point us toward meaning—in theoretical modeling, you have to say *how* you think different factors are related quantitatively, and *why* changes in one factor will lead to changes in another factor over time.

Phil: And this requires particular types of mathematics, right? For instance, covariance-based statistics, like correlation matrices, tell you about associations but not about the nature of relations between factors. On the other hand, calculus is all about the study of change. It can at least tell you about trends and changes in trends, which are often a necessary component of explanatory and predictive theories. So, it makes sense that researchers are trying to apply calculus to developmental phenomena.

Ed: Calculus? I thought we were supposed to *reduce* complexity in order to learn anything. Aren't those the "details [that] are confusing"?

Phil: Speaking of confusing, we seem to have found a choice point—see more of the museum or go back out into the sunshine?

Ed: Onward!

*The three move deeper into the museum . . .*

Phil: So we have a desire for simplicity on one hand and confusing details on the other. Isn't that why we moved things into the laboratory in the first place—to reduce complexity?

Mira: But that doesn't simplify the *child*—you still have the same rambunctious 5-year-old; you've just limited the ways he can bounce off the walls! Experiments simplify behavioral complexity by stripping down the environment; models can go one step further. Think of it as reducing the internal complexity of the child—this way, we are controlling the details of the environment *and* the organism. When we test mathematical models in simulations, then, we can actually know *everything* that's going on—not just in the environment, but also in our simplified "child."

Phil: This sounds similar to the way physics makes use of simple systems like a pendulum or hypothetical frictionless ramp to understand more complex systems.

Mira: Exactly! The reason physicists—or chemists or biologists—could make progress based on those simple systems is because they moved to a more complex level of analysis. All of these sciences did that by using advanced mathematics, like calculus, to describe and analyze complex systems. Once you get to a certain level of complexity in the system that you're studying—whether it's a physical system, a biological system, or a social system—you can no longer rely on verbal concepts or linear statistics alone. Analysis of complex systems requires the use of complex tools.

Ed: But I'm not sure complex systems absolutely *require* complex mathematics. Think about Darwin: Biologists have learned a lot through observing and cataloging the natural world.

Phil: Well, that depends on exactly what you want to know. Are you satisfied with cataloging children's different behaviors, or do you want something beyond that—overarching principles or common explanations that tie development together across tasks and domains?

Mira: Right, think back to the model I showed you. Someone could catalog the looking times at each painting as people go through the museum, but if you want to generalize beyond those people, those paintings, and this museum, you need to understand something more general about why people look at paintings the way they do.

Phil: I think this provides an important analogy with other sciences: we can arrive at general explanations, and therefore understanding, only if we can synthesize results across different experiments and areas of research.

Mira: Right, modeling approaches are tested most stringently by taking data from a variety of studies and integrating the findings into a single model. This is what I see as the third benefit of modeling in developmental science—organization and integration of empirical findings.

Phil: History backs you up. As each science has collected more accurate and detailed observations, their particulars eventually force more complex models. For example, in astronomy, the meticulous data collection by Galileo and Kepler demanded better models, which eventually required the development of calculus.

Mira: And now, as we said, calculus has much broader applications, allowing scientists to make detailed predictions about development. From the motion of stars to the motion of children . . .

Phil: That's beautiful!

Mira: Thank you, Phil.

Phil: Sorry, I was referring to that painting . . .

Mira: Oh . . . well, are we ready to head out of here?

Phil: Yeah, let's go.

*The three walk out of the museum and take a nearby path toward Old Fort Marcy Park overlooking Santa Fe. The conversation resumes:*

Phil: It sounds as though you have some good arguments for how computational models can be useful in developmental science, and there are already some "success stories" in the field. So why haven't models been more influential thus far?

Ed: Frankly, one reason is because they are often based on equations that aren't familiar to most developmental scientists. Even with all my statistics training, I still don't understand most equations used in computational models.

Mira: I know. Without some mathematical background, it's hard to understand models beyond a superficial level. This creates a cycle: Each person who hasn't been taught how to read and understand modeling equations goes on to train students in the same way, and so on down the line. There has been a breakdown in graduate training that will have to be addressed if we

are going to bridge the divide between modelers and experimentalists.

Phil: There must be some examples of accessible papers for nonmodelers to read. Lack of communication can't be the only reason modeling isn't more accepted.

Ed: I'll admit, one reason I'm suspicious of models is that, from what I've seen, if your model isn't working, you can just add a new parameter, or scale the outputs, or whatever you need to make it work. It just seems so arbitrary!

Mira: Of course, there are examples where modeling is done poorly, giving modeling a bad name . . .

Ed: Ah, so you modelers aren't perfect after all!

Mira: . . . but *any* tool can be used poorly: There are also lots of bad experiments on children's behavior, or bad observational studies, but no one would consider dumping those methods as a result!

Phil: Maybe Ed's resistance to modeling can tell us something about which modeling approaches are more convincing than others. From what he's been saying, it seems like it's important to clarify why your model is set up the way it is, and to justify any changes that need to be made. Beyond that, how are we to determine what makes a good model?

Ed: Well . . . I suppose one standard that matters to me is whether the model can tell us something new, either about a specific behavior or about some underlying concept or assumption.

Phil: Sure, new insights are central to scientific progress. But what about when there are multiple models of the same phenomenon that make the same type of predictions? How do we judge models against one another when they can differ on so many levels?

Mira: In these situations, models are typically evaluated on how well they fit the pattern of data, and maybe how many free parameters they have. Although that is an important component of evaluating models, in my opinion it puts too much emphasis on mimicking data sets—which may have problems of their own.

Ed: That's one aspect of modeling that really turns me off—tweaking parameters to get a better fit. How does that further our knowledge?

Mira: Some amount of "tweaking" will always be necessary because you have to specify *everything* in a model. Some things aren't known, but you still have to put some values into your formula, and to begin with, you might just have to guess what values will work.

Phil: And I'm guessing with those arbitrary values to start, you don't always get good performance out of your model.

Mira: Right, so you have to test out different options to figure out what works. Careful modelers will show how their model works with a range of parameter values—and even figure out where the model's behavior breaks down. That can give us new ideas, for example, about developmental disabilities.<sup>2</sup>

Ed: Interesting, I never thought about how a model's failure could be informative.

Phil: I would argue that the failure of a model is actually what is *most* informative. Science could never progress with only confirmatory evidence.

Mira: Right, you often learn the most from what makes your model fail. But it's also important to specify how modeling approaches differ. That is, what assumptions do they make about how information is sampled, integrated, and processed? What concepts are embodied in their variables and formulas? Does the model generalize to new phenomena, predict new things, and integrate a range of empirical observations?

Phil: So . . . you're saying that a good model is like a good theory.

Ed: That sounds better, but it also sounds tricky, particularly when it isn't clear exactly how a model is tied to specific phenomena.

Mira: That's another reason models are viewed with some skepticism. As models become more complex, they can lose their connection to real behavior, not only in the lab but outside the lab as well.

Phil: Sounds to me as though you need something, or perhaps *someone*, else!

Mira: I don't follow . . .

Ed: You need *me*! I can help you keep it real.

Mira: You're right—you know the way back to the hotel! Seriously, now I'm with you: What was so great about that symposium was not just the modeling; it was the integration of modeling with empirical work. The best work integrates the most careful empirical and behavioral work with the most neurologically, and behaviorally, realistic models.

Ed: Finally, I feel useful again!

Mira: An important dimension in model evaluation is how well modeling frameworks connect with the details of behavior. We need more process-based models that can handle the task-specific nature of behavior but also connect with known processes or neural mechanisms of learning and development.

Ed: If you had a modeling framework that could give me new hypotheses to design experiments in my lab—that would be really innovative.

Phil: Could that be . . . a dialogue?

*Walking onward, the three enter the park and sit down to continue their discussion. A family with two young children plays nearby.*

Phil: It sounds as though you two at least agree now that both sides could benefit from more interaction. How do you think we can get there as a science?

Mira: Maybe the first step is to integrate modeling, or at least the mathematical foundations, into graduate training programs. No matter what your area of study, a better understanding of mathematical techniques will generate shared understandings—a foundation for formulation and evaluation, using the most general representational system available—giving students a tool to broaden their perspectives on different theories and theory testing. This would also provide a richer

foundation for understanding other sciences, like neuroscience and genetics.

Phil: You might need to go further than that, with more rigorous mathematics requirements for undergraduate psychology majors to help prepare them for that training. Every other science expects students to be proficient in advanced mathematics, so maybe psychology should, too. After all, the brain is the most complicated system we know of!

Ed: That's a good point. Better mathematical training could only be a good thing. And having a stronger math background would help students use more advanced statistical techniques as well. But not all psychology programs have faculty who do modeling, so who is going to teach those courses?

Mira: There are other ways to learn about specific modeling perspectives. For instance, the Cognitive Science Society and other organizations regularly have preconference workshops and tutorials on all types of modeling approaches. These offer great training opportunities.

Phil: I've also seen more intensive "summer school" programs, where you spend a week or so learning about the history and applications of particular modeling perspectives.

Mira: Right, or if you just want to know more about the math behind the models, you can look for courses in math, engineering, or computer science departments to give you the relevant foundational training.

Ed: It's good to know there are resources out there. But it sounds as though all of these require a huge time investment. Why can't I just pick up a journal article and make sense of a computational model?

Phil: Maybe that's what the modeling side should work on—finding the best way to communicate with a nonmodeling audience. There are ways to describe the math behind your model that can be accessible, even if the audience doesn't really understand the equations.

Mira: A few modelers have done just that, writing articles or books designed to be primers on a particular type of modeling.<sup>3</sup> They often include software or information on the Web to help guide a novice through using the model.

Phil: That still requires an investment from the reader, but I suppose that's unavoidable. It sounds as though it doesn't include many styles of modeling, though. Hopefully these attempts at connecting with new potential modelers will become more common and accessible.

Ed: I'm glad you think there's still work to be done from the modeling side. It seems as though modelers are happy to just keep working with no real connection back to our experiments.

Mira: Well, that may be the case for some, but I think most modelers want their work to inform other researchers' experiments. One challenge is that people expect to get the full "take" on the model from a single paper. I don't think that's an attainable goal. If you want to start understanding modeling papers, you have to change your strategy a little and track the arguments across several papers.



Ed: Truthfully, that's the case for empirical papers as well—to really understand the details, you have to ground yourself in the different ways people study the phenomena of interest.

Phil: Whether you use models or experiments, science is just hard! But isolating yourself from progress doesn't make it any easier. I mean, take that family over there. Neither of you have told me how you are going to get to that level of complexity. Lewis Thomas once wrote an essay listing the seven wonders of the modern world—it was a challenge put to him in the form of a dinner invitation which he, interestingly, declined.<sup>4</sup> Do you know what his Seventh Wonder of the modern world was?

Ed: No, what? Lasagna?

*Mira and Phil look annoyed . . .*

Ed: Sorry, I'm getting hungry.

Phil: Thomas's Seventh Wonder was the development of a human child. How, my brilliant companions, are you going to explain *that*?

*Phil points to the children playing nearby.*

Mira: It's like O'Keeffe said: Details are confusing.

Ed: They certainly are, but let's not drown in the details. That's what experiments are all about: controlling some of those details, making the situation reliable, reproducible . . . analyzing behavior.

Mira: Well, that's what modeling is all about too—controlling the details, testing your assumptions, recreating behavior to analyze what makes those children tick.

Phil: There's some common ground—a similar strategy for handling a huge challenge. That's your basis for working together.

Mira: You're right. When we get back, you can tell me about some of your new data, Ed. I've been working on a model that might be relevant to your area of research.

Ed: That sounds good; maybe you can help me design some new experiments—I'm a bit stuck on a few effects that don't make sense.

Phil: Now that's what I call dialogue—modeling and empirical approaches are strongest when they work together to “select, eliminate, and emphasize” to “get at the real meaning of things.”

Ed: Great. Now I need to act on this hunger!

Mira: And I need a beer! Let's head back to the O'Keeffe museum. I hear there's a good restaurant there . . . and clearly there are one or two things we might learn from her.

*The three walk back down the path into town.*

## RECOMMENDED READINGS

- Churchland, P. S., & Sejnowski, T. J. (1992). *The computational brain*. Cambridge, MA: MIT Press.
- Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press.
- Mareschal, D., Johnson, M. H., Spratling, M., Thomas, M. S. C., & Westermann, G. (2007a). *Neuroconstructivism: Vol. I. How the brain constructs cognition*. Oxford, UK: Oxford University Press.

Mareschal, D., Sirois, S., Westermann, G., & Johnson, M. H. (Eds.). (2007b). *Neuroconstructivism: Vol. II. Perspectives and prospects*. Oxford, UK: Oxford University Press.

Mareschal, D., & Thomas, M. S. C. (2007). Computational modeling in developmental psychology. *IEEE Transactions on Evolutionary Computation*, *11*, 137–150.

McClelland, J. L., & Jenkins, E. (1991). Nature, nurture, and connections: Implications of connectionist models for cognitive development. In K. VanLehn (Ed.), *Architectures for intelligence: The twenty-second Carnegie Mellon symposium on cognition* (pp. 41–73). Hillsdale, NJ: Erlbaum.

<sup>3</sup>McClelland, J. L., & Rumelhart, D. E. (1988). *Explorations in parallel distributed processing*. Cambridge, MA: MIT Press.

McClelland, J. L., Spencer, J. P., & Thomas, M. S. C. (Eds.). (2009). *Toward a unified theory of development. Connectionism and dynamic systems theory re-considered*. New York: Oxford University Press.

<sup>3</sup>McLeod, P., Plunkett, K., & Rolls, E. T. (1998). *Introduction to connectionist modelling of cognitive processes*. Oxford, UK: Oxford University Press.

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.

<sup>2</sup>Richardson, F. M., & Thomas, M. S. C. (2006). The benefits of computational modelling for the study of developmental disorders: Extending the Triesch et al. model to ADHD. *Developmental Science*, *9*, 151–155.

Schöner, G. (2009). Development as change of system dynamics: Stability, instability, and emergence. In J. L. McClelland, J. P. Spencer, & M. S. C. Thomas (Eds.) *Toward a unified theory of development: Connectionism and dynamic systems theory re-considered*. (pp. 25–47). New York: Oxford University Press.

<sup>3</sup>Shultz, T. R. (2003). *Computational developmental psychology*. Cambridge, MA: MIT Press.

Shultz, T. R. (in press). Computational models in developmental psychology. In P. D. Zelazo (Ed.), *The Oxford handbook of developmental psychology*. New York: Oxford University Press.

Shultz, T. R., & Sirois, S. (2008). Computational models of developmental psychology. In R. Sun (Ed.), *The Cambridge handbook of computational psychology* (pp. 451–476). New York: Cambridge University Press.

Spencer, J. P., & Schöner, G. (2003). Bridging the representational gap in the dynamic systems approach to development. *Developmental Science*, *6*, 392–412.

Sun, R. (Ed.). (2008). *The Cambridge handbook of computational psychology*. New York: Cambridge University Press.

<sup>4</sup>Thomas, L. (1983). *Late night thoughts on listening to Mahler's Ninth Symphony*. New York: Viking.

<sup>2</sup>Thomas, M. S. C., & Karmiloff-Smith, A. (2002a). Are developmental disorders like cases of adult brain damage? Implications from connectionist modeling. *Behavioral and Brain Sciences*, *25*, 6.

<sup>2</sup>Thomas, M. S. C., & Karmiloff-Smith, A. (2002b). Modeling typical and atypical cognitive development: Computational constraints on mechanisms of change. In U. Goswami (Ed.), *Blackwell handbook of childhood cognitive development* (pp. 575–599). Malden, MA: Blackwell.

<sup>2</sup>Thomas, M. S. C., & Karmiloff-Smith, A. (2003). Connectionist models of cognitive development, atypical development and individual differences. In R. J. Sternberg, J. Lautrey, &

- T. Lubart (Eds.), *Models of intelligence for the next millennium* (pp. 133–150). Washington, DC: American Psychological Association.
- <sup>2</sup>Triesch, J., Teuscher, C., Deak, G., & Carlson, E. (2006). Gaze following: Why (not) learn it? *Developmental Science*, *9*, 125–147.
- <sup>1</sup>Urbain-Jean-Joseph, L. V. (2007). *Encyclopedia Britannica*. Retrieved May 18, 2010, from <http://www.britannica.com/EBchecked/topic/333461/Urbain-Jean-Joseph-Le-Verrier>
- Westermann, G., Sirois, S., Shultz, T. R., & Mareschal, D. (2006). Modeling developmental cognitive neuroscience. *Trends in Cognitive Sciences*, *10*, 227–232.