Chapter 3

Computational Modeling of Learning and Teaching

Michael S. C. Thomas and Diana Laurillard

Introduction

In this chapter, we consider how computer models are being used to advance our understanding of learning within education. By the notion of 'model', we mean a simplified representation and implementation of a phenomenon that captures the key theoretical principles of its operation. In the case of computer models, the implementation of the model is in the form of a computer program.

In educational neuroscience, computers are employed as model systems in two different but related ways. First, computers are used to understand the *cognitive mechanisms* that underlie the learning process. Second, computers are used as *teaching tools* that model the interaction of the teacher with the learner; these tools can take the form of *intelligent tutoring systems* or *adaptive microworlds*. Cognitive models and computer-based teaching tools are related in that both require an understanding of the learner and the way in which particular sets of experiences and kinds of feedback can advance the learner's knowledge. Indeed, some computer programs have been employed both as cognitive models and as the basis of intelligent tutoring systems (see the later example of ACT-R). The two uses differ in that for cognitive models the target of the model is the learning process taking place within the child or adult, while for the teaching tools the target of the model is the behavior of the human teacher and his or her interactions with the learner.

Educational Neuroscience, First Edition. Edited by Denis Mareschal, Brian Butterworth and Andy Tolmie.

Both approaches are concerned with *individual differences*. Cognitive modeling seeks to understand the causal mechanisms – be they of genetic or environmental origin – that lead to differences in learning outcomes. For the teaching tools, intelligent tutoring systems seek to tailor the knowledge and tasks that the individual learner must revisit to acquire a specific domain, given the successes and failures he or she has exhibited on previous tasks (an *instructivist* pedagogic model). Adaptive microworlds contain a model of how a user can interact with a domain to construct an object or event; the program then adapts the difficulty of the task goal to the individual learner's performance on the last task (a *constructionist* pedagogic model).

In terms of the dialogue between education and neuroscience, cognitive modeling demonstrates one active interface between the disciplines. *Artificial neural networks* are used to build models of learning based on the computational principles observed in actual neural circuits. To the extent that these cognitive models are successful, they will uncover the nature of knowledge representations in the learner, as well as the sequence of knowledge development. In the future, they may thus inform computer-based teaching programs.

Our review of computational methods is structured as follows. In the first half of the chapter, we discuss the use of computers as cognitive models. We discuss why building explicit, implemented models is an effective way to advance our understanding of the nature of learning, and we summarize the general principles and aims of building such models. We discuss two of the main approaches to building models of the cognitive system, the *symbolic* and *subsymbolic* approaches. The subsymbolic approach makes widespread use of artificial neural network models. We consider the way in which such models have been inspired by neuroscience research into the principles of computation in the brain and how artificial neural networks have certain properties that make them well suited to modeling cognitive mechanisms of learning. By way of illustration, we describe research in which a large population of artificial neural networks is used to simulate individual differences in rates of language development, and which addresses the specific educational issue of why language delay, when diagnosed early in children, sometimes disappears of its own accord, but other times persists and requires intervention. We finish the first half of the chapter with a look to the future, in how computational models of learning may advance beyond current limitations to enrich our understanding of the plethora of phenomena that represent the educational experience.

In the second part of the chapter we discuss the use of computers as teaching tools. We summarize the key properties of educational models of teaching and learning. We then describe the components of two different digital learning environments, intelligent tutoring systems and adaptive microworlds. In each case, we make reference to the pedagogical theories that they embody.

We illustrate the second of these approaches with an example of an adaptive microworld designed to address weaknesses in the understanding of number sense found in children with dyscalculia, a specific deficit in the learning of mathematical skills.

Finally, although the approaches of cognitive modeling and computer-based teaching tools are conceptually related, it will become apparent that they sometimes use terms and ideas in different ways, and approach their account of the process of learning from different directions. We use the example of *feedback* to highlight where the approaches line up and where they do not. We finish with a summary of the main points of the chapter.

Computational Models of Cognition

The use of models to understand mechanisms of learning

At the heart of education lies the concept of learning – facilitating change in knowledge and abilities over time. Computational systems can provide models to understand the cognitive processes underlying learning. Such models are formal systems that track the changes in information processing that take place as a behavior or skill is acquired. Models are generally implemented as psychologically constrained computer simulations, which learn tasks such as reasoning, concept formation, and language and literacy skills.

To date, models have mainly been applied to the study of cognitive development, focusing in particular on how transitions are achieved from one level of competence to the next via experience and/or maturation. Models have been used to probe questions such as how much 'preprogrammed' or innate knowledge exists in the infant mind, and how the sophistication of reasoning can increase in children with age and experience. Education differs from cognitive development with respect to the assumed learning environment. Whereas development reflects what the learner discovers through interaction with his or her natural physical and social environments, formal education concerns the acquisition of knowledge and skills accumulated by a culture across many generations. The extended, structured learning environments provided by education are powerful enough to sculpt new brain systems, such as those involved in reading and mathematics. Education is built on the foundation of cognitive development, and, with respect to the innate component of cognitive development, education both takes advantage of inherited mechanisms of brain plasticity and is also constrained by the types of knowledge representation that the mind can support.

Computer models have proved invaluable tools to help developmental psychology shift from a descriptive science into a mature explanatory science

(Mareschal & Thomas, 2007). A descriptive science deals in summaries of observations of what happens – in the case of development, the abilities that children exhibit at different ages. An explanatory science deals in mechanism – revealing the underlying causal processes that give rise to the observed behavior. The construction of computer models has aided the shift to an explanatory science of development because, when researchers have to translate their underlying theories into explicit computer models, they must specify precisely what is meant by the various terms in the causal theory. Terms such as *representations*, *symbols*, *variables*, and *learning* must have an exact definition to allow implementation. The degree of precision required to construct a working computer model avoids the possibility of arguments arising from the misunderstanding of imprecise verbal theories. For example, the idea of "attention" conveniently summarizes a cluster of human behaviors. Yet it is another thing to build a processing system that has the ability to select certain sorts of information for enhanced processing. How does the system select what information to attend to, and how is the processing of attended (and unattended) information altered? There is no longer room for vagueness when building a working model of the process.

General principles and aims of computational models of learning

A cognitive computational model of the learning process usually comprises the following four elements. First, there is a *computational system*, of which there are different varieties. A computational system is a mechanism that acquires, stores, and manipulates information, which it can use to drive behavior. Second, there are *representations of information*. Some of these representations are specified by the modeller and correspond to the information supplied to the system by its simulated environment, along with the output responses that are required to generate the requisite behavior (e.g., behaviors such as naming a word, giving the number that is the answer to a mathematic problem, or inferring the intended meaning of an analogy). Other representations of information may be developed by the computational system itself during the learning process – this is the information the system needs in order to generate the appropriate behavior, given the input with which it is supplied. Third, the model has a *learning algorithm*. This is a process by which the system alters its internal structures to improve its performance on the target problem, given feedback on its current performance (see the later section on Feedback for more detail on types of feedback in computer learning systems). Fourth, there is a *training set*, corresponding to the problem domain that the model must learn. Training sets can be supplied externally by the modeler. Alternatively, the modeler can construct an artificial microworld in which the model generates its own training set by its behavior and subsequent experiences in the artificial microworld. The microworld may contain other individuals (or *agents*) with which the model can interact. In addition, modellers sometimes test the learning system on a *generalization set*, corresponding to novel problems within the target domain. Success on a generalization set ensures that the model has learnt the general principles of the problem domain, rather than just memorizing the individual items in the training set.

There is a wide range of possible computational systems that can serve as models for learning in the cognitive system (see, e.g., Mitchell, 1997, and Sun, 2008, for introductions to different types of machine learning system). These include concept learning, decision tree learning, artificial neural networks, Bayesian (probabilistic) learning, instance-based learning, genetic algorithms, and reinforcement learning. One broad distinction that has characterized computational models of development in particular is that between *symbolic* and *subsymbolic* models. Researchers using *symbolic* models maintain that cognition is best characterized as a rule-governed physical symbol system, such as a conventional computer program. In this view, cognitive development consists in the construction and modification of mental rules. By contrast, researchers using *subsymbolic* models view cognition in terms of a highly interactive dynamic system, such as an artificial neural network. An artificial neural network contains simple processing units, each with an activation level analogous to the firing rate of a neuron. The processing units are wired together in networks with weighted connections. The strength of the connection between any two units determines how much the activity of one unit can affect the subsequent activity of the other. Networks can learn tasks as transformations between different activation states. They do so by gradually changing the strength of the connection weights to produce the appropriate activation states, based on information from the environment. In this type of system, the causal entities are not rules but continuous activation states distributed across the network, states that sometimes cycle over time. Such networks do not operate as physical symbol systems, or at best approximate them in certain narrow circumstances. That is, the networks sometimes show *rule-following* behavior, without being *rule driven*. In a subsymbolic framework, both learning and development consist in the continuous tuning of the underlying parameters of the cognitive system (for a network, these parameters are the connection weights), in order to bring the responses of the system closer to the desired behaviors.

In symbolic models, encoded knowledge can be clear and transparent. In some cases, it corresponds to a rule-based description of the observed behavior (e.g., the rules for addition and subtraction in a system that performs arithmetic). Rules can be very powerful in producing a range of complex behaviors (e.g., the rules of grammar can be used to generate an infinite variety of sentences). In subsymbolic models, it can be less obvious what knowledge is encoded, since dynamic patterns of activation across a network of simple processing units may not be directly relatable to behavior. Given the task presented at input, the activation patterns must simply serve to provide the correct answer at output. Most symbolic models have emphasized the transparency of knowledge representations involved in cognitive development at the expense of implementing mechanisms for the acquisition of new knowledge and abilities. That is to say, it has proved difficult to understand how systems that run according to rules learn new sets of rules, particularly rules that are more sophisticated than those previously operating in the system. In contrast, most subsymbolic models have emphasized the specification of a learning mechanism for incrementally improving behavior on a problem domain, at the expense of the transparency of the knowledge representations that the system acquires. In short, symbolic models have more transparent workings but do not readily capture the process of learning, while subsymbolic models are good at learning but their internal functioning is more opaque.

For researchers in psychology who investigate the nature of the mind and the nature of development, the different learnabilities of, respectively, rules and activation patterns has led to a conundrum. To the extent we think that the human mind needs complicated and densely structured mental representations to deliver a cognitive skill, it is hard to fathom how these representations are learned. There are three ways out of this conundrum (Mareschal & Thomas, 2006). *Either* complex behavior is generated by representations that are in large part innate (i.e., not learned at all; Chomsky's theory of universal grammar would provide one example of such an approach in the field of language development – this is the idea that humans easily acquire complex language because we are born with a blueprint for grammar), *or* we do not yet understand the full repertoire of learning mechanisms available to the human mind (that is, somehow complex mental representations are learnable in a way we have not yet understood), *or* we are currently overestimating the complexity of the representations that the human mind needs to generate its complex behaviors.

How does a researcher tell if he or she has constructed a good model of a certain cognitive process? Certainly the model should reproduce the behaviors observed in people. A model of learning should be able to learn the skills that people can learn, be unable to learn the skills that people cannot learn, and should exhibit the same trajectory of learning, including the same kinds of error that people make when they are going through the process of learning.

However, the evaluation of a model can be more nuanced. A model must be constrained by empirical psychological evidence. No "unrealistic" components or processes should be included to make the model work. For instance, a model of learning to read should encode written words (orthography) in the way we believe children do, and encode spoken words (phonology) in the way we believe children do. This does not necessarily mean it should have a visual system to recognize the written words and a mouth to pronounce the spoken words. It does mean that the information it receives about written words should be similar to what we know the visual system extracts from the page, and the information it outputs should respect the structure of spoken sounds, in terms of articulatory features. Next, the model should be exposed to the same kinds of learning experience that children are, such as the juxtaposition of written and spoken letters and words. Moreover, the model should contain a learning mechanism that we think could plausibly operate in the mind. However, a model will necessarily contain simplifications (such as the lack of eyes to read words and a mouth to speak them). It is, after all, a model intended to capture the key principles of the process under study, rather than duplicating the system in every regard.

A cognitive model should be evaluated according to several criteria. (i) Does the model simulate human behavior in the target domain? (ii) Does it help *explain* why the human behavior occurs? (This requires that the modeller understands why the model works!) (iii) Is the model successful in simulating the target behavior due to its key design principles – the theory that the model embodies – or due to its design simplifications (fixes that the researcher has used to get the model to work)? If the answer is the former, the model can be viewed as a demonstration of the viability of the theory it embodies. (iv) Does the model explain a range of behaviors rather than just one (i.e., is it parsimonious)? Finally, (v) can the model generate any new behaviors observable under different conditions (e.g., in novel situations, or perhaps when the model is damaged in certain ways); that is, does the model generate novel predictions that can be corroborated by subsequent psychological experiments with people?

Examples of symbolic and subsymbolic cognitive models: ACT-R and artificial neural networks

One widely used symbolic model of cognitive processing is called ACT-R (Anderson, 2007; Anderson & Lebiere, 1998). ACT-R stands for Adaptive Control of Thought – Rational. The system is intended as an overarching cognitive architecture, capturing how the whole mind works. The key components of ACT-R are inspired by processing distinctions observed in the brain, and in particular the distinction between declarative memory (explicit facts and knowledge) and procedural memory (implicit knowledge and skills). The system has different specialized components that reflect this distinction.

At the heart of ACT-R lie procedural IF–THEN rules. If a certain set of conditions hold, then a certain behavior is produced. The system has working memories or buffers of current knowledge, reflecting both new inputs and its previous processing states. The system's library of rules competes to find the rule that most closely matches the current state of the buffers. The winning rule then produces the subsequent behavior. When the rule is executed, it may then alter the knowledge in the buffers, which triggers the next winning rule, and so forth. In ACT-R, cognition proceeds as a succession of rule operations. ACT-R is symbolic in the sense that it contains discrete variables and syntactic rule-based operations, although the performance of the system may also depend on what is referred to as 'subsymbolic quantities', such as the strength of productions, and the base level of activation for a chunk of knowledge (Anderson & Schunn, 2000). ACT-R has been used to model a variety of cognitive processes including memory, attention, executive control, language, and problem solving.

Two further points are of note. First, as we saw in the previous section, cognitive models that rely on rule-based representations struggle to find a ready means to learn new rules. In keeping with this, ACT-R has found relatively little application to modeling cognitive processes involved in learning and development. Second, as we shall see in the second half of the chapter, ACT-R has nevertheless been successfully used as the engine on which intelligent tutoring systems are based.

One widely used subsymbolic model of cognitive processing is the artificial neural network (Rumelhart & McClelland, 1986; Spencer, Thomas, & McClelland, 2009; Thomas & McClelland, 2008). As we have seen, artificial neural networks are abstractions that capture some of the key properties of computation carried out in neural circuits. The brain comprises a large number of neurons that electrically signal to each other via highly connected networks. Artificial neural networks contain simple processing units, each with an activation value, and a network of connections through which the activity of each processing unit can influence the activity of other processing units. A key property of neural systems is that they are adaptive. The strengths of the connections between units can be altered incrementally to bring the network's output closer to the desired behavior given its inputs. Many of the artificial neural network learning algorithms are based on *Hebbian learning* (Hebb, 1949), the principle that "units that fire together should wire together". In other words, if two units are firing at the same time, they are probably both involved in performing the same computation; therefore, the connection between them should be strengthened so that they encourage each other to fire when the units receive equivalent inputs the next time around.

As they learn, artificial neural networks are able to develop their own internal representations of knowledge over their banks of processing units. Artificial neural networks are subsymbolic in the sense that information is encoded as continuous patterns of activation. As we saw in the previous section, these representations of knowledge are not necessarily transparent in what information they contain. When used as cognitive models, artificial neural networks therefore emphasize the learning of new abilities at the expense of the transparency of the exact knowledge that is being acquired. Artificial neural networks have been used to model a wide range of developmental phenomena, including perceptual learning and object-oriented behaviors in infants, language and literacy acquisition in children, and the development of reasoning in children (Elman et al., 1996; Mareschal & Thomas, 2007). In addition, these models have begun to provide a platform to understand development and individual differences within the same explanatory framework: that is, why children of the same age should differ in their abilities, and the respective role of genetic variation and environmental variation in generating these differences (see, e.g., Thomas, Baughman, Karaminis, & Addyman, 2012a; Thomas, Forrester, & Ronald, in press; Thomas, Knowland, & Karmiloff-Smith, 2011). In the next section, we outline an example of an artificial neural network model of language development that illustrates these points, as well as more general principles of the construction and evaluation of cognitive models.

An example of cognitive modeling in educational neuroscience: individual differences in language development

One central concern of cognitive modeling within educational neuroscience is the issue of *individual differences*. What are the genetic and/or environmental causes of differences in learning outcomes? An understanding of these causes may help us optimize learning outcomes for children with different abilities or from different backgrounds. So why do children learn at different rates? While it is long established that individual differences can have both environmental and genetic causes, there is a lack of detailed cognitive modeling that stipulates how these influences unfold in generating behavior. For example, there has been a recent renewal of interest in how *socioeconomic status* (SES) affects children's development and their educational outcomes (Hackman & Farah, 2009). However, SES is associated with many differences in children's physical and social environments, and it is unclear which causal pathways are responsible for the observed variation in developmental outcomes. One possibility considered within the field of language development is that SES is associated with differences in the quality and quantity of the *information* (in this case, language input) to which the child is exposed (see Chapter 6, this volume, for a wider review of relevant issues in language development). In lower-SES families, there is simply less language directed towards the child (Hart & Risley, 1995). Targeting the role of language input more precisely, Huttenlocher, Vasilyeva, Cymerman, and Levine (2002) found that the proportion of complex sentences produced by teachers

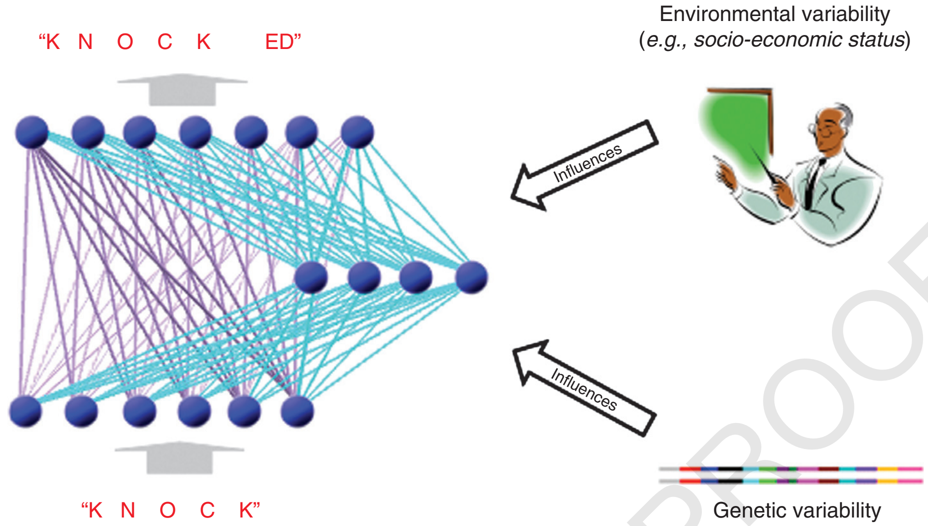

Figure 3.1 An example of a cognitive model of one aspect of language acquisition, based on an artificial neural network. The model learns to form the past tense of English verbs. The model simulates a population of learners who show individual differences in their learning due to variations in learning abilities and environments.

predicted 18% of the variance in the improvement in children's performance on a syntax comprehension task over a year of preschool. Differences in language input appear to be an important (if not sole) contributor to differences in rates of language development across the SES range of developed countries.

Thomas et al., (in press) built a cognitive model of the acquisition of one aspect of English grammar, the English past tense. The model was designed to capture the range of developmental trajectories of a large population of simulated children, and incorporated individual differences from both intrinsic sources (i.e., the power of the learning mechanism each child had) and extrinsic sources (the quality of the environment to which the child was exposed, by hypothesis influenced by SES). Population-level modeling is a relatively recent innovation, which has become possible through increases in computational power that allow thousands of models to be run rather than just a few (Thomas et al., 2012a). The Thomas, Forrester, and Ronald model is illustrated in Figure 3.1. It comprised an artificial neural network, with a phonological representation of the English verb stem at input, along with information about the verb's meaning, and a phonological representation of the past tense at the output. The network was exposed to verb stem–past tense pairs for both regular and irregular English verbs (e.g., knock–knocked; think–thought), and underwent an extended developmental trajectory as it acquired this aspect of grammar. The model aimed to capture empirical data from Bishop (2005), which reported the effects of SES on the acquisition of English past tense for a sample of 300 six-year-old children. For these children, regular past tenses were produced more accurately than irregular past tenses. SES explained around 1% of the variation in children's regular-verb performance but around 5% of the variation in irregular-verb performance.

The model succeeded in capturing the predictive power of SES that was observed in the empirical data, and in particular the greater predictive power of SES on irregular than regular verbs. The model suggested that the empirical data were best captured by relatively wide variation in learning abilities of children and relatively narrow variation in (and good quality of) environmental information. The model served as a demonstration of the viability of the theory that variations in language input are one causal pathway through which SES may operate. In addition, the model generated a novel prediction not previously considered by any researcher: it predicted that SES should reliably predict *gifted* performance in children (e.g., whether a child would fall in the top 10% of the population) but not *delayed* performance (e.g., whether the child would fall in the bottom 10%). This surprising prediction was subsequently borne out by the Bishop (2005) data set.

In a follow-up paper, Thomas and Knowland (submitted) used the model to focus on an important current issue in the field of developmental language disorders. It is optimal to diagnose language delay in children early on (say, at three or four years of age) in order to optimize chances for effective intervention. However, of the children diagnosed with delay at this young age, over half subsequently have their delay resolve of its own accord without the need for intervention. Early intervention is therefore optimal but risks treating children in whom (potentially costly) intervention is unnecessary. Thomas and Knowland used the population-modeling technique to focus on the issue of the outcome of early-diagnosed delay. Figure 3.2 shows how simulated children were diagnosed with delay (here defined as falling more than one standard deviation below the population mean) at five different developmental time points. Figure 3.3 summarizes the number of simulated children with delay at each time point – and confirms that in the model, too, the number of cases of delay fell by over half from the first to the fifth time point. Thomas and Knowland then examined the cases of persisting early delay and resolving early delay in more detail, tracing individual trajectories. Several of these trajectories are displayed in Figure 3.4. There were in fact four patterns of development: typical development, persisting delay, resolving delay with low–average final outcome, and resolving delay with good final outcome.

At this point, the investigation focused on what properties differed between these four groups, in terms of both the environmental conditions and the learning properties of the artificial neural networks. Artificial neural networks have a

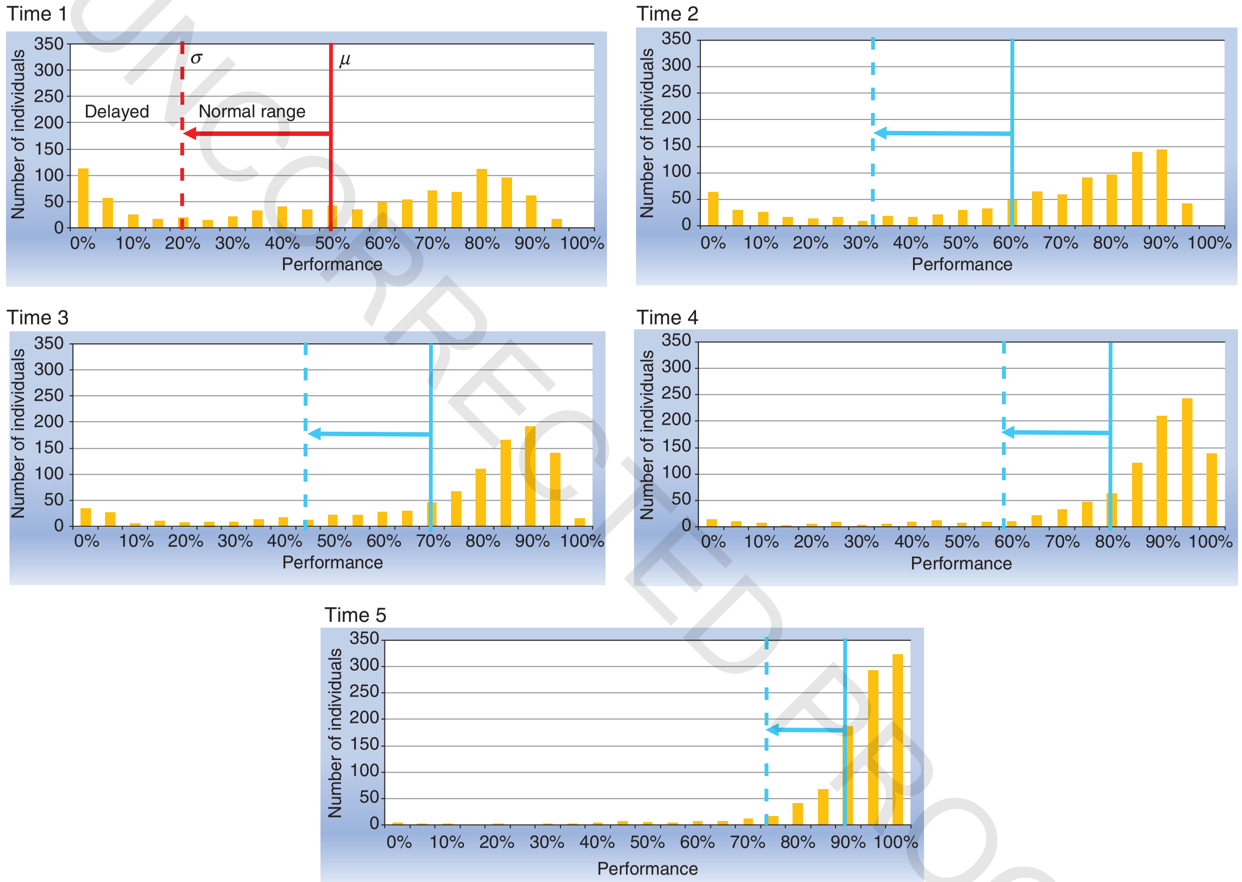

Figure 3.2 Frequency distributions of the number of simulated children performing at each level of accuracy, at five points in development. Individuals exhibiting delay are diagnosed at each time point as falling more than one standard deviation (σ) below the mean (μ). The proportion of delayed individuals reduces over time, suggesting that some cases of delay must resolve.

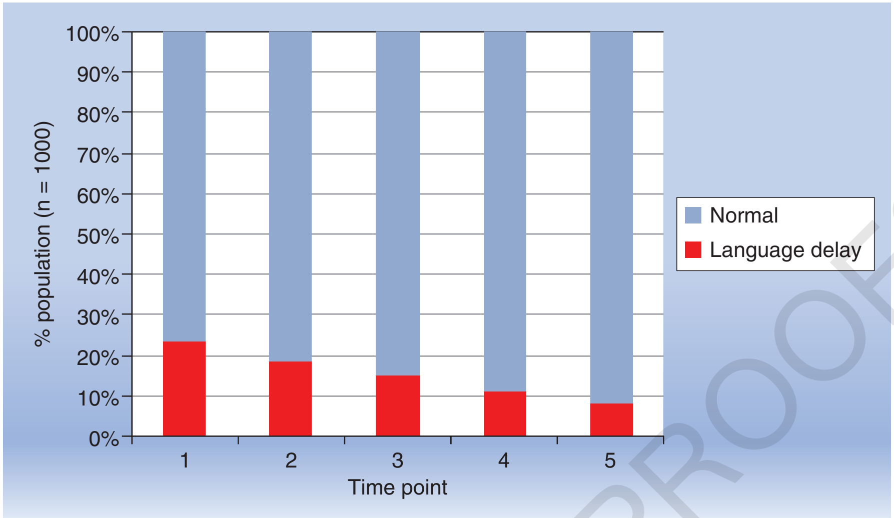

Figure 3.3 The proportion of the simulated population showing developmental delay at each time point.

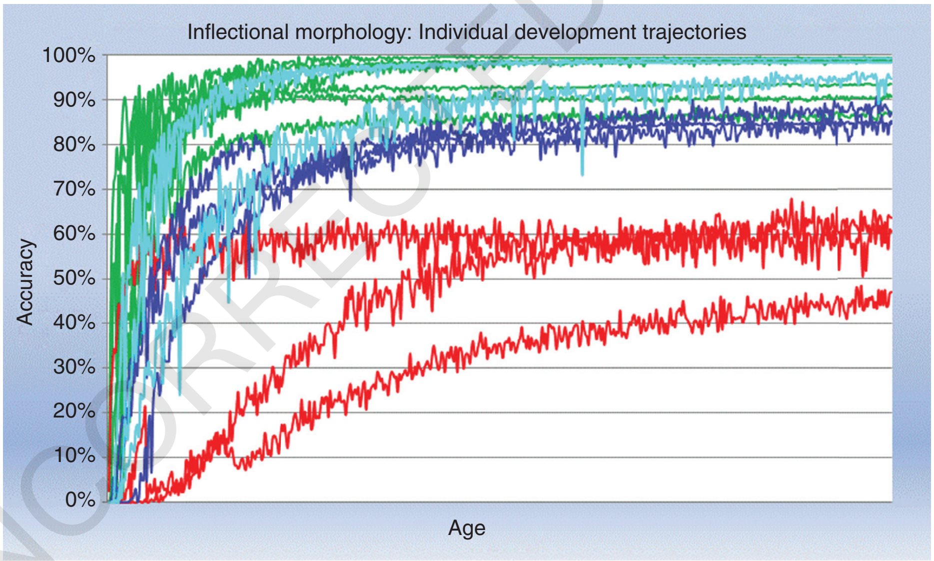

Figure 3.4 Example developmental trajectories showing four different patterns: typical development (green); persisting developmental delay (red); resolving delay with low–average final outcome (turquoise); and resolving delay with good final outcome (dark blue).

range of properties that can alter how much information they can learn and how quickly they can learn it. We can refer broadly to these properties as the *capacity* and *plasticity* of the system. Capacity is affected by, for example, how many units and connections there are inside the network. Plasticity is affected by how quickly

Table 3.1 Computational causes of the four types of developmental trajectory.

		Computational Plasticity	*Computational Capacity*	*Environment*
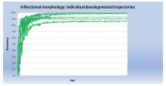	Normal	Okay	Okay	Okay
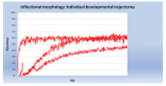	Persistent deficit	Okay / low	Poor	Okay
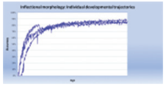	Resolving low-normal	Low	Okay	Poor
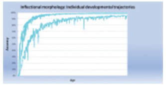	Resolving normal	Low	Okay	Good

connection weights can alter their strength when the network is required to change its behavior. Combined with the quality of the learning environment, the properties of capacity and plasticity allowed Thomas and Knowland to distinguish between the four groups.

The results are shown in Table 3.1. When development was delayed, this could arise from limitations in capacity or plasticity. Persisting delay was associated with limitations in capacity, while environmental conditions seemed unimportant. Resolving delay was associated with limitations in plasticity, i.e., it took these learners longer to adapt and they learned at a slower rate. Crucially, for this second group, the quality of the environment then predicted the final level of performance after the delay had resolved. A rich environment (high SES) was associated with good final outcome, while poorer environments (lower SES) were associated with low–average final outcome. Once more this was a prediction that had not arisen from any previous theory, and once more the prediction was confirmed in the data set of Bishop (2005), for a large sample of British children who were diagnosed as at risk of language delay aged four, and whose past-tense abilities were then tested at age six. The next step of this research program is to isolate behavioral or neural markers that can distinguish low capacity from low plasticity in the early diagnosis of language delay, and so narrow the focus of language interventions.

This example of a cognitive model illustrates several of the design principles we introduced earlier. The model was aimed towards capturing the acquisition of a specific task domain. It comprised a computational system – an artificial neural network; representations of information – phonological encodings of English verb stems and past tenses, along with information about their meaning; a learning algorithm – in this case backpropagation, a supervised learning algorithm (see later, in "Feedback"), which is itself a variant of Hebbian learning (Thomas & McClelland, 2008); and a training set – pairs of English verb stems and their associated past tenses. The model was evaluated according to how well it simulated real empirical data – children's ability to learn regular and irregular English past tenses, and the influence of SES on individual differences in this ability. It was evaluated by the extent to which it achieved this success via its design principles rather than simplifications – in this case, the theory being implemented was the idea that SES corresponded to differences in the richness of the language information in the environment to which the child was exposed, against a background of individual differences in learning ability. The model was further evaluated against its ability to explain a range of phenomena – in this case, both normal language development, gifted language development, and delayed language development. Moreover, the model was evaluated against its ability to produce novel empirical predictions, which were then borne out by real empirical data – in this case, the model predicted differential effects of SES on gifted versus delayed language development, and on different types of delayed development, both of which were subsequently confirmed.

The broader perspective: neuroconstructivism and education

In the previous section, we described a cognitive model aimed at capturing the influence of variations in the environment on children's language development. This model used an artificial neural network as its basic computational system. Such networks embody principles derived from neuroscience, and in this way cognitive-level modeling provides a link between neuroscience and the overt aspects of children's behavior that are the central concern of education. Nevertheless, this is only one model, targeting a fairly circumscribed aspect of language acquisition. It is also important to consider the broader perspective that this theoretical approach implies, and its potential impact on educational theories.

The idea that neuroscience principles should influence cognitive-level theories of learning amounts to the proposal that the way a cognitive system (the "mind") is implemented in the brain makes certain ways of thinking and learning easier and others harder. One cannot, therefore, derive a theory of cognition without reference to how the brain delivers cognition. *Neuroconstructivism* is one theoretical approach that has recently attempted to flesh out this idea (Elman et al., 1996; Mareschal et al., 2007). In particular, neuroconstructivism builds on the Piagetian view that development corresponds to the progressive elaboration in the complexity of mental representations via experience-dependent processes, enabling new competences to develop based on earlier, simpler ones (*constructivism*). *Neuro*-constructivism also incorporates recent theories of functional brain development, proposing that the increase in representational complexity is realized in the brain by a progressive elaboration of functional cortical structures (see Sirois et al., 2008; Thomas et al., 2008; Westermann et al., 2007; Westermann, Thomas, & Karmiloff-Smith, 2010).

One might well ask, then, which principles of brain function should influence the formation of cognitive theory? Here are five such principles. (1) The brain uses *partial representations of knowledge*: whole concepts are rarely used, only the dimensions of knowledge required to drive particular behaviors relevant to the current context of action. Whole concepts may, indeed, be rarely acquired. (2) *Contextualisation*: mechanisms always act in context – genes operate in the context of other genes, neurons operate within the context of a neural network, brain regions operate within the context of a set of brain regions, the brain operates within the context of the body, and the individual operates within the context of a culture and society. (3) *Timing*: the timing of developmental events can be crucial, so that the same event happening at different times can have different consequences. (4) *Emergent specialization* (and brain localization): systems become more specialized with development, tuning their function to particular domains depending on experience. For example, within vision, dedicated systems for face recognition and written word recognition are experience-dependent specializations of an initially more general object recognition system. (5) Developmental events in the brain must be construed within the wider framework of *evolutionary developmental biology*: an adaptive framework informs the functions established during brain development. What has evolution designed the system to do, and what are the neural constraints fashioned into the structure of the brain that allow the individual to achieve that goal when the child is raised in a normal environment? How can these constraints respond to novel environments, such as the evolutionary novel (cultural environment) of literacy and numeracy?

These ideas are recent enough that their implications for educational theory have not yet been fully explored. In respect of *timing*, for example, research has begun to focus on what sensitive periods in brain development may mean for the timing of the delivery of educational curricula (e.g., Thomas, 2012; Thomas & Knowland, 2009). In some cases, this work has once more relied on the use of computational modeling to connect neuroscience principles to high-level behavior (Thomas & Johnson, 2006). However, it is likely that there are more deep-seated implications for education to be derived from the neuroconstructivist thesis. For example, the notion of *partial representations* of knowledge suggests that different dimensions of a concept are activated according to context. This means that knowledge may be intrinsically bound by context, including during its acquisition.[1] In turn, this implies that the acquisition of a full, abstract concept requires exposure to all contexts of its usage. To give a concrete example, a child may learn that 5 is a number that falls between 4 and 6; that 5 is the result of summing 1 and 4; that 5 is the result of dividing 10 by 2. But each of these reflects the use of the number 5 in a given context. The ultimate goal of learning is to acquire the decontextualized concept: to learn that 5 is just 5. Constrained by modes of brain function, the child will always begin by acquiring concepts in a perceptual and contextually bound fashion. This predicts that multiple contexts of presentation must be deployed to liberate concepts from the shackles of context and the sensorimotor conditions of their acquisition, in order to construct the abstract idea. The abstract idea is then applicable across a range of situations including ones that the child is yet to encounter.

The future of cognitive modeling in education

Although cognitive modeling is a powerful method to advance our theories of learning and development, there are a number of reasons why this approach is currently somewhat limited with respect to education. This is because many of the central phenomena in education are among the most psychologically complex – involving the social context of the classroom, the dynamics of the interaction between learner and teacher, and the combination of knowledge and motivation. Current models are limited for several reasons: because they are insufficiently complex – though, as we saw earlier, models must retain some degree of simplicity to serve their explanatory goals – and because models currently target individual cognitive systems or components within that system, which makes it hard to capture the dynamics of learner–teacher interaction, or the community phenomenon of a classroom. Finally, there also remain unresolved debates within the study of cognition: what do representations of high-level conceptual knowledge look like? How does meta-cognition work? How do emotions, rewards, and motivation mediate learning?

Nevertheless, one can sketch out a picture of how the cognitive modeling approach could contribute to education in the future. Its ultimate aim will be to optimize the timing, regimes, and contexts of learning by understanding

[1] See Thomas, Purser, & Mareschal (2012b) for a computational modeling treatment of this idea, and in particular the proposal that the importance of language in problem-solving is that it allows the individual to bring to bear information that is not suggested by the individual's immediate context.

mechanistic principles of how the brain acquires, consolidates, and abstracts knowledge. It will contribute an understanding of how representations of knowledge form in the learner, how learners interact to develop a shared understanding in a classroom context, the role of attention and motivation in this process, how other factors may affect the learning properties of the brain (such as the role of sleep in consolidating memories, or of aerobic fitness in modulating brain plasticity), and the factors that may alter changes in brain plasticity with age, in order to optimize learning across the lifespan. Cognitive modeling will also contribute an understanding of cognitive mechanisms in the teacher: how the teacher represents the current state of the learner, how the teacher uses this knowledge to present information relevant to the task domain to the advance learner's knowledge, how the teacher generates feedback that is meaningful to the learner given the current state of knowledge, and how the teacher's emotional and motivational states modulate these processes. Finally, cognitive modeling will contribute an understanding of how each of these processes can vary across individuals, from the least to most gifted.

Cognitive modeling is a computational approach in neuroscience to understanding the process of learning and development. We now consider computational approaches in education that use our, albeit partial, understanding of the process of learning to develop a computational model of teaching. Here, there are two distinct approaches: intelligent tutoring systems and adaptive microworlds.

Computers as Teaching Systems

Educational models of teaching and learning

From the educational point of view, any teaching–learning environment includes a set of properties that must be present to make it possible for students to learn. The terminology is different, but the properties found in the educational literature all have their counterparts in the neuroscience account of learning developed above. These are referenced in parentheses in the following educational account, where a teaching–learning environment must specify a *learning outcome* (*requisite/target behavior*), a *method of assessment* of achievement of the learning outcome (*generalization set*, i.e., the test that learning has occurred), and a *set of task activities* (*training set*) the learner is to work through in order to achieve the outcome, where each activity consists of *learner actions* (*output responses*) to achieve a *goal* (*target problem*), and *meaningful feedback* (*feedback*) on the actions in relation to the goal. The teaching–learning environment may also give learners access to *peer learners* (*agents*). These properties are common both to conventional human/physical learning environments and to digital teaching–learning environments that attempt to emulate the human teacher. The digital version is a computer program in which the role of the teacher, the task activities, the goal, and the feedback all take place through the learner interacting with a either an "intelligent teaching system" or an "adaptive microworld".

The remaining correspondence is between the *computational system* that models learning in a cognitive system, and the *teacher's model of learning*, which is meant to correspond to the way their learners learn. In a digital teaching–learning environment in education, the model of learning draws on one or more theories of learning, such as constructivism, social constructivism, constructionism, and conceptual, experiential, collaborative learning (Laurillard, 2012), none of which have clear equivalents in the computational system models of learning above: concept learning, decision tree learning, artificial neural networks, Bayesian (probabilistic) learning, instance-based learning, genetic algorithms, and reinforcement learning. This is probably because the computational system models are relevant for learning the elements of knowledge or skills (such as the link between the verb stem and the past tense), whereas the educational theories operate at a different level of description of the curriculum (such as "the different forms of verb conjugation", "the laws of motion", or "the causes of the first world war"). However, the educational models of learning are much less well specified in terms of clearly agreed parameters and mechanisms, and lack synergy (Bransford et al., 2006).

One advantage of trying to construct a digital teaching–learning environment is that, just as with the cognitive modeling discussed above, the process demands specificity about exactly what the model of learning consists in. By combining the expectations on teacher and learner of all the current educational models of learning, it is possible to derive an explicit model, the "conversational framework," for describing the teaching–learning process in education (Frederickson, Reed, & Clifford, 2005; Laurillard, 2002). This model defines the process as a continual iteration between teacher and learner, between learner and peer, and between each participant's concepts and actions. Figure 3.5 illustrates the relationships between teacher, learner, and peer, and with the practice environment, real or virtual.

The iterations between the teacher's conceptual knowledge (TC) and the learner's conceptual knowledge (LC) are mediated by forms of representation such as language, symbols, diagrams, animations, and so on, through reading, listening, watching, debating, discussing, and so on, at the conceptual level. At the practice level, the teacher generates a modeling environment that emulates the world (TME), such as exercises, labs, fieldwork, and so on, in which the learner can use his or her practice repertoire (LP), in the form of goal-oriented actions, feedback, and revised actions, at the action level. Peer learning is represented in terms of discussions with peers about their concepts (PC) and exchanges of their practice outputs (PP). The within-participant iterations represent the generation of

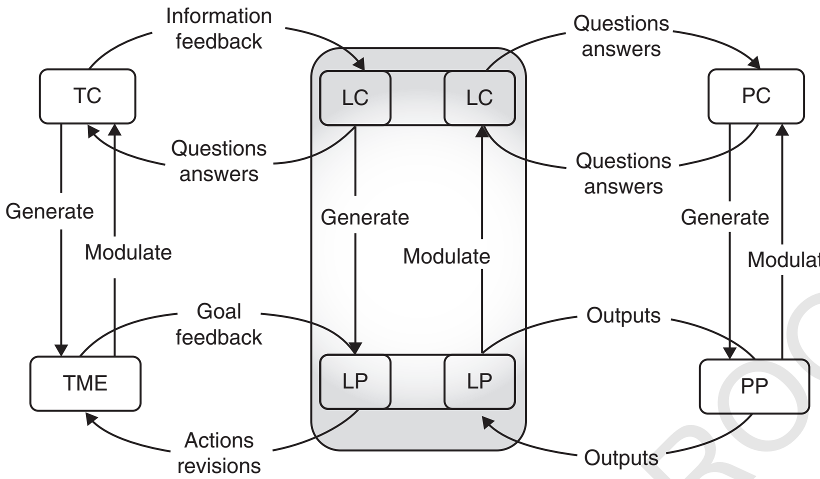

Figure 3.5 The conversational framework for individual and peer learning.

actions in the light of their current concepts and the modulation of their concepts in the light of feedback on actions. This includes the teacher generating a practice environment in the light of learners' discussions and questions, and modulating his or her own discourse in the light of learners' actions. The framework attempts to capture the dynamics of teacher–learner–peer interactions both within and beyond the classroom. The theories of learning currently used in education can each be mapped onto all or part of the framework.

In terms of the two types of model discussed above, the iterative nature of this model and the adaptive nature of the processes of generating actions and modulating concepts make it closest to the learning mechanism for incrementally improving behavior that is based on *subsymbolic* models. In the conversational framework, the knowledge representations that constitute the learner's concepts and actions are similarly opaque – they can only be detected in terms of what the learner produces as conceptual representations in their interactions with the teacher and peers, or the actions performed at the practice level.

In the next sections we look at the two main approaches to computational modeling in education: intelligent tutoring systems and adaptive microworlds.

Computational modeling of teaching and learning: intelligent tutoring systems

Intelligent tutoring systems (ITSs) use a symbolic model, and derive from theories of human information processing. A computational model of cognitive processing, such as ACT-R, enables an intelligent tutoring system to make inferences about what and how students are learning, as it monitors their outputs on a set of activities provided by the system (Sawyer, 2006).

The intelligent tutoring system has three components that are critical for aligning the student and teacher.

The *learning model* is the network of declarative chunks and production rules that generate responses to the problem set, i.e., the equivalent of the *computational system* in neuroscience, and the *teacher's model of learning* in education.

The *diagnosis* of the student's current needs is carried out by the system monitoring the student's behavior (*learner actions* or *output responses*) and comparing it with the behavior predicted by the model, in order to deduce which declarative knowledge chunks and production rules are being used. For example, if a student is making mistakes, it deduces the erroneous knowledge the student is using that would generate such mistakes.

The *teacher feedback* deals with the discrepancies diagnosed between the actual and predicted behavior, and is provided by the system in the form of help, scaffolding, and "dynamic instruction to repair the holes in their knowledge" (Anderson & Schunn, 2000, p. 19).

This account does not make any explicit reference to the other components of the cognitive model, although they are present in an intelligent tutoring system: the *task activities* or *training set* take(s) the form of the actions the student has to take to achieve the *goals* set by the system. So there is a good correspondence between the cognitive model and the computational modeling offered by the intelligent tutoring system.

Research on the intelligent tutoring system approach has some features in common with cognitive modeling, therefore, but was overtaken in the 1990s by the explosion of alternative forms of computer-based learning activities such as web resources, multimedia, user-generated content methods, and online communications technologies (Laurillard, 2010), and it has not progressed to having any major mainstream impact.

Computational modeling: adaptive microworlds

By contrast, an adaptive microworld has no explicit model of learning, being closer to the subsymbolic model of learning, and is built on Papert's ideal of "learning without being taught" (diSessa, 2001; Papert, 1980). A microworld is an interactive computational model of an aspect of the world, with its own constraints and assumptions, in which learners can experience the relevant concepts by using the program "to engage tasks of value to them, and in doing so … come to understand powerful underlying principles" (diSessa, 2001). It is adaptive (i) when it responds to the learner by showing the result of their actions in that world, and (ii) when it is designed to adjust the difficulty of the task in the light of the learner's current performance.

The microworld approach has fared better, and the fundamental idea of "constructionism" as a model of how learning can succeed is still current. Pioneered by Seymour Papert at MIT, and influenced by Piagetian psychology, "constructionism" embodies the theory that we learn complex concepts and ideas best by constructing representations that use them (Papert, 1980; Papert & Harel, 1991). The idea was applied to curriculum topics in science, but had most impact in school maths in the form of "Logo" for learning geometry, in many different countries (Hoyles & Noss, 2003). The fundamental concept of learning through construction is applicable across a wide range of discipline areas, at all levels of learning. The concept has now been implemented as "NetLogo", a modeling tool that enables learners to set up and investigate models of the behavior of systems such as population growth, electrical circuits, and climate change, wherever a computational model is possible (Gilbert & Troitzsch, 2005).

The computational modeling of pedagogy

To provide a teaching–learning environment, a computer program must include all the properties defined above, which a simple modeling environment such as Logo or NetLogo does not. A simple modeling environment is adaptive to students' actions, but not to their level of performance. Here, it is the teacher who monitors and sets up the task activities. By contrast, an *adaptive microworld* combines the task model with rules for monitoring student performance and adapts the difficulty of the next task according to the learner's needs. Figure 3.6 shows how the two contrasting computational environments for teaching and learning can be mapped onto the framework.

In Figure 3.6(a) the ITS uses the model of the teacher's modeling environment (TME) to generate the task goal (1, 2); the learner uses his or her concept knowledge to generate an action to achieve the goal (3, 4); the ITS monitors the learner's action and modulates the extrinsic feedback or guidance (5, 6), allowing the learner to adapt his or her concept knowledge to generate a revised action (7, 8).

In Figure 3.6(b) the teacher's modeling environment is a microworld (TME) that sets the task goal (1); the learner uses his or her concept knowledge to generate an action to achieve the goal (2, 3); the microworld models and shows the result of the action (4); the learner uses the feedback to modulate his or her concept knowledge (5) and generate a revised action (6, 7); an adaptive microworld monitors the learner's actions and modulates the selection of the teacher's concept (8) to generate a more, or less, challenging task goal (9, 10).

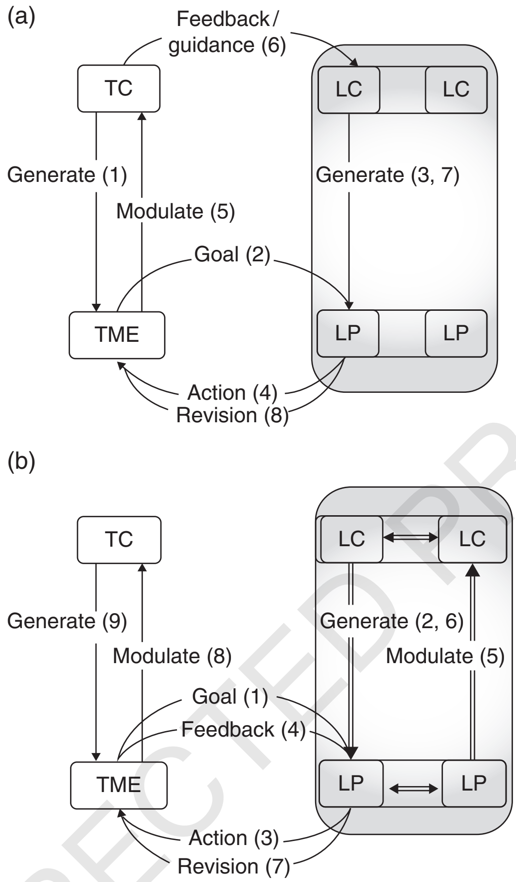

Figure 3.6 (a) An intelligent tutoring system, and (b) an adaptive microworld mapped to the conversational framework.

The point of adapting to the learner's performance is to make the learning situation challenging, to keep the learners in the "zone of proximal development" (Vygotsky, 1978), where inevitably they will make errors. We know that the response to errors is critical in the neural basis of learning. The brain mechanisms try to reduce the difference between the organism's response and the correct or optimal outcome through *prediction error learning* (Dayan & Abbott, 2001). Similarly, the "constructionist" pedagogy relies on the learner being able to interpret the nature of the error (i.e., that the feedback is meaningful to him or her, and then constructing the correct response, thereby recruiting the "predication error learning" mechanisms to dealing with the task set). For this to be possible, the task set must be within the current repertoire, so it is important for the program to monitor the *learner actions* or *output responses*, provide *meaningful feedback* so he or she can change the action or response until it matches the *goal*, and if still having difficulty, adapt the *goal* or *target problem* by changing the *task activity* or *training set*.

An example of an adaptive microworld

As we have seen, neuroscience can help us understand the process of learning by providing cognitive models such as *prediction error learning*, or learning algorithms based on machine learning systems, which can inform the approach a human teacher or a tutoring program might take and are compatible with educational models such as constructionism. It can also help with identifying the type of knowledge that must be targeted. An example of this is the identification of "dyscalculia" as a particular type of neural deficit, sometimes referred to as a "lack of number sense" (see Chapter 8). Children and adults who are dyscalculic need to spend time making sense of how numbers work, and teachers of special needs (SEN) classes have developed materials and techniques to help tackle this specific deficit (Butterworth & Yeo, 2004). Working with, for example, rods of different lengths to represent numbers (such as Cuisenaire rods), learners work on tasks such as constructing the relations between sets (e.g., finding any two rods that make up the length 10, or finding which two identical rods make 10). The same approach has now been implemented as a digital environment by modeling the teacher in the form of an adaptive microworld, which embodies the assumption that the learner is using the cognitive mechanism of *prediction error learning*. One such example is a program for learning the number bonds of ten. Its properties are as follows:

learning outcome – able to compute, e.g., 3 + ? = 10;
method of assessment – tasks such as 3 + ? = 10, etc.
set of task activities – find the correct number bond for 10 for a given number (rods fall within a 10 unit wide column; the sequence of tasks progresses from rods with colour + length, to length only, colour + length + digit, length + digit, digit only; rods fall more slowly if performance is poor);
learner actions – select a rod to fit the column from the pile of 10 rods;
goal – select the rod that fits;
meaningful feedback – rods overlap or show a gap or wriggle into place.

This program provides no peer interaction, and is designed for the individual learner working without a teacher (see www.number-sense.co.uk for other examples). The teacher's model of learning embodied in the program is the idea of "constructionism", that by trying to construct a pairing that fits, and by seeing the result of this action, and then attempting to improve it, the learner will begin to make sense of the relationship between the cardinalities of the numbers 0 to 10 in

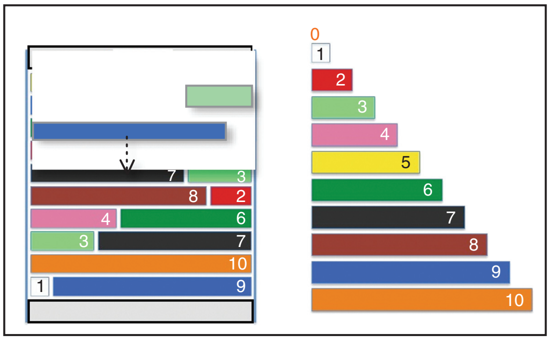

As a rod falls, the learner must select the appropriate rod to make a length of 10. The feedback shows an overlap, gap, or fit, and if incorrect, the same rod falls again, so the learner can improve their response in the light of the feedback.
Task difficulty varies according to whether the objects display length, colour, digit, in different combinations, or, at the highest level, only digits.

Figure 3.7 Learning the number bonds of 10.

terms of their representation as lengths, and eventually in terms of their representation as digits (see Figure 3.7 for a representation of the interface[2]). Like the human teacher, the program adapts the next task item to the pace and accuracy of the response, and the next task set to the learner's ability, proposing that the learner repeat the same task if he or she was very slow or very inaccurate (although the learner can override this). Tests with learners in SEN classes (the *method of assessment*, or *generalization set*) in primary and secondary school (ages 8–13) show that, for example, (i) their performance improves over the short term (two weeks), and (ii) the number of tasks completed is significantly more than in an SEN class of three learners: 4–11 trials per minute were completed by individual learners, while only 1.4 trials per minute were completed on average during ten-minute observations of the classes (Butterworth, Varma, & Laurillard, 2011).

Like the cognitive modeling example above, and the intelligent tutoring system in this section, therefore, this approach also addresses the essential properties for a teaching–learning system. The principal differences between the computational modeling in neuroscience and in education are that in the former the learning process is being modelled by the system, the output of which is to be the same as that of the human learner, whereas in the latter there is only an assumed model of how the human learner is learning, and it is the human teacher whose behavior is modelled by the systems. To examine this in more detail we consider in the next section how the concept of "feedback" is used in the two types of system.

Contrasting perspectives on "feedback"

Computational modeling in neuroscience and education sometimes uses similar terminology, but the concepts behind this terminology do not always directly line up. The notion of feedback offers an instructive example. Feedback is crucial

[2] Now available in the App Store as 'Number Bonds by Thinkout'.

in education, to improve performance on a given task and thereby build an understanding of the topic. Based on the student's current performance on a set of activities, the teacher offers a form of feedback, of which there are two types – "extrinsic", where the teacher interprets what the student needs to be told in order to improve their performance, and "intrinsic", where the environment provides information about the result of their action in relation to their intended goal. In both cases, feedback must be meaningful to the learners if they are to make use of it. The distinction is important because classroom research strongly suggests that intrinsic motivation and reward are more effective than extrinsic (Deci, Koestner, & Ryan, 2001).

Intelligent tutoring systems endeavour to capture and automate the role of the teacher in providing the appropriate *extrinsic* feedback, given the actions of the learner in the task environment of the system. For example, in a system tutoring multiplication, the student attempts a set of problems. Based on the characteristics of the errors displayed, the tutoring system will infer the student's current (erroneous) understanding of the multiplication procedure and provide direct *extrinsic* feedback to repair the gaps in knowledge and rules, and then set further problems to enable the student to improve his or her performance.

Adaptive microworlds create a practice environment that provides *intrinsic* feedback on the learner's actions on a task, based on its model of the world, although by modeling a specific aspect of the real world it focuses the learner's attention on the concepts and skills relevant to the learning outcome. It then selects further tasks, according to the learner's level, to enable him or her to improve in performance.

Cognitive modeling considers how feedback operates inside learning mechanisms. Models are generally addressed to mechanisms for acquiring particular behaviors (such as learning to read, or as we saw above learning aspects of English grammar or aspects of number). Acquisition occurs through exposure to the problem domain. For artificial neural networks, feedback is usually construed as falling into three different classes. In *self-organizing learning* mechanisms, the goal is for the system to develop categories that capture the key dimension of the problem domain, without necessarily generating overt behavior. Self-organizing systems require no external feedback, but instead attempt to optimize some property of their internal knowledge representations as they are exposed to more and more examples of the problem domain. One such property would be how concise or parsimonious the representations are.

In *supervised learning*, very detailed feedback is given to the learning mechanism to improve its performance. For a given input (i.e., example from the problem domain), the mechanism must learn to output a given response (i.e., the right answer, or a step toward the right answer to be fed into another mechanism in the wider cognitive system). Learning occurs in the following way. For the given input, the mechanism outputs its current "best guess" of the appropriate response. This is compared with the actual desired response. The disparity between the two is used as an *error signal* to adjust the connection strengths of the network in such a way that the next time the network encounters this problem its output will be closer to the desired answer. Through repeated exposure to examples, along with detailed feedback, the network gradually acquires the required knowledge. From the point of view of the learning mechanism, the source of the desired response is simply viewed as external to the mechanism – it might originate either from another part of the cognitive system, or as informational feedback from the environment, or from a teacher. As a type of learning algorithm, therefore, supervised learning would be neutral as to whether the feedback was intrinsic or extrinsic.

The third type is called *reinforcement learning*, and falls in between self-organized learning (no feedback) and supervised learning (detailed feedback). In reinforcement learning, the mechanism offers its best guess as to the required response, but the feedback it receives is much more vague. It is similar to a game of locating a hidden object, where one is told "warmer, warmer, cooler, cooler" to encourage looking in one place but discourage looking in other places. This type of learning mechanism is less powerful for learning detailed knowledge. However, researchers have used this type of mechanism to build models of the development of decision making and behavioral control, based on whether the child finds the outcome of each decision to be good or bad; and models of how decision-making abilities can differ in disorders such as attention deficit hyperactivity disorder (see, e.g., Williams & Dayan, 2005). Because this type of learning operates by attempting to minimize the disparity between the expected reward of an action (e.g., whether you will be told "warmer" or "cooler" in the find-a-hidden-object game) and the actual reward, it is sometimes called *prediction error learning*.

The example of feedback shows that all forms of computational modeling discussed here require very detailed specification of the information provided to the learner to improve performance on a task. In the case of intelligent tutoring, providing feedback requires a model of the learner that enables the system to match the output to the presumed input, and thereby direct the feedback to the behavior that produced the erroneous output. In the case of adaptive microworlds the model of the task interaction must provide appropriate informational feedback on the action that enables the learner to interpret how to improve the action. In the case of cognitive models, feedback considers ways in which information is directed to specific mechanisms to alter their knowledge and thereby incrementally improve subsequent behavior.

Conclusion

This chapter has reviewed recent work on the use of computational models in educational neuroscience, in the related methods of cognitive modeling and digital teaching tools. We highlighted the strength of the modeling method, in making explicit theories of learning and teaching. We demonstrated that there are several points of similarity between the approaches of cognitive modeling and digital teaching tools, even though they have different goals and different theoretical origins. Cognitive modeling aims to understand the learning process in the brain; within the teaching tools, intelligent tutoring systems aim to model the learner's knowledge as a version of the domain knowledge in order to generate appropriate remedial tutoring; adaptive microworlds aim to use our understanding of the learning process in the brain to model an environment in which the learner can use this process to learn formal concepts and skills. The theoretical origins of the contrasting approaches draw primarily on neuroscience, information processing, and constructionism, respectively.

We have shown that the two approaches identify a common set of properties that a learning environment must have. This is encouraging for the future of the interdisciplinary field of educational neuroscience. Despite their very different goals and theoretical origins, the computational models have identified something like the essence of what it takes to learn, and formalized it as the set of conditions that make learning possible. From this analysis we can begin to envisage the transfer of a learning model, or task activity, or feedback type from one discipline to another; or the transfer of a finding from one to being tested in another; as well as the contesting of terminology and precise definitions that could help to advance the disciplines of education and neuroscience separately while beginning to bind them together. In short, we are moving towards the potential of integrating educational, neuroscience, and psychological approaches in our developing understanding of learning.

We conclude that the two types of modeling provide the basis for the constructive interdisciplinary dialogue that can now take place between neuroscience-informed cognitive models and education-informed teaching systems.

Acknowledgements

This work was supported by ESRC grant RES-062-23-2721. We would like to thank Kaska Porayska-Pomsta and Ed Hubbard for helpful discussions during the writing of this chapter, as well as the helpful comments of three anonymous reviewers.

References

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York: Oxford University Press.

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Erlbaum.

Anderson, J. R., & Schunn, C. D. (2000). Implications of the ACT-R learning theory: No magic bullets. In R. Glaser (Ed.), *Advances in instructional psychology: Educational design and cognitive science* (Vol. 5) (pp. 1–34). Mahwah, NJ: Erlbaum.

Bishop, D. V. M. (2005). DeFries–Fulker analysis of twin data with skewed distributions: Cautions and recommendations from a study of children's use of verb inflections. *Behavior Genetics*, *35*(4), 479–490.

Bransford, J., Vye, N., Stevens, R., Kuhl, P., Schwartz, D., Bell, P., Meltzoff, A., Barron, B., Pea, R., Reeves, B., Roschelle, J., & Sabelli, N. (2006). Learning theories and education: Toward a decade of synergy. In P. Alexander & P. Winne (Eds.), *Handbook of educational psychology* (2nd ed.) (pp. 209–244). Mahwah, NJ: Erlbaum.

Butterworth, B., Varma, S., & Laurillard, D. (2011). Dyscalculia: from brain to education. *Science*, *332*, 1049–1053.

Butterworth, B., & Yeo, D. (2004). *Dyscalculia guidance: Helping pupils with specific learning difficulties in maths*. London: nferNelson.

Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience: Computational and mathematical modeling of neural systems*. Cambridge, MA: MIT Press.

Deci, E. L., Koestner, R., & Ryan, R. M. (2001). Extrinsic rewards and intrinsic motivation in education: Reconsidered once again. *Review of Educational Research*, *71*(1), 1–27. DOI: 10.3102/00346543071001001

diSessa, A. A. (2001). *Changing minds: Computers, learning and literacy*. 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.

Frederickson, N., Reed, P., & Clifford, V. (2005). Evaluating web-supported learning versus lecture-based teaching: Quantitative and qualitative perspectives. *Higher Education Research and Development*, *50*, 645–664.

Gilbert, N., & Troitzsch, K. G. (2005). *Simulation for the social scientist*. Maidenhead, UK: McGraw-Hill.

Hackman, D. A., & Farah, M. J. (2009). Socioeconomic status and the developing brain. *Trends in Cognitive Sciences*, *13*(2), 65–73.

Hart, B., & Risley, T. R. (1995). *Meaningful differences in the everyday experience of young American children*. Baltimore, MD: Brookes.

Hebb, D. O. (1949). *The organization of behavior*. New York: John Wiley & Sons, Inc.

Hoyles, C., & Noss, R. (2003). What can digital technologies take from and bring to research in mathematics education? In A. J. Bishop, M. A. Clements, C. Keitel, J. Kilpatrick, & F. K. S. Leung (Eds.), *Second international handbook of research in mathematics education* (pp. 323–349). Dordrecht: Kluwer.

Huttenlocher, J., Vasilyeva, J., Cymerman, E., & Levine, S. (2002). Language input and child syntax. *Cognitive Psychology*, *45*, 337–374.

Laurillard, D. (2002). *Rethinking university teaching: A conversational framework for the effective use of learning technologies* (2nd ed.). London: RoutledgeFalmer.

Laurillard, D. (2010). Effective use of technology in teaching and learning in HE. In Peterson, P., Baker, E., & McGaw, B. (Eds.), *International encyclopedia of education* (Vol. 4, pp. 419–426). Oxford: Elsevier.

Laurillard, D. (2012). *Teaching as a design science: Building pedagogical patterns for learning and technology*. New York: Routledge.

Mareschal, D., Johnson, M., Sirios, S., Spratling, M., Thomas, M. S. C., & Westermann, G. (2007). *Neuroconstructivism: How the brain constructs cognition*. Oxford: Oxford University Press.

Mareschal, D., & Thomas, M. S. C. (2006). How computational models help explain the origins of reasoning. *Computational Intelligence Magazine, IEEE*, *1*(3), 32–40.

Mareschal, D., & Thomas M. S. C. (2007) Computational modeling in developmental psychology. *IEEE Transactions on Evolutionary Computation (Special Issue on Autonomous Mental Development)*, *11*(2), 137–150.

Mitchell, T. M. (1997). *Machine learning*. New York: McGraw-Hill.

Papert, S. (1980). *Mindstorms: Children, computers, and powerful ideas*. Brighton, UK: Harvester.

Papert, S., & Harel, I. (1991). Situating constructionism. In I. Harel & S. Papert (Eds.), *Constructionism: Research reports and essays, 1985–1990*. Norwood, NJ: Ablex.

Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing. Volume 1*. Cambridge, MA: MIT Press.

Sawyer, K. (Ed.). (2006). *Cambridge handbook of the learning sciences*. Cambridge: Cambridge University Press.

Sirois, S., Spratling, M., Thomas, M. S. C., Westermann, G., Mareschal, D., & Johnson, M. H. (2008). Précis of neuroconstructivism: How the brain constructs cognition. *Behavioral and Brain Sciences*, *31*, 321–356.

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

Sun, R. (2008). *Cambridge handbook of computational cognitive modelling*. Cambridge: Cambridge University Press.

Thomas, M. S. C. (2012). Brain plasticity and education. *British Journal of Educational Psychology – Monograph Series II: Educational Neuroscience*, *8*, 142–156.

Thomas, M. S. C., Baughman, F. D., Karaminis, T., & Addyman, C. (2012a). Modelling development disorders. In C. Marshall (Ed.), *Current issues in developmental disorders* (pp. 93–124). Hove, UK: Psychology.

Thomas, M. S. C., & Johnson, M. H. (2006). The computational modelling of sensitive periods. *Developmental Psychobiology*, *48*(4), 337–344.

Thomas, M. S. C., & Knowland, V. (2009). Sensitive periods in brain development: Implications for education policy. *European Psychiatric Review*, *2*(1), 17–20.

Thomas, M. S. C., & Knowland, V. C. P. (submitted). Modelling mechanisms of persisting and resolving delay in language development. Manuscript submitted for publication.

Thomas, M. S. C., Knowland, V. C. P., & Karmiloff-Smith, A. (2011). Mechanisms of developmental regression in autism and the broader phenotype: A neural network modeling approach. *Psychological Review*, *118*(4), 637–654.

Thomas, M. S. C., & McClelland, J. L. (2008). Connectionist models of cognition. In R. Sun (Ed.), *Cambridge handbook of computational cognitive modelling* (pp. 23–58). Cambridge: Cambridge University Press.

Thomas, M. S. C., Purser, H. R. M., & Mareschal, D. (2012b). Is the mystery of thought demystified by context-dependent categorisation? Towards a new relation between language and thought. *Mind and Language*, *27*(5), 595–618.

Thomas, M. S. C., Forrester, N. A., & Ronald, A. (in press). Modelling socio-economic status effects on language development. *Developmental Psychology*.

Thomas, M. S. C., Westermann, G., Mareschal, D., Johnson, M. H., Sirois, S., & Spratling, M. (2008). Studying development in the 21st century. *Behavioral and Brain Sciences*, *31*, 345–356.

Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Cambridge, MA: Harvard University Press.

Westermann, G., Mareschal, D., Johnson, M. H., Sirois, S., Spratling, M. W., & Thomas, M. S. C. (2007). Neuroconstructivism. *Developmental Science*, *10*(1), 75–83.

Westermann, G., Thomas, M. S. C., & Karmiloff-Smith, A. (2010). Neuroconstructivism. In U. Goswami (Ed.), *Blackwell handbook of child development* (2nd ed.) (pp. 723–748). Oxford: Blackwell.

Williams, J. O. H., & Dayan, P. (2005). Dopamine, learning and impulsivity: A biological account of ADHD. *Journal of Child Adolescent Psychopharmacology*, *15*(2), 160–179.

Further Reading

Laurillard, D. (2012). *Teaching as a design science: Building pedagogical patterns for learning and technology*. New York: Routledge.

Mareschal, D., Johnson, M., Sirios, S., Spratling, M., Thomas, M. S. C., & Westermann, G. (2007). *Neuroconstructivism: How the brain constructs cognition*. Oxford: Oxford University Press.

Sawyer, K. (Ed.). (2006). *Cambridge handbook of the learning sciences*. Cambridge: Cambridge University Press.

Thomas, M. S. C., Baughman, F. D., Karaminis, T., & Addyman, C. (2012). Modelling development disorders. In C. Marshall (Ed.), *Current issues in developmental disorders* (pp. 93–124). Hove, UK: Psychology