How Educational Theories Can Use Neuroscientific Data

Daniel T. Willingham[1] and John W. Lloyd[2]

ABSTRACT—Although many articles have addressed the relationship of neuroscience and education at a theoretical level, none has considered as a practical matter how one integrates neuroscientific data into a behavioral theory that uses hypothetical constructs. We describe 4 techniques by which researchers may do so: (a) direct observation of hypothetical constructs in the brain, (b) validation of hypothetical constructs through brain imaging, (c) using neural architecture to infer a behavioral architecture, and (d) using well-developed knowledge of brain function to select among competing behavioral theories. A detailed examination of these 4 techniques indicates that they are amenable to educational research and, indeed, have been used to inform research and theory. We argue that these techniques are not applicable to all educational research problems, but rather they are appropriate to relatively low-level behaviors (e.g., reading) rather than complex behaviors (e.g., classroom management).

HOW CAN BRAIN IMAGING HELP EDUCATION RESEARCH?

Interest in the relationship of neuroscience and education is clearly on the increase. In the past 2 years, a spate of articles have examined how the fields relate (e.g., Ansari & Coch, 2006; Goswami, 2006; Katzir & Paré-Blagoev, 2006), although the idea that neuroscientific data might inform educational theory and practice is by no means new (e.g., Berninger & Abbott, 1992; Gaddes, 1983; Iran-Nejad, Hidi, & Wittrock, 1992).

The present article focuses on the practical use of neuroscientific data. Educational theory operates at a level of description different than that used by neuroscientists. How,

[1]Department of Psychology, University of Virginia
[2]Curry School of Education, University of Virginia

Address correspondence to Daniel T. Willingham, Department of Psychology, University of Virginia, Box 400400, Charlottesville, VA 22904; e-mail: willingham@virginia.edu.

then, might educational theories use neuroscientific data? To address this question, we use cognitive neuroscience as a model. Researchers in cognitive neuroscience have been using neuroscientific data to inform models of cognitive function for 25 years or more, so it provides a working analogy to understand potential contributions to education. We then consider whether these particular methods are applicable to educational research. We conclude that they are—and indeed, researchers have been using these methods—but the nature of educational research brings different opportunities and challenges that must be considered.

THE NATURE OF COGNITIVE THEORY

What sort of neuroscientific data might be applied to psychological theory? We will focus on localization—measures that associate the location of brain activity with psychological functions. There are, of course, other measures of central nervous system activity, including neuromodulatory (i.e., those measuring chemical changes, such as microdialysis, microinjection, or perfusion) or temporal (i.e., those measuring the timing of neural activity, such as electroencephalography [EEG] or magnetoencephalography). We focus on localization due partly to space restrictions and partly to our sense that these data currently receive the most interest from educational researchers. Several techniques provide localization data, but three of these are most often used: single-cell recording, lesions, and functional imaging.

How do localization data contribute to cognitive theory? The need for internal representations for a successful theory of behavior was the basis of the cognitive revolution in experimental psychology during the 1950s and 1960s (e.g., Newell, Shaw, & Simon, 1958) and has remained central to cognitive theory since that time. Internal *representations* function as symbols (Markman, 2002). They stand for objects in the physical world, as well as the properties and relationships of these objects, much as symbols on a map stand for geographical entities in the physical world. It is also assumed that humans have mental *processes* by which they manipulate

internal representations. For example, processes might combine concepts to create new ones. Properly conceived, a theory would propose a parsimonious number of representations and processes that would simultaneously be flexible enough to account for a broad variety of behavioral data and specific enough to make testable predictions.

This point about testable predictions is critical. Internal representations and processes are, of course, not directly observable. Critics charged that such theories were therefore inherently unscientific (e.g., Skinner, 1963), but proponents responded that using unobservable constructs in a theory is justifiable if their use brings order to a broad enough array of data (e.g., Simon, 1998). A logical conclusion of this justification is that the proposed theories must be sufficiently detailed to make specific predictions and the behavioral data collected must be precise enough to test them. Without this specificity of theory and data, it is possible to propose literally an infinite number of theories that can account for the same set of observations (Anderson, 1978; Hunt, 1999).

Cognitive psychologists have most often dealt with this problem in two ways, although there are others. First, a subset of them have turned to quantitative modeling to ensure that their theories make specific predictions about outcomes. Second, many have used response–time data. It is true that one can create an unlimited number of models to account for the same set of behaviors, but those models will differ in the efficiency with which they carry out the different tasks, depending on the processes and representations that they employ. Cognitive psychologists have used response times as a proxy for efficiency, and efficiency is the metric used to differentiate theories.

It might seem that brain localization would offer a third way to differentiate theories; one could simply look at the brain to see which theory is right. After all, internal representations must be realized in the physical substrate of neural tissue, and perhaps, scientists can examine the brain and deduce what it represents. In the next section, we will argue that brain localization techniques can be used to select among competing psychological theories, but the strategies for doing so are not simple.

HOW LOCALIZATION DATA CAN DIFFERENTIATE AMONG THEORIES

In this section, we describe four common strategies that cognitive neuroscientists have employed to integrate neuroscientific data and cognitive theories. Before we describe these methods, we call the reader's attention to Figure 1, which depicts levels of analysis for neural and for behavioral researchers, and in which we introduce two new terms.

Most behavioral researchers assume that the mind is organized hierarchically. Representations and the processes
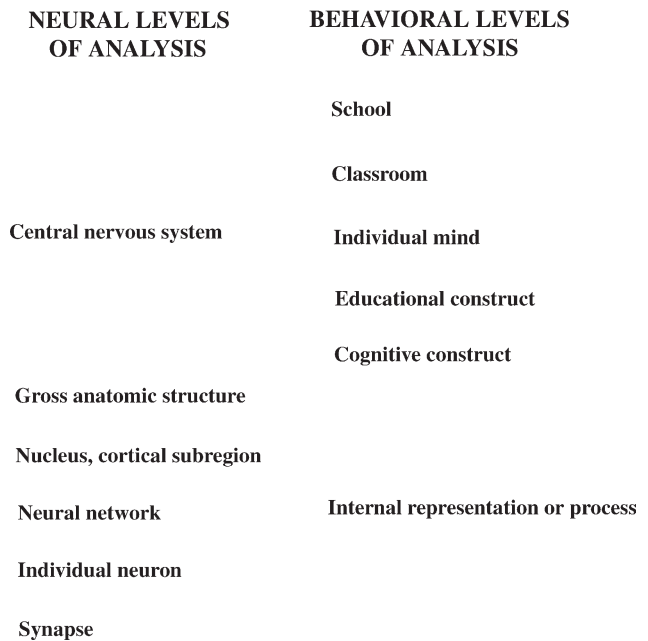
NEURAL LEVELS OF ANALYSIS	BEHAVIORAL LEVELS OF ANALYSIS
	School
	Classroom
Central nervous system	Individual mind
	Educational construct
	Cognitive construct
Gross anatomic structure	
Nucleus, cortical subregion	
Neural network	Internal representation or process
Individual neuron	
Synapse	

Fig. 1. Levels of neural and behavioral analysis. The vertical dimension reflects spatial extent, but on an ordinal, not an interval scale. The horizontal dimension is meant to reflect the approximate match between neural and behavioral levels, for example, an individual mind is reflected in the entire central nervous system. A "neural network" is composed of hundreds of thousands or millions of neurons.

that manipulate them lie at the bottom of the hierarchy. They serve as building blocks for the theoretical entities at the next level, which we will call *cognitive constructs*. Each of these serves a single cognitive function or a group of closely related functions. An example would be "working memory" (Baddeley, 2003). Cognitive constructs make use of multiple representations and processes. We refer to theoretical entities at the next level of description as *educational constructs*. Each of these serves a particular function in educational theory and will typically have multiple cognitive constructs embedded in it. An example of an educational construct would be "mathematical reasoning," which would draw on cognitive constructs such as "attention," "working memory," "long-term memory," and so forth.

The terms "cognitive construct" and "educational construct" are imprecise. Some constructs would be difficult to classify as one or the other (e.g., "reading"), but precision is not important for our purposes. The important features of the behavioral analysis are (a) its hierarchical organization, with lower order constructs embedded in higher order constructs and (b) the fact that educational constructs will generally be of a higher order than cognitive constructs.

Each level of behavioral or neural analysis has its own integrity because a researcher might propose a theory of how the theoretical entity works without specifying the detailed operation of those entities embedded in it. Thus, a cognitive researcher might propose a theory of working memory

that describes subcomponent processes and representations functionally but not mechanistically; that is, the theorist describes what the subcomponents do but not the particulars of how they do it. So too, an education researcher might propose a theory of reading that describes a function for working memory, a function for phonological decoding, and so forth, but without articulating a complete theory of these or other component processes. In just the same way, a neuroscientist may propose a theory of the operation of a neuron without describing the workings of each dendrite and synapse.

The independence of these levels of analysis has an important implication. Each level is suitable to support a theory. Although the levels of analysis for neuroscience and for behavior are hierarchically organized, one level of theoretical description is not fully reducible to the next. There are two strong reasons to believe that this independence of levels of analysis is sensible. First, it is known that, at least for some levels of analysis, the whole is more than the sum of the parts. That is, there are properties that arise not from individual components and their interactions but from all the components acting together. This phenomenon is often called *emergence.* Classroom effects provide an example. Knowing the properties of the students and teachers as individuals is not enough to predict these effects. Second, there is a practical consideration. Suppose a neuroscientist sought to describe the neural events associated with hunger. A theory at a low level of description might account for hunger in terms of the firing patterns of tens of millions of synapses. Even if emergence were not an issue—that is, the description of the hunger were complete—an account entailing millions of synaptic events would not be comprehensible to the human mind. If a key function of scientific theory were to bring order and comprehensibility to the world, such an account would fail to serve this function.

Bearing in mind the hierarchical nature of these levels of analysis (the vertical dimension of Figure 1) and the manner in which the neural and behavioral levels are comparable (the horizontal dimension of Figure 1) will prove useful as we consider the techniques by which neural data may inform behavioral theory.

Strategy 1: Direct Observation of an Internal Representation

This first strategy comes closest to the idea of simply looking at the brain. Single-cell recording data, in which a probe records the activity of a neuron or a small group of neurons, are often interpreted as reflecting what a neuron codes. If one knows what a neuron codes, then a behavioral researcher might be wise to include that representation in his or her theory. For example, Georgopoulos, Schwartz, and Kettner (1986) showed that neurons in the monkey primary motor cortex are sensitive to movement direction, not the target location or end point of the movement. Following this important work,

other researchers proposed behavioral models of movement that used direction as a primitive representation (e.g., Takeda & Funahashi, 2004). Researchers interested in reading and reading disorders have used single-cell recording data in a similar fashion, although infrequently. The work of Stein and Walsh (1997) provides one example. They hypothesized that, among people with dyslexia, the magnocellular or large-cell visual system that processes low-contrast, fast-moving objects, interacts poorly with the parvocellular (small-cell) visual system that processes colors and, especially, patterns. This hypothesis is rooted in single-cell recording work that first described properties of the magnocellular and parvocellular streams of processing (e.g., Derrington, 1984). The evidence for the connection to dyslexia comes from several lines of studies, including one showing abnormalities in the magnocellular layer (but not the parvocellular layer) of the lateral geniculate nucleus of the thalamus in individuals with dyslexia (Livingstone, Rosen, Drislane, & Galaburda, 1991). Although the evidentiary bases for the theory have been challenged subsequently (e.g., Hutzler, Kronbichler, Jacobs, & Wimmer, 2006; Johannes, Kussmaul, Münte, & Mangun, 1996), this work illustrates how single-cell recording studies of neurons can inspire a theoretical account of behavior.

Strategy 2: The Reality of Cognitive Constructs

Functional imaging techniques are interpreted as measures of brain activity. Positron emission tomography and functional magnetic resonance imaging (fMRI) actually provide indirect measures of blood flow. The well-founded assumption is that increased neural activity creates a demand for increased oxygen or glucose at the active site, and so increased blood supply is shunted to the neurally active region.

Experiments that localize brain activity with functional imaging techniques are rarely designed to investigate a single process or representation; they are designed to investigate a cognitive construct or a subcomponent thereof, such as long-term memory retrieval (e.g., Cabeza, Prince, & Daselaar, 2004), lexical access (e.g., Petersen, Fox, Posner, Mintun, & Raichle, 1988), or maintenance in working memory (e.g., D'Esposito, Postle, Ballard, & Lease, 1999). Localizing the brain activity associated with a cognitive construct appears to provide compelling evidence for its reality and importance. It is one thing to propose that there is a process in experienced readers by which phonemes (basic units of sound in speech) are derived from graphemes (letters or letter combinations corresponding to phonemes); it is another thing to show, by a series of careful experiments, exactly where that process happens in the brain (e.g., Petersen, Fox, Snyder, & Raichle, 1990).

But why is it compelling? There is a spurious sense that an unobservable cognitive construct is made more scientifically rigorous if it is associated with observable brain activity. For example, one might think that, if we compare the

brain activity associated with considering a moral dilemma and that associated with solving other problems (Greene, Sommerville, Nystrom, Darley, & Cohen, 2001), we can "see" morality in the brain. That interpretation is misleading.

The problem is that any mental activity leads to some pattern of brain activation. The pattern may even be reliable across participants, but that does not make it theoretically meaningful. To take an extreme example, suppose someone proposes a theory of classroom management that includes a *monitor selection module*, an educational construct describing a mental process dedicated to the job of selecting a student to monitor classroom behavior when a teacher steps out for a moment. Researchers carefully design an experiment conducted in an fMRI magnet that requires experienced teachers to observe a videotape of a classroom for 5 minutes and then to select a classroom monitor. Researchers use the proper control conditions and find a network of eight brain regions that, in every participant they test, is reliably associated with selecting a monitor. Have they provided evidence to support the reality of the *monitor selection module* by localizing it in the brain?

Likely no one would be impressed. A better interpretation of the data is that selection of the monitor requires a set of processes (e.g., attention, memory, social comparisons) that are briefly recruited to solve a problem. Monitor selection may show reliable localization, but it is not a useful educational construct because it is applicable to a single behavior. How, then, can we know whether a construct is useful? We make that judgment based on behavioral testing. As noted above, the justification for using an unobservable construct in a theory is that the resultant theory can account for a good deal of data. The construct must carry its freight. An unobservable construct such as "grapheme to phoneme conversion" carries its freight; "monitor selection" does not. Brain imaging adds nothing if all it shows us is the localization of the construct because every construct that we can dream up must be localized somewhere in the brain.

So why localize? We just said that cognitive or educational constructs must be broad and flexible, that is, applicable to a variety of tasks. At the same time, the theorist must be able to predict when the construct will be invoked. If the set of constructs in a theory are *too* broad and flexible, there may be many ways the theory can account for performance in a given task, and the theorist finds himself or herself unable to predict which constructs will be engaged.

Brain localization is one tool that a researcher can use to show that both desiderata—flexibility and specificity—have been satisfied. Localization provides a dependent measure that can show that the theorist understands the relationship of the construct and behavioral tasks. Thus, if a theorist proposes that storage of spatial material in working memory is subserved by a network that includes the prefrontal cortex and the dorsal parietal cortex (Wager & Smith, 2003), then

a researcher should be able to predict when those areas will be active and when they will not. If the researcher can make such predictions, one should be persuaded that the relationship between the unobservable construct and overt behavior are sufficiently specific. Localization can also prove useful in demonstrating the flexibility of the construct. The researcher should be able to devise a wide variety of tasks, some of which may appear dissimilar, but that the researcher predicts all draw on the construct, and therefore all lead to activation in the predicted locations.

This method of applying neuroscience to behavioral research has been employed in educational theory. For example, a theory of dyslexia that contends that reading problems are associated with specific language problems would benefit from evidence that poor performance in specific aspects of reading occur in the same areas of the brain across languages (e.g., Eden & Moats, 2002; Shaywitz, 2003). In fact, although factors such as the fundamental print form of a language affect location of deficit processes (Siok, Perfetti, Jin, & Tan, 2004), in alphabetic languages there is remarkable similarity in the location of atypical activity among individuals who learn to read in different languages, including German, Italian, and English (Paulesu et al., 2001). This consistency in the localization of dyslexic abnormalities across varied languages increases one's confidence that the impaired reading processes have been aptly described.

Strategy 3: Separate Systems From Apparently Unitary Functions

Brain localization has often been used to delineate separate cognitive systems. If researchers observe that brain structures X, Y, and Z are consistently associated with one set of tasks, and brain structures P, Q, and R are consistently associated with another set of tasks, then they might assume that different cognitive systems support task Set 1 and task Set 2. In turn, that implies that different theories should be required to account for the tasks in Set 1 and Set 2.

For example, brain imaging and lesion data indicate that the premotor, posterior parietal, and prefrontal cortices, along with the medial temporal lobe, are crucial for learning a sequence of locations in space (Eliassen, Souza, & Sanes, 2001). A different set of brain regions (supplementary motor area, anterior cingulate, posterior parietal cortex, and striatum) supports learning in a quite similar task in which participants are not asked to learn the sequence but simply press buttons in response to spatial locations. Sequence learning is inferred from decreasing response times (e.g., Willingham, Salidis, & Gabrieli, 2002). Many researchers interpret that anatomic dissociation as strong evidence that the two types of tasks are qualitatively different (e.g., Poldrack & Willingham, 2006) despite their apparent similarity.

Note, however, that this conclusion is not logically compelled. One could argue that the real test of separability for a cognitive psychologist is cognitive separability, and anatomic separability may not be a foolproof guide. To use an analogy, two files created by Microsoft Word may reside in different places on your computer's hard drive, but the representations they employ and the processes needed to use the files are not distinguishable. In the same way, it is possible that two varieties of tasks may use the same rules at a cognitive level of description, even though they are implemented in different parts of the brain. Thus, when researchers collect data indicating that two behavioral tasks are supported by different neural structures, that anatomic separation does not prove beyond doubt that the two tasks are supported by different systems at a cognitive level of description. Nevertheless, most researchers take anatomic separability as strongly indicating cognitive separability.

Educational theorists have also made good use of this technique. For example, Shaywitz and colleagues (Shaywitz, Lyon, & Shaywitz, 2006; Shaywitz et al., 2003) observed differences in the pattern of activation associated with three groups of readers: persistently poor readers, accuracy (but not fluency)–improved readers, and nonimpaired readers. During real-word reading, the accuracy-improved readers showed underactivation of left-posterior regions of the brain compared to nonimpaired readers. In contrast, the persistently poor readers showed robust activation of these posterior regions.

The results from the Shaywitz group correspond with results from other studies of dyslexia. For example, Castles, Datta, Gayan, and Olson (1999) examined the heritability of reading deficits in dyslexic twins categorized on the basis of their reading of pseudo-words and irregularly spelled words. Although there was a genetic contribution to the reading performance of the children in both groups, genetic factors contributed more to the reading performance of the group of children whose scores on the phonological processing measure were higher than their scores on the orthographic processing measure in comparison to the group of children whose orthographic processing measure was higher than their phonological processing measure. Taken together, these findings suggested to researchers that there are two subgroups of participants with reading difficulties: one group with a primarily genetic etiology and the other reflecting environmental influences to a greater degree. The differences in anatomic localization of reading responses form a cornerstone of this hypothesis.

Strategy 4: Reliable Knowledge of the Brain Guides Cognitive Theory

In this technique, researchers accept that they do not have a full understanding of the cognitive processes that contribute to a task, and they use brain imaging to inform cognitive theory.

A classic example of this approach may be drawn from the visual imagery debate. During the 1970s, there was a lively and sometimes acrimonious debate between researchers (e.g., Kosslyn, 1976) who believed that the cognitive system used two forms of representation (linguistic and analog, i.e., quasi-pictorial) and other researchers (e.g., Pylyshyn, 1973) who believed that linguistic representations could account for memorial abilities and that analog representations were an unnecessary complication. The tasks that proved critical in the controversy were those that participants claimed to solve by using mental imagery (e.g., "what shape are a German Shepherd's ears?"). Settling this issue through behavioral data alone was difficult (Anderson, 1978) in part because mental imagery is not observable. Fortunately, each theoretical position made a straightforward prediction about the neural basis of imagery tasks. If imagery tasks are supported by linguistic representations, then imagery tasks should be localized in parts of the brain known to support language (at a very rough cut, the superior temporal lobe). But if imagery tasks are supported by quasipictorial representations, they should be localized in visual areas like the primary and secondary visual cortex. Data from lesion studies (Farah, 1988) and from brain imaging (Kosslyn, Alpert, Thompson, & Maljkovic, 1993) showed that imagery tasks are supported, in part, by visual cortices and that language regions of the temporal lobe play no role, decisively supporting the dual representation theories.

Naturally, this strategy is only applicable when fairly good knowledge of the brain is available. It was effective in the case of imagery because the predictions concerned brain regions that were well characterized. The number of brain regions for which researchers can confidently assign functions is limited, but growing. Despite this limitation, educational researchers have made use of this technique. Most notably, researchers have observed lowered activation in the brains of dyslexic readers in regions known from other work to support phonological coding (e.g., Paulesu et al., 1996). That finding complements behavioral findings that implicate phonological processing in dyslexia (e.g., Brady & Shankweiler, 1991; Wagner & Torgesen, 1987).

HOW EDUCATION DIFFERS

The foregoing section described four different techniques by which neuroscientific data have been used to inform cognitive theory. We have also provided examples of uses of each technique in educational research, focusing on dyslexia. Is it then the case that these four techniques could generally be used to integrate neuroscientific data and educational theory?

Educational theory differs from cognitive theory in several important ways, but two of these strike us as especially relevant to the topic at hand. One leads to an opportunity for greater use of neuroscientific data, whereas the other presents a limitation.

Neuroscience and Diagnosis of Learning Disabilities

An important difference between cognitive and educational research is that study of dysfunction is commonly taken to be part of the latter, but not the former. Neuroscientific data offers, in our view, considerable potential to be useful in this effort, specifically through diagnosis.

There seems ample reason to be optimistic that some learning disorders have a detectable neural basis and will be reliably diagnosed via neuroscientific measures in the near future. We focus here on dyslexia as an example, and indeed, it is the disability for which we believe researchers are closest to realizing this goal.

It is well established that at least some forms of dyslexia are due to disorders of auditory processing (e.g., Breier et al., 2001). This problem is likely one of brain connectivity and in some cases is at least partly attributable to genetic anomalies (e.g., Meng et al., 2005). It has been shown that auditory processing problems are detectable with EEG (Stoodley, Hill, & Stein, 2006), even in infants less than 1 year old (Benasich et al., 2006). Failures of auditory processing have signature patterns of EEG response, and the infants with this pattern later showed delayed language learning (e.g., Guttorm, Leppänen, & Poikkeus, 2005), which is also known to be associated with difficulties in learning to read.

It is already known that dyslexic children and adults show different EEG patterns than those of unimpaired readers (Molfese, Key, & Kelly, 2006), but that difference may be due to faulty instruction or any of a host of other factors. The question is whether there might be a signature EEG pattern indicative of likely difficulty in learning to read that would be observable before the child ever sets foot in a classroom. That would be of tremendous importance, given that intervention is known to be more effective if it is earlier rather than later (e.g., Foorman, Francis, Shaywitz, Shaywitz, & Fletcher, 1997). Although several labs are actively pursuing a solution to this problem (e.g., Espy, Molfese, & Molfese, 2004; Lyytinen, Guttorm, & Huttunen, 2005), a reliable marker has not yet been identified. Still, it seems likely that, with a vigorous effort, this discovery is not more than 5 or 10 years in the future.

Levels of Analysis

We earlier drew attention to the fact that education and cognition operate at different levels of behavioral analysis, as shown in Figure 1. Ultimately, educational researchers are interested in behaviors that bear on educational outcomes—reading, scientific problem solving, synthesizing historical documents, and so on. They are interested in cognitive constructs—memory, working memory, attention—insofar as they shed light on educational constructs. We asserted that educational constructs typically have two or more cognitive constructs embedded within them. The hierarchical nature of behavioral levels of analysis, coupled with the relationship between behavioral and neural analyses, has important implications for the promise of a neuroscience of education.

A significant problem in trying to map behavioral and neuroscientific analyses to one another is the presence of interactions among the constituents of a cognitive construct. That is, a cognitive construct operates differently, depending on what other cognitive constructs are doing. In one striking example, Price et al. (1997) compared brain activation in two conditions: silently reading a word and silently saying "okay" when the word appeared. They observed (among other findings) that the left extrastriate cortex was more active in the silent reading than in the okay condition. The extrastriate cortex subserves vision. Why would visual processing change depending on whether one silently read a word or silently said okay in response to it? The authors concluded that this activity might represent top-down modulation of visual cortical processing. In other words, it represents an interaction. Visual processing operates differently if one is naming objects versus simply observing them.

The concern, then, is that cognitive constructs operate differently depending on what other constructs are doing. Even low-level constructs, such as visual processing of relatively simple stimuli, show these context effects. A neuroscientific analysis is just as likely to be complicated by interactions, based on anatomic connectivity. Figure 2 shows a diagram of the connectivity of macaque visual areas. It is not quite the case that "everything is connected to everything" as some mournful students have commented, but there is ample room for complex interactions, recursion, and so forth. In addition, it is quite possible or even probable that functionally separate circuits are interdigitated within a single cortical area (e.g., Callaway, 1998).

We highlight the reality of interactions in behavioral analysis and in neural analysis to emphasize the difficulty of the task at hand. Developing a behavioral or a neuroscientific theory on its own is difficult. It is that much more difficult to map one complex system to another complex system.

But things are still harder in the case of education. Note in Figure 1 that levels of behavioral and neural analysis rarely align horizontally. A cognitive construct rarely maps directly to a single entity in a neural analysis; constructs like "long-term memory retrieval" involve multiple neural sites. This problem is compounded in the case of education. Education constructs are one level removed from cognitive constructs.
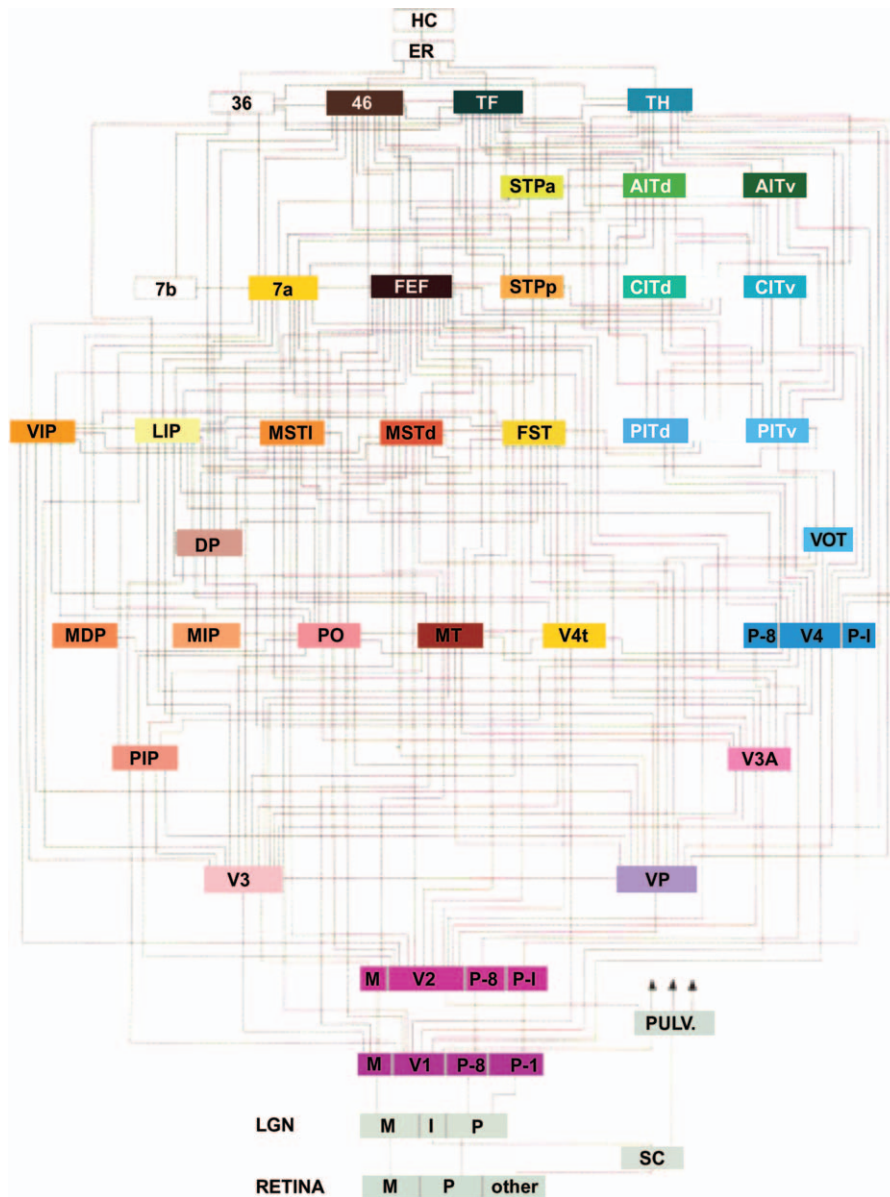
Fig. 2. Connectivity of macaque visual cortical (colored) and subcortical (gray) structures. From Van Essen, Anderson, and Felleman (1992). Reprinted with permission from AAAS.

Thus, there is even *greater* opportunity for complex interactions in the behavioral analysis, making it still more difficult to map behavior to brain.

Figure 1 highlights another limitation of the potential applicability of neuroscientific data to education. There are levels of analysis in behavior that are known to be important to educational outcomes (effects that occur at the level of the individual mind, the classroom, or the school) for which there are not parallel neuroscientific levels of analysis. This limitation does not diminish the importance of neuroscientific data at other levels of analysis, but it does serve as a reminder that

there are important behavioral effects that cannot be directly informed by neuroscientific data.

CONCLUSIONS

We have specified four ways that neuroscientific data can inform psychological theory. In so doing, we have illustrated that all these methods have been used by educational researchers in the study of dyslexia. We believe that, as a topic of study, dyslexia has two features that made this effort successful.

First, reading is a relatively low-level educational construct. That means that it is more tractable in a neuroscientific analysis because it uses a small number of cognitive constructs. When we try to analyze many educational constructs of interest—for example, the effect of a classroom's emotional atmosphere on learning (e.g., Rimm-Kaufman, La Paro, Downer, & Pianta, 2005) or how the expectation of negative social interactions affects adolescents' choices to use alcohol and drugs (e.g., Allen, Leadbeater, & Aber, 1994)—we will be observing most of the brain at work.

The second reason that neuroscientific data have been profitably applied to theories of dyslexia is that educational researchers had already developed sophisticated theories of reading and of dyslexia based on behavior. These theories guided the interpretation of the neuroscientific data. For example, some theories of reading predicted the potential importance of phonological decoding, so when parts of the brain associated with phonology were implicated in dyslexia, there was a ready interpretation.

A careful reading of the literature shows relative agreement on the potential for neuroscientific data to inform educational research. One finds only occasional naysayers who argue that neuroscience offers nothing to education (Davis, 2004), and some irrationally exuberant (and inexact) boosters of neuroscience in education may be found in the popular press (e.g., Connell, 2002; Jensen, 2005).

Scholarly treatments have been positive about the prospects, but more sober, and most have taken a position that is broadly consistent with ours. They argue that neuroscience has been and will continue to be helpful to education—indeed, recent reviews show beyond doubt that this is true (e.g., Katzir & Paré-Blagoev, 2006)—but they argue that data from neuroscience must be funneled through a behavioral level of analysis (e.g., Bruer, 1997, 1998; Hirsh-Pasek & Bruer, 2007) or that neuroscience should be part of a broader approach to research in education, not the sole savior (e.g., Ansari & Coch, 2006; Byrnes & Fox, 1998; Fischer et al., 2007; Geake & Cooper, 2003). Naturally, there is variation in how this general orientation is expressed. Byrnes and Fox, for example, suggest that findings from neuroscience can be used as constraints on educational theory. Szűcs and Goswami (2007) offer a provocative vision of educational neuroscience centered on the development of mental representations, instantiated as neural networks.

The view offered here differs from others in the specific methods by which we argue that advances can be made and the guidelines we have offered for a research agenda. We argue that educational researchers would do well to leverage methods developed in cognitive neuroscience. If they do so, however, neuroscientific data will be most fruitfully applied to educational issues that entail relatively low-level behavioral constructs and for which there are already well-developed behavioral theories.

REFERENCES

Allen, J. P., Leadbeater, B. J., & Aber, J. L. (1994). The development of problem behavior syndromes in at-risk adolescents. *Development and Psychopathology*, 6, 323–342.

Anderson, J. R. (1978). Arguments concerning representations for mental imagery. *Psychological Review*, 85, 249–277.

Ansari, D., & Coch, D. (2006). Bridges over troubled waters: Education and cognitive neuroscience. *Trends in Cognitive Sciences*, 10, 146–151.

Baddeley, A. D. (2003). Working memory: Looking back and looking forward. *Nature Reviews: Neuroscience*, 4, 829–839.

Benasich, A. A., Choudhury, N., Friedman, J. T., Realpe-Bonilla, T., Chojnowska, C., & Gou, Z. (2006). The infant as a prelinguistic model for language learning impairments: Predicting from event-related potentials to behavior. *Neuropsychologia*, 44, 396–411.

Berninger, V. W., & Abbott, R. D. (1992). The unit of analysis and the constructive processes of the learner: Key concepts for educational neuropsychology. *Educational Psychologist*, 27, 223–242.

Brady, S., & Shankweiler, D. (1991). *Phonological processes in literacy*. Hillsdale, NJ: Erlbaum.

Breier, J. I., Gray, L., Fletcher, J. M., Diehl, R. L., Klaas, P., Foorman, B. R., et al. (2001). Perception of voice and tone onset time continua in children with dyslexia awith and without attention deficit/hyperactivity disorder. *Journal of Experimental Child Psychology*, 80, 245–270.

Bruer, J. T. (1997). Education and the brain: A bridge too far. *Educational Researcher*, 26, 4–16.

Bruer, J. T. (1998). Let's put brain science on the back burner. *NASSP Bulletin*, 82, 9–19.

Cabeza, R., Prince, S. E., & Daselaar, S. M. (2004). Brain activity during episodic retrieval of autobiographical and laboratory events: An fMRI study using a novel photo paradigm. *Journal of Cognitive Neuroscience*, 16, 1583–1594.

Callaway, E. M. (1998). Local circuits in primary visual cortex of the macaque monkey. *Annual Review of Neuroscience*, 21, 47–74.

Castles, A., Datta, H., Gayan, J., & Olson, R. (1999). Varieties of developmental reading disorder: Genetic and environmental influences. *Journal of Experimental Child Psychology*, 72, 73–94.

Connell, D. (2002). Left brain, right brain. *Instructor*, 112, 28–32, 89.

Davis, A. (2004). The credentials of brain-based learning. *Journal of the Philosophy of Education*, 38, 21–35.

Derrington, A. M. (1984). Spatial and temporal contrast sensitivities of neurons in lateral geniculate nucleus of macaque. *Journal of Physiology*, 357, 219–240.

D'Esposito, M., Postle, B. R., Ballard, D., & Lease, J. (1999). Maintenance versus manipulation of information held in working memory: An event-related fMRI study. *Brain and Cognition*, 41, 66–86.

Eden, G. F., & Moats, L. (2002). The role of neuroscience in the remediation of students with dyslexia. *Nature Neuroscience*, 5, 1080–1084.

Eliassen, J. C., Souza, T., & Sanes, J. N. (2001). Human brain activation accompanying explicitly directed movement sequence learning. *Experimental Brain Research*, 141, 269–280.

Espy, K. A., Molfese, D. L., & Molfese, V. J. (2004). Development of auditory event-related potentials in young children and relations to word-level reading abilities at age 8 years. *Annals of Dyslexia*, 54, 9–38.

Farah, M. J. (1988). Is visual imagery really visual? Overlooked evidence from neuropsychology. *Psychological Review*, *95*, 307–317.

Fischer, K. W., Daniel, D. B., Immordino-Yang, M. H., Stern, E., Battro, A., & Koizumi, H. (2007). Why *Mind, Brain, and Education*? Why now? *Mind, Brain and Education*, *1*, 1–2.

Foorman, B. R., Francis, D. J., Shaywitz, Í. E., Shaywitz, B. A., & Fletcher, J. M. (1997). The case for early reading intervention. In B. A. Blachman (Ed.), *Foundations of reading acquisition and dyslexia: Implications for early intervention* (pp. 243–264). Mahwah, NJ: Erlbaum.

Gaddes, W. H. (1983). Applied educational neuropsychology: Theories and problems. *Journal of Learning Disabilities*, *16*, 511–515.

Geake, J., & Cooper, P. (2003). Cognitive neuroscience: Implications for education? *Westminster Studies in Education*, *26*, 7–20.

Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, *233*, 1416–1419.

Goswami, U. (2006). Neuroscience and education: From research to practice. *Nature Reviews Neuroscience*, *7*, 2–7.

Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science*, *293*, 517–523.

Guttorm, T. K., Leppänen, P. H. T., & Poikkeus, A.-M. (2005). Brain event-related potentials (ERPs) measured at birth predict later language development in children with and without familial risk for dyslexia. *Cortex*, *41*, 291–303.

Hirsh-Pasek, K., & Bruer, J. T. (2007). The brain/education barrier. *Science*, *317*, 1293.

Hunt, E. (1999). What is a theory of thought? In R. J. Sternberg (Ed.), *The nature of cognition* (pp. 3–49). Cambridge, MA: MIT Press.

Hutzler, F., Kronbichler, M., Jacobs, A. M., & Wimmer, H. (2006). Perhaps correlational but not causal: No effect of dyslexic readers' magnocellular system on their eye movements during reading. *Neuropsychologia*, *44*, 637–648.

Iran-Nejad, A., Hidi, S., & Wittrock, M. C. (1992). Reconceptualizing relevance in education from a biological perspective. *Educational Psychologist*, *27*, 407–414.

Jensen, E. (2005). *Teaching with the brain in mind* (2nd ed.). New York: ASCD Press.

Johannes, S., Kussmaul, C. L., Münte, T. F., & Mangun, G. R. (1996). Developmental dyslexia: Passive visual stimulation provides no evidence for a magnocellular processing deficit. *Neuropsychologia*, *34*, 1123–1127.

Katzir, T., & Paré-Blagoev, J. (2006). Applying cognitive neuroscience research to education: The case of literacy. *Educational Psychologist*, *41*, 53–74.

Kosslyn, S. M. (1976). Can imagery be distinguished from other forms of internal representation? Evidence from studies of information retrieval times. *Memory & Cognition*, *4*, 291–297.

Kosslyn, S. M., Alpert, N. M., Thompson, W. L., & Maljkovic, V. (1993). Visual mental imagery activates topographically organized visual cortex: PET investigations. *Journal of Cognitive Neuroscience*, *5*, 263–287.

Livingstone, M. S., Rosen, G. D., Drislane, F. W., & Galaburda, A. M. (1991). Physiological and anatomical evidence for a magnocellular defect in developmental dyslexia. *Proceedings of the National Academy of Sciences of the United States of America*, *88*, 7943–7947.

Lyytinen, H., Guttorm, T. K., & Huttunen, T. (2005). Psychophysiology of developmental dyslexia: A review of findings including studies of children at risk for dyslexia. *Journal of Neurolinguistics*, *18*, 167–195.

Markman, A. B. (2002). Knowledge representation. In H. Pashler & D. Medin (Eds.), *Steven's handbook of experimental psychology (3rd ed.), Vol. 2: Memory and cognitive processes* (pp. 165–208). New York: Wiley.

Meng, H., Smith, S. D., Hager, K., Held, M., Liu, J., Olson, R. K., et al. (2005). DCDC2 is associated with reading disability and modulates neuronal development in the brain. *Proceedings of the National Academy of Sciences of the United States of America*, *102*, 17053–17058.

Molfese, D. L., Key, A. F., & Kelly, S. (2006). Below-average, average, and above-average readers engage different and similar brain regions while reading. *Journal of Learning Disabilities*, *39*, 352–363.

Newell, A., Shaw, J. C., & Simon, H. A. (1958). Elements of a theory of human problem solving. *Psychological Review*, *65*, 151–166.

Paulesu, E., Démonet, J.-F., Fazio, F., McCrory, E., Chanoine, V., Brunswick, N., et al. (2001). Dyslexia: Cultural diversity and biological unity. *Science*, *291*, 2165–2167.

Paulesu, E., Frith, U., Snowling, M., Gallagher, A., Morton, J., Frackowiak, R. S. J., et al. (1996). Is developmental dyslexia a disconnection syndrome? Evidence from PET scanning. *Brain*, *119*, 143–157.

Petersen, S. E., Fox, P. T., Posner, M. I., Mintun, M., & Raichle, M. E. (1988). Positron emission tomographic studies of the cortical anatomy of single-word processing. *Nature*, *331*, 585–589.

Petersen, S. E., Fox, P. T., Snyder, A. Z., & Raichle, M. E. (1990). Activation of extrastriate and frontal cortical areas by visual words and wordlike stimuli. *Science*, *249*, 1041–1044.

Poldrack, R. A., & Willingham, D. T. (2006). Skill learning. In R. Cabeza & A. Kingstone (Eds.), *The handbook of functional neuroimaging* (2nd ed., pp. 113–148). Cambridge, MA: MIT Press.

Pylyshyn, Z. W. (1973). What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin*, *80*, 1–24.

Rimm-Kaufman, S. E., La Paro, K. M., Downer, J. T., & Pianta, R. C. (2005). The contribution of classroom setting and quality of instruction to children's behavior in the kindergarten classroom. *Elementary School Journal*, *105*, 377–394.

Shaywitz, B. A., Lyon, G. R., & Shaywitz, S. E. (2006). The role of functional magnetic resonance imaging in understanding reading and dyslexia. *Developmental Neuropsychology*, *30*, 613–632.

Shaywitz, S. (2003). *Overcoming dyslexia: A new and complete science-based program for reading problems at any level*. New York: Knopf.

Shaywitz, S. A., Shaywitz, B. A., Fulbright, R. C., Skudlarskic, P., Mencla, W. E., Constable, R. T., et al. (2003). Neural systems for compensation and persistence: Young adult outcome of childhood reading disability. *Biological Psychiatry*, *54*, 25–33.

Simon, H. A. (1998). Discovering explanations. *Minds and Machines*, *8*, 7–37.

Siok, W. T., Perfetti, C. A., Jin, Z., & Tan, L. H. (2004). Biological abnormality of impaired reading is constrained by culture. *Nature*, *431*, 71–76.

Skinner, B. F. (1963). Behaviorism at fifty. *Science*, *140*, 951–958.

Stein, J., & Walsh, V. (1997). To see but not to read; the magnocellular theory of dyslexia. *Trends in Neurosciences*, *20*, 147–152.

Stoodley, C. J., Hill, P. R., & Stein, J. F. (2006). Auditory event-related potentials differ in dyslexics even when auditory psychophysical performance is normal. *Brain Research*, *1121*, 190–199.

Takeda, K., & Funahashi, S. (2004). Population vector analysis of primate prefrontal activity during spatial working memory. *Cerebral Cortex, 14*, 1328–1339.

Wager, T. D., & Smith, E. E. (2003). Neuroimaging studies of working memory: A meta-analysis. *Cognitive, Affective, & Behavioral Neuroscience, 3*, 255–274.

Wagner, R. K., & Torgesen, J. K. (1987). The nature of phonological processing and its causal role in the acquisition of reading skills. *Psychological Bulletin, 101*, 192–212.

Willingham, D. B., Salidis, J., & Gabrieli, J. D. (2002). Direct comparison of neural systems mediating conscious and unconscious skill learning. *Journal of Neurophysiology, 88*, 1451–1460.