Part IV

An Examination of Specific Learning and Motivational Issues

12 Learning and Constructivism

Walter Kintsch *University of Colorado*[*]

In this chapter, I try to make explicit some issues that have been somewhat overlooked in the debate over the "Failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching" (Kirschner, Sweller, & Clark, 2006). The tendency has been to lump all these methods under the term "constructivist" and hence to identify constructivism with minimal guidance in instruction. This practice is quite general, but it obscures another meaning of the term constructivism: that learning is an active process, that knowledge is constructed. This is a very important point, about which there is considerable agreement in the research literature. In rejecting "constructivism" we do not want to revert to a view of learning as passive knowledge acquisition. The active role that the learner plays in acquiring knowledge must be clearly understood. Learners are not simply receiving information or acquiring knowledge by osmosis, but must be actively engaged in knowledge building. The role of instruction is to constrain and guide their activities. The question of how much guidance is optimal for learning is a separate issue.

This is not to say that Kirschner et al. (2006) advocate a view of learning as passive information reception. They are quite clear and explicit about this: the goal of instruction is to alter long-term memory, and long-term memory is not a passive repository of information; knowledge in long-term memory must be constructed. Thus, what I am discussing here is nothing new. Nevertheless, it is an issue that could use further clarification. Although the terminological confusion in the term constructivism is clearly recognized by Kirschner et al. (2006) as well as the authors who replied to their article, it may easily be misunderstood by some readers. Therefore, I would like to elaborate on how knowledge is constructed, on the differences between novices and experts, and on the role of guidance in instruction. I shall do this from a viewpoint that is a little different. The discussion so far has focused on problem solving, whereas I propose to view the issue of constructivism from the viewpoint of comprehension, specifically text comprehension. There is a rich literature in this area that, as I shall show, complements the literature on problem solving in sometimes illuminating ways. Furthermore, text comprehension is not a well-structured domain such as problem solving in mathematics or physics, and extending the discussion beyond such

* I thank Gerhard Fischer and Eileen Kintsch for their helpful comments.

domains would be useful, as Schmidt, Loyens, van Gog, and Paas (2007) have suggested.

My goal here is to distinguish clearly the constructive aspects of learning, the process of knowledge construction, from the question of how much guidance is optimal for learning. Although minimal guidance and discovery learning have frequently been advocated by constructivists, minimal guidance does not necessarily follow from a constructivist view of learning. Instructional methods are most effective when they respect the view of learning as an active (and, indeed, often effortful) process, with the right amount of guidance determined by the characteristics of the learner and the to-be-learned material—which is not necessarily minimal guidance. Again, there is nothing new about this claim: Kirschner et al. (2006), as well as Hmelo-Silver, Duncan, and Chinn (2007) and Schmidt et al. (2007) explicitly agree that the level of guidance for optimal learning must be adapted to the learner and the material they are supposed to master (although they might disagree on what constitutes a minimal and optimal level of guidance). However, considering how this issue plays out in the domain of text comprehension might help us to obtain a better grasp of it.

Learning as an Active Process

What do we mean when we say "Learning is an active process"? We first need to specify the term "learning," as there are many types of learning. Pavlovian conditioning, operant conditioning, associative learning, skill learning, rote memorization, learning by doing, and leaning from text differ in important ways. The focus here is on school learning, that is, the processes whereby students acquire knowledge and skills in school settings. Indeed, for clarity and specificity, I shall limit this discussion to a particular type of school learning—learning from texts.

To see why learning must be regarded as a constructive process, consider the input and end result of that process (see Kintsch, 1998, for more detail). The input is a text, that is, a series of written words, organized into sentences, paragraphs, and higher-order discourse units. The end result is a situation model that faithfully represents the meaning of that text, both at a local and global level, and integrates it with the reader's prior knowledge and learning goals. Turning the written text into a situation model in the reader's mind requires going beyond the written word. Even constructing a decent representation of the text itself—a textbase—requires active processing, for texts are never fully explicit. Inferences of several kinds are required from the reader—referents have to be identified, coherence gaps have to be bridged, the macrostructure of the text must be mentally represented. A well-written text gives the reader all kinds of cues on how to go about textbase construction, but it is up to the reader to infer the discourse entity a pronoun refers to, to come up with the right bridging inference linking two seemingly unrelated sentences, or a suitable high-level generalization to characterize a macro-unit of the text. The passive reader, who does not perform this required activity, will end up with an inadequate textbase. But the activity required for the construction of a textbase is much less of a problem for most readers than that required to construct a good situation model. After all, the text usually cues the reader on how to construct a textbase, but for the construction of the situation model the reader is on his/her own. It is their specific background knowledge that matters, their particular interests and reading goals that have to be integrated with the text, and the text cannot provide detailed guidance for every reader, since knowledge and goals differ widely among readers. Thus, a major problem in school learning is the student's failure to construct a situation model at all, or the inability to construct an adequate one.

What kind of situation model will be constructed depends, inter alia, on the reader's goals. Readers whose goal is to prepare for a test emphasizing fact retrieval will focus on different aspects of the text than readers who try to understand the text in preparation for a class discussion. Similarly, reading for appreciation or reading for doing will give rise to different situation models. Thus, an important aspect of a situation model is how students perceive the learning environment, which depends on how that environment is implemented by the teacher. Hence the kind of expectations the teacher creates in a classroom play a large role in fostering either superficial reading or deep understanding. In our experiments on text recall with college students we regularly find that they faithfully reproduce whatever names and numbers there are in the text, because that is the sort of thing they are often asked about in tests, but fail to generate inferences that would result in a deeper understanding of the text, even when they are able to recall all the premises.

Since situation models link an individual's background knowledge and personal experience, goals and purposes with information from the text, they differ among individuals more than textbases, which generally hew closely to the text. Nevertheless veridical situation models have much in common, because they must be constrained by the text (van Dijk & Kintsch, 1983; Trabasso & Suh, 1993; Zwaan & Radvansky, 1998; Tapiero, 2007; Therriault & Rinck, 2007). The situation model includes not only verbal or propositional information like the textbase, but may also include sensory imagery (either retrieved from prior knowledge or constructed on the basis of the text itself), emotional markers, and action plans. Importantly, situation models are cumulative: as one reads more and more on a given topic, the situation model changes, not only by accretion, but also by reorganization and error correction. The situation model is the product of the learning process. Bereiter and Scardamalia, in talking about learning in general, not just learning from texts, have used the term knowledge building for this process (Bereiter, 2002; Bereiter & Scardamalia, 2003; Scardamalia & Bereiter, 2003).

The goal of instruction is to make knowledge building possible. Two aspects of knowledge building are critical for instruction. First, in the words of Harel and Papert (1991), "(the building of knowledge structures) happens especially felicitously in a context when the learner is consciously engaged in constructing a public entity, whether it's a sand castle on the beach or a theory of the universe," or, we might add, when the learner summarizes a report or writes a critical essay. Second, knowledge objects do not stand alone but are grounded on a shared, cultural knowledge base (Hirsch, 1987, 2006). Not only is cultural knowledge necessary for understanding objects in a culture (such as texts); common knowledge also assures that different individuals in that culture build situation models that share important features, thereby becoming members of a cultural community. It must be emphasized, however, that the way that knowledge is acquired, and the way it is used in building new knowledge, is an active, constructive process. Cultural objects become building blocks for the construction of knowledge, not something that can be absorbed through passive reading. The question of how much guidance is optimal to facilitate knowledge building will be considered below, but first we need to discuss knowledge building in more detail.

It is important to understand the difference between how experts and novices go about constructing a situation model (Ericsson & Kintsch, 1995; Kintsch, 1998). In both cases, the text is processed in cycles corresponding more or less to sentence units. The information about the sentence currently being read is held in working memory and processed in various ways. For instance, the inferences required to form a coherent textbase are performed, often automatically (as in the case of most bridging inferences). but sometimes they depend on conscious search and reasoning. Novices are undoubtedly not as good as experts at these tasks, but there is really no qualitative difference in what they do. However, processing is not restricted to the level of the textbase, since at the same time a situation model is also being constructed. Here, experts and novices differ qualitatively. Domain experts, reading texts within their domain of expertise, have available retrieval structures that link the information in working memory—the current portion of the text they are working on—automatically with relevant information in their long-term memory. Thus, as they read, a situation model is formed, largely without conscious effort. Reading goals come into play in much the same way, via retrieval structures that favor goal-relevant portions of long-term memory. Thus, their reading results in the automatic generation of a situation model that integrates the textbase with their prior knowledge and that is structured in a way that reflects their interests and goals.

Novices are in a very different situation with regard to text comprehension and learning from texts. By the upper grades in school many students have successfully acquired expert adult reading strategies enabling them to form an adequate textbase more or less automatically and effortlessly. These strategies receive a great deal of practice during the school years, and to some extent there is transfer from general comprehension strategies used in listening, which are practiced even more throughout life. Thus, many young adults are no longer novices in this sense—they are practiced readers, even experts, so long as their reading material concerns everyday matters for which they have adequate background knowledge. However, there are many readers in middle school and high school, and even in college, for whom that is not the case. For these, the formation of a textbase remains a task that involves conscious effort, even when they are reading about familiar topics. They must employ explicit strategies to assure comprehension, strategies that must be directly taught. Thus the novice reader has to learn to consciously search for relevant prior knowledge, since it will not be automatically activated; he must learn to ask himself what the author meant with a particular sentence, or why the author said it, or how it relates to what was said before; and he must learn to discern what is the main argument and what are ancillary points, where the author presents evidence and where she makes claims.

But even if students are expert in reading general, familiar texts (such as stories or newspapers), this is seldom the case with the instructional texts that they read in school. They are lacking the retrieval structures that ensure smooth comprehension for domain experts when they read these texts. Retrieval structures are links between items held in working memory (roughly conscious awareness) and relevant associated knowledge in long-term memory that is thereby activated without overloading the limited-capacity working memory. Knowledge activation with retrieval structures is automatic and effortless, and hence characterizes expert knowledge and reading behavior (Ericsson & Kintsch, 1995). However, expert retrieval structures are the product of extended practice in knowledge building, the kind of deep comprehension that results in a well-grounded situation model. Constructing a situation model for novices requires conscious, effortful memory searches to retrieve relevant background knowledge and the use of explicit comprehension strategies to compensate for their lack of automatic retrieval structures. Novices must problem solve their way through the text by identifying places that call for elaboration and clarification. They must paraphrase and re-explain text passages in their own words to explicate the relation between the new information in the text and what they already know.

Thus, we can distinguish four types of readers: readers with good general comprehension strategies and expert domain knowledge; good readers without domain knowledge; poor readers without domain knowledge; and, finally, poor readers with high domain knowledge (such as the soccer experts studied by Schneider, Körkel, & Weinert, 1989). Most students fall into either one of the two middle categories. For these students, comprehension is an active, effortful, resource-demanding construction process as described above. The role of instruction is to support this process (Brown & Campione, 1994; King, 1997; Palincsar & Brown, 1984; Pearson & Fielding, 1991; Scardamalia, Bereiter, & Lamon, 1994).

A number of instructional implications for readers lacking domain knowledge follow from such an active view of learning. Two examples that we have explored in our work will be briefly discussed here.

Deep comprehension, and hence learning from text, is not possible unless there is at least some background knowledge present. Thus, educators try to assign instructional texts that are attuned to the students' level of prior knowledge. The texts need to be within a "zone of proximal learning" to make knowledge acquisition possible (Kintsch, 1994; Wolfe et al., 1998). Just how much one must know about a topic before one can learn more about it, and what one does not have to know but can learn on demand is by no means a straightforward problem, however.

Another implication of the constructivist view of learning is the role of metacognition. If meaning construction is an effortful, demanding process, readers may try to get by without the effort. The easiest way to do so is by not thoroughly analyzing the level of comprehension that is being achieved. Since a superficial level of understanding is easy enough to attain, students need to learn that this is not sufficient. By superficial, I mean understanding at the textbase level, without forming a reasonable situation model. Such understanding is good enough to reproduce the gist of the text and some of its detail, but it remains inert knowledge, unconnected to a person's store of knowledge and hence it is easily forgotten and unusable in novel situations (Kintsch, 1998, Chapter 9).

One way to make readers aware of their lack of comprehension is to problematize instruction (Reiser, 2004). An experiment that nicely illustrates the possibilities and pitfalls of such an approach was reported by McNamara, Kintsch, Songer, and Kintsch (1996). This study employed readers with good background knowledge and readers with inadequate background knowledge with respect to a particular science text they were asked to learn. Two versions of the text were used: one was well-written, well-structured, explicit, and in general provided the reader with all the support that was possible. The other version was deliberately lesioned: there were gaps in the text requiring bridging inferences, undefined terms, the organization of the text was obscured, and in general the text was made hard to read.

The results of this study were instructive. When only superficial comprehension was required (the ability to recall part of the text), the well-written version was always superior, whether students had high or low background knowledge. However, when the text required deeper understanding, that is, a well-worked-out situation model (as assessed by inference questions and a problem-solving task), students with high background knowledge performed better with the poorly written text than they did with the well-written version. With the well-written text, they easily formed a good textbase, which made them think they understood what there was to understand, but they never did the processing that was required for building a situation model (in spite of their background knowledge, they were far from being domain experts who would have understood this text without further effort). When the text was difficult, they realized they did not understand it well and were forced into processing at a deeper level, with beneficial results for their understanding.

Our high-knowledge students had just enough background knowledge to make it possible for them to draw the required inferences and construct an adequate situation model. For the students with little background knowledge of the domain, the situation was different. When the text was well-written, they could at least come up with a good textbase, and hence were able to reproduce the text. But they could not construct the required situation model, for the text lay outside their zone of proximal learning. When they were given the difficult text to study, they were lost: lacking the knowledge to fill in coherence gaps or identify the referents, they became confused by the lack of a clear organization in the text, and hence they could neither recall it well, nor understand it at a deeper level. Problematizing a text is fine, but you need to make sure that the reader can solve the problem!

The claim so far has been that successful learning from texts requires the construction of a good situation model, which, unless the reader is already a domain expert, is a resource-demanding process involving conscious effort. So what are the conditions that are necessary for the success of this effort?

Cognitive-load theory (Sweller, 1988) provides some useful answers. The reader's resources are finite and they have to be used purposefully. We have seen that increasing the cognitive load under certain circumstances can improve learning (the high-knowledge readers with the poorly written text in the McNamara et al. experiment). In that case the increased cognitive load arose from activities that were directly relevant to the learning process: in overcoming the difficulties of the text, the readers were able to reach a deeper level of understanding. Increases in cognitive load are not beneficial, however, when the activities involved are extrinsic to the learning process. Thus, reading a text in a foreign language may overload working memory because the text has to be translated to be understood, with the result that very little information about that text reaches long-term memory.

Motivating learners to expend the required effort to construct a situation model is just as important as considerations of cognitive load. Project-based learning and related approaches that engage the student in problems that are interesting and relevant to the student offer great possibilities in this respect. As Kuhn (2007) has pointed out, learners make their own choices about how to construct knowledge. They need to have some reason to be interested in what they are doing (Hidi & Renninger, 2006). In our work, we have used a software program in several hundred classrooms in Colorado to teach middle school students how to write summaries. The program was demonstrably successful in doing so, but not in all classrooms. One crucial variable was whether the teacher gave the students a good reason for the activity of summary writing (e.g., preparing for a class discussion or a presentation, or a particular project) or whether the program was introduced as just another decontextualized activity, in which case students typically failed to learn anything (Caccamise et al., in preparation). Learning is an active process, it is the student who must be active, and instruction must provide reasons for the active effort, which can be done by engaging the student's interests and motivation. Problem-based learning, project-based learning, or scientific inquiry appear to be effective means toward that end. It is important, however, that the problem-solving activity does not become a goal in itself when the real goal is to learn the science involved. The project is not the important outcome—science knowledge is.

Learning and Learning-to-Learn

Learning from texts often has two goals that operate simultaneously. Suppose I ask students to study a text on the theory of plate tectonics. I want them to learn about plate tectonics—what are the claims, the data, the controversies, etc. To do so, the students must form a good situation model that integrates the text they have just read with their prior knowledge about geology and geography, as discussed above. But I also have a second goal: I want my students to become not exactly experts in geology, but more expert-like in their ability to read science texts.

Simon's estimate that it takes 10,000 hours or 10 years of deliberate practice to become an expert in any area of science, sport, or the arts is widely accepted today (Simon, 1996). Schools, generally, do not produce real experts, but strive to move students a bit closer to expertise and provide them with the tools to develop further on their own. Therefore, it is worthwhile to look at the literature on expertise for hints about how to become an expert (Ericsson, Charness, Hoffman, & Feltovich, 2006). What strikes one first is the sheer amount of practice necessary for expertise, not just practice but deliberate, guided practice. The important lesson for classroom instruction and learning is the need to provide the opportunity for guided practice of the skill that is to be learned, including the skill of text comprehension.

Guided practice is best illustrated by what a sports coach is doing, say a ski instructor (Fischer, Brown, & Burton, 1984). On the one hand, the instructor provides feedback about the student's current performance, and on the other she selects new, more advanced tasks for the student that are within his proximal zone of learning. This is a tricky business, for too much challenge just scares the skier off the slopes, while without challenge he will be consigned to the groomed slopes forever, and soon get bored with the sport. Csikszentmihalyi (1990) has discussed this dilemma in terms of maintaining the flow experience, which is threatened on the one hand by anxiety when the learner is over-challenged, and tedium when not challenged enough, in reading comprehension as in skiing. But while the pleasure of flow can play an important motivating role, learning is the result of deliberate practice, "in which individuals engage in (typically planned) activity aimed at reaching a level just beyond the currently attainable level of performance by engaging in full concentration, analysis after feedback, and repetitions with refinement" (Ericsson & Ward, 2007, p. 349). Thus the flow is not the goal of instruction, learning is, which is hard work, but the flow may provide motivation to engage in the hard work of deliberate practice.

Instruction, therefore, must provide students with ample opportunity for guided practice. If we want them to acquire expert-like strategies for reading science texts, just making them read (and maybe take a test afterwards) will not suffice. We must provide feedback that allows the student to assess her current level of understanding, hints about what to do when her understanding is inadequate, and we must carefully select new texts to be studied that afford the student opportunities to learn more advanced strategies. If we want to teach students how to summarize, we must give them feedback about what they have written, hints on how to improve it, and the opportunity to work on more and more difficult tasks. Franzke, Kintsch, Caccamise, Johnson, and Dooley (2005) provide an example of such an approach to teaching summarizing, embedded in a computer-based tool, called *Summary Street*. Using the content-based feedback delivered by the system, students not only improved the quality of their summaries, but the benefits persisted over time, even when the students summarized without the support. Especially telling is the fact that students in the control group, who summarized an equivalent number of texts without the guidance from *Summary Street*, did not improve at all. They made the same errors after practice as they did before. Similarly, mere activity in a flight simulator does not improve the performance of pilots, but guided practice does (see the discussion in Ericsson & Ward, 2007).

It is clear that one cannot become an expert without guided practice, a great deal of guided practice, in fact. But what kind of guided practice? Two approaches have been suggested, one involving the teaching of general thinking skills, the other focusing instead on domain-specific strategies. The first modern theories of problem solving emphasized general problem-solving strategies (Newell & Simon, 1972). Instruction, accordingly, should focus on teaching "students to use their minds well,... skills of inquiry and skills of argument" (Kuhn, 2007). However, it soon became apparent that effective problem solving tends to be domain specific: expert problem solving is characterized by the use of domain-specific strategies rather than general problem-solving skills (Chase & Simon, 1973; Schunn & Anderson, 2001). Experts develop retrieval structures that link particular contents in working memory with relevant knowledge in long-term memory. Both the patterns in working memory that trigger the operation of retrieval structures and the contents of long-term memory are highly domain specific. Problem-solving skills are situated and in general do not transfer across situations. Being an expert in one domain does not make one an expert in a different domain.

This domain specificity of expertise poses a serious dilemma for schools. Schools (preceding law school or medical school) are not expected to produce experts, but well-rounded general problem solvers who can function in many different environments and are capable of becoming experts in some environment with further training. Thus, school learning is designed to un-situate, decontextualize knowledge and skills, so as to make it flexible and usable in a variety of situations. As is well-known, that is a tricky task, because there is a delicate balance between knowledge that is so situation-bound that it is usable only in that very situation and knowledge that is so decontextualized that it becomes inert knowledge, usable in no situation whatever. School learning is always in danger of producing inert knowledge, but nevertheless its goal remains to provide students with general knowledge that is broadly usable and is not tied to the context of its acquisition (Bereiter, 1997). Schools need to teach students to construct knowledge at the right level of abstraction, knowledge that is neither limited to concrete situations nor completely decontextualized, but rather linked to abstract, generalizable features of situations.

Irrespective of the importance of domain-specific strategies and domain knowledge, general-purpose strategies also play a role in thinking and problem solving. First of all, it has been shown that expert problem solving is not limited to domain-specific strategies but typically employs a mixture of general and specific reasoning methods (e.g., Greeno, 1983; Duncan, 2007). Second, there are at least two genuine skills that are generalizable across domains: metacognitive strategies (Flavell & Wellman, 1977) and reading-comprehension strategies (Perfetti, 1989). By the time we are young adults, most of us are expert comprehenders: we have had many years of practice with spoken-language comprehension, as well as more than a decade of practice (not always deliberate) with reading comprehension. We read fluently and comprehend automatically—but only as long as we read familiar texts (Ericsson & Kintsch, 1995), like true experts who rely upon automatic retrieval structures. But a curious thing happens when we read texts in an unfamiliar domain, say about string theory or meiosis and mitosis, rather than the daily newspaper or an airplane novel. To read such texts we need two kinds of expertise: general reading skills as well as domain expertise.

Students, when they read for learning, may be good readers, but they still do not comprehend because of their lack of domain knowledge. Reading-comprehension strategies can help them deal with this situation. Where the domain expert would rely on automatic retrieval structures to construct a valid situation model, the reader who is not a domain expert must consciously and intentionally go through the many steps required in this process—make inferences to fill gaps in the text, retrieve relevant background knowledge, identify the structure of an argument, and so on. The teacher can model the required behaviors, teaching in effect comprehension strategies that students can use to achieve their goals. There exists a great deal of evidence that such strategies can be very helpful to students, that they smooth the path towards expert comprehension (e.g., Palincsar & Brown, 1984; McNamara, 2007). Thus, reading-comprehension strategies play an important role in instruction in helping the learner to build new knowledge from the instructional texts they read in school.

There are some open questions about how comprehension strategies should be taught. Some computer tutors (such as *Summary Street*, Caccamise et al., in preparation; Franzke et al., 2005) do not teach students specific comprehension strategies, but guide their practice through judicious feedback so that they learn to adopt suitable strategies. In contrast, other successful systems (such as reciprocal teaching, Palincsar & Brown, 1984; questioning the author, Beck & McKeown, 2006; or *iSTART*, McNamara, 2007) explicitly teach relevant comprehension strategies, providing students with a set of consciously available comprehension tools. It is not clear at this point what combination of these approaches is most effective, for which students, and at what stages of learning.

An example of how learning from a text can be combined with training in general comprehension strategies can be found in the teacher and student manuals of Hampton (2007) for Grades 7 and 8. She has students working on strategies for constructing a faithful textbase (such as pronoun identification, sentence connectives, vocabulary, text structure) as well as for building a situation model (think-aloud and teacher modeling, discussion, summarizing, and essay writing). There is nothing new about these strategies; what is different here is how they are embedded within the textbase–situation model framework and, crucially, that this is all done within one specific knowledge domain that is relatively familiar to students to begin with and that is systematically expanded during the course of 30 lessons. Thus, students learn at the same time a set of general reading strategies and build a cumulative situation model about an important concept in biology. For the students, the strategy knowledge remains implicit; for the teachers, a general understanding of how comprehension works provides a meaningful framework for their activities.

Levels of Guidance for Learning

Learning from text is by its very nature a constructive process, guided through feedback. Learning effective comprehension strategies requires a great deal of deliberate practice, which also implies some kind of guidance. The guidance can come from the teacher, the nature and organization of the instructional texts, or it can be self-guidance through metacognitive control. At one extreme we have direct instruction, where the teacher and the learning materials firmly guide the learning process, leaving little to the discretion of the student. At the other extreme would be completely unguided discovery learning, with, of course, numerous shades of gray in between. The eventual goal is to have a self-guided learner, but what is the best road to that goal is not so clear, which is one of the things the "constructivist" controversy is all about. Other chapters in this volume speak to this general issue. Here, I am merely exploring the implications of the literature on text comprehension with respect to that complex question.

One area where the guidance issue has been extensively explored is hypertext. The familiar linear text guides the reader by sequencing the text in the order the author thought would be optimal. Sometimes, however, the way the author has ordered an expository text does not mesh well with a reader's goals. For instance, if a reader is looking for certain pieces of information, these may not be easy to locate in a linear text. A well-organized hypertext with a transparent structure and proper navigation aids can be searched much more efficiently. So, for instance, when students are given a particular problem to solve for which purpose they need to find relevant information from textual sources, hypertext can be very helpful (Dillon & Gabbard, 1998).

A different question is whether hypertext is a good alternative to linear text for promoting learning, when the goal is not to find a specific piece of information needed to solve a problem, but to acquire knowledge about some domain. Originally, the expectation was very much that hypertext would be helpful in this respect, too (McKnight, Dillon, & Richards, 1993). However, the burgeoning literature soon disappointed these expectations (e.g., Unz & Hesse, 1999). Hypertext users tend to use three different strategies for choosing which node to follow (Salmerón, Kintsch, & Cañas, 2006): (a) they choose the node that promises to provide the most coherent continuation of what they have just read; (b) they choose the node that looks like it would be of most interest to them; or (c) they follow some superficial strategy, like selecting the node printed at the top of the screen. If they follow the coherence strategy, they learn quite well (Foltz, 1996), but not otherwise. Specifically, low-knowledge readers do poorly at situation-model-level comprehension unless the text is presented in a coherent order (Salmerón, Cañas, Kintsch, & Fajardo, 2005) or follow a coherence strategy in selecting their own order (Salmerón et al., 2006). High-knowledge readers, on the other hand, can do well even if they follow an interest strategy in selecting nodes.

Giving a hypertext to high-knowledge readers is one of the techniques that can be used to ensure active reading and construction of a proper situation model—they cannot read superficially, because at each choice point they must select a good continuation. Since they have enough knowledge to either find a coherent continuation or, if they follow their own interests, are capable of forming coherence links with other parts of the text, this activity is beneficial for their comprehension, just as the need to fill coherence gaps and identify referents was beneficial for high-knowledge readers in the McNamara et al. (1996) studies with linear texts reported earlier. Increasing the cognitive load for these readers is more than balanced by the benefits of active processing. The situation is different for low-knowledge readers. If they do not follow a coherence strategy, they end up with a disorganized and fragmentary situation model; if they do, the increased cognitive load may not leave them with enough resources for processing other aspects of the text (Sweller, 1988). Thus, using hypertexts for learning problematizes comprehension, which can be beneficial, but carries considerable risk. One cannot claim that the level of guidance provided by linear text is necessarily superior to letting the reader make his own choices in a hypertext, but the conditions under which hypertext is superior for learning are narrowly circumscribed.

It seems not unreasonable to generalize the conclusions about the role of guidance arrived at above in the discussion of learning from hypertext. Minimally guided learning (hypertext is not unguided—you cannot just go anywhere, but must follow a given set of links, and typically you have available overviews and other navigation aids) can be as good as or better than guided instruction, but the potential risks must be carefully thought out and weighed for different kinds of learners.

What I have stressed here is that learning is a constructive activity, an active, intentional process that may demand considerable commitment and effort from the learner. It is difficult to state in general terms what the appropriate level of guidance is for the learning process. As we have seen, the amount of guidance needed differs, depending on the nature of the material, the background of the learner, as well as the stage of learning. The level of guidance should support the goal of keeping the learner actively engaged; it must motivate the learner, by challenging him or by interesting him, to engage in the laborious task of comprehension.

Conclusions

The title of this volume refers to the success or failure of "constructivist instructional theory." In these notes I have tried to forestall a possible terminological confusion about constructivism. At issue here is the effectiveness of minimally guided instructional methods, such as discovery, problem-based, experiential, project-based, and inquiry-based teaching, which are commonly labeled "constructivist." Constructivism, however, is also a theory of comprehension and learning. The central idea of this theory is that meaning must be constructed, that knowledge building is an active process on the part of the learner, not a passive process of information absorption. Just about every current learning theory is constructivist in that sense. Minimal guidance is a separate issue, but if we do not clearly and explicitly distinguish between these two uses of the term "constructivism" we invite confusion.

I have also tried to bring to the discussion a fresh viewpoint and a novel set of evidence. The discussion so far has been framed mostly in terms of problem solving in domains like math and physics. Here I have examined results from the field of text comprehension and learning from text, which nicely complement the literature on problem solving. Text comprehension is an ill-structured domain, unlike problem solving in formal disciplines. I have discussed the ways in which comprehension and learning from text are considered constructive processes. Central to this argument is the need to construct situation models on the basis of texts. The primary instructional problem is to get learners to construct adequate situation models and not be satisfied with a superficial understanding.

The cumulative construction and elaboration of situation models is a form of knowledge building. Since situation models always build on a foundation of prior knowledge, the process of situation-model construction is very different for domain experts and domain novices, smooth and automatic for the former and effortful and intentional for the latter. The crucial role of deliberate practice in becoming an expert was discussed.

As to the central question of the present volume—how much guidance is optimal for learning—the literature on text comprehension suggests a nuanced answer. Minimal guidance, such as in unconstrained discovery learning, is not generally effective, because it makes demands that easily exceed the resources of the learner, especially learners who lack appropriate background knowledge. However, maximal guidance, as in forms of instruction that reduce the learner to a passive information recipient, can also be counterproductive when it prevents the learner from the active, deep processing of the text that is required for the construction of adequate situation models. It is difficult to state in general terms what the optimal level of guidance is for the learning process. As we have seen, optimal levels of guidance differ, depending on the nature of the material, the background of the learner, as well as the stage of learning. The level of guidance should support the goal of keeping the learner focused on the topic of interest and actively engaged.

Question: Spiro and DeSchryver. *One of the great strengths of your chapter is the careful, nuanced demonstration that the answer to the question "How much guidance is optimal for learning?" depends on "the nature of the material, the background of the learner, as well as the stage of learning." Can you extend your conclusions beyond "how much" (levels or amounts of support) to questions of kind of guidance. Part of what we argued in our chapter is that the nature of optimal guidance shifts in more ill-structured domains from the kind that has been shown to be most beneficial in studies of predominantly well-structured domains. If this might be so, it would be useful in the constructivism–direct instruction debate to be able to specify where this shift occurs.*

You present text comprehension and learning from text as ill-structured domains, and point to kinds of guidance that have been shown empirically to be effective (e.g., those used in Reciprocal Teaching and Questioning the Author). To what extent do the kinds of strategies advocated in such approaches fall on the other side of a qualitative divide from the Kirschner et al. criteria? That is, would successful strategies for *learning from text be able to be characterized as fully explaining essential procedures? Or in domains like learning from text is the support necessarily of a vaguer kind (e.g., "Look for connections to prior knowledge")?*

Reply: Kintsch. You raise an important point. How much guidance is needed is only part of the story; what kind of guidance is needed is equally important, and it depends, too. But having agreed with you on this point, I must plead ignorance about how to frame an adequate answer to your question. Research like your own will surely provide more detailed answers, but at present we know little more than the bare outlines of this problem.

For well-formed problems, Kirschner et al.'s emphasis on essential information and full explanation seems right. An example from the field of text comprehension would be arithmetic word problems, and the kind of technological support that can guide the construction of situation models for students, as I have discussed in my answer to another question. But, as you suggest in your question, learning from text can make quite different demands. In some domains (you cite biology—not exactly an ill-formed domain, there are correct answers!) the distinction between essential and inessential information is difficult to make, and it is not even clear what would count as a full explanation. Content-free prompts can be as effective as explanation and feedback—but just where the boundaries are must await further research.

Reaction to Kintsch's reply: Spiro and DeSchryver. *We find ourselves happily in agreement with your response to our question, but we do have one point of exception, which we think is worth discussion: you contest our claim that* aspects *of biology are ill-structured. There are two important points here. First, we don't claim that domains are entirely well- or ill-structured, though some may be predominantly more one than the other. It's clear, for example, that Newtonian mechanics is predominantly well-structured and the concept of "period" in art history is predominantly ill-structured, according to our use of these notions. All domains have* both *well- and ill-structured aspects. Similarly, we don't claim that biology is predominantly ill-structured, just that some important aspects of it are. For example, the mechanisms of* adaptation *are both complex in individual instances and, more importantly, evince considerable conceptual irregularity across instances in the way conceptual features are instantiated and configured, making generalizations and abstractive reductions problematic. Further, understanding of adaptation instances often involves interpretive processes that it would be a stretch to refer to as "correct answers."*

The second important point is that we argue that instruction and support are a function of the degree of well-structuredness of a given aspect, not the overall, predominant pattern of the domain, and that the ill-structured aspects require a different kind of instruction and instructional support than the well-structured aspects. So, for example, we would expect there to be qualitative differences in the nature of instructional support for the more determinate microbiological arena than for some topics in the macrobiological realm (like adaptation). By the way, we take no stand on whether the ill-structured aspects of biology are that way in principle, or, instead, *are a reflection of limitations of current knowledge which might be remedied at some future time. The implications for instruction at this time would be the same.*

Question: Schwartz et al. *In our read of the chapter, we had a small confusion that we think has important ramifications. The confusion involved a separation between students learning "content knowledge" and students learning "process knowledge." It is an important question whether content knowledge should be taken as separate from the processes and contexts associated with its acquisition. However, our question is simpler. We did not understand why, on the one hand, you stated,*

> *Problem-based learning, project-based learning, or scientific inquiry appear to be effective means towards that end [motivation]; it is important, however, that the problem-solving activity does not become a goal in itself when the real goal is to learn the science involved. The project is not the important outcome—science knowledge is.*

But, on the other hand, in the context of reading instruction, you stated that you wanted "students to become not exactly experts in geology, but more expert-like in their ability to read science texts." Did you mean to imply that people should learn the content of the text and the process of reading well, whereas for science, the goal is learning science content but learning the process of inquiry is irrelevant except for its motivational value?

Reply: Kintsch. I expressed myself badly about the role of general problem-solving strategies, both in science learning and in reading instruction. Let me see whether I can get it straight. I do not think one can separate content and process in science learning. When I said "science knowledge," I meant knowledge both of content and process. In science, the process of inquiry is closely tied to the content: strategies are domain specific, as the problem-solving literature shows; general problem-solving strategies that could be used in all science domains play a minor role. The process of inquiry is sufficiently different in different branches of science, so that an expert in one domain has little advantage in another domain. Thus, science instruction must involve both content and process at the same time. When I objected to the emphasis on projects, I had in mind the fancy packaging and presentation that one sometimes observes in science fairs and class projects. This may have some motivational value, but should not be the primary focus of a science project.

With respect to reading instruction, things are a little different. I cited in my chapter a paper by Perfetti who claims that reading strategies are in fact general problem-solving strategies. Thus, teaching students general reading strategies and giving them plenty of opportunity for guided practice is indeed important. The eventual goal in developing reading expertise is the automatic use of these strategies.

Question: Schwartz. *Your distinction between the text-base and the situation model is powerful, well-supported, and it adds nuance to the otherwise flat claim* *that all knowledge is constructed. For readers who have not read your chapter yet, the distinction might be coarsely characterized as the difference between "parsing" the text and "understanding" the text. Have you or others applied your framework successfully to other domains where the primary input is also symbolic? For instance, when confronted with a mathematics word problem or equation, people might make a "symbol base" of the quantities in the problem and then a mental model of the situation to which the quantities could apply. Alternatively, the mental model may drive the construction of the symbol base. There are other possible applications of your framework, for example, when it comes to parsing graphs and constructing a mental model of their referents. It would be interesting to hear your thoughts on the construction–integration framework when applied to domains other than reading.*

Reply: Kintsch. I enthusiastically agree with your suggestion: there is a great deal that could be done along those lines. In fact, some time ago we worked out a model for understanding and solving word problems that made explicit the distinction between different levels of representations and their interdependence (Kintsch & Greeno, 1985) and designed a software program (Nathan, Kintsch, & Young, 1992) that guided students' efforts to solve algebra word problems by showing them the consequences at the situation-model level of what their equations implied. When they saw the faster plane leave before the slower one in an overtake problem, they immediately realized the mistake they had made! Much has been learned since then that would force us to modify some of the details of this work, but I am convinced that this remains a very promising approach where modern technology could have a direct impact on instruction, not only in mathematics, but in the other areas you mention too.

References

Beck, I. L., & McKeown, M. G. (2006). *Improving comprehension with Questioning the Author: A fresh and expanded view of a powerful approach.* New York: Scholastic.

Bereiter, C. (1997). Situated cognition and how to overcome it. In D. Kirshner & J. A. Whitson (Eds.), *Situated cognition: Social, semiotic, and psychological perspectives* (pp. 281–300). Hillsdale, NJ: Erlbaum.

Bereiter, C. (2002). *Education and mind in the knowledge age.* Mahwah, NJ: Erlbaum.

Bereiter, C., & Scardamalia, M. (2003). Learning to work creatively with knowledge. In E. De Corte, L. Verschaffel, N. Entwistle, & K. van Merriënboer (Eds.), *Powerful learning environments* (pp. 55–68). Amsterdam: Elsevier.

Brown, A. L., & Campione, J. C. (1994). Guided discovery in a community of learners. In K. McGilly (Ed.), *Classroom lessons: Cognitive theory and classroom practice* (pp. 229–272). Cambridge, MA: MIT Press.

Caccamise, D., Snyder, L., Johnson, N., Allen, C., DeHart, M., Kintsch, E., et al. (in preparation). *Summary Street: Scale-up and evaluation.*

Chase, W. G., & Simon, H. A. (1973). The mind's eye in chess. In W. G. Chase (Ed.), *Visual information processing.* New York: Academic Press.

Csikszentmihalyi, M. (1990). *Flow: The psychology of optimal experience.* New York: HarperCollins.

Dillon, A., & Gabbard, R. (1998). Hypermedia as an educational technology: a review of the quantitative research literature on learner comprehension, control and style. *Review of Educational Research, 68*, 322–349.

Duncan, R. G. (2007). The role of domain-specific knowledge in generative reasoning about complicated multileveled phenomena. *Cognition and Instruction, 25*, 271–336.

Ericsson, K. A., Charness, N., Hoffman, R. R., & Feltovich, P. J. (2006). *The Cambridge handbook of expertise and expert performance*. New York: Cambridge University Press.

Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review, 102*, 211–245.

Ericsson, K. A., & Ward, P. (2007). Capturing the naturally occurring superior performance of experts in the laboratory. *Current Trends in Psychological Science, 16*, 346–350.

Fischer, G., Brown, J. S., & Burton, R. (1984). Analysis of skiing as a success model of instruction: Manipulating the learning environment to enhance skill acquisition. In B. Rogoff & J. Lave (Eds.), *Everyday cognition: Its development in social context*. Cambridge, MA: Harvard University Press.

Flavell, J. H., & Wellman, H. M. (1977). Metamemory. In R. V. Kail & J. W. Hagen (Eds.), *Perspectives on the development of memory and cognition*. Hillsdale, NJ: Erlbaum.

Foltz, P. W. (1996). Comprehension, coherence, and strategies in hypertext and linear text. In J.-F. Rouet, J. J. Levonen, A. Dillon, & R. J. Spiro (Eds.), *Hypertext and cognition* (pp. 106–136). Mahwah, NJ: Erlbaum.

Franzke, M., Kintsch, E., Caccamise, D., Johnson, N., & Dooley, S. (2005). Summary Street®: Computer support for comprehension and writing. *Journal of Educational Computing Research, 33*, 53–80.

Greeno, J. G. (1983). Conceptual entities. In D. Gentner & A. Stevens (Eds.), *Mental models* (pp. 227–252). Mahwah, NJ: Erlbaum.

Hampton, S. (2007). *Foundations: Comprehending texts*. Washington, DC: America's Choice.

Harel, I., & Papert, S. (Eds.). (1991). *Constructionism*. Norwood, NJ: Ablex Publishing Corporation.

Hidi, S., & Renninger, K. A. (2006). The four-phase model of interest development. *Educational Psychologist, 41*, 111–127.

Hirsch, E. D. (1987). *Cultural literacy: What every American needs to know*. Boston, MA: Houghton Mifflin.

Hirsch, E. D. (2006). *The knowledge deficit*. Boston, MA: Houghton Mifflin.

Hmelo-Silver, C. E., Duncan, R., & Chinn, C. A. (2007). Scaffolding and achievement in problem-based and inquiry learning: A response to Kirschner, Sweller, and Clark (2006). *Educational Psychologist, 42*, 99–107.

King, A. (1997). ASK to THINK-TEL WHY: A model of transactive peer tutoring for scaffolding higher-level complex learning. *Educational Psychologist, 32*, 221–235.

Kintsch, W. (1994). Text comprehension, memory, and learning. *American Psychologist, 49*, 294–303.

Kintsch, W. (1998). *Comprehension: A paradigm for cognition*. New York: Cambridge University Press.

Kintsch, W., & Greeno, J. G. (1985). Understanding and solving word arithmetic problems. *Psychological Review, 92*, 109–129.

Kirschner, P. A., Sweller, J., & Clark, R. E. (2006). Why minimal guidance during instruction does not work: An analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching. *Educational Psychologist, 41*, 75–86.

Kuhn, D. (2007). Is direct instruction an answer to the right question? *Educational Psychologist, 41*, 109–113.

McKnight, C., Dillon, A., & Richards, J. (Eds.). (1993). *Hypertext: A psychological perspective*. Mahwah, NJ: Erlbaum.

McNamara, D. S. (Ed.). (2007). *Reading comprehension strategies*. New York: Erlbaum.

McNamara, D. S., Kintsch, E., Songer, N., & Kintsch, W. (1996). Are good texts always better? Text coherence, background knowledge, and levels of understanding in learning from text. *Cognition and Instruction, 14*, 1–43.

Nathan, M. J., Kintsch, W., & Young, E. (1992). A theory of word algebra problem comprehension and its implications for the design of learning environments. *Cognition and Instruction, 9*, 329–389.

Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Palincsar, A. S., & Brown, A. L. (1984). Reciprocal teaching of comprehension-fostering and monitoring strategies. *Cognition and Instruction, 1*, 117–175.

Pearson, P. D., & Fielding, L. (1991). Comprehension instruction. In R. Barr, M. L. Kamil, P. B. Bosenthal, & P. D. Pearson (Eds.), *Handbook of reading research* (Vol. 2, pp. 8115–8860). White Plains, NY: Longman.

Perfetti, C. A. (1989). There are generalized abilities and one of them is reading. In L. B. Resnick (Ed.), *Knowing, learning, and instruction* (pp. 307–336). Hillsdale, NJ: Erlbaum.

Reiser, B. J. (2004). Scaffolding complex learning: The mechanisms of structuring and problematizing student work. *The Journal of the Learning Sciences*, *13*, 273–304.

Salmerón, L., Cañas, J. J., Kintsch, W., & Fajardo, L. (2005). Reading strategies and hypertext comprehension. *Discourse Process, 40*, 171–191.

Salmerón, L., Kintsch, W., & Cañas, J. J. (2006). Reading strategies and prior knowledge in learning from hypertext. *Memory & Cognition, 34*, 1157–1171.

Scardamalia, M., & Bereiter, C. (2003). Knowledge building. In *Encyclopedia of education* (2nd ed., pp. 1370–1373). New York: Macmillan.

Scardamalia, M., Bereiter, C., & Lamon, M. (1994). The CSILE Project: Trying to bring the classroom into World 3. In K. McGilly (Ed.), *Classroom lessons* (pp. 201–228). Cambridge, MA: MIT Press.

Schmidt, H. G., Loyens, S. M. M., van Gog, T., & Paas, F. (2007). Problem-based learning is compatible with human cognitive architecture: Commentary on Kirschner, Sweller, and Clark (2006). *Educational Psychologist, 42*, 91–97.

Schneider, W., Körkel, J., & Weinert, F. E. (1989). Domain-specific knowledge and memory performance: A comparison of high- and low-aptitude children. *Journal of Educational Psychology, 81*, 306–312.

Schunn, C. D., & Anderson, J. R. (2001). Acquiring expertise in science: Exploration of what, when, and how. In K. Crowley, C. D. Schunn, & T. Okada (Eds.), *Designing for science: Implications from everyday, classroom, and professional settings* (pp. 83–114). Hillsdale, NJ: Erlbaum.

Simon, H. A. (1996). *The sciences of the artificial* (3rd ed.). Cambridge, MA: MIT Press.

Sweller, J. (1988). Cognitive load during problem solving. *Cognitive Science, 12*, 257–285.

Tapiero, I. (2007). *Situation models and levels of coherence*. Mahwah, NJ: Erlbaum.

Therriault, D. J., & Rinck, M. (2007). Multidimensional situation models. In F. Schmalhofer & C. A. Perfetti (Eds.), *Higher level language processes in the brain*. Mahwah, NJ: Erlbaum.

Trabasso, T., & Suh, S. (1993). Understanding text: Achieving explanatory coherence through on-line inferences and mental operations in working memory. *Discourse Processes, 16*, 3–34.

Unz, D. C., & Hesse, F. W. (1999). The use of hypertext for learning. *Journal of Educational Computing Research, 20*, 279–295.

Van Dijk, T. A., & Kintsch, W. (1983). *Strategies of discourse comprehension.* New York: Academic Press.

Wolfe, M. B., Schreiner, M. E., Rehder, R., Laham, D., Foltz, P. W., Landauer, T. K., et al. (1998). Learning from text: Matching reader and text by Latent Semantic Analysis. *Discourse Processes, 25*, 309–336.

Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin, 123*, 162–185.