The Knowledge-Learning-Instruction Framework: Bridging the Science-Practice Chasm to Enhance Robust Student Learning

Kenneth R. Koedinger,[a] Albert T. Corbett,[b] Charles Perfetti[c]

[a]*Human-Computer Interaction Institute and Department of Psychology, Carnegie Mellon University*
[b]*Human-Computer Interaction Institute, Carnegie Mellon University*
[c]*Learning Research & Development Center and Department of Psychology, University of Pittsburgh*

Abstract

Despite the accumulation of substantial cognitive science research relevant to education, there remains confusion and controversy in the application of research to educational practice. In support of a more systematic approach, we describe the Knowledge-Learning-Instruction (KLI) framework. KLI promotes the emergence of instructional principles of high potential for generality, while explicitly identifying constraints of and opportunities for detailed analysis of the knowledge students may acquire in courses. Drawing on research across domains of science, math, and language learning, we illustrate the analyses of knowledge, learning, and instructional events that the KLI framework affords. We present a set of three coordinated taxonomies of knowledge, learning, and instruction. For example, we identify three broad classes of learning events (LEs): (a) memory and fluency processes, (b) induction and refinement processes, and (c) understanding and sense-making processes, and we show how these can lead to different knowledge changes and constraints on optimal instructional choices.

Keywords: Instructional principles; Learning principles; Knowledge representation; Cognitive modeling; Education; Experimentation; Cognitive task analysis

1. Introduction

A substantial base of highly refined and extensively tested theories of cognition and learning (e.g., Anderson & Lebiere, 1998; McClelland & Cleeremans, 2009; Newell, 1990;

Correspondence should be sent to Ken Koedinger, Human-Computer Interaction Institute, Carnegie Mellon University, Pittsburgh, PA 15213. E-mail: koedinger@cmu.edu

Sun, 1994) provide broadly useful, but limited guidelines for instruction. Theoretical frameworks and design methods have been proposed that are directly relevant to instructional decision making (Bruner, 1966; Engelmann & Carnine, 1991; Gagne, 1985; Sweller & Chandler, 1994; van Merrienboer & Kirshner, 2007). However, we need instructional applications that are more grounded in cognitive theories that explicate the role in learning of basic cognitive science concepts, such as memory, chunking, analogical transfer, reasoning, and problem solving. To bring instructional decision making into closer contact with the science of learning, we propose a *theoretical framework* that builds on the computational precision of cognitive science while addressing instruction, not as an additional consideration, but as part of its basic conception. Our framework embodies research-based principles and theoretical propositions, while demonstrating pathways to instructional decision making. In doing so, the framework supports ''rigorous, sustained scientific research in education,'' as urged by the National Research Council report (Shavelson & Towne, 2002).

In the following sections, we first develop the basic background that motivates a general theoretical framework for improving student learning. We then define and explain the basic concepts that shape the theoretical framework. We follow with elaboration and examples across academic domains, and exemplary instructional principles. We conclude by suggesting how the framework can aid a sustained research effort and generate new hypotheses.

1.1. Educational context and the need for instructional principles

Heated debates in education (''education wars'') have been most visible in reading (e.g., Rayner, Foorman, Perfetti, Pesetsky, & Seidenberg, 2001) and math (e.g., Schoenfeld, 2004). Research has informed these debates, particularly in reading. However, its impact is blunted by larger contextual factors, including assumptions of both researchers and policy makers about educational values, child development, and research standards. For example, advocates in these debates tend to differ on the kinds of scientific support they value. Some emphasize rigor (e.g., internal validity, randomized controlled experiments) and others, relevance (e.g., ecological and external validity and appropriate settings and assessments). Our approach addresses this divide by embedding rigorous experimental research methods in the context of real learning, with real students, in real contexts (see the Pittsburgh Science of Learning Center at learnlab.org).

Our emphasis on connecting learning research to instructional principles is not unique. Research in the learning sciences has led to statements of principles, including the influential NRC report ''How People Learn'' (Bransford, Brown, & Cocking, 2000), instructional design principles (e.g., Clark & Meyer, 2003), and learning principles drawn from psychology research (http://www.psyc.memphis.edu/learning/principles/; Graesser, 2009). Progress in the ''reading wars'' was marked by a succession of evidence-based recommendations (National Reading Panel, 2000; Snow et al., 1996). Broad learning principles have been directed specifically at teachers and educational decision makers (e.g., http://www.institute-forlearning.org/; Pashler et al., 2007). We join the tradition of extracting principles from research because we believe that finding the right level to communicate evidence-based generalizations is important for guidance in instruction. However, we also are committed to

the idea that principles must be based on sound evidence that supports their *generality* and identifies their limiting conditions.

1.2. Levels of analysis

Finding the right level of analysis or grain size for theory is a major question. In all theories, including physical theories, the theoretical entities vary in the levels of analysis to which they apply, both in terms of the grain of the unit (micro to macro) and its functional level (neurological, cognitive, behavioral, social). We refer to these two together as ''grain size'' and follow past researchers (Anderson, 2002; Newell, 1990) in using the time course of cognitive processes to distinguish levels of analysis. The requirement of having an instructionally relevant theory requires a grain size that is intermediate between existing theoretical concepts in education and cognitive psychology.

Education theories have tended to use macro levels of analysis with units at large grain sizes. For example, situated learning, following its origins as a general proposition about the social-contextual basis of learning (Lave & Wenger, 1991), has been extended to an educational hypothesis (e.g., Greeno, 1998). It tends to use a rather large grain size including groups and environmental features as causal factors in performance and to focus on rich descriptions of case studies. These features work against two important goals that we have: (a) identifying mechanisms of student learning that lead to instructional principles, and (b) communicating instructional principles that are general over contexts and provide unambiguous guidelines to instructional designers. We assume that learning is indeed situated, and understanding learning environments, including social interaction patterns that support learning, is an important element of a general theory. But we take the interaction between a learner and a specific instructional environment as the unit of analysis.

Learning theories have small grain sizes as well. At the neural level, variations on Hebb's idea (Hebb, 1949) of neuronal plasticity refer to a basic mechanism of learning, often expressed simplistically as ''neurons that fire together, wire together.'' This idea has been captured by neural network models of learning mechanisms that use multi-level networks (e.g., O'Reilly & Munakata, 2000). Micro-level theories also give accounts of elementary causal events using symbols, rules, and operations on basic entities expressed as computer code, as in the ACT-R theory of Anderson (1993). Although initially developed without attention to biology, ACT-R has been tested and extended using brain imaging data (e.g., Anderson et al., 2004). It is important for the learning sciences that such theories demonstrate the ability to predict and explain human cognition at a detailed level that is subject to empirical testing. The small grain size of such theories and the studies that support them leave them mostly untested at the larger grain size of knowledge-rich academic learning. Thus, they tend to be insufficient to constrain instructional design choices.

The level of explanation we target is intermediate to the larger and smaller grain sizes exemplified above. This level must contain propositions whose scope is observable entities related to instructional environments and learner characteristics that affect learning. These propositions must be testable by experiments and allow translation, both downward to micro-level mechanisms and upward to classroom practices. The theoretical framework we

describe in section 2 uses this intermediate cognitive level, which we refer to as the *Learning Event level*. First, we briefly describe a research context for theory development.

1.3. Generalizations and constraints

Among many factors that can limit the application of even principles backed by solid evidence are two that are especially important for the application of principles to educational settings: variations among subject matter domains and among students. The content of a domain is a specific challenge to a general learning principle. Physics is not the same as history, and neither is the same as language. Our approach is to develop a domain-independent framework for characterizing knowledge structures that in turn captures differences in the structure of content knowledge across domains. Thus, the KLI framework is domain-independent in its general structure and domain-sensitive at the level of the knowledge components (KCs) that are central to its deeper analysis of learning. This allows the framework to cover multiple academic content areas so that generalizations and limits on generalizations become visible.

Variation in student characteristics is a more complex challenge. As studies must be done with students in specific school and social-geographical settings, strong generalizations demand new studies beyond the original samples. We draw primarily from studies at the junior high school through college levels in K12 Schools and University settings. However, every theoretical proposition and instructional principle derived from this research is an empirically testable hypothesis for other populations.

We address our goal of grounding our framework in application settings by emphasizing research within actual courses and their students. We call this in vivo experimentation (e.g., Koedinger, Aleven, Roll, & Baker, 2009; Salden, Aleven, Renkl, & Schwonke, 2008; see also learnlab.org/research/wiki). We address challenges of different domains and student characteristics in terms of differences in knowledge demands (section 3) and stage of learning (sections 4 and 6).

2. The knowledge-learning-instruction framework

Many efforts at instructional ''theory'' are really frameworks (e.g., Bloom, 1956; Gagne, 1985; van Merriënboer & Sweller, 2005; Sweller & Chandler, 1994) because they do not lead directly to precise predictions. Nevertheless, a theoretical framework does entail a hypothesis-testing research agenda:

> *Frameworks* are composed of the bold, general claims … They are sets of constructs that define important aspects of [interest] … Frameworks, however, are insufficiently specified to enable predictions to be derived from them, but they can be elaborated, by the addition of assumptions, to make them into *theories*, and it is these theories that generate predictions. A single framework can be elaborated into many different theories. (Anderson, 1993, p. 2)

The propositions within the KLI framework can help generate research questions within specific domains and instructional situations that, with further work, yield precise and falsifiable predictions. However, our main goal here is to identify the broad constructs and claims that serve more specific instantiations. We pursue this goal by specifying three taxonomies, kinds of knowledge, kinds of learning processes, and kinds of instructional choices, and dependencies between them. We show how kinds of knowledge constrain learning processes and how these processes constrain which instructional choices will be optimal in producing *robust student learning*. Learning is robust when it lasts over time (long-term retention), transfers to new situations that differ from the learning situation along various dimensions (e.g., material content, setting, cf., Barnett & Ceci, 2002; Chen & Klahr, 2008), or accelerates future learning in new situations (Bransford & Schwartz, 1999).

2.1. Learning events, instructional events, assessment events

The KLI framework relates a set of observable and unobservable events: learning events (LEs), instructional events (IEs), and assessment events (AEs), as illustrated in Fig. 1. Instructional and AEs are observable activities in or changes to the instructional environment controlled by an instructor, instructional designer, or experimenter. *IEs*, which are intended to produce learning (they cause LEs), can be observed in a lesson from a 10 s episode on a computer to a series of didactic moves by a teacher or computer tutor. *AEs* involve student responses that are evaluated. Some are instructional; some are not. Although AEs are usually test items, they can also be embedded in the context of instruction, for instance, tracking whether a student is correct on their first attempt in a tutored problem-solving scenario (e.g., Aleven, Roll, McLaren, & Koedinger, 2010; Feng, Heffernan, & Koedinger, 2009). *LEs* are essentially changes in cognitive and brain states that can be inferred from
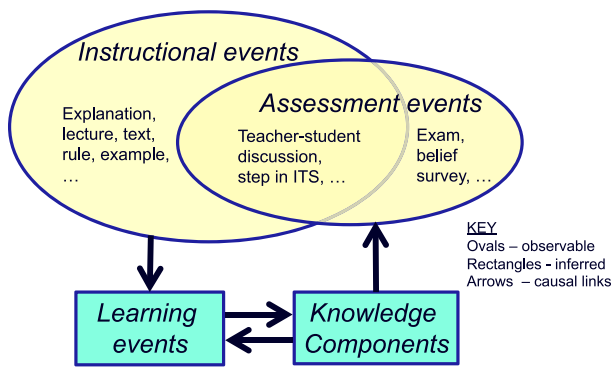

Fig. 1. *Instructional events* are (usually planned) variations in the learning environment that cause *learning events (LEs)*. LEs are unobservable processes that cause unobservable changes in *knowledge components (KCs)* and can be effected by existing knowledge. KCs cause student performances, which can be observed in *assessment events (AEs)*. Examples of instructional and AEs are shown, with some being only instructional (e.g., an explanation), some only assessment (e.g., an exam), and some both (e.g., a step taken by a student in an intelligent tutoring system).

data but cannot be directly observed or directly controlled. Learning processes and knowledge changes are inferred from assessments at both immediate and remote time points (long-term retention, future learning) and in tasks (transfer, future learning) that may differ from those during instruction. Each of the three central events is decomposable, both temporally and structurally. Temporally, instructional sequences contain event segments that vary from less than 1 s (e.g., exposure to a printed word) and a little more than a second (e.g., a fact retrieval task with instructional feedback) to a few minutes (e.g., turns in a classroom dialog around the concept of an integer). Structurally, we decompose the cognitive changes that arise from these *IEs* into *KCs*. LEs produce KCs and the acquisition and modification of these components is the KLI framework's explanation for consistency in student performance and transfer across related AEs.

The arrows in Fig. 1 represent inferences about causal links to the unobservables (knowledge changes and the LEs that produce them) from the observables (Instructional Events and AEs). Unobservable KCs can be inferred by comparing performances exhibited by different kinds of AEs (''tasks'' or ''items''). For example, by contrasting student performance on assessment tasks that presented algebra problems in words versus matched equations, Koedinger and Nathan (2004) identified components of knowledge that students particularly struggle to acquire. In contrast to teachers' beliefs, beginning algebra students performed worse on equations than word problems and this contrast led the researchers to infer that students have greater knowledge deficits for comprehending algebraic expressions than corresponding English sentences. In cases like these, instructors may incorrectly treat knowledge and its acquisition as directly observable, not requiring empirically based inference, and conclude equations are easier than word problems because it looks like equations are simpler.

The typical instructional experiment, in KLI terms, explores how variations in IEs affect performance on subsequent AEs. The interpretation of such experiments may involve inferences about mediating LEs and KC changes. As an example, we can explain results of Aleven and Koedinger (2002) in these terms. They found that adding prompts for self-explanation to tutored problem solving practice (without adding extra instructional time) produced greater explanation ability *and* conceptual transfer to novel problems while maintaining performance on isomorphic problems. In KLI terms, changing instruction from pure practice to practice with self-explanation (kinds of IEs) engaged more verbally mediated explanation-based learning in addition to non-verbal induction (kinds of LEs) and thus produced more verbal declarative knowledge in addition to non-verbal procedures (kinds of KCs). KC differences were inferred from the observed contrast in student performance on isomorphic problem solving test items compared with conceptual transfer test items (kinds of AEs). These results require an explanation that appeals to KCs: The groups perform the same when assessment tasks allow either kind of knowledge (verbal declarative or non-verbal procedural), but only the self-explainers do well on the transfer tasks that require just verbal declarative knowledge.

We generalize such examples for linking knowledge to learning and instruction by specifying taxonomies of knowledge (section 3), learning (section 4), and instruction (section 5). Taxonomies have value when classification supports scientific analysis. In the case of the KLI framework, taxonomies are used to make more visible dependencies between the kinds

of knowledge to be learned, the learning processes that produce knowledge changes, and the instructional options that affect robust learning.

Examples of possible dependencies have been suggested in the literature. Rohrer and Taylor (2006) summarize a meta-analysis of spacing effects (Donovan & Radosevich, 1999) with ''the size of the spacing effect declined sharply as conceptual difficulty of the task increased from low (e.g., rotary pursuit) to average (e.g., word list recall) to high (e.g., puzzle).'' Wulf and Shea (2002) suggest that ''situations with low processing demands [simple skills] benefit from practice conditions that increase the load and challenge the performer, whereas practice conditions that result in extremely high load [complex skills] should benefit from conditions that reduce the load to more manageable levels.'' These indications of dependencies between knowledge acquisition goals and effective learning and instruction illustrate the gain in conceptual clarity that may come from a systematic articulation of kinds of knowledge, learning, and instruction that provides the basis for exploring how they interrelate.

It is useful to consider whether such taxonomies can help explain why opposing recommendations for optimal instruction remain in the literature. One such opposition contrasts recommendations for instruction that increases demands on students so as to produce ''desirable difficulties'' (e.g., Cepeda, Pashler, Vul, Wixted, & Rohrer, 2006; Roediger & Karpicke, 2006; Schmidt & Bjork, 1992) with recommendations for instruction that decreases demands on students so as to reduce extraneous ''cognitive load'' (e.g., van Merriënboer & Sweller, 2005). KLI can be used to help generate hypotheses for how this opposition may be resolved based on differences in the kinds of KCs, kinds of LEs, or kinds of AEs on which these research paradigms focus. Research supporting desirable difficulties has tended to focus on fact knowledge (a kind of KC), memory processes (a kind of LE), and long-term retention tests (a kind of AE) to argue in favor of instructional approaches, such as increased test-like practice (Roediger & Karpicke, 2006) and mixed or spaced practice (Schmidt & Bjork, 1992). Cognitive load research has tended to focus on rule or schema knowledge (kinds of KCs), induction or compilation processes (kinds of LEs), and transfer tests (a kind of AE) to argue in favor of instructional approaches, such as increased study of examples (Renkl, 1997). KLI provides a frame for (a) noticing contradictions in instructional recommendations, and (b) pursuing whether the resolution lies within differences in the kinds of knowledge being addressed, in the kinds of learning processes being evoked (LEs), or in the nature of the assessment strategies (AEs) being employed.

Articulating potential dependencies among kinds of knowledge, learning, and instruction requires a stipulation of those kinds. Such articulation suggests a variety of open research questions that go beyond the open issues that motivated our creation of the KLI framework.

3. Knowledge: Decomposing task complexity and transfer

Others have argued for the importance for educational design of analyzing learning goals into components of knowledge (e.g., Minstrell, 2001; diSessa, 1993). Anderson and Schunn (2000) suggested, ''there is a real value for an effort that takes a target domain, analyzes it

into its underlying KCs, … communicates these components, and monitors their learning.'' Although cognitive task analysis to design instruction has been demonstrated to be effective in a number of training domains (e.g., Clark, Feldon, van Merriënboer, Yates, & Early, 2007; Lee, 2003), it has not been a common approach for designing academic instruction. Cognitive task analysis remains as much an art as a science, in part because of the unobservable nature of knowledge and limited scientific tools for characterizing it at a useful level of analysis. Thus, we think an effort toward defining a taxonomy of kinds of KCs is worthwhile and can be useful even without costly implementation of computational models that has been the traditional approach of cognitive science.

We define a *knowledge component (KC)* as an *acquired* unit of cognitive function or structure that can be inferred from performance on a set of related tasks. These tasks are the AEs of the KLI framework (see Fig. 1). As a practical matter, we use ''knowledge component'' broadly to generalize across terms for describing pieces of cognition or knowledge, including production rule (e.g., Anderson & Lebiere, 1998; Newell, 1990), schema (e.g., Gick & Holyoak, 1983; van Merriënboer & Sweller, 2005), misconception (e.g., Clement, 1987), or facet (e.g., Minstrell, 2001), as well as everyday terms, such as concept, principle, fact, or skill (cf. Bloom, 1956). Many KCs describe mental processes at *about* the *unit task* level within Newell's (1990) time scales of human action (see Table 1). Unit tasks last about 10 s and are essentially the leaf nodes or smallest steps in the decomposition of a reasoning task—that is, the application of a single operator in a problem solving space (e.g., applying a theorem in a geometry proof). Unit tasks are at the interface between what Newell called the ''cognitive band'' and ''rational band.'' Scientific investigation at these time scales is critical to make productive bridges between neuroscience research within the biological band, where attention is on millisecond changes, and educational research within the social band, where attention is on changes occurring over months.

A range of unit task times are illustrated in Fig. 2, which shows a sample of three learning curves of student performance in language and math domains. Each curve displays the

Table 1
Newell's time scales of human action

Scale (s)	Time Units	System	World (Theory)
10^7	Months		
10^6	Weeks		Social band
10^5	Days		
10^4	Hours	Task	
10^3	10 min	Task	Rational band
10^2	Minutes	Task	
10^1	10 s	Unit task	
10^0	1 s	Operations	Cognitive band
10^{-1}	100 ms	Deliberate act	
10^{-2}	10 ms	Neural circuit	
10^{-3}	1 ms	Neuron	Biological band
10^{-4}	100 μs	Organelle	

**Correct Time Learning Curves
for Different Kinds of Knowledge Components**

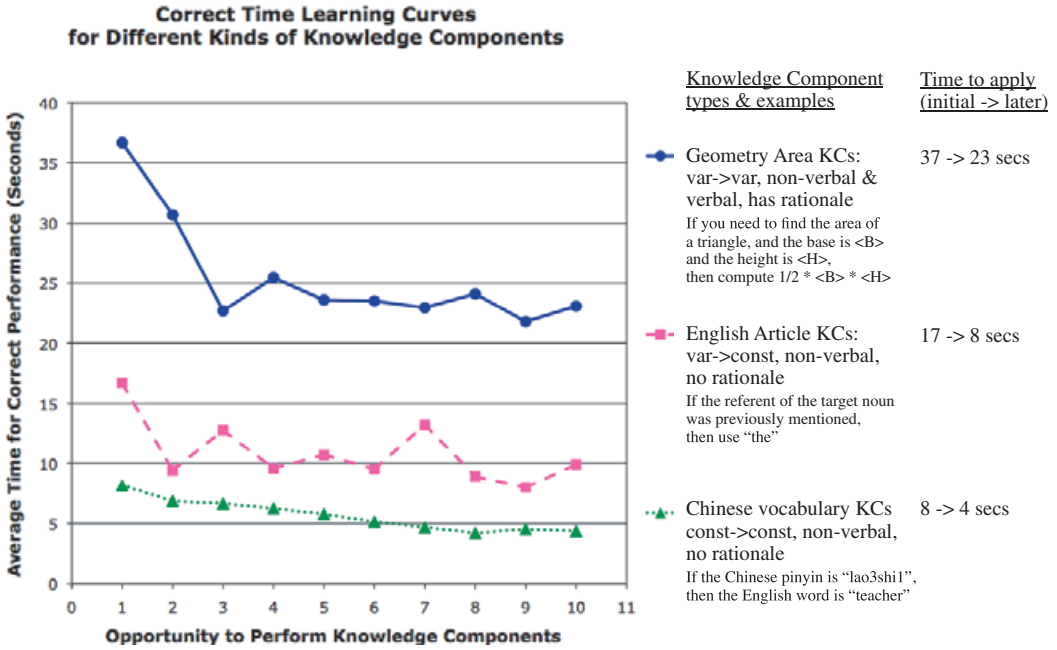|	Knowledge Component types & examples	Time to apply (initial -> later)
Geometry Area KCs:	var->var, non-verbal & verbal, has rationale	37 -> 23 secs
	If you need to find the area of a triangle, and the base is <B> and the height is <H>, then compute 1/2 * <B> * <H>	
English Article KCs:	var->const, non-verbal, no rationale	17 -> 8 secs
	If the referent of the target noun was previously mentioned, then use "the"	
Chinese vocabulary KCs	const->const, non-verbal, no rationale	8 -> 4 secs
	If the Chinese pinyin is "lao3shi1", then the English word is "teacher"	

Fig. 2. Performance time learning curves for three different kinds of knowledge components (KCs) show how more complex KCs take longer to perform than less complex KCs. The data come from student use of computer tutors for High School Geometry, College English as a Second Language, and College Chinese and are available from DataShop (http://pslcdatashop.org). The curves show the time for a correct performance of a KC (averaged across students and KCs) on successive opportunities to perform and learn that KC (the *x*-axis). Each opportunity is an assessment event because the tutor observes and record students' success, and an instructional event because students get correctness feedback and as-needed instruction.

average time it takes students to correctly apply a single KC across successive opportunities during tutored problem solving. Application time for each KC systematically decreases as the KC is further refined and strengthened. Comparing the three curves illustrates the variability in average time to correctly apply KCs across the examples from three different domains—Chinese vocabulary, English articles, and geometric areas. However, rather than assuming such variability is intrinsic to a whole domain, the working assumption of our framework is that such variability is associated with specific KCs. The prototypical KCs in Fig. 2 are drawn from different types, as elaborated in Tables 2 and 3 and discussed below. Distinguishing properties of KCs are part of a deep analysis of domain knowledge, and thus a powerful source for innovative instructional design.

KCs that function in a given task are hierarchical, thus posing a problem of how to target the right components for a given learning situation. For example, sentence comprehension relies on word identification, which relies on letter recognition and trigonometry relies on geometry, which relies on multiplication. As the KLI framework targets the analysis of academic learning, our strategy is to focus on knowledge that is to be acquired by students in a given academic course. This general strategy leads to a more specific one, to focus on

Table 2
Examples of different kinds of knowledge components in second language learning, mathematics, and science

Domain/Knowledge Category	Application Conditions	Response	Example
Second Language			
Chinese vocabulary	Constant	Constant	Chinese radical ''日'' => ''sun'' in English
English determiners	Variable	Constant	Choose article for \<Noun\> & it is definite => use ''the''
Plural English spellings	Variable	Variable	\<Noun\> ending \<consonant\> ''y'' => remove ''y'' from \<Noun\> & add ''ies''
Mathematics			
Rule from name	Constant	Constant	''circle area formula'' => ''$A = \pi*r^2$''
Rule from context	Variable	Constant	Find area of \<Fig\> & it is a circle => use ''$A = \pi*r^2$''
Solution planning	Variable	Variable	Find area of irregular shape made up of regular shapes \<S1\> and \<S2\> => find area \<S1\> and \<S2\> and combine
Science			
Content fact	Constant	Constant	''gravitational acceleration'' => ''$g = 9.8/s^2$''
Categorize situation	Variable	Constant	Find unit for \<quantity\>, a small volume => use ''milliliters''
Experimental design	Variable	Variable	Test hypothesis \<Var1\> causes \<Var2\> => run an experiment that varies only \<Var1\> and nothing else and measure differences in \<Var2\>

the component level at which the novice student makes errors. Thus, within a hierarchy of components, a knowledge analysis for a particular course may focus only on a single level that lies just above the level at which novices have achieved success and fluency. KC descriptions at the target level (e.g., word identification or geometry) can treat lower levels (e.g., letter identification or multiplication) as atomic under the empirical constraint that the target population has mastered the lower level.[1]

3.1. Kinds of knowledge components

KCs can be characterized in terms of various properties that cut across domains to capture their functioning in LEs. Time scale is one property we have already noted. Conditions of applications and the student response to an AE are others. Table 2 illustrates the properties of application conditions and responses for examples across three domains.[2] Some KCs are applied under unvarying, constant conditions, while others are applied under variable conditions. A KC has a constant condition when there is a single unique pattern to which the KC applies, and a variable condition when there are multiple patterns to which it applies (e.g., a feature or features that can take on many different values). Paired associates (e.g., Roediger & Karpicke, 2006) and ''examples'' (e.g., Anderson, Fincham, & Douglass, 1997) are constant condition KCs, whereas categories, rules, schemas, or principles are variable condition KCs. Similarly, the response of a KC can be a single value or constant

Table 3
Some basic knowledge component categories

Application Conditions	Response	Relationship	Rationale	Labels
Constant	Constant	Non-verbal	No	Association
Constant	Constant	Verbal	No	Fact
Variable	Constant	Non-verbal	No	Category
Variable	Constant	Verbal	No	Concept
Variable	Variable	Non-verbal	No	Production, Schema, Skill
Variable	Variable	Verbal	No	Rule, Plan
Variable	Variable	Verbal	Yes	Principle, Rule, Model

(e.g., a category label) or it can vary as a function of the variable information extracted in the condition.[3]

The condition-response format emphasizes that knowing is not just about knowing what or knowing how, but knowing when. Indeed, learning the conditions of application of knowledge, the ''when,'' may be more difficult than learning possible responses, the ''what'' (cf., Chi, Feltovich, & Glaser, 1981; Zhu, Lee, Simon, & Zhu, 1996).

For purposes of first approximation, these properties define a useful (if perhaps incomplete) notion of ''complexity.'' KCs that have variable conditions of application, variable conditions of response, and operate on a longer time scale can be said to be more complex than KCs that are constant in application or response, and at a shorter time scale. To illustrate a relatively simple KC, consider learning the meaning of the Chinese radical 日. An English language learner acquires the connection between this form and the English word ''sun.'' The implementation of this KC then is that given 日 and the goal of generating an English translation (the conditions), produce ''sun.'' While knowledge may well be non-symbolic in its brain-based implementation, for purposes of scientific analysis, we describe KCs in a symbolic format (e.g., English or computer code).

As illustrated in Table 3, our knowledge taxonomy makes four key distinctions (the first two have been introduced above): the generality of the conditions of application, the generality of the response whether the KC is verbal or not, and the extent to which the KC can be rationalized. In the Chinese KC example above, the application condition (the Chinese radical 日) is a constant and the response (sun) is a constant. The relationship is verbal (i.e., students express it in words), and while a few radicals may have iconic value, the relationship generally has no rationale, instead being a writing convention.[4] Non-language examples of ''constant-constant'' KCs are shown in Table 2 as are examples of KCs with variable conditions and/or variable responses. As indicated in the last column of Table 3 (second row), constant-constant KCs are commonly called ''facts.'' The labels in Table 3 (facts, categories, rules, etc.) are rough mappings to commonly used terms (cf. Gagne, 1985) and not necessarily one-to-one with the cells in this taxonomy.

3.1.1. Constant-constant KCs

Whether the conditions and response of KCs are constants or variables is associated with the broad kinds of learning processes (e.g., memory, categorization, induction) that

have been studied. Learning research that emphasizes memory (e.g., Cepeda et al., 2006; Roediger & Karpicke, 2006; Schmidt & Bjork, 1992) has primarily focused on constant-constant KCs and these are often assessed with ''paired associate'' tasks. For example, instructional recommendations derived from research on spacing and testing effects (Pashler et al., 2007) have tended to emphasize constant-constant KCs with tasks, such as math fact recall (e.g., $8 \times 5 = 40$ in a study of spaced practice; Rea & Modigliani, 1985).

These relatively simple ''associations'' or ''facts'' (see Table 3) are pervasive in academic learning, such as vocabulary facts in second language learning.[5] Similar KCs are essential in middle school, high school, and post-secondary mathematics and science, particularly, definitions of terms. For instance, ''Pi (Π) is the ratio of a circle's circumference to its diameter'' represents a constant-constant KC in math that relates a term to its definition. These tend to be relatively minor components in post-elementary mathematics, but they can be significant barriers to learning. For example, student performance in the ASSISTment computer tutoring system (Feng et al., 2009) reveals that errors in simple fraction multiplication word problems (e.g., ''What is 3/4 of 1/2?'') are sometimes not about the math, but about vocabulary—follow-up questions indicate that students make more errors on ''What does 'of' indicate?'' (where multiply is one of four operator choices) than on ''what is ½ times ½?''

We emphasize that all KCs, even these simple facts, are not directly observable and multiple AEs are needed to infer whether or not a student has robustly acquired a KC. Variability in the timing of AEs and the kinds of performance metrics used are important. Success on an AE immediately after instruction provides evidence of some initial acquisition of a KC, but is weaker evidence of robust learning than success on a delayed AE (cf. Cepeda et al., 2006). Accuracy in performance does not guarantee acquisition of fluency, but AEs that provide timing data or assess accuracy in time-limited contexts would (cf., De Jong & Perfetti, 2011).

3.1.2. Variable-constant KCs

While memory research has emphasized constant-constant KCs, research on conceptual and perceptual category learning (Medin & Schaffer, 1978) and artificial grammar rule learning (Frank & Gibson, 2011) has primarily focused on variable-constant KCs, which are essentially category-recognition rules with many-to-one mappings. English article selection is an example of this type: ''to construct a noun phrase with a unique referent, use the article 'the' '' (e.g., ''The moon …''). Such categorical or classification knowledge exists in mathematics as well, for instance, any expression that indicates the quotient of two quantities is a fraction.[6] Other examples of variable-constant KCs are given in Table 2. That some knowledge can be used in (or generalizes across) a variety of different situations is reflected in KLI by KCs that have variable conditions.

3.1.3. Variable-variable KCs

Research that emphasizes more complex rule or schema structure learning and transfer (e.g., Gick & Holyoak, 1983; Sweller & Chandler, 1994) has primarily focused on variable-variable KCs. These KCs map one relational structure (Gentner et al., 2009) to another and

variables are used to express the many possible arguments to those relations. An example in second language learning is a rule for generating an English plural: To form the plural of a singular noun <N> ending in an ''s'' sound or a ''z'' sound, form the word <N> ''es''. In mathematics and science, KCs that apply formulas to solve problems have a variable-variable structure: To find the area of a triangle with height <H> and base <B>, multiply <H> * <B> * 1/2. Other math and science examples of variable-variable KCs are shown in Table 2.

Variety in task contexts is needed to infer acquisition of variable condition KCs from AEs. Just because a second language English student correctly selects ''an'' in ''[a∕an] orange'' does not ensure the student has learned a (variable-constant) KC with the right generality (Wylie, Koedinger, & Mitamura, 2009). They may or may not have (implicitly) induced a KC for which ''the noun begins with a vowel'' is a condition. Variety in assessment tasks, such as ''[a∕an] honor,'' is needed to infer that the student has learned the correct condition, ''the noun begins with a vowel *sound.*'' Aleven and Koedinger (2002) illustrate the use of a variety of assessments to disambiguate KCs with incorrectly generalized conditions (''angles that look equal => are equal'') from KCs with correctly generalized conditions (e.g., ''base angles of an isosceles triangle => are equal'').

3.1.4. Non-verbal versus verbal KCs

Some KCs represent associations, perceptual categories, skills, or procedures that cannot be readily verbalized (cf., Alibali & Koedinger, 1999; Dienes & Perner, 1999). Other KCs represent concepts, procedures, principles, or theories that learners can readily verbalize. The KLI distinction between verbal and non-verbal is similar (but not identical[7]) to the ACT-R distinction between declarative and procedural knowledge (Anderson & Lebiere, 1998), which has been influential in instructional design (e.g., Koedinger & Aleven, 2007). The ACT-R distinction is about unobservable cognitive mechanisms, emphasizing whether or not knowledge can be accessed by other knowledge (declarative can, procedural cannot). The KLI distinction is about observable behavior, emphasizing whether students can ''do'' but not explain (indicating non-verbal knowledge), explain but not do (indicating ''inert'' verbal knowledge), or do *and* explain (indicating both non-verbal and verbal knowledge).

One reason to emphasize this distinction derives from the observation that much of what experts know is only in non-verbal or ''implicit'' form (Posner, in press). Much language knowledge (e.g., English article selection) is non-verbal—English speakers can converse effectively and fluently, but most either cannot explain their choices at all or cannot do so to a reasonable standard of coherence. Non-verbal knowledge is common in math too; for example, students may be able to accurately recognize which expressions are fractions but not articulate how they do so. Aleven and Koedinger (2002) provide evidence of non-verbal procedural knowledge in the performance of geometry students who were more correct on making geometric inferences than they were on matched explanation items. A mathematical model of their data provides evidence that non-verbal procedural and verbal declarative knowledge of the same content can co-exist and that different mixtures of such knowledge can yield identifiable patterns of performance across multiple assessments.

As experts cannot directly access or articulate many aspects of what they know, they are susceptible to ''expert blind spot'' (Koedinger & Nathan, 2004). Instructors and designers may thus underemphasize or even completely overlook non-verbal knowledge in their instructional delivery or design. The pervasiveness of non-verbal visual or procedural knowledge is one reason why Cognitive Task Analysis is so powerful in improving instruction (Clark et al., 2007).

Verbal instruction is intrinsically bound to learning verbal KCs. Moreover, even in tasks involving substantial non-verbal knowledge, the verbal articulation of knowledge (by teacher, textbook, tutor, or student) can enhance learning. It is not well enough understood when explanations (intended to produce verbal KCs) should be given to or elicited from students to best enhance their learning (cf. Renkl, Stark, Gruber, & Mandl, 1998). Much research on ''self-explanation'' indicates that prompting students to provide verbal explanations of actions (whether their own or ones given in a worked example) can often aid robust learning (Aleven & Koedinger, 2002; Chi, de Leeuw, Chiu, & LaVancher, 1994; Graesser et al., 2005; Hausmann & VanLehn, 2007). Although verbalization during mainly perceptual tasks can sometimes interfere with learning (Schooler, Fiore, & Brandimonte, 1997), verbalization in such tasks can also be supportive (Fox & Charness, 2009). In some situations, the support provided for verbalization is offset by the additional instructional time it requires (Wylie et al., 2009). These somewhat ambiguous conclusions may reflect variations in the specific IEs and the target learning. The KC taxonomy may help differentiate positive and negative results of the effects of verbalization, including prompting for self-explanations. The negative results tend to involve variable-constant KCs in domains where non-verbal acquisition is sufficient, whereas the positive results involve variable-variable KCs in domains where verbal KCs (as well as non-verbal KCs) are learning objectives (e.g., students are expected to be able to express mathematical theorems and scientific principles).

One reason why educators in some domains value verbal articulation of rules and principles is that, once expressed, they are available for rational argument. How much rationalization can be done around a KC is a topic we discuss next.

3.1.5. Connective tissue—KCs with and without rationales

While some ''rules'' or ''principles'' clearly have a rationale, such as theorems in mathematics, other ''rules'' may reflect regularities of seemingly arbitrary conventions, as in some English spelling rules. The availability of rationales should be considered as graded rather than all or none and dependent on the depth of theory development in a domain.

The rules for creating plurals are relatively difficult to rationalize insofar as different languages have different ways to create plurals. These rules gain some rationale within the constraints of a given language, but they remain largely opaque to everyday usage.[8] Other KCs have rationales grounded in nature. The formula for the area of a triangle is a provable regularity of Euclidean spaces that are approximated in the real world. However, just as a skilled language user will not know the rationales for plurals, the successful geometry student may not understand the rationale for the area of a triangle.

The rationale feature of KCs is relevant in considering whether certain forms of instruction, such as collaborative argumentation or discovery learning, will be effective for a particular kind of KC. The rationale of a KC can support sense-making strategies and be used to reconstruct a partially forgotten KC, adapt it to unfamiliar situations, or even construct a KC from scratch. Instruction that involves students in explicitly discovering KCs from data or deriving KCs through argumentation may be productive for KCs with a rationale, but not for ones without.

3.2. Integrative knowledge components and other complexity factors

Beyond the basic KC taxonomy illustrated in Tables 2 and 3 are some additional features of KCs that have significance for learning and instruction. Integrative knowledge (e.g., Case & Okamoto, 1996; Slotta & Chi, 2006), prerequisite conceptual and perceptual knowledge (e.g., Booth & Koedinger, 2008), probabilistic knowledge (e.g., Frishkoff et al., 2008), and shallow or incorrect knowledge including misconceptions (e.g., Aleven & Koedinger, 2002; Chi et al., 1981; Clement, 1987; Minstrell, 2001) are among the important knowledge complexities that have been made visible in learning research.

KCs sometimes are not inferable from a single behavioral pattern, but only from behavioral patterns across task situations varying in complexity. We call such a component an ''*integrative knowledge component*'' because it integrates or must be integrated (or connected) with other KCs to produce behavior. Descriptions of integrative KCs make reference to internal mental states either in their condition (e.g., a deep feature produced by another KC) or in their response (e.g., a subgoal for another KC to achieve). A typical strategy for inferring an integrative component uses a subtraction logic that takes the differences between two tasks of overlapping complexity as implicating an integrative KC. For example, Heffernan and Koedinger (1997) found students were significantly worse at translating two-step algebra story problems into expressions (e.g., 800–40 ×) than they were at translating two closely matched one-step problems (with answers 800-y and 40 ×). They hypothesized that many students are missing an integrative KC that is necessary to solve the two-step problems, namely a recursive grammar rule indicating expressions (e.g., 40 ×) can be embedded within other expressions (e.g., 800–40 ×). Instruction specifically designed to address learning of this KC significantly improved student performance (Koedinger & McLaughlin, 2010).

For integrative knowledge, non-verbal forms of instruction (e.g., example study, repeated practice) may not be optimal for effective induction and refinement. Providing and eliciting explanations may be critical to help learners break down or externalize the complex inference needed when processing negative feedback to fully identify all the KCs and integrative KCs in a reasoning chain and to revise any KCs that are not correct as part of that chain (cf. MacLaren & Koedinger, 2002).

3.2.1. Estimating and measuring KC complexity

Informally, complexity reflects the condition encoding requirements and the response requirements of the task (e.g., how many coding operations on perceptual stimulus are

needed; how many response operations). However, the knowledge taxonomy captures additional possibilities: KCs with variable conditions or responses tend to be more complex than those with constant conditions or responses. Knowing a KC both verbally and non-verbally is more complex than knowing it just non-verbally. And knowing the rationale of a KC as well as the KC itself is more complex than knowing a KC without a rationale.

Beyond these informal approximations, we have two ways to ground and estimate KC complexity. The first is a simple heuristic: Simply put, the more complex the description of the KC the more complex is the KC, following a general definition of complexity as description length (Rissanen, 1978).[9] The ideal description language for KCs is a formal cognitive modeling language, such as an ACT-R production system (Anderson & Lebiere, 1998) or structure mapping theory (Gentner et al., 2009). But KC descriptions in English may also serve this purpose particularly when closely guided by empirical cognitive task analysis methods, such as think alouds or difficulty factors assessments (Koedinger & Terao, 2002; Rittle-Johnson & Koedinger, 2001). Neither approach is guaranteed, but reasonable predictions are possible. Employing description length with the KC descriptions in Table 2 suggests that constant-constant KCs in the three domains tend to be simpler (9, 6, and 6 words) than variable-constant KCs (11, 12, and 10 words), which are simpler than variable-variable KCs (12, 21, and 21 words).

A second grounding of KC complexity is empirical: the difficulty students have in applying the KC, one measure of which is the time it takes for students to correctly execute a KC. In general, simpler KCs can be executed more quickly as suggested by the learning curves displayed in Fig. 2. In these curves the *y*-axis is the average time for correct entries in computer-based tutoring systems for second language and math learning. The average is for each student on each KC in a particular unit of instruction. The *x*-axis shows the number of opportunities that students have had to practice and learn these KCs. Each opportunity is both an AE, whereby the system measures accuracy (whether or not the student is correct on first attempt without a hint) and latency (how long to perform the action), and an IE, whereby a correct action gets positive feedback, incorrect actions get substantive negative feedback, and successively more detailed hints are provided at request. In general, students learn with more opportunities and this relation is reflected in Fig. 2 where the curves go down, indicating faster correct performance as opportunities increase.

Fig. 2 shows learning curves from a Chinese vocabulary unit, an English grammar unit on use of articles, and a Geometry unit on area of figures. The Chinese Vocabulary unit involves constant-constant KCs, for example, the Chinese character for ''lao3shi1'' translates as ''teacher.'' Students' correct performance of these KCs, which involves retrieving the correct response and typing it in, takes about 3–6 s on average (see Fig. 2a). The English Article unit involves variable-constant KCs, such as ''if the referent of the target noun was previously mentioned, then use 'the.' '' Students' correct performance of these KCs, which involves retrieving or reasoning to a correct response and selecting it from a menu, takes about 6–10 s (see Fig. 2b). The Geometry Area unit involves variable-variable KCs, such as ''if you need to find the area of a circle with radius <R>, then compute 3.14 * <R> ^ 2.'' These typically involve retrieving or constructing the needed mathematical operations or formula and typing

in the arithmetic steps to be taken (e.g., 3.14*3^2) or the final result (e.g., 28.26). Correct reasoning and entry takes about 10–14 s (see Fig. 2c). These examples illustrate the possibility of supporting theoretically derived hypotheses about KC complexity (e.g., from KC category or description length) with empirical results (e.g., time to execute the KC)—the variable condition or action KCs with longer descriptions correspond with slower execution times.

3.3. Kinds of knowledge components guide assessment and instructional decisions

3.3.1. Kinds of KCs drive assessment event choices

The AEs, as we have mentioned, need variability to support reliable inferences of student learning of KCs. First, AEs need to assess long-term as well as short-term retention, and fluency as well as accuracy. Second, variety in task contexts is needed to assess whether generality has been achieved in variable condition KCs. Third, factorial variations in AEs are needed to infer the presence and level of acquisition of *integrative* KCs (which are usually variable-variable KCs).

Assessment of variable condition KCs requires multiple AEs to determine whether the condition acquired by the student is accurate in level of generality. Correct learning depends not only on the nature of the correct target KC but also on the context of alternative interpretations that students acquire from instruction (e.g., worked examples or text descriptions of rules). Students may acquire overspecialized conditions, like the statistics students in Chang (2006) who (implicitly) induce rules, such as ''if you want to display demographic data, use a pie chart'' rather than ''if you want do display data <V> and <V> is categorical, then use a pie chart.'' Multiple assessment tasks that vary surface details can detect KCs with overspecialized conditions.

Alternatively, students may acquire incorrectly generalized conditions, as in our previous example of a student who learns that ''an'' rather than ''a'' is used when the first *letter* of the following noun is a vowel. This rule is consistent with most cases, but fails when applied to ''honor.'' In geometry, many students (implicitly) induce a rule that when angles look equal in a diagram, they are equal (Aleven & Koedinger, 2002). There are no sure fire methods to anticipate likely student misconceptions like these (simulated students notwithstanding, Matsuda, Cohen, Sewall, Lacerda, & Koedinger, 2008) and hence a sampling of a wide variety of assessment tasks are needed along with detailed error analysis.

Assessment of integrative KCs requires a different form of variety. As integrative KCs work with other supporting KCs, they need to be assessed a combination of harder tasks that require the hypothesized integrative KC and easier tasks that require the supporting KCs but *not* the integrative KC. The logic of inference in its simplest form assumes that the probability of knowing the integrative KC (P_i) is independent of the probability of knowing the supporting KC (P_s). In that case, P_i can be found dividing success rate on the hard task ($P_i * P_s$) by success rate on the easier task (P_s). This essential logic and a logistic regression generalization of it have been successfully employed to identify integrative KCs for symbolic solution composition (Koedinger and McLaughlin, 2010) and for problem decomposition planning (Stamper & Koedinger, 2011).

3.3.2. Kinds of KCs drive instructional event choices

The differentiation of knowledge types has implications for the effectiveness of an instructional principle. A key motivation of the knowledge taxonomy is that its distinctions may help resolve apparent contradictions among instructional recommendations. A related KLI assumption is that many learning processes and instructional design decisions are not restricted to a domain as whole but are determined by the type of KCs being learned (cf. Sweller & Chandler, 1994; Wulf & Shea, 2002). While some KCs (e.g., constant-constant KCs) may be prevalent in certain domains (e.g., second language vocabulary learning), instructional principles should refer to KCs rather than to domains. A hypothetical principle "drill and practice is not effective for mathematics" is at the wrong level of analysis because it does not describe what it is about mathematics that makes drill and practice unsuitable. There are also aspects of math learning (e.g., non-verbal knowledge of the grammar of algebra, Koedinger and McLaughlin, 2010) where pure practice may be an ideal choice. Similarly, while language learning often benefits from repetition (e.g., vocabulary), other aspects of second language learning (e.g., pragmatics or intercultural competence) may benefit from a combination of dialog, explanation, and practice (cf. Ogan, Aleven, & Jones, 2009).

Rather than associate optimal instructional choices with domains (as disparate literatures on math education, physics education, reading, and second language learning are wont to do), the KLI framework suggests that instructional choices depend on the kinds of KCs being targeted. Some of these will be domain specific and some will not. We offer below some tentative hypotheses that link instructional choices to our taxonomy of KCs. Each hypothesis has some support in the research literature and the complete set reflects an overarching hypothesis that the optimal complexity of student behavior in an IE is correlated with the complexity of the underlying learning event (see section 5.4). But all these hypotheses require further testing.

- Simpler constant-constant KCs (e.g., historical facts, vocabulary, see Table 2) and non-verbal, probabilistic variable condition KCs (e.g., perceptual categories, simple concepts and skills) may imply instructional approaches that emphasize recall (Roediger & Karpicke, 2006), spacing of practice (Cepeda et al., 2006), tutored practice (e.g., Corbett & Anderson, 2001), and optimized scheduling of practice (Pavlik, 2007).
- More complex variable-variable KCs (e.g., designing a controlled experiment, see Table 2) imply instruction that includes comparison/blocking (Gentner et al., 2009; Gick & Holyoak, 1983) and more worked example study (Sweller & Cooper, 1985).
- Integrated variable-variable KCs to be learned in both non-verbal procedural and verbal declarative form (e.g., math and science principles) imply instruction that prompts for self-explanations (Aleven & Koedinger, 2002).
- KCs with rationales (i.e., are not conventions, but reflect discoverable principles) imply instruction that includes classroom argumentation and instructional dialogs (Michaels, O'Connor, & Resnick, 2008).

The various instructional recommendations seem mutually incompatible without the taxonomy. They would reflect "education wars": More worked example study is at odds with more testing of recall, blocked comparison of examples is at odds with spacing, and pure

non-verbal practice is at odds with prompts for self-explanation and extended classroom dialog and argumentation. The KC taxonomy does not resolve these debates and apparent contradictions but provides a path toward resolution suggesting possible knowledge-by-instruction interactions that researchers can explore experimentally and theoretically. We turn now to explaining why the taxonomy might provide guidance for instructional choices. The key idea is that KC categories are tuned to LEs that must be supported by instruction aligned with those LEs.

4. Learning: Toward a taxonomy of processes for knowledge acquisition and improvement

For a simple taxonomy of LEs, we propose three very broad types of learning processes as a starting point:

1. *Memory and fluency-building processes*: Non-verbal learning processes involved in strengthening memory and compiling knowledge, producing more automatic and com-posed (''chunked'') knowledge. Fluency building can be conceived as making the link between the condition and response of a KC more direct, more consistent, and more resistant to interference as well as making the response execution faster.

2. *Induction and refinement processes*: Non-verbal learning processes that improve the accuracy of knowledge. They include perception, generalization, discrimination, clas-sification, categorization, schema induction, and causal induction. (We classify these as non-verbal because although they are often supported by verbalizations, they do not require it.) These processes modify the conditions (e.g., which specific conditions satisfy the variable condition in a variable-constant or variable-variable KC) or response (e.g., which responses satisfy a variable response type KC). These processes refine a KC, making it more accurate, appropriately general, and discriminating.

3. *Understanding and sense-making processes*: Explicit, verbally mediated learning in which students attempt to understand or reason. This includes comprehension of verbal descriptions, explanation-based learning, scientific discovery, and verbal rule-medi-ated deduction. Sense making can be conceived as linking non-verbal with verbal forms of knowledge or a KC with its rationale.

Fig. 3 illustrates how these different learning processes can lead to different kinds of knowledge changes and ultimately to measurable robust learning outcomes.

The knowledge changes that result in robust learning are produced through the memory and fluency building, induction and refinement, and understanding and sense-making pro-cesses we have described above. There is no one-to-one mapping between these learning processes and robust learning outcomes. For example, it is not the case that only sense mak-ing leads to transfer. Accelerated future learning can include three very different learning processes: (a) learning how to learn, that is, acquiring general learning strategies that can be used for more effective learning in a new domain, (b) acquiring deep concepts or founda-tional skills that facilitate learning in a new domain, and (c) increasing ''cognitive head

room'' through fluency with core knowledge that leaves more capacity for using that knowledge in new learning. An example of this cross-domain accelerated learning comes from Hausmann and VanLehn (2007), who found that prompting students to self-explain in a Physics unit on electricity led to accelerated learning in a later magnetism unit. This effect may reflect (a) acquisition of a learning strategy (self-explanation), (b) deeper learning of electricity concepts (e.g., electrical field principles) allowing better learning of similar magnetism concepts (e.g., magnetic field principles), or (c) greater fluency with core concepts and skills (e.g., elements of field equations) allowing more head room for learning magnetism. In a follow-up study, Hausmann et al. (2009) contrasted two types of explanation prompts in electrodynamics problems, justification-based prompts that focused on the physics principle that justifies a step and meta-cognitive prompts that focused on how each step relates to the student's existing knowledge, and found that justification-based prompts supported greater learning. This study did not examine transfer across topics, but it suggests that the transfer effects in the prior study reflected the acquisition of deep concepts more than meta-cognitive learning or increased fluency.

4.1. Memory and fluency-building processes

The brain is continuously engaged in creating and strengthening connections between the conditions and responses of KCs in use. Thus, memory strengthening operates throughout learning from the initial formation of a KC to each time it is subsequently used. It operates on all kinds of KCs from the simplest non-verbal constant-constant associations to the most complex variable-variable schemas with verbal descriptions and rationales. Especially relevant for instruction is that memory improves with increased frequency of exposure (implying practice—Anderson & Lebiere, 1998), is more reflective of robust learning when recall rather than recognition is required (implying a ''testing effect''—Roediger & Karpicke,

Learning processes	*Knowledge changes*		
Understanding & Sense-making	Verbal KCs with rationales facilitate rederivation	Verbal KCs with rationales can be adapted to new needs	Verbal KCs with rationales can derive new knowledge
Induction & Refinement	Dual-coded KCs support retrieval	More general KCs apply more broadly	More general KCs are better building blocks
Memory & Fluency building	Stronger KCs are better retained	Stronger KCs leave cognitive headroom for adaptation	Stronger KCs leave cognitive headroom for learning
	⬇ Longer retention	⬇ Further transfer	⬇ Better future learning

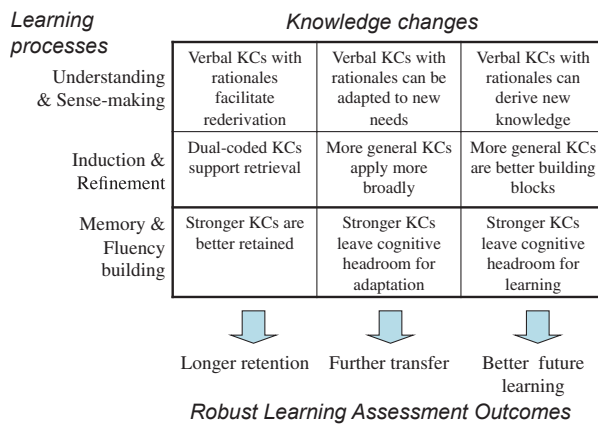*Robust Learning Assessment Outcomes*

Fig. 3. How different classes of learning processes change knowledge to yield different robust learning outcomes.

2006), and when practice is more widely distributed in time (implying a ''spacing effect''—Cepeda et al., 2006).

Even after a KC is learned enough to produce accurate responses, memory, and fluency-building processes continue, leading to fast and effortless performance. Two principal processes underlying fluency gains are knowledge compilation (e.g., Anderson & Lebiere, 1998) and memory strengthening (e.g., Logan, 1988). In compilation, an initial declarative encoding of KCs is proceduralized into a directly executable form, and chains of small KCs may be composed into a single larger KC and producing more automatic processing (Schneider & Shiffrin, 1977). In strengthening, KCs become more accessible with repeated use, resulting in faster and more reliable *retrieval* of KCs. Fluency is often considered only as a matter of specific skill acquisition. However, it is important to test its possible role in accelerating future learning (see the lower right cell in Fig. 3).

4.2. Induction and refinement processes

Induction and refinement processes modify KCs, especially the conditions that control the retrieval and application of knowledge. Like memory and fluency-building processes, induction and refinement processes function in both the initial construction of a KC and its subsequent revision. These processes are relevant to variable-constant and variable-variable KCs, modifying the condition part of KCs by adding missing relevant features (a ''discrimination'') or removing irrelevant features (a ''generalization''). For instance, from examples in which equal angles look equal in a diagram (e.g., an isosceles triangle ABC is shown with base angle B = 70° and it is shown that base angle C = 70°), a geometry novice may *induce* the KC: ''If angles look equal, then they are equal'' (Aleven & Koedinger, 2002). (This induction may be done with little or no deliberate awareness.) Although this KC can yield correct answers, it is incorrect and in need of refinement. Refinement can occur through learning processes that remove the irrelevant feature ''angles look equal'' or add a relevant feature, such as ''angles opposite each other in crossing lines'' or ''angles that are base angles of an isosceles triangle.''

Many specific kinds of induction and refinement learning mechanisms can be found in the cognitive science and machine learning literatures. These include perceptual chunking (e.g., Gobet, 2005; Servan-Schreiber & Anderson, 1990), rule or schema induction and analogy (e.g., Frank & Gibson, 2011; Gentner et al., 2009; Gick & Holyoak, 1983), generalization (Hummel & Holyoak, 2003; Shepard, 1987), discrimination (e.g., Chang, 2006; McClelland, Fiez, & McCandliss, 2002; Richman, Staszewski, & Simon, 1995), error-driven learning (Ohlsson, 1996), classification and categorization (e.g., Blum & Mitchell, 1998; Medin & Schaffer, 1978; Quilici & Mayer, 1996), and non-verbal explanation-based learning (Mitchell, Keller, & Kedar-Cabelli, 1986).

While memory and fluency-building processes involve core mechanisms of the cognitive architecture, induction and refinement processes make use of existing knowledge, as more elemental KCs become part of the condition or response of a larger KC. These processes draw on cognitive resources and take time to execute. Unlike the understanding and sense-making processes, induction and refinement processes are non-verbal (although

verbally mediated sense making may work in service of these processes). Learning from examples or ''by experience'' may result in feature inductions or refinements that students cannot verbalize. For instance, first language learners acquire the features for correct choice of articles, such as ''a'' and ''the,'' without being able to articulate the explicit rules for article choice, although their learning is produced by verbal events. Even *second* language learners, as well as math and science learners, engage in such non-verbal feature induction and refinement (cf. Koedinger & Roll, in press; Posner, in press). Non-verbal induction can lead to a situation in which students can correctly perform mathematics (e.g., find a correct angle in a geometry problem by subtracting from 180°) that they cannot explain (e.g., indicating that the target and given angle form a line; Aleven & Koedinger, 2002).

4.3. Understanding and sense-making processes

Understanding and sense-making processes are robust learning strategies by which students engage in higher level language-mediated thinking to create knowledge. They involve explicit reasoning and include comprehension strategies, self-explanation, and social argumentation. While sense making can focus on application conditions and responses, it can also convert non-verbal relationships into verbal ones, thereby transforming constant-constant associations into facts, variable-constant categories into concepts, and variable-variable productions into rules (see Table 3). Similarly, when sense-making focuses on the rationale for a feature-response relationship, it may transform variable-variable rules into principles.

Understanding and sense-making processes include explicit comprehension strategies (e.g., Graesser et al., 2005; Kintsch, 1998; Palincsar & Brown, 1984), verbally mediated self-explanation (e.g., Ainsworth & Loizou, 2003; Chi, Bassok, Lewis, Reimann, & Glaser, 1989; Lewis, 1988; VanLehn, 1999), explicit hypothesizing and scientific discovery processes (Klahr & Dunbar, 1988; Levine, 1966), deductive proof (e.g., Stylianides & Stylianides, 2009), and explicit argumentation or collaborative discourse (e.g., Asterhan & Schwarz, 2009). Although understanding and sense making may operate in the early formation of knowledge, many competencies (not just first language) emerge from non-verbal induction and only later *may* be ''understood'' and articulated verbally. In contrast with the first two categories of learning processes, understanding and sense-making processes are more deliberate, occurring when a student chooses to engage in them. Unlike induction and refinement learning, for which verbal explanations are largely inaccessible, understanding and sense making are explicitly supported by language or external symbols (whether subvocalized, written, or spoken in social dialog).

4.4. Knowledge and learning process dependencies

A taxonomy of learning processes helps provide causal links between instructional methods and changes in student knowledge. The KLI framework suggests that there are likely to be important dependencies between kinds of knowledge, learning processes, and choices of

most effective instructional methods. For example, fluency-building processes may be most important for learning simple constant-constant components without a rationale, whereas sense-making processes may be most important for learning more complex variable-variable components that have a rationale. Fluency-building processes may also be relevant for more complex variable condition components that may become inaccessible in long-term memory without appropriate repetition. Similarly, refinement processes are also relevant for the kinds of complex integrated and interconnected KCs (with rationales) produced by sense-making processes. These observations suggest a potential asymmetry, whereby simpler learning processes (fluency and refinement) may support complex knowledge, but complex learning processes (e.g., argumentation) may fail to support simple knowledge (e.g., arbitrary constant-constant associations).

Table 4 relates three kinds of KCs to the three broad learning process categories. It suggests an asymmetry in the relevance of learning processes to kinds of KCs. While all three learning processes may be important for robust (and efficient) learning of more complex principles and mental models (integrated verbal and non-verbal KCs with rationales), for the simplest facts (constant-constant paired associates without rationales) memory and fluency processes are more important than the other learning processes. We refer to this idea as the asymmetry hypothesis, which we elaborate in the next section. For such KCs, there are no variable conditions (generalizations) that need to be induced and refined and there are no rationales to use to engage in understanding or sense making. Thus, two minus signs appear in the ''Facts'' column of Table 4. Similarly, for intermediate complexity rules (variable condition KCs which need not be verbalized and do not have rationales), induction and refinement are important, but understanding and sense making may not be (cf. Wylie et al., 2009)—thus the minus sign in the ''Rules'' column of Table 4.

On the other hand, humans sometimes fail to learn rules not because of failures in induction but failures in memory—as was nicely demonstrated recently by Frank and Gibson (2011) in artificial grammar learning tasks. In other words, memory processes remain necessary (and should be supported in optimal instruction) even for rules (see the single plus in Table 4). The other single plus signs (in the last column) indicate hypotheses that memory and fluency processes and induction and refinement processes are also important for learning principles—at least for lasting retention of the verbal form of those principles and for accurate and fluent performance of actions (generated by non-verbal procedures) corresponding with those principles (e.g., quickly find the area of a circle as well as state the formula). These have single rather than double plus signs in Table 4 because research is insufficient and there are opposing views. For example, ''memorization'' (as might be supported by spacing of practice) may not be (as) important for principles because interconnections with other knowledge may sufficiently support retrieval (cf. Rohrer & Taylor, 2006).

In addition to the kind of KC, the stage of learning may also have implications for what learning processes (and therefore what instructional choices) are relevant. For example, in learning of a rule, inductive processes are critical for initial formation of the rule and refinement for improvement, but then memory and fluency process are important to improve retrieval reliability and application speed.

Table 4
Which learning processes are effective for which kinds of knowledge components (KCs)

		Facts (constant–constant KCs)	Rules (variable condition KCs)	Principles (verbal KCs with a rationale)
Learning Processes (simpler on bottom)	Understanding & sense making	– nothing to explain	– implicit rule learning is more efficient	++ principles must be understood
	Induction & refinement	– no generalization needed	++ rules must be induced	+ principles can be inert without associated rules
	Memory & fluency building	++ facts must be memorized	+ rules (& instances) must be remembered (dual paths may help)	+ principles must be remembered, but can be reconstructed

Knowledge Components
(simpler on left)

We turn next to examples of instructional principles based on experimental results that support hypothesis-based instructional interventions in producing robust learning efficiency outcomes. In section 6, we provide a more detailed analysis of one of these principles. Finally, in section 7, we return to the issue of dependency and, in particular, hypothesize a general relationship between KC types and instructional methods.

5. Instructional principles and hypotheses about their effectiveness

The main question for this section is, what kinds of IEs yield robust learning in an efficient way? We are concerned not only about robust learning outcomes but the instructional time required. We seek principles that achieve greater robust learning outcomes without taking more time, that achieve equivalent robust learning in less time, or, most generally, that increase the rate at which robust learning occurs. Table 5 shows a list of example instructional principles from simplest to most complex, where by ''principle'' we mean a kind of IE for which there is substantial evidence that it enhances robust learning efficiency. The simplest, Spacing, Testing, and Optimized Scheduling, have tended to be used with the simplest kinds of KCs, constant-constant facts. The most complex, Accountable Talk, has tended to be used with the most complex kinds of KCs, verbal principles with rationales.

Our taxonomy of instruction follows directly from the taxonomy of learning and includes three major kinds: (a) memory and fluency enhancing instruction, (b) induction and refinement enhancing instruction, and (c) understanding and sense making enhancing instruction. We describe some example principles in these three categories in more detail.

5.1. Memory and fluency enhancing instruction

5.1.1. Spacing and testing effects

The spacing recommendation in the ''Organizing Instruction'' practice guide (Pashler et al., 2007) suggests that spacing practice over longer time intervals leads to better long-term retention than massing practice over shorter intervals (e.g., Cepeda et al., 2006). The testing recommendation (''use quizzing to promote learning'') from that same guide suggests that long-term retention is enhanced by practice in recalling target material more than by repeated studying of the same material (e.g., Roediger & Karpicke, 2006).

Table 5
Some examples of instructional principles in learning process categories roughly ordered from simpler to more complex

Learning Processes	Instructional Principle	Description	Example References
Memory and fluency	Spacing and testing	Long-term retention of knowledge components (KCs) is enhanced with longer intervals between practice and when active recall is required (''tests'' or ''problems'').	Cepeda et al. (2006)
	Optimized scheduling	Selection of practice instances based on prior statistics and on each student's experience with each target KC.	Pavlik (2007)
Induction and refinement	Timely feedback	Providing an evaluative response (e.g., correct or incorrect) soon after a student's attempt at task or step.	Corbett and Anderson (2001)
	Feature focusing	Instruction leads to more robust learning when it guides the learner's attention (''focuses'') to valid or relevant features of target KCs.	Dunlap et al. (2011)
	Worked examples	Students learn more efficiently and more robustly when more frequent study of worked examples is interleaved with problem solving practice as opposed to practice that is all problem solving.	Sweller and Cooper (1985)
Understanding and sense making	Prompted self-explanation	Encouraging students to explain to themselves parts of instruction (steps in worked example or sentences in a text) yields more robust learning than not prompting or providing such explanations to students.	Chi et al. (1994); Hausmann and VanLehn (2007)
	Accountable talk	Teacher use of ''talk moves,'' particular question and response patterns which encourage students to be accountable to accurate knowledge, rigorous reasoning, and the classroom community, leads to more robust learning.	Michaels et al. (2008)

5.1.2. Optimized scheduling

Optimizing scheduling is a more specific version of the spacing principle and builds on Pavlik's (2007) observation that much past research has not controlled for time on task and thus has underestimated the benefit of shorter practice intervals early in KC acquisition. This principle involves applying an instructional schedule that has been ordered to optimize robust learning. More precisely, what is optimized is instructional efficiency, that is, gains in robust learning per instructional time. Optimization is achieved mathematically by deriving when a student should repeat practice of a KC. The time interval between practice opportunities of a KC is optimal (neither too short nor too long) when it best balances the benefit of enhanced memory strength, a benefit higher at a long interval (spaced practice), with the cost of time to retrain owing to retrieval failure, a cost higher at a long interval. Mathematical models may be used to produce optimized schedules by computing the KC that will be most efficiently learned, if practiced next.

5.2. Induction and refinement enhancing instruction

5.2.1. Feature focusing

This principle asserts the value of attending to cues or features that are valid for the targeted KCs. Focus on key features may help students to learn more quickly those KCs most important for the goals of learning. More generally, focusing may also result in students spending more time during a learning event on a particular KC and thus increase its strength.

An example of feature focusing comes from learning to read Chinese, whose characters often are compounds, consisting of two components. Often these components (radicals) provide cues to pronunciation and meaning. For example, consider the compound character 晴, which is translated as fair weather. On the left is a semantic radical 日 that means ''sun'' and on the right is a phonetic radical 青 that is pronounced ''qing.'' Knowing that 日 means ''sun'' is useful in learning the meaning of this compound and others that contain it. Feature focus directs attention to the form for ''sun'' in association with is meaning. Standard Chinese reading instruction tends not to do this, emphasizing instead the meaning of the character as a whole. However, the research indicates that focusing on the feature of component form-meaning association supports learning of characters (Taft & Chung, 1999) and that a short instruction to focus on the semantic radical brings improvement that is dramatic and immediate (Wang, Liu, & Perfetti, 2004). In one implementation of the feature focus principle, Dunlap, Perfetti, Liu, and Wu (2011) found that highlighting the semantic component as a student moves the computer mouse over it improves learning the character.

A second example comes from science education. Chen and Klahr (1999) found that when elementary school students were provided with instruction that drew their attention to the specific features that distinguish confounded and unconfounded experiments, there were significantly greater gains in children's ability to design good experiments when compared with a condition in which children were only asked questions but not instructed about such specific features.

5.2.2. Worked examples

In a worked example, students are given a problem description along with a step-by-step solution to the problem and are asked to study or self-explain the solution. Sweller and Cooper (1985) demonstrated that alternating worked examples with standard problems leads to more efficient and more effective learning than having students to solve all the problems in the set: Students work through the interleaved worked examples and problems more quickly and perform better on corresponding problems in a posttest. Subsequent studies have refined the worked example technique. Renkl, Atkinson, and Maier (2000) introduced a ''fading'' method that yields even better learning outcomes. In this fading technique, a succession of problems is presented. The first is completely solved and the rest incrementally replace a solved step with a request for a student solution until ultimately the student is providing the solution for all steps in a problem. Students are also encouraged to self-explain the solved solution steps in this fading method and it yields consistent evidence of more robust learning than problem solving alone. Section 6 provides an expanded analysis of this principle.

5.3. Understanding and sense-making enhancing instruction

5.3.1. Prompted self-explanation principle

When students are given a worked example or text to study, prompting them to self-explain each step of the worked example or each line of the text usually causes higher learning gains than having them study the material without such prompting (e.g., Aleven & Koedinger, 2002; Chi et al., 1994; Renkl et al., 1998), but exceptions are discussed below. Hausmann and VanLehn (2007) found that prompting students to self-explain while solving a physics problem produced more learning than providing them with high-quality explanations to study. When it comes to explaining physics examples, it appears it is better to do it yourself, even if you get it wrong, than to study someone else's explanations.

This principle provides a good example of how the effectiveness of a principle may be dependent on the nature of the target KCs. When both non-verbal and verbal versions of knowledge are objectives of instruction (i.e., we want both fluent doing and deliberate explaining), prompting self-explanation can pay off doubly by both: (a) strengthening verbal forms of knowledge, and (b) providing redundant support (co-training) for acquiring non-verbal forms of knowledge. The redundant support idea relies on the assumption that self-explanation is a verbal explication process that functions in addition to non-verbal learning processes. Verbal knowledge is an instructional objective for much of math and science, where students are expected to be able to state principles (e.g., Newton's laws) and provide explanations of solutions. In contrast, in language, a fluent speaker is not expected to explain principles that might underlie grammatical choices, such as English article selection.

Prompting students to self-explain may not aid learning for some kinds of KCs. One such KC type is grounded in perceptual learning, for which learner verbalization can serve as an inhibitor (Schooler et al., 1997). Something similar may occur when a to-be-learned KC is relatively simple, that is, non-verbal, lacking a deep rationale, probabilistic, and/or

dependent primarily on perceptual information. Hints in that direction come from a study of learning the English double dative construction (Frishkoff et al., 2008). English allows speakers to say either ''John gave the book to Mary'' or ''John gave Mary the book.'' However, rather than a set of rules, native speakers' choices about which noun to put first are subject to a web of factors that are implicitly weighted (Bresnan et al., 1997). These factors can be translated into instructional heuristics. When Frishkoff et al. (2008) implemented these heuristics into problem examples, they found that English second language learners subsequently made choices more in line with those of native speakers. However, when they provided these heuristics as feedback following some degree of implicit example-only learning (with correctness-only feedback), they found it was not helpful. This might imply that verbalization—in this case not by the learner—can be helpful or harmful as a function of specifically how it connects to learning that is proceeding implicitly. Self-explanation also might function within a similar space, supportive of learning in general, but sometimes interfering with learning processes that are essentially non-verbal (e.g., Schooler et al., 1997).

Finally, a more general qualification on the principle of Self-Explanation is that its successful application, as is the case with sense-making activities generally, depends on having access to relevant knowledge or to learning strategies that can support the explanation process. McNamara (2004) found that low knowledge learners reading a science text were unable to benefit from prompted self-explanation without training. However, with training in strategies to support comprehension, including self-explanation strategies, low knowledge learners were able to show gains in comprehension of the science text.

5.4. Linking knowledge analysis and instructional principles

Our examples so far have illustrated how instructional methods may align with the learning processes. The existence of these alignments suggests that instructional principles will not apply universally but will be dependent on the kind of KCs that are the targets of instruction. Across a wide variety of instructional experiments, simpler IE types (involving less time, less feedback, less verbalization and reasoning) tend to be associated with simpler KCs (involving less time, less complex conditions and responses, less integration with related KCs). Table 6 illustrates this relationship.

The apparent association in Table 6 suggests that the complexity of instruction should be aligned with the complexity of the knowledge goals (the *alignment hypothesis*). Thus, (1) for simple kinds of knowledge, complex forms of instruction are not needed and (2) for complex kinds of knowledge, simple forms of instruction are not needed. However, our analysis of which learning processes are relevant to which kinds of knowledge (shown in Table 4) suggests an alternative in which simple kinds of knowledge become embedded in more complex kinds, leading to an asymmetry. Instruction for complex KCs (columns to the right of both Tables 4 and 6) may include simpler IEs, whereas instruction for simpler KCs (columns to the left in both tables) would not benefit from more complex IEs (the *asymmetry hypothesis*). For simple KCs, as in our vocabulary examples, memory processes may be sufficient. Somewhat more complex KCs, such as learning English articles, require induction of a category structure

Table 6
A possible correlation between the complexity of knowledge components (KCs) and the complexity of the instruction that best produces such knowledge

	Instructional Principles (simpler on bottom)	Chinese vocab	French articles	English articles	Algebra eq	Geometry rules	Chemistry rules	Help seeking skls	Physics prinpls	Chem models	Pressure concept
Understanding and Sense Making	Accountable Talk										+
	Collaboration			0							+
	Self-explanation			0	+	+			+		
	Worked examples				+	+	+				
Induction and Refinement	Diagram coordination					+				+	
	Feature Focusing	+	+								
	Feedback		+			+		+			
Mem & Fluency	Optimal Scheduling	+				+					
		Facts			Rules				Principles		

Knowledge Components
(simpler on left)

Note. ''+'' indicates the corresponding instructional principle has been experimentally demonstrated to enhance robust learning efficiency for the corresponding kind of KCs. ''0'' indicates an experiment found no difference. An empty cell indicates absence of relevant experiments.

as well as memory. Learning articles involves forming generalizations, for instance, that ''the'' is used across the variable set of situations in which the referent of the target noun was previously mentioned. But, because the acquired category structure must also be remembered, instructional methods that effectively engage memory processes are also important (cf. Frank & Gibson, 2011). This line of reasoning argues against the alignment hypothesis. It implies that simpler memory enhancing instructional methods, such as spacing, testing, or optimized scheduling, are also effective for more complex KCs.

What about more complex instructional methods for simpler KCs? To learn an arbitrary constant-constant association (a fact), a generalization process like that implemented by category induction is unneeded. Furthermore, for an arbitrary association (e.g., that a Chinese character that does not contain the ''ching'' phonetic radical is nevertheless pronounced ''ching'') an explanation structure cannot be used to generate or re-derive a KC. Re-derivation is functional only for knowledge that has an underlying rationalization

(complex knowledge). These observations suggest the asymmetry hypothesis is more nearly correct.

One counter argument to the asymmetry hypothesis is that more complex instructional methods, such as accountable talk or even self-explanation prompting, can indirectly achieve the same robust learning efficiency outcomes (e.g., memory enhancement) that simpler methods (e.g., optimized scheduling) achieve. Engaging basic memory processes may not be necessary to the extent that complex forms of instruction help students form an integrated network of knowledge (e.g., combinations of principles and reasoning strategies) that can be used to regenerate or re-derive forgotten knowledge. Supporters of mnemonic or memory elaboration strategies might take this argument further to suggest that robust learning of simple paired associates can be enhanced by engaging learners in an explanation process in that the mnemonic is essentially an explanation.

To the extent the instructional goal is robust learning *efficiency*, outcomes must not only last and transfer but also be achieved with less time or, at least, without extra time. Too many theoretical analyses and experimental studies do not address the time costs of instructional methods. Practically, use of more complex instructional strategies may not always be worth the extra time they tend to require.

Others have suggested that the generality of instructional principles may be bounded by the complexity of targeted knowledge (e.g., Rea & Modigliani, 1985; Sweller & Chandler, 1994; Wulf & Shea, 2002), but more research is needed to clarify boundaries. The KLI taxonomies provide a conceptual space to guide alternative hypothesis formation, associated experimentation, and theoretical interpretation.

6. Using the KLI framework to illuminate and compare instructional principles

We provide an example of how the KLI framework can be used to analyze instructional principles, showing how they operate in LEs at the KC-level. We focus on the *worked example principle* and make connections to the testing and self-explanation principles. As described in section 5.2, integrating worked examples with problem solving is more efficient and effective learning for novices than conventional problem solving alone. Our analysis sheds light on alternative theoretical accounts and demonstrates how the application of the KLI framework helps specify boundary conditions on the kinds of knowledge and the types of students for which worked examples are effective.

6.1. Theoretical interpretations

The theoretical framework that drives much of the worked example research is *Cognitive Load Theory* (van Merriënboer & Sweller, 2005). Human processing capacity is limited, and this theory argues that much of the cognitive load in novice problem solving practice is extraneous to robust learning. In problem solving, novices rely heavily on means-ends analysis and trial-and-error strategies. While these processes are useful for solving novel problems, they require cognitive resources that cannot then be used (or be used as much) to

engage in induction or sense making (e.g., in reflecting upon or self-explaining generated solution steps). Thus, novices are less likely to develop a deep understanding of the domain while engaged in pure problem-solving practice.

An alternative, but possibly complementary, theory proposes that problem solving yields poorer learning outcomes not because students' cognitive resources are depleted, but because there is less environmental information to support students in filling *KC gaps* (cf., McNamara, Kintsch, Songer, & Kintsch, 1996; VanLehn, 1999). As a complete and correct solution is not available in problem solving, students can apply incorrect knowledge without realizing it, resulting in the refinement and strengthening of incorrect KCs. In addition, when students do recognize knowledge gaps, they are more likely to induce appropriate KCs with worked examples because a worked example provides more input (KC conditions *and* responses) than a problem statement that has only the initial problem (KC conditions only). This additional input allows induction processes to build better solution knowledge.

6.2. Knowledge component level analysis

In the earliest worked example research, the basic unit of analysis was the complete solution: Students were asked to study full worked examples and in problem solving were given feedback and answers based on full solutions. The KLI framework suggests that examining instruction at the level of KCs, rather than whole problems, can illuminate a deeper theoretical analysis that provides more precision in defining the critical novice state. Subsequent computational modeling and empirical research that analyzes the worked example principle at the level of KCs is beginning to tease apart the competing explanations.

On the empirical front, Renkl and colleagues shifted the analytic focus for worked examples to the grain size of KCs when they compared interleaved worked examples and problem solving with an example "fading" condition (Atkinson, Renkl, & Merrill, 2003; Renkl, Atkinson, Maier, & Staley, 2002; Renkl et al., 2000). (See section 5.2) They found that fading was more effective, yielding better learning outcomes with the same learning time.

Jones and Fleischman (2001) subsequently examined a formal KC-level explanation of worked examples by applying a computational model of learning called Cascade (VanLehn, 1999) to a set of Newtonian physics problems, asking it both to explain worked examples and solve problems. These simulations, which model information available in the environment but not cognitive load nor motivation, replicated the benefits of faded examples over complete worked examples on subsequent problem solving. They provide a theoretical argument that the Knowledge Gap theory is sufficient (although it may not be necessary) to produce the worked example effect, and Cognitive Load Theory is not necessary (although it may be sufficient). In addition, the authors found that Cascade learns more effectively when the specific problem-solving steps that are faded are ones that introduce new KCs.

Finally, Renkl et al. (2002) began to examine the ordering of faded steps to compare the predictions of the Cognitive Load and KC gap-filling theories. They found that *backward* fading (fading the last step, then the last two steps, etc.) is more efficient and effective than

forward fading (fading the first step, then fading the first two, etc.). This result appears to support Cognitive Load theory under the assumption that a student's cognitive load is lower at the end of a problem, so the student is better able to learn from faded steps at the end of a problem. However, Renkl, Atkinson, and Große (2004) note that this result is ambiguous because the KC content of faded steps is typically confounded with the position of faded steps. In two studies, they varied the position of a faded step independently of the specific KC required for the faded step. The results were consistent with the KC gap-filling theory rather than the Cognitive Load theory. Students' test performance was reliably predicted by the specific KC that governs a faded step but not by the position of the faded step.

6.3. Defining novices: A knowledge component analysis

In contrast to the positive results of worked examples for novice students, a number of studies (e.g., Kalyuga, Ayres, Chandler, & Sweller, 2003; Kalyuga & Sweller, 2005) have demonstrated that for more experienced students, straight problem solving yielded better learning than interleaved worked examples and problem solving. This result has been dubbed the ''expertise reversal effect.''

Expertise reversal raises the question of defining the novice-expert boundary. In most worked-example research, novices are students who have just received instruction and are just beginning to apply it; the distinction between novices and advanced students is defined at the grain size of broad topics within the domain. However, beginning students can vary greatly in their prior knowledge. Kalyuga and Sweller (2005) addressed this variability by developing pretests that categorize students as novices or advanced students, and this showed that channeling the novices into worked examples and the advanced students directly into problem solving yielded more efficient learning outcomes.

In the KLI framework the novice/advanced student distinction is that novices are primarily engaged in KC induction and sense making, whereas advanced students are primarily engaged in refinement and fluency building. Novices may benefit from seeing examples of solution steps to induce variable-variable KCs and from seeing the entire solution structure to make sense of the role of each step so as to construct integrated KCs for generating plans and subgoals (cf. Catrambone, 1996). Problem solving, in contrast, offers advanced students the opportunity to build memory and fluency through active retrieval opportunities (as in the testing effect) and to refine the conditions of application through feedback on incorrect solution attempts (Matsuda et al., 2008). A similar case, in which knowledgeable students benefit from active retrieval, may arise in the reverse cohesion effect in learning from text (McNamara, 2001; McNamara et al., 1996). In this reversal effect, readers with high knowledge in a domain do better in learning from a text when the text has cohesion gaps, which presumably stimulate more active knowledge retrieval and inference making.

Within the KLI framework the expertise reversal effect marks a boundary between two seemingly contradictory principles, the worked example principle (being shown answers is more effective than retrieving answers) and the testing effect (retrieving answers is more effective than being shown answers). In addition to the KC-based boundary suggested above

(i.e., worked examples are better for variable condition KCs, testing for constant-constant KCs), an LE boundary is that worked examples are better early to support induction and testing is better later to support memory and fluency.

The optimal grain size for these boundaries is at the level of individual KCs, rather than the broad topics within a domain, and the ideal goal is to monitor each student's progression from sense making and induction to refinement and fluency building for each KC. A recent study by Salden, Aleven et al. (2010) pursued this KC-based approach and demonstrated that adaptive fading of examples led to greater robust learning than fixed-fading and no fading conditions (see the description of example fading in section 5). In the adaptive fading condition, the Cognitive Tutor's built-in student modeling algorithm was used to monitor each student's accuracy in generating explanations for each of the KCs in the cognitive model. When the student modeling algorithm estimated a student could correctly generate explanations for the application of a KC, the student was judged to be transitioning from induction/sense making to refinement/fluency building and in subsequent problems, solution steps employing that KC were faded (i.e., presented as problem steps rather than example steps).

6.4. Summary

We illustrated the use of the KLI framework to analyze an instructional principle, the worked example effect. This analysis (a) finds support for the KC Gap theory as a viable alternative to Cognitive Load theory, (b) suggests that ''novice'' is not a relation between a student and a domain, but between a student and a KC, (c) proposes boundary conditions between this principle and a related one, the testing effect, and (d) illustrates how knowledge-level analysis can be combined with a general instructional principle (faded worked examples) to produce a student-adaptive version of that principle that enhances robust learning.

7. Conclusion

Our goal has been to put forward a theoretical framework (KLI) that links knowledge, learning, and instruction to organize the development of instructional theory at a grain size appropriate for guiding the design, development, and continual improvement of effective and efficient academic course materials, technologies, and instructor practices. This goal reflects the purpose of our Pittsburgh Science of Learning Center, ''to leverage cognitive theory and computational modeling to identify the conditions that cause robust student learning.'' It is also consistent with broader calls for cumulative theory development in education that is supported by ''rigorous, sustained scientific research'' (e.g., Shavelson & Towne, 2002). The need for cumulative theory is illustrated both by the general lack of consensus around educational practices that work and by the limitations of large-scale randomized control trials, which are great tests of instructional practices but do not generate new theory or practices.

In elaborating the KLI framework, we proposed three fundamental taxonomies of kinds of knowledge, learning processes, and instructional principles. We outlined potential interdependencies between categories in these taxonomies and illustrated how the framework can be used to generate new research questions and frame alternative hypotheses.

In developing the KLI framework, we emphasized the importance of KCs as opposed to domains (Geometry; English). In contrast to Bloom's well-known taxonomy (Bloom, 1956), which is expressed in terms of instructional objectives, our taxonomy focuses on the knowledge needed to achieve those objectives and is expressed in cognitive process terms. It is at a more abstract and coarse-grain level than the representations used in computational models of cognition (e.g., Anderson & Lebiere, 1998; McClelland & Cleeremans, 2009; Newell, 1990; Sun, 1994). Knowledge, in our account, is decomposable into units that relate some input characteristics or features of the student's perceived world or mental state (the conditions) to some output in the student's changeable world or mental state (the response). Unlike production rules in theories of cognitive architecture (Anderson & Lebiere, 1998; Newell, 1990), which are implicit components outside a students' awareness, the KCs in KLI include explicit, verbalizable knowledge. Given the prominence of comprehension, reasoning, dialog, and argumentation in more complex forms of instruction (e.g., prompted self-explanation, accountable talk), the KLI knowledge taxonomy distinguishes between kinds of KCs that have accessible rationales, such that students can effectively reason and argue about them, and some that do not, such that explicit reasoning and argumentation may be of little value for learning.

Learning occurs as unobservable events that can be inferred from performance and can be appropriately attributed to instruction events under circumstances of experimental control. The processes of learning include both simple associations and more complex, reflective processes that result in KC changes of three broad types: (a) memory and fluency building, (b) induction and refinement, and (c) understanding and sense making. These learning processes can proceed more or less independently or in some synchrony.

Instructional principles emerge from research that is sufficiently convergent to support generalization. Instructional principles are intended to be widely applicable across domains and situations but in fact are likely constrained in their applicability by the kinds of KCs to be learned and students' stage of learning. Table 5 summarizes seven such principles, which have broad experimental support. We noted a general trend for a correspondence between the complexity of the principle and complexity of the KCs being targeted in those studies. But our learning process analysis suggests an asymmetric relationship could turn out to be more nearly correct: Simple instructional principles are generally relevant, but more complex principles are only relevant for the most complex kinds of knowledge.

We developed the case of worked examples in enough detail to illustrate the richness of applying the KLI framework to a single question that has been the focus of much research. Using the KLI framework and specifically the analysis of KCs led to experiments that found that instruction that was individualized to specific KCs led to more robust learning than instruction that was not (Salden et al., 2008). Such studies not only validate the

basic assumption that KCs are the functional unit of analysis for learning, they suggest instructional procedures that can be the object of further research, leading at some point to a broad instructional principle.

The strategy of creating a student-adaptive version of a principle by applying a knowledge analysis could be applied not only to adaptive fading of worked examples as described above but also to other principles as Pavlik (2007) has demonstrated for the spacing effect. For example, it may be productive to employ student-adaptive fading from blocked practice to random practice (cf. Wulf & Shea, 2002), from comparison (Gentner et al., 2009) to sequential spaced presentation (Rohrer & Taylor, 2006), from low content variability to high variability (cf. Paas & Van Merrienboer, 1994). Regarding the testing effect, the Salden et al. (2008) study indicates that adaptive fading from study trials (''examples'' in this context) to test trials (''problems'') is sometimes more effective. The knowledge taxonomy provides a guide for hypothesizing how far the results may generalize outside the Geometry content that was investigated. Instead of making a domain-general claim about applicability, we suggest that fading of example study to problem/test should work for content with variable-variable KCs with rationales. For constant-constant KCs, adaptive fading may be wasted effort and simply providing study trials (examples) only after failures on test/problem attempts may be optimal, as is done in the most effective conditions in the testing effect experiments.

Other researchers have recognized the importance for effective instructional design of a detailed analysis of domain content into the components of knowledge that students bring to a course, and those we would like them to take away. Nevertheless, Anderson and Schunn (2000) expressed concern that ''detailed cognitive analyses of various critical educational domains are largely ignored by psychologists, domain experts and educators.'' They noted the tendency for psychologists to value domain-general results, domain experts (e.g., mathematicians and linguists) to value new results in their domain, and educational researchers to value more holistic explanations. Some progress has been made (e.g., Clark et al., 2007; Lee, 2003), but careful cognitive task analysis of domain knowledge is not a standard research practice in any discipline. Such analysis needs to become a more routine part of instructional design for new instructional domains as well as for existing ones.

Finally, we emphasize that the KLI framework implies a broad range of empirical studies that can disconfirm as well as strengthen some of its propositions. The framework is not a set of frozen taxonomies but an interconnected set of theoretical and empirical propositions that imply hypothesis-testing experiments. As with any theoretical framework, its utility is tested by whether it stimulates sufficient work to lead to its revision, abandonment, or enrichment through an increasingly well-targeted set of research results from the learning and educational sciences.

Notes

1. The level of KC analysis is a function of the target content *and* the target student population. It may not have a tight age correspondence. The relevant level for adults

learning to read Chinese as second language may be lower than second graders learning to read English as a first language. Reading errors of adults learning Chinese could lead to a focus on KCs at the level of the stroke patterns that form the graphemes (characters) that make up words, whereas reading errors of second graders learning English may lead to a focus on KC relevant for word identification that are above the level of recognizing graphemes (letters).

2. While our notion of KC shares similarity in condition-action structure with production rules in theories of cognitive architecture (Anderson & Lebiere, 1998; Newell, 1990) and with situation-action schemas in theories of analogical transfer (e.g., Gick & Holyoak, 1983), it is a more general idea that encompasses both and includes explicit, verbalizable knowledge. The inclusion of a condition component, even for verbal, declarative knowledge, represents a KLI assumption that learning *when* to do or think something is a critical part of learning.

3. In formal computational modeling, paired associates (constant-constant mappings) can be represented in a schema or production rule without the use of any variables, whereas categories (variable-constant) and general rules or principles (variable-variable) require the use of variables. Or in structure mapping (Gentner, Loewenstein, Thompson, & Forbus, 2009), relations are variables and the primitive objects (at the leaves of a structure map) are the constants.

4. The sun radical, one can argue, reflects the evolution of an iconic non-arbitrary representation of the sun to a more abstract graph that continues to be non-arbitrary; however, from a reader's point of view graphs in general have at most an opaque, seemingly arbitrary, connection to meaning, even in Chinese and certainly in alphabetic writing systems.

5. Vocabulary KCs for words with explicit morphological markers (e.g., past tense of regular verbs in English, like ''jumped'') are not members of this constant-constant fact category, but of the variable-variable rule category (e.g., To form the past tense of <verb>, produce <verb> followed by ''ed'').

6. Note that knowledge of recognizing expressions as fractions (e.g., saying that ''3/4'' is a fraction but that ''3–4'' is not) is not the same as being able to state the definition of a fraction, which is a constant-constant KC (mapping ''fraction'' to ''quotient of two quantities'').

7. Knowledge of visual images (e.g., someone's face) is in declarative memory in ACT-R, but people cannot always verbalize such knowledge so it is non-verbal in KLI.

8. For example, while the plural of *tip* is pronounced *tip*[s], the plural of *tub* is *tub*[z]. The difference between these two pronunciations is handled by a natural phonotactic constraint that makes it easier to assimilate the voicing feature from the preceding consonant (/p/and/s/are unvoiced,/b/and/z/are voiced), thus producing *tub*[z], than to shift from voiced to unvoiced to produce *tub*[s]. This constraint is not arbitrary because its source is the human speech mechanism.

9. In a KC analysis for a particular course, a KC should be unpacked into smaller KCs when incoming student performance on tasks assessing those smaller KCs is

sufficiently lower than a high threshold success rate (e.g., 95%), or, in cases emphasizing fluency, slower than a desired reaction time.

Acknowledgments

References

Ainsworth, S., & Loizou, A. Th. (2003). The effects of self-explaining when learning with text or diagrams. *Cognitive Science*, *27*, 669–681.

Aleven, V., & Koedinger, K. R. (2002). An effective metacognitive strategy: Learning by doing and explaining with a computer-based cognitive tutor. *Cognitive Science*, *26*(2), 147–149.

Aleven, V., Roll, I., McLaren, B. M., & Koedinger, K. R. (2010). Automated, unobtrusive, action-by- action assessment of self-regulation during learning with an intelligent tutoring system. *Educational Psychologist*, *45*(4), 224–233.

Alibali, M. W., & Koedinger, K. R. (1999). The developmental progression from implicit to explicit knowledge: A computational approach. (Commentary on Z. Dienes & J. Perner, A theory of implicit and explicit knowledge.). *Behavioral and Brain Sciences*, *10*, 755–756.

Anderson, J. R. (1993). *Rules of the mind*. Mahwah, NJ: Lawrence Erlbaum.

Anderson, J. R. (2002). Spanning seven orders of magnitude: A challenge for cognitive modeling. *Cognitive Science*, *26*(1), 85–112.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, *111*(4), 1036–1060.

Anderson, J. R., Fincham, J. M., & Douglas, S. (1997). The role of examples and rules in the acquisition of cognitive skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 932–945.

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

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

Asterhan, C. S. C., & Schwarz, B. B. (2009). The role of argumentation and explanation in conceptual change: Indications from protocol analyses of peer-to-peer dialogue. *Cognitive Science*, *33*, 373–399.

Atkinson, R. K., Renkl, A., & Merrill, M. M. (2003). Transitioning from studying examples to solving problems: Effects of self-explanation prompts and fading worked-out steps. *Journal of Educational Psychology*, *95*, 774–783.

Barnett, S. M., & Ceci, S. J. (2002). When and where do we apply what we learn? A taxonomy for far transfer. *Psychological Bulletin*, *128*, 612–637.

Bloom, B. S. (Ed.) (1956). *Taxonomy of educational objectives. Handbook 1: Cognitive domain*. New York: McKay.

Blum, A., & Mitchell, T. (1998). Combining labeled and unlabeled data with co-training. In S. Ben-David & P. Long (Eds.), *Proceedings of Eleventh Annual Conference on Computational Learning Theory (COLT)* (pp. 92–100). New York: ACM Press.

Booth, J. L., & Koedinger, K. R. (2008). Key misconceptions in algebraic problem solving. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 64–70). Austin, TX: Cognitive Science Society.

Bransford, J. D., Brown, A. L., & Cocking, R. R. (Eds.). (2000). *How people learn: Brain, mind, experience, and school*. Washington, DC: National Academy Press.

Bransford, J. D., & Schwartz, D.. (1999). Rethinking transfer: A simple proposal with multiple implications. In A. Iran-Nejad, & P. D. Pearson (Eds.), *Review of Research in Education* (Vol. 24, pp. 61–100). Washington, DC: American Educational Research Association.

Bresnan, J., Cueini, A., Nikitina, T., & Baayen, R. H. (2007). Predicting the dative alteration. In G. Bouma, I. Kramer, & J. Zwarts (Eds.), *Cognitive foundations of interpretation*. Chicago, IL: University of Chicago Press.

Bruner, J. S. (1966). *Toward a theory of instruction*. New York: W. W. Norton & Co.

Case, R., & Okamoto, Y. (1996). The role of central conceptual structures in the development of children's thought. *Monographs of the Society for Research in Child Development*, *61*(1-2, Serial No. 246).

Catrambone, R. (1996). Generalizing solution procedures learned from examples. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 1020–1031.

Cepeda, N. J., Pashler, H., Vul, E., Wixted, J. T., & Rohrer, D. (2006). Distributed practice in verbal recall tasks: A review and quantitative synthesis. *Psychological Bulletin*, *132*, 354–380.

Chang, N. M. (2006). *Learning to discriminate and generalize through problem comparisons*. Unpublished doctoral dissertation, Pittsburgh PA: Carnegie Mellon University.

Chen, Z., & Klahr, D. (1999). All other things being equal: Children's acquisition of the control of variables strategy. *Child Development*, *70*(5), 1098–1120.

Chen, Z., & Klahr, D., (2008). Remote transfer of scientific reasoning and problem-solving strategies in children. In R. V. Kail (Ed.) *Advances in Child Development and Behavior, Vol. 36* (pp. 419–470). Amsterdam: Elsevier.

Chi, M. T. H., Bassok, M., Lewis, M., Reimann, P., & Glaser, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science*, *13*, 145–182.

Chi, M. T. H., de Leeuw, N., Chiu, M. H., & LaVancher, C. (1994). Eliciting self-explanations improves understanding. *Cognitive Science*, *18*(3), 439–477.

Chi, M. T. H., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, *5*, 121–152.

Clark, R. E., Feldon, D., van Merriënboer, J., Yates, K., & Early, S. (2007). Cognitive task analysis. In J. M. Spector, M. D. Merrill, J. J. G. van Merriënboer, & M. P. Driscoll (Eds.), *Handbook of research on educational communications and technology* (3rd ed., pp. 577–593). Mahwah, NJ: Lawrence Erlbaum Associates.

Clark, R. C., & Meyer, R. E. (2003). *e-Learning and the science of instruction*. San Francisco: Jossey-Bass/Pfeiffer.

Clement, J. (1987). Overcoming students' misconceptions in physics: Fundamental change in children's physics knowledge. *Journal of Research in Science Teaching*, *28*, 785–797.

Corbett, A. T., & Anderson, J. R. (2001). Locus of feedback control in computer-based tutoring: Impact on learning rate, achievement and attitudes. *Proceedings of ACM CHI'2001 Conference on Human Factors in Computing Systems*, (pp. 245–252). New York: ACM Press.

De Jong, N., & Perfetti, C. A. (2011). Fluency training in the ESL classroom: An experimental study of fluency development and proceduralization. *Language Learning*, *61*(2), 533–568.

Dienes, Z., & Perner, J. (1999). A theory of implicit and explicit knowledge. *Behavioral and Brain Sciences*, *22*(5), 735–808.

Donovan, J. J., & Radosevich, D. J. (1999). A meta-analytic review of the distribution of practise effect: Now you see it, now you don't. *Journal of Applied Psychology*, *84*, 795–805.

Dunlap, S., Perfetti, C. A., Liu, Y., & Wu, S. M. (2011) Learning vocabulary in Chinese as a foreign language: Effects of explicit instruction and semantic cue reliability. http://www.pitt.edu/~perfetti/PDF/Dunlap Learning vocabulary.pdf

Engelmann, S., & Carnine, D. (1991). *Theory of instruction: Principles and applications*. Eugene, OR: ADI Press.

Feng, M., Heffernan, N. T., & Koedinger, K. R. (2009). Addressing the assessment challenge in an online system that tutors as it assesses. *User Modeling and User-Adapted Interaction: The Journal of Personalization Research (UMUAI)*, *19*(3), 243–266.

Fox, M. C., & Charness, N. (2009). How to gain eleven IQ points in ten minutes: Thinking aloud improves Raven's Matrices performance in older adults. *Aging, Neuropsychology, and Cognition*, *17*(2), 191–204.

Frank, M. C., & Gibson, E. (2011). Overcoming memory limitations in rule learning. *Language, Learning, & Development*, *7*, 130–148.

Frishkoff, G. A., Pavlik, P., Levin, L., & de Jong, C. (2008). Providing optimal support for robust learning of syntactic constructions in ESL. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the Annual Meeting of the Cognitive Science Society* (CogSci08) (pp. 64–70). Austin, TX: Cognitive Science Society.

Gagne, R. M. (1985). *The conditions of learning and theory of instruction*, 4th ed. New York: Holt, Rinehart, and Winston.

Gentner, D., Loewenstein, J., Thompson, L., & Forbus, K. D. (2009). Reviving inert knowledge: Analogical abstraction supports relational retrieval of past events. *Cognitive Science*, *33*, 1343–1382.

Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, *15*, 1–38.

Gobet, F. (2005). Chunking models of expertise: Implications for education. *Applied Cognitive Psychology*, *19*, 183–204.

Graesser, A. C. (2009). Cognitive scientists prefer theories and testable principles with teeth. *Educational Psychologist*, *44*(3), 193–197.

Graesser, A. C., McNamara, D. S., & VanLehn, K. (2005). Scaffolding deep comprehension strategies through Point&Query, AutoTutor, and iSTART. *Educational Psychologist*, *40*, 225–234.

Greeno, J. G. (1998). The situativy of knowing, learning, and research. *American Psychologist*, *53*(1), 5–26.

Hausmann, R. G. M., Nokes, T. J., VanLehn, K., & Gershman, S. (2009). The design of self-explanation prompts: The fit hypothesis. In N. Taatgern & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 2626–2631). Austin, TX: Cognitive Science Society.

Hausmann, R. G. M., & VanLehn, K. (2007). Explaining self-explaining: A contrast between content and generation. In R. Luckin, K. R. Koedinger, & J. Greer (Eds.), *Proceedings of AI in Education* (pp. 417–424). Amsterdam: IOS Press.

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

Heffernan, N., & Koedinger, K. R. (1997). The composition effect in symbolizing: The role of symbol production versus text comprehension. In M. G. Shafto, & P. Langley (Eds.), *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society* (pp. 307–312). Hillsdale, NJ: Erlbaum.

Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review*, *110*, 220–264.

Jones, R. M., & Fleischman, E. S. (2001) CASCADE explains and informs the utility of fading examples to problems. In J. D. Moore & K. Stenning (Eds.), *Proceedings of the 23rd Annual Conference of the Cognitive Science Society* (pp. 459–464). Mahwah, NJ: Erlbaum.

Kalyuga, S., Ayres, P., Chandler, P., & Sweller, J. (2003). The expertise reversal effect. *Educational Psychologist*, *38*, 23–32.

Kalyuga, S., & Sweller, J. (2005). Rapid dynamic assessment of expertise to improve the efficiency of adaptive e-learning. *Educational Technology Research and Development*, *53*, 83–93.

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

Klahr, D., & Dunbar, K. (1988). Dual space search during scientific reasoning. *Cognitive Science*, *12*, 1–48.

Koedinger, K. R., & Aleven, V. (2007). Exploring the assistance dilemma in experiments with cognitive tutors. *Educational Psychology Review*, *19*(3), 239–264.

Koedinger, K. R., Aleven, V., Roll, I., & Baker, R. (2009). *In vivo* experiments on whether supporting metacognition in intelligent tutoring systems yields robust learning. In D. J. Hacker, J. Dunlosky, & A. C. Graesser (Eds.), *Handbook of metacognition in education*. (pp. 383–412). New York: Routledge.

Koedinger, K. R., & McLaughlin, E. A. (2010). Seeing language learning inside the math: Cognitive analysis yields transfer. In S. Ohlsson & R. Catrambone (Eds.), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society* (pp. 471–476). Austin, TX: Cognitive Science Society.

Koedinger, K. R., & Nathan, M. J. (2004). The real story behind story problems: Effects of representations on quantitative reasoning. *The Journal of the Learning Sciences*, *13*(2), 129–164.

Koedinger, K. R., & Terao, A. (2002). A cognitive task analysis of using pictures to support pre-algebraic reasoning. In C. D. Schunn, & W. Gray (Eds.), *Proceedings of the Twenty-Fourth Annual Conference of the Cognitive Science Society* (pp. 542–547). Mahwah, NJ: Lawrence Erlbaum Associates.

Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge, UK: Cambridge University Press.

Lee, R. L. (2003). *Cognitive task analysis: A meta-analysis of comparative studies*. Unpublished doctoral dissertation, Los Angeles, California: University of Southern California.

Levine, M. (1966). Hypothesis behavior by humans during discrimination learning. *Journal of Experimental Psychology*, *71*, 331–338.

Lewis, C. H. (1988). Why and how to learn why: Analysis-based generalization of procedures. *Cognitive Science*, *12*, 211–256.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, *95*(4), 492–527.

MacLaren, B., & Koedinger, K. R. (2002). When and why does mastery learning work: Instructional experiments with ACT-R ''SimStudents.'' In S. A. Cerri, G. Gouarderes, & F. Paraguacu (Eds.), *Proceedings of the 6th International Conference on Intelligent Tutoring Systems* (pp. 355–366). Berlin: Springer-Verlag

Matsuda, N., Cohen, W., Sewall, J., Lacerda, G., & Koedinger, K. R. (2008). Why tutored problem solving may be better than example study: Theoretical implications from a simulated-student study. In B. Woolf, E. Aimeur, R. Nkambou, & S. P. LaJoie (Eds.), *Proceedings of the 9th International Conference of Intelligent Tutoring Systems* (pp. 111–121). Berlin: Springer-Verlag.

McClelland, J. L., & Cleeremans, A. (2009). Connectionist models. In T. Byrne, A. Cleeremans, & P. Wilken (Eds.), *Oxford companion to consciousness*. New York: Oxford University Press.

McClelland, J. L., Fiez, J. A., & McCandliss, B. D. (2002). Teaching the/r/-/l/discrimination to Japanese adults: Behavioral and neural aspects. *Physiology & Behavior*, *77*, 657–62.

McNamara, D. S. (2001). Reading both high and low coherence texts: Effects of text sequence and prior knowledge. *Canadian Journal of Experimental Psychology*, *55*, 51–62.

McNamara, D. S. (2004). SERT: Self-explanation reading training. *Discourse Processes*, *38*(1), 1–30.

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

Medin, D., & Schaffer, M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207–238.

van Merriënboer, J. J. G., & Sweller, J. (2005). Cognitive load theory and complex learning: Recent developments and future directions. *Educational Psychology Review*, *17*(1), 147–177.

Michaels, S., O'Connor, C., & Resnick, L. B. (2008). Deliberative discourse idealized and realized: Accountable talk in the classroom and in civic life. *Studies in the Philosophy of Education*, *27*(4), 283–297.

Minstrell, J. (2001). Facets of students' thinking: Designing to cross the gap from research to standards-based practice. In K. Crowley, C. D. Schunn, & T. Okada (Eds.), *Designing for Science: Implications for Professional, Instructional, and Everyday Science* (pp. 415–444). Mahwah, NJ: Lawrence Erlbaum Associates.

Mitchell, T. M., Keller, R. M., & Kedar-Cabelli, S. T. (1986). Explanation-based generalization: A unifying view. *Machine Learning*, *1*, 47–80.

National Reading Panel (2000). *Report of the National Reading Panel. Teaching children to read: An evidence-based assessment of the scientific research literature on reading and its implications for reading instruction (00-4769)*. Washington, DC: National Institute of Child Health & Human Development.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

Ogan, A., Aleven, V., & Jones, C. (2009). Advancing development of intercultural competence through supporting predictions in narrative video. *International Journal of Artificial Intelligence in Education*, *19*(3), 267–288.

Ohlsson, S. (1996). Learning from performance errors. *Psychological Review*, *103*(2), 241–262.

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

Paas, F. G. W. C., & Van Merrienboer, J. J. G. (1994). Variability of worked examples and transfer of geometrical problem-solving skills: A cognitive-load approach. *Journal of Educational Psychology*, *86*, 122–133.

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

Pashler, H., Bain, P., Bottge, B., Graesser, A., Koedinger, K., McDaniel, M., & Metcalfe, J. (2007). *Organizing instruction and study to improve student learning (NCER 2007–2004)*. Washington, DC: National Center for Education Research, Institute of Education Sciences, U.S. Department of Education.

Pavlik., P. I. Jr (2007). Understanding and applying the dynamics of test practice and study practice. *Instructional Science*, *35*, 407–441.

Posner, M. I. (in press). The expert brain. In J. Staszewski (Ed.), *Expertise and skill acquisition: The impact of William G. Chase*. New York: Psychology Press.

Quilici, J. H., & Mayer, R. E. (1996). Role of examples in how students learn to categorize statistics word problems. *Journal of Educational Psychology*, *88*, 144–161.

Rayner, K., Foorman, B. R., Perfetti, C. A., Pesetsky, D., & Seidenberg, M. S. (2001). How should reading be taught. *Psychological Science*, *2*(2 suppl.), 31–74.

Rea, C. P., & Modigliani, V. (1985). The effect of expanded versus massed practise on the retention of multiplication facts and spelling lists. *Human Learning: Journal of Practical Research & Applications*, *4*, 11–18.

Renkl, A. (1997). Learning from worked-out examples: A study on individual differences. *Cognitive Science*, *21*, 1–29.

Renkl, A., Atkinson, R. K., & Große, C. S. (2004). How fading worked solution steps works—a cognitive load perspective. *Instructional Science*, *32*, 59–82.

Renkl, A., Atkinson, R. K., & Maier, U. H. (2000). From studying examples to solving problems: Fading worked-out solution steps helps learning. In L. Gleitman, & A. K. Joshi (Eds.), *Proceeding of the 22nd Annual Conference of the Cognitive Science Society* (pp. 393–398). Mahwah, NJ: Erlbaum.

Renkl, A., Atkinson, R. K., Maier, U. H., & Staley, R. (2002). From example study to problem solving: Smooth transitions help learning. *Journal of Experimental Education*, *70*, 293–315.

Renkl, A., Stark, R., Gruber, H., & Mandl, H. (1998). Learning from worked-out examples: The effects of example variability and elicited self-explanations. *Contemporary Educational Psychology*, *23*, 90–108.

Richman, H. B., Staszewski, J. J., & Simon, H. A. (1995). Simulation of expert memory with EPAM IV. *Psychological Review*, *102*, 305–330.

Rissanen, J. (1978). Modeling by shortest data description. *Automatica*, *14*, 465–471.

Rittle-Johnson, B., & Koedinger, K. R. (2001). Using cognitive models to guide instructional design: The case of fraction division. In J. D. Moore, & K. Stenning (Eds.), *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society* (pp. 857–862). Mahwah, NJ: Erlbaum.

Roediger, H. L., & Karpicke, J. D. (2006). The power of testing memory: Basic research and implications for educational practice. *Perspectives on Psychological Science*, *1*, 181–210.

Rohrer, D., & Taylor, K. (2006). The effects of overlearning and distributed practice on the retention of mathematics knowledge. *Applied Cognitive Psychology*, *20*, 1209–1244.

Salden, R. J. C. M., Aleven, V., Renkl, A., & Schwonke, R. (2008). Worked examples and tutored problem solving: Redundant or synergistic forms of support? In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 589–594). Austin, TX: Cognitive Science Society.

Schmidt, R. A., & Bjork, R. A. (1992). New conceptualizations of practice: Common principles in three paradigms suggest new concepts for training. *Psychological Science*, *3*, 207–217.

Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search and attention. *Psychological Review*, *84*, 1–66.

Schoenfeld, A. (2004). The math wars. *Educational Policy*, *18*(1), 253–286.

Schooler, J. W, Fiore, S., & Brandimonte, M. A. (1997). At a loss from words: Verbal overshadowing of perceptual memories. *Psychology of Learning and Motivation: Advances in Research and Theory*, *37*, 291–340.

Servan-Schreiber, E., & Anderson, J. R. (1990). Learning artificial grammars with competitive chunking. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*, 592–608.

diSessa, A. A. (1993). Toward an epistemology of physics. *Cognition and Instruction*, *10*(2 & 3), 105–225.

Shavelson, R. J., & Towne, L. (Eds.) (2002). *Scientific research in education*. Washington, DC: National Academy Press.

Shepard, R. N. (1987). Toward a universal law of generalization of psychological science. *Science*, *237*, 1317–1323.

Slotta, J. D., & Chi, M. T. H. (2006). The impact of ontology training on conceptual change: Helping students understand the challenging topics in science. *Cognition and Instruction*, *24*, 261–289.

Snow, C., Burns, M., & Griffin, P. (Eds.) (1998). *Preventing reading difficulties in young children*. Washington, DC: National Academy Press.

Stamper, J., & Koedinger, K. R. (2011). Human-machine student model discovery and improvement using data. In J. Kay, S. Bull, & G. Biswas (Eds.), *Proceedings of the 15th International Conference on Artificial Intelligence in Education* (pp. 353–360). Berlin: Springer.

Stylianides, G. J., & Stylianides, A. J. (2009). Facilitating the transition from empirical arguments to proof. *Journal for Research in Mathematics Education*, *40*(3), 314–352.

Sun, R., 1994. *Integrating rules and connectionism for robust commonsense reasoning*. New York: John Wiley and Sons.

Sweller, J., & Chandler, P. (1994). Why some material is difficult to learn. *Cognition and Instruction*, *12*, 185–233.

Sweller, J., & Cooper, G. A. (1985). The use of worked examples as a substitute for problem solving in learning algebra. *Cognition and Instruction*, *2*, 59–89.

Taft, M., & Chung, K. (1999). Using radicals in teaching Chinese characters to second language learners. *Psychologia*, *42*, 243–251.

VanLehn, K. (1999). Rule learning events in the acquisition of a complex skill: An evaluation of cascade. *Journal of the Learning Sciences*, *8*(1), 71–125.

Wang, M., Liu, Y., & Perfetti, C. A. (2004). The implicit and explicit learning of orthographic structure and function of a new writing system. *Scientific Studies of Reading*, *8*(4), 357–379.

Wulf, G., & Shea, C. H. (2002). Principles derived from the study of simple skills do not generalize to complex skill learning. *Psychonomic Bulletin & Review*, *9*, 185–211.

Wylie, R., Koedinger, K. R., & Mitamura, T. (2009). Is self-explanation always better? The effects of adding self-explanation prompts to an English grammar tutor. In N. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 1300–1305). Amsterdam: The Netherlands.

Zhu, X., Lee, Y., Simon, H. A., & Zhu, D. (1996). Cue recognition and cue elaboration in learning from examples. *Proceedings of the National Academy of Sciences*, *93*, 1346–1351.