Chapter 5

Knowledge Representation and Acquisition

ARTHUR C. GRAESSER, ANNE M. LIPPERT, & KEITH T. SHUBECK

University of Memphis

This chapter discusses how knowledge is represented in our minds when we learn about new topics in school and life. How do we encode and think about subject matters in fields as diverse as psychology, literature, art, history, biology, physics, mathematics, and computer technology? The knowledge representations and reasoning in these fields often differ (Goldman et al., 2016). In psychology and physics, we think like a scientist. We think about hypotheses and how to test them by collecting data in experiments. In mathematics, we puzzle over formulas and proofs. In literature, we construct imaginary worlds in our mind that may or may not correspond to anything in the real world. In computer technology, we think about procedures for running programs that perform some practical task. The representations and ground rules for thinking are quite different in these different disciplines.

There are multiple ways to represent experiences and topics of interest. Popular music is a great example of this. Consider how people represent music when they listen to songs such as *Hey Jude* by the Beatles, *Crazy in Love* by Beyoncé, or *Yankee Doodle*. Some have representations that focus on the melody, others the lyrics, others the emotions, others visual images, and others the rhythm and meter that inspire dance or other forms of physical motion. Most of us have mental representations with some combination of these dimensions. There is no right or wrong representation, but memory for the songs is influenced by the nature of the representations that people construct (Rubin, 1995). Psychologists in the learning sciences investigate the nature of the representations that we construct when we learn new topics and use the knowledge when performing tasks.

Mental representations of what we perceive are not perfect copies of the world out there. The mental representations we construct about the world are simplifications that often have errors and distortions. As an interesting exercise, draw from memory a floorplan of your home, with the various doors, windows, and pieces of furniture. Then compare the sketch with your actual home and note the differences. Or if you prefer, sketch your town with the streets and landmarks. Although you have experienced your home and town for hundreds of thousands of days, there are still distortions. Psychologists in the cognitive sciences investigate theories about the properties of these mental representations and conduct experiments to test the theories.

This chapter identifies some of the theories of representation that cognitive and learning scientists have developed. Their goal is to explain how children and adults represent knowledge during learning. The focus of this chapter is on learning when adults acquire subject matters in schools, the workforce, and their personal lives. In contrast, Chapter 4 ("Concepts: Structure and Acquisition") and Chapter 17 ("Development of Human Thought") take on the development of representations in infants and children. Our emphasis is also on deeper lev-

https://doi.org/10.17885/heiup.470.c6668

els of comprehension and learning (Millis, Long, Magliano, & Wiemer, 2019). A recent report by the National Academy of Sciences, Engineering and Medicine on *How People Learn* (volume 2, 2018) contrasts six basic types of learning: habit formation and conditioning, observational learning, implicit pattern learning, perceptual and motor learning, learning of facts, and learning by making inferences from mental models. This chapter emphasizes the learning of facts and making inferences from mental models, although the other types of learning are sometimes very relevant.

Instructional media and technology will play an important role in this chapter because they dominate the world we live in today. Media and technology shape how we think and represent information. For example, a few decades ago it would have taken days to find an answer to a question as people walked to libraries, to card catalogues, to stacks of books, and searched pages and paragraphs for an answer. The same question can now be answered in seconds on the computer. We expect swift answers to questions and get irritated by delays. A decade ago students submitted essays for grading and waited for days or weeks for a grade. Now essays can be graded immediately with validity comparable to experts (Foltz, 2016). We now live in a world of intelligent tutoring systems that tailor learning to the individual student (Graesser, Hu, & Sottilare, 2018) and computer environments where groups of people can learn and solve problems together (Fiore & Wiltshire, 2016). We now live in a world where facts need to be checked for misinformation and contradictions (Rapp & Braasch, 2014) and technology has the only major capacity to do so. We live in a world of media, games, and adutainment. These seductions appeal to our motivational and emotional seductions and run the risk of competing with the learning of important subject matter. All of these advances in media and technology influence how we represent and acquire knowledge.

5.1 Knowledge Components

This first approach to representing subject matter knowledge consists of a list of **knowledge components**. A knowledge component is much like a sentence that expresses a particular idea that is important to know about a topic. Example knowledge components in psychology can be captured in such expressions as “absence makes the heart grow fonder” (as the opposite to “out of sight, out of mind”), “team members in groups may not respond because they expect other members to respond”, or “correlation does not imply causation.” An example in physics is “force equals mass times acceleration” whereas an example in mathematics is “the circumference of a circle is pi times the diameter.” Some knowledge components are if-then rules with contingencies: “If a person has XX chromosomes, they are female; if a person has XY chromosomes, they are male.” The subject matter on a topic may consist of a long list of dozens to hundreds of knowledge components. As students learn a subject matter, students and teachers do not know how well the performance on these knowledge components is progressing. However, computers can track this progress for individual students in intelligent tutoring systems (Graesser, 2016; Koedinger, Corbett, & Perfetti, 2012) and for individuals and groups in team learning (von Davier, Zhu, & Kyllonen, 2017). When the computer determines that enough of the knowledge components have been learned by the student, the system then decides that the student has mastered the topic.

How does the student, instructor, or computer know whether a knowledge component (KC) has been mastered? The answer is debatable. Consider once again the knowledge component “team members in groups may not respond because they expect other members to respond.” How would one know whether this KC has been mastered by a learner? There are many possible operational definitions. Can the learner recite the KC in words that have the same meaning as the KC? Does the learner send important requests to individuals rather than groups in social communication systems (knowing that there may be diffusion of responsibility in groups)? Mastery of some KC’s may be reflected in a number of cognitive measures, such as response times to requests, eye movements, and neuroscience indicators (see Chapter 3, “Methods for Studying Human

Thought"). Individual learners may differ in how they behaviorally show mastery of a particular KC. They may exhibit mastery in words, drawing figures, gestures, problem solving, or other actions.

Mastery of knowledge components improves over time if there is knowledge acquisition. Computers can track this. Suppose a computer tracks whether or not a student on a KC has a successful response (1) or an unsuccessful response (0) over 8 episodes of being assessed. The following sequence would reflect successful learning on assessment episode number 4: 00011111. The sequence 01010101 shows no learning because the number of 1's is the same for the first four episodes and the second four. Probabilistic learning is reflected in 00101011 because there is only one 1 among the first four episodes but three 1's in the last four episodes. Mastery of a topic is achieved when many of the KCs are mastered in performance assessments.

5.2 The Representation of Knowledge Components

The mastery of a knowledge component depends how it is represented and how picky one is as to whether it is mastered. A precise standard for a verbal representation would be an exact match between the expected knowledge component and the student's language. However, it is important to match on meaning rather than precise language (Kintsch, 1998). There are many ways to articulate "team members in groups may not respond because they expect other members to respond" in particular contexts, such as "there is diffusion of responsibility in the group", "tell John personally because he expects others on the team to handle the task", or "the likelihood of a team member completing an assigned task is lower than when an individual is assigned the task." How can one determine whether these answers match the KC when they are worded so differently? Computers have made major advances in evaluating the accuracy of semantic matches in a field called computational linguistics (Jurafsky & Martin, 2008), but they are far from perfect. Expert human judges have moderate agreement on whether two sentences have the same or different meanings, but they also do not always agree.

Multiple levels of language and discourse need to be considered when deciding whether two verbal expressions have the same meaning (Pickering & Garrod, 2004; McNamara, Graesser, McCarthy, & Cai, 2014). We need to consider whether the words have the same or similar meaning. For example, the phrase "team members in groups" is very similar in meaning to "people in groups" in the example KC but not to "sports in groups." Syntax and word order matter when interpreting meaning. The meaning of the phrase "team members in groups" is quite different in meaning than "to members group in teams" and the nonsensical expression "groups team in members." The discourse context also needs to be considered when deciding whether two sentences have the same meaning. The expression "absence makes the heart grow fonder" makes sense in a psychology class when debating whether a romance will survive after two lovers part for a few months. It does not make sense when a student tries to explain to an instructor why an exam was missed.

Mastery of a knowledge component is manifested in its meaning rather than the precise surface structure (i.e., wording and syntax). People tend to remember in long-term memory the meaning of ideas rather than the surface structure (Craik & Lockhart, 1972). Surface structure is normally short-lived, a minute or less, whereas the semantic meaning lasts a long time. Therefore, verbal memory assessments of how well a student has mastered a subject matter need to consider the meaning of the KCs rather than the exact wording. An essay test that taps meaning is superior to a test on reciting texts verbatim.

Mastery of a knowledge component is often manifested nonverbally. Actions, facial expressions, eye movements, pointing gestures, and other behaviors can signal mastery. Consider a KC that "some chemical sprays from groundkeepers cause people to sneeze." When someone starts sneezing, this KC is likely to have been mastered if the person gets up and looks out the window, glares in contempt at the groundkeeper, points to the groundkeeper, closes the window, and/or puts on an allergy mask. There is no need to articulate the KC in words.

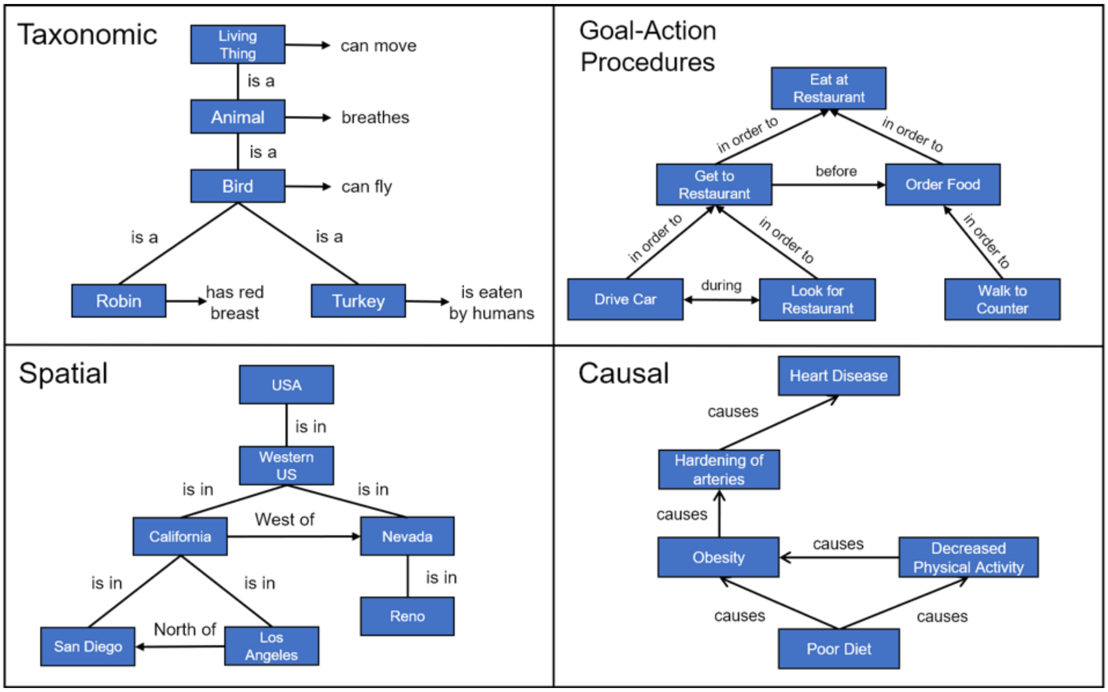

Figure 5.1: Four different types of knowledge structures: Taxonomic, spatial, causal, and goal-action procedures.

5.3 Knowledge Structures

Our description of the knowledge component representation does not take into consideration the structural relations between ideas. This section emphasizes these relational connections. Four types of structures are being discussed here to illustrate the importance of relations. These are shown in Figure 5.1: Taxonomic, spatial, goal-action procedures, and causal structures. There are many other types of **knowledge structures**, such as organizational charts of positions in a corporation and the lineage in family trees. All of these knowledge structures emphasize how knowledge is interconnected and that ideas close to each other in the structure are more conceptually related than ideas far away. When an idea is activated during learning, it tends to activate its nearby neighbors in the structure more than neighbors far away (Collins & Loftus, 1975).

There is a terminology that researchers use to talk about these knowledge structures. **Nodes** are basic ideas that can be expressed in a word, phrase or sentence. As explained above, however, it is the meaning rather than the surface structure that captures the essence of a node. Nodes are sometimes assigned to epistemic categories, such as concept, state, event, process, goal, or action. An *arc* is a connection between two nodes. An arc is directed (forward, backward, or bidirectional) and often assigned to categories (such as is-a, has-as-parts, property, contains, cause, reason). A **graph** consists of a set of nodes connected by **arcs**. Below we describe some different kinds of graphs that are depicted in Figure 5.1.

5.3.1 Taxonomic Structures

Taxonomic structures represent the concepts that were discussed in Chapter 4, “Concepts: Structure and Acquisition”. The concepts are organized in a hierarchical structure that is connected by *is-a* arcs. A robin is-a bird, a turkey is-a bird, a bird is-a animal, an animal is-a living thing. These is-a arcs that are directly represented in the graph, but others can be inferred by the principle of transitivity: a robin is an animal, a turkey is an animal, a robin is a living thing, a turkey is a living thing, and a bird is a living thing. Each of these concept nodes have distinctive *properties*, such as a robin has a red breast, a turkey is eaten by humans, a bird can fly, an animal breathes, and living things can move. These properties can be inherited by transitive inference, such as the following expressions: a robin can fly, a robin breathes, a robin can move, a bird can move, and so forth. There is some evidence that these inferred expressions take a bit more time to judge as true or false than the direct expressions (Collins & Loftus, 1975).

5.3.2 Spatial Structures

Spatial structures have a hierarchy of regions that are connected by *is-in* arcs (or the *inverse* contains relation). As shown in Figure 5.1, Los Angeles is-in California, San Diego is-in California, Reno is-in Nevada, California is-in the western US, Nevada is-in the western US, and the western US is-in the USA. From these, we can derive via transitivity the following inferences: Los Angeles is in the western US, San-Diego is in the western US, Reno is in the western US, Los Angeles is in the USA, and so on. The locations within each region can also be connected by relational arcs that specify north, south, east, and west. We see in Figure 5.1 that Los Angeles is north-of San-Diego and California is west-of Nevada. We can infer by transitivity that San Diego is west of Reno. Most of these transitive inferences are correct when we look at actual maps. However, these inferences are not always correct (Stevens & Coupe, 1978). For example, San Diego is actually east of Reno rather than west of Reno according to an actual map. Similarly, Seattle is actually north of Toronto and El Paso is actually west of Denver. Knowledge structures and these transitive inferences are often accurate, but sometimes generate some interesting errors. The knowledge structures also can to some extent predict biases in distance. For example, distances between cities within a region can also, to some extent, seem closer than distances between cities from different regions. The distance from Memphis to Jackson, Tennessee seems closer than to Jackson, Mississippi, yet the actual distance is the opposite.

5.3.3 Goal-action Procedures

Goal-action procedural structures are organized into a hierarchy of nodes connected by "in order to" arcs. The nodes refer to goals or desired states that are organized hierarchically and that guide a sequence of actions that achieve the goals if the procedure is successfully performed. Imagine you have a goal of eating at a restaurant. The structure in Figure 5.1 shows how this could be accomplished. In order to eat at the restaurant, you need to get to the restaurant and order your food. In order to get to a restaurant, you need to drive your car and look for the restaurant. This specific knowledge structure in Figure 5.1 does not require careful deliberation to plan and execute. The procedure becomes a routine through experience and repetition. It would be exhausting to plan through problem solving for each step of every goal-action procedure you carry out throughout the day. However, such problem solving (see Chapter 9) is needed when a person visits another country.

The structure in Figure 5.1 is taken from the perspective of one person who needs food. However, there are other people who have their own agenda, such as the cook and the person at the counter. A **script** is a structure that considers all of the people who participate in the organized activity of a restaurant (Bower, Black, & Turner, 1979). The cook, the person at the counter who collects money, and the customer all have their own goal structures and perspectives. The script also has taxonomic structures (cook → employee → person) and spatial structure (table → restaurant → building).

These goal-action procedures and script structures explain a number of psychological phenomena. Each goal-action node is broken down into subordinate nodes that become much more detailed in the activity. People tend to forget the lower-level details of the actions and procedures (Bower et al., 1979), which are often automatized from repetition and experience (see Chapter 13, "Expertise"). People tend to notice obstacles to goals being accomplished and may become frustrated, as everyone who has waited for many minutes trying to order food at a counter knows. When people visually observe scripts being enacted, they tend to notice event boundaries (i.e., junctures, separations) after a goal is achieved/interrupted, when there is a new spatial setting, and when a new person enters a scene (Zacks, Speer, & Reynolds, 2009). When people read stories, sentences take more time to read when they introduce new goals, spatial settings, and characters (Zwaan & Radvansky, 1998). These structures also explain answers to questions. When asked, "Why do you go to a restaurant?", a good answer would go up the structure (in order to eat food) but not down the structure (in order to drive). When asked "how do you go to a restaurant?", a good answer would be down the structure (you drive) but not up the structure (you eat). Organized structures like these explain a large body of data involving neuroscience, cognition, behavior, emotion, and social interaction.

5.3.4 Causal Networks

Causal networks can be used to answer the question, "What causes something to occur?" For example, one could use causal networks to show the chain of events that cause a volcanic eruption, cancer, the winner of an election and other phenomenon in physical, biological, and technological systems (van den Broek, 2010). In a causal network, nodes represent events (or states, or processes) whereas arcs point from one node to another if an event causes or enables another event. For example, in Figure 5.1, we have a causal network showing how heart disease can be a result of a causally driven chain of events. Some of these events are inspired by sociological factors (getting a divorce) and psychobiological factors (smoking), whereas other events are entirely products of biological systems (hardening of the arteries). The events in the causal system that are linked through *enables* arcs convey a weak sense of causality, while the *causes* arcs indicate a stronger sense of causality. Causal networks are complex. They are not strictly hierarchical or follow a linear order but can have many paths of connections and loops.

The structures in Figure 5.1 are very systematic, organized, and conceptually precise. The mental structures are not that neat and tidy. One approach to help people learn is to have them construct such graphs during or after they comprehend text, digital environments on the internet, conduct an experiment, or perform some other activity. The activity of constructing these conceptual graphs can help them learn a subject matter even though they are not likely to generate neat and tidy structures. Available research has also revealed that nodes that are more central in the structure (i.e., many arcs radiate from them) are more important and better remembered (Bower et al., 1979; van den Broek, 2010).

5.4 Associative Representations of Knowledge

According to classical associationism, ideas vary in how strongly associated they are with each other. That is no doubt true, but the deep secret lies on what can predict the strength of association. A word like "evil" has likely strong associations to words like "bad" (a functional synonym), "good" (an opposite), "Halloween" (an event), "Knieval" (part of the phrase evil Knievel, the dare devil), and "devil" interesting etymology), but not to words like "smooth", "birthday", and "Michael Jordan."

What makes associations strong versus weak? Strength of repetition is clearly one factor. The strength of association between ideas increases with the frequency of the ideas occurring together at the same time and location. Another prediction is the similarity of the ideas. The strength of association between two ideas is stronger to the extent they are similar in meaning. Positive outcomes is yet another prediction: two ideas have stronger association to the extent that they lead to positive outcomes (a reward, a solution) rather than negative outcomes (punishment, failure). In summary, repetition, similarity, and reinforcement are major predictions of the strength of association between two ideas.

These principles of associationism have been known for at least two centuries. They are deeply entrenched in modern cognitive models of perception, categorization, memory, judgment, and other auto-

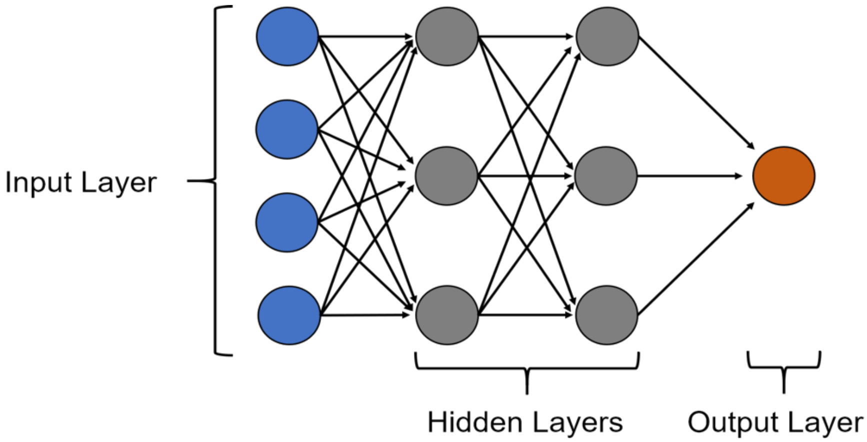

Figure 5.2: A neural network with an input layer, two hidden layers, and an output node.

mated processes of cognition. **Neural networks** are a noteworthy class of models that implement associationism (McClelland & Rumelhart, 1987). Figure 5.2 presents an example of a neural network. A neural network is a structure of nodes (analogous to neurons) in multiple layers that are interconnected by directed, weighted arcs that potentially activate the nodes (positive weights) or inhibit the nodes (negative weights). A node is fired (all-or-none) if the arcs that feed into it receive enough activation, with the sum of the activation being stronger than the inhibition.

In order to illustrate the mechanisms of a neural network, consider a neural network that detects whether or not a person's face shows confusion. The input layer of nodes would correspond to states, events, or processes on parts of the face at particular positions. For example, the right eyelid opens wide, the mouth opens wide, or the left corner of the lip contracts. Ekman and his colleagues developed a facial action coding system that defines these features for those who investigate facial expressions (Ekman & Rosenberg, 2005). The output node is activated if the set of activated input node features show a pattern of confusion, but otherwise it is not activated. There may also be one or more hidden layers of nodes that refer to intermediate states, events, or processes. Exactly what these hidden nodes refer to is not necessarily clear-cut and easy to interpret. They could refer to higher order categories, such as the overall amount of movement, positive versus negative emotions, upper face parts versus lower face parts, or angle of perspective. The hidden layers and nodes within these layers are statistically derived characteristics that depend on a long history of experiences that the individual person has had. It is important to emphasize that these neural networks learn from experience. The nodes and arcs are strengthened or otherwise altered with each experience. The networks capture the associationist principles of repetition, similarity, reinforcement, and contiguity of events in time and space.

Today neural networks are frequently used in machine learning and artificial intelligence to enable computers to perceive people, objects, events, and scenes, to guide robots in completing routine tasks, and to solve some types of problems. In this "deep learning" revolution, massive amounts of experiences are fed into the computer during training of the neural network, far more than a single person would ever receive. As a consequence, the computer outperforms humans in precisely defined tasks. This has the potential to threaten the workforce for some jobs that humans traditionally perform (Elliot, 2017). These neural networks can handle only specific tasks, however. A neural network for detecting confusion would not be of much use to detect surprise or boredom – they cannot generalize and transfer to other tasks. Nevertheless, it is widely acknowledged that generalization and transfer are also very difficult for humans to accomplish (Hattie & Donoghue, 2016). Perhaps the human mind is little more than a large collection of these specialized neural networks. This is a debate in the cognitive and learning sciences.

Another example of associative knowledge representations is latent semantic analysis, LSA (Landauer, McNamara, Dennis, & Kintsch, 2007). LSA is a statistical representation of word knowledge and world knowledge that considers what words appear together in documents, such as articles in books, speeches, conversations, and other forms of verbal communication. According to LSA, the meaning of a word depends on the other words that accompany it in real-world documents. The word *riot* often occurs in the company of other particular words in documents, such as *crowd*, *dangerous*, *protest*, *police*, and *run*. These words do not always occur with the word riot of course, but they do with some co-occurrence probability. These probabilities of words with other words define a word's meaning, which is very different than word meanings in a dictionary or thesaurus. LSA has been found to predict data in many cognitive tasks such as priming (a word automatically activates another word), judgments of sentence similarity, inferences, and summarization of text (Landauer et al., 2007). LSA has also been used in computer systems that automatically grade student essays (Foltz, 2016) and tutor them in natural language (Graesser, 2016).

5.5 The Body in Cognition

Proponents of **embodied cognition** believe that mental representations are shaped and constrained by the experience of being in a human body. Our bodies influence what we perceive, our actions, and our emotions. These embodied dimensions are often incorporated in representations when we comprehend text (Zwaan, 2016) and influence how we learn (Glenberg, Goldberg, & Zhu, 2011). Embodied representations are constructed, for example, when you read a novel and get lost in the story world. There is a rich mental model of the spatial setting, the actions performed by characters, and their emotions. Your experience is similar to watching a movie or acting the parts yourself. Mental representations are often colored with perceptual images, motoric actions, and visceral emotions rather than being abstract conceptualizations. The meaning of abstract concepts (such as love) is often fortified by these dimensions of perception, action, and emotion, such as visual image of a wedding cake, a dance, or a first kiss (Barsalou & Wiemer-Hastings, 2005). There is substantial evidence that memory is improved for verbal material when learners construct visual images in their mind (Clark & Paivio, 1991) or they perform actions associated with the content.

The importance of embodied cognition in comprehension is obvious when you go someplace new and ask for directions to a specific location, such as the city hall. When you ask a stranger, "Where is the city hall?" the helpful stranger nearly always points in the right direction and launches several sentences with landmarks, paths, and left-right-straight comments, typically accompanied by hand gestures. You get confused by the second sentence but politely nod. Then you follow the suggested direction and soon ask the next person. The problem is that there is very little shared knowledge between you and the stranger so you have no foundation for constructing a precise embodied path to the destination. Embodied representations are necessary for precise comprehension of important messages about the physical, social, and digital worlds.

The importance of embodied representations on reading comprehension has been confirmed in the *Moved by Reading* program (Glenberg, Goldberg, & Zhu, 2011). Readers who struggle with reading comprehension experience difficulty constructing an embodied representation of the text. Suppose that students read a text about events that occur at a tea party. This would be difficult to imagine if they had no knowledge or experience with tea parties. In *Moved by Reading*, the student is presented with an image of a tea set on a computer screen and then asked to act out a story on the content by pouring tea, sipping tea, and performing other actions conveyed in the story. Students are also later asked to imagine acting out the story so they will internalize the strategy of constructing a mental model of the text. When compared to students who were asked to simply reread the text, the students who were asked to imagine manipulating the objects showed large gains in comprehension and memory. One of the interesting research questions is whether it is better to physically perform the actions compared to digitally moving images on a computer screen or to imagine performing actions in the mind.

5.6 Conversations

People have learned by observing and participating in conversations throughout most of the history of personkind, especially prior to the invention of the printing press and computer technologies. The secrets of family life and a person's livelihood were learned by holding conversations with members of a family, a tutor, a mentor, a master, or a group of people participating in the practical activities. Knowledge representations are to some extent shaped by these conversations that are observed, enacted, remembered, or otherwise internalized in the mind (Vygotsky, 1978). Texts that are written in the style of stories and oral conversation are read faster, comprehended better, and remembered better than technical text that is distant from conversation.

There is also solid evidence that one-on-one human tutoring helps to learn subject matter in courses more than simply listening to lectures or reading texts (Cohen, Kulik, & Kulik, 1982; VanLehn, 2011). The individual tutor can find out the problems the learner is facing, provide hints or direct assertions on helping them, and answer their questions. Researchers have developed intelligent tutoring systems that simulate human tutors (VanLehn, 2011), including some systems like AutoTutor that hold conversations with the student in natural language (Graesser, 2016). These systems help students learn subject matters like computer literacy, physics, and

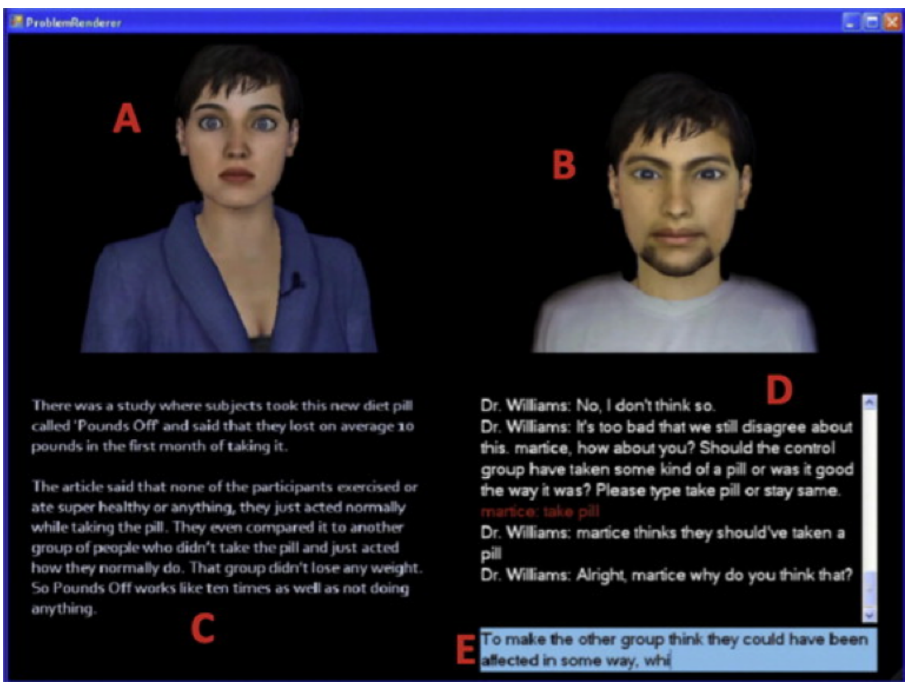

Figure 5.3: This is a screenshot showing pedagogical agents used in an intelligent tutoring system (D'Mello, Lehman, Pekrun, & Graesser, 2014). In this example, the tutor agent, Dr. Williams is on the left of the screen, and the peer agent, Chris, is on the right of the screen. Reprinted from Learning and Instruction, 29, D'Mello, S., Lehman B., Pekrun, R., & Graesser, A.C. Confusion can be beneficial for learning. 153-170. ©(2014), with permission from Elsevier.

scientific reasoning about as good as human tutors, both of which are better than conventional training methods like reading texts and listening to lectures.

A promising approach to establish deeper knowledge representations is to plant contradictions and information that clashes with prior knowledge to the point of the learner experiencing cognitive disequilibrium. Cognitive disequilibrium occurs when people face obstacles to goals, interruptions, contradictions, incongruities, anomalies, impasses, uncertainty, and salient contrasts. Cognitive conflicts can provoke information-seeking behavior, which engages the learner in inquiry, reasoning and deep learning. Learning environments with computer agents have been designed to stage contradictions and debates, thereby inducing cognitive disequilibrium (D'Mello, Lehman, Pekrun, & Graesser, 2014). These studies had tutor and peer agents engage with the student in conversational trialogues while critiquing research studies in psychology, biology, and chemistry. An example screenshot is shown in Figure 5.3. Most of the research studies had one or more flaws with respect to scientific methodology. For example, one case study described a new pill that purportedly helps people lose weight, but the sample size was small and there was no control group. During the course of the three-way conversation, the agents periodically expressed false information and contradictions. Disagreements between the agents and with what the student believed tended to create cognitive disequilibrium, confusion, and disagreement. During the course of the trialogue conversation, the agents periodically asked students for their views (e.g., "Do you agree that the control group in this study was flawed?"). The students' responses were coded on correctness and also the vacillation in making decisions when asked a question multiple times throughout a conversation. There were also measures of confusion. The correctness and confusion scores confirmed that the cognitive disequilibrium that resulted from contradictions improved learning, particularly among the students who had enough knowledge and thinking to be confused. That is, the experience of confusion, a signal of thinking, played an important role in the deep learning.

Table 5.1: Key affordances of learning technologies (National Academy of Sciences, Engineering, and Medicine, 2018). ©National Academies Press. Reprinted with permision. https://www.nap.edu/catalog/24783/how-people-learn-ii-learners-contexts-and-cultures

1. *Interactivity.*	The technology systematically responds to the actions of the learner.
2. *Adaptivity.*	The technology presents information that is contingent on the behavior, knowledge, or characteristics of the learner.
3. *Feedback.*	The technology gives the learner information about the quality of their performance and how it could improve.
4. *Choice.*	The technology gives learners options on what to learn and how to regulate their own learning.
5. *Nonlinear access.*	The technology allows the learner to select or receive learning activities in an order that deviates from a set order.
6. *Linked representations.*	The technology provides quick connections between representations for a topic that emphasizes different conceptual viewpoints, media, and pedagogical strategies.
7. *Open-ended learner input.*	The technology allows learners to express themselves through natural language, drawing pictures, and other forms of open-ended communication.
8. *Communication with other people.*	The learner communicates with one or more people or agents.

5.7 Importance of Media and Technology in Knowledge Representation and Learning

Theories of distributed cognition assume that the mind is shaped and constrained by the physical world, technologies, and other people in their environment (Dror & Harnad, 2008; Hutchins, 1995). An expert problem solver in a distributed world needs to assess whether a technology, a social community, the external physical world, or his/her own analytical mind is best suited for achieving particular steps in solving challenging problems. Judgments are involved in the decisions you make when you decide whether to trust your own analytical judgment, the output of a computer program, or a decision of a group. There are questions such as "Should I write down on a piece of paper the groceries I need to buy or try to memorize them?"; "Should I compute this square root by hand or use a calculator?"; "Should I ask my friends where to on vacation or decide that for them?" These are decisions in a distributed world.

Media and technology play a central role in shaping cognitive representations in a distributed world. It is important to take stock of how they do so. Old-school media consisted of listening to lectures, watching video presentations, and reading books. For these media, the learners passively ob-

Table 5.2: Mayer's (2009) Principles to Guide Multimedia Learning. Adapted from NAESM (2018). With permission from National Academy of Sciences, Engineering, and Medicine, 2018. ©National Academies Press. https://www.nap.edu/catalog/24783/how-people-learn-ii-learners-contexts-and-cultures

1. *Coherence Principle*	People learn better when extraneous words, pictures and sounds are excluded rather than included.
2. *Signaling Principle*	People learn better when cues that highlight the organization of the essential material are added.
3. *Redundancy Principle*	People learn better from graphics and narration than from graphics, narration and on-screen text.
4. *Spatial Contiguity Principle*	People learn better when corresponding words and pictures are presented near rather than far from each other on the page or screen.
5. *Temporal Contiguity Principle*	People learn better when corresponding words and pictures are presented simultaneously rather than successively.
6. *Segmenting Principle*	People learn better from a multimedia lesson is presented in user-paced segments rather than as a continuous unit.
7. *Pre-training Principle*	People learn better from a multimedia lesson when they know the names and characteristics of the main concepts.
8. *Modality Principle*	People learn better from graphics and narrations than from animation and on-screen text.
9. *Multimedia Principle*	People learn better from words and pictures than from words alone.
10. *Personalization Principle*	People learn better from multimedia lessons when words are in conversational style rather than formal style.
11. *Voice Principle*	People learn better when the narration in multimedia lessons is spoken in a friendly human voice rather than a machine voice.
12. *Image Principle*	People do not necessarily learn better from a multimedia lesson when the speaker's image is added to the screen.

serve or linearly consume the materials at their own pace. However, the learning environments in today's world require learners to be more active by strategically searching through hypermedia, constructing knowledge representations from multiple sources, performing tasks that create things, and interacting with technologies or other people (Chi, 2009; Wiley et al., 2009). From the standpoint of technology, it is worthwhile taking stock of the characteristics of learning environments that facilitate active, constructive, interactive learning environments. Table 5.1 shows some of these characteristics that were identified by the National Academy of Sciences, Engineering, and Medicine in the second volume of *How People Learn* (NASEM, 2018). It is important to consider these characteristics when selecting technologies to support the acquisition of knowledge representations in different subject matters, populations, and individual learners. All of these characteristics have been implemented in learning technologies and have shown some successes in improving knowledge representations and learning.

Unfortunately, there is an abundance of commercial technologies that are not well designed, are not based on scientific principles of learning, and have no evidence they improve learning. There are many bells and whistles of multimedia in so many products (a lot of razzle dazzle), but under the hood there is no substance in helping people learn and build useful knowledge representations. We live in a world replete with games and social media that contribute to shallow rather than deep knowledge representations.

It is important to consider the characteristics of the learning technologies that support deeper knowledge representations and learning (Millis et al., 2019; NASEM, 2018). Mayer (2009) has also identified 12 principles of multimedia learning that improve knowledge representation and acquisition (see Table 5.2). These principles are all based on psychological theories and confirmed by data collected in experiments.

The hope is that stakeholders and policy makers in education encourage learning environments which support knowledge representations needed in the 21st century. Citizens in the 21st century are faced with complex technologies, social systems and subject matters (National Research Council, 2012; Levy & Murnane, 2006). Mastery of facts and routine procedures are necessary, but not sufficient for participation in a world that demands deeper comprehension of technical material and more complex problem solving, reasoning, information handling and communication. Understanding the nature of knowledge representations will be extremely important in meeting this challenge.

Summary

1. People construct mental representations when they experience the social, physical, and digital world. Our perceptions are not exact copies of the world, but are simplified with errors and missing information. Learning and performance on tasks are influenced by how our knowledge is represented.

2. This chapter has reviewed the different types of representations that have been proposed by researchers in the cognitive and learning sciences who investigate adult learning of different subject matters. The types of representations include (1) ensembles of knowledge components, (2) knowledge structures, (3) associationistic neural networks, (4) embodied perceptions, actions, and emotions, (5) conversation, and (6) distributed cognition with diverse multimedia and technologies.

3. Knowledge of a specific subject matter is represented by a set of knowledge components which express ideas relevant to the topic. Knowledge structures consist of nodes, which represent concepts, states, events, goals or processes, and arcs that connect the nodes with

different types of relations (e.g., is-a, has-a, contains, causes). Four example knowledge structures were discussed: taxonomic, spatial, causal, and goal-action procedures.

4. Neural networks model associationistic representations with neuron nodes connected by associative weights. The strengths of the associations are determined by repetition, similarity, how often nodes co-occur in time, and positive versus negative outcomes.
5. Knowledge representations and acquisition are influenced by our human experience and how we interact with our environment. Embodied representations capture perception, action, and emotion. Conversational representations include the social discourse we observe and enact with families, tutors, mentors, and groups.
6. Digital technologies will continue to shape and constrain the mental representations and influence how people learn. These technologies are making information about topics more distributed across people, times, locations, and media sources.

Review Questions

1. Sketch a map of your town or city, including major landmarks and streets, based on your memory. Try to be as detailed as possible. After you finish, compare your sketch with an actual map. What did you get right, what did you miss, and what errors did you make in your mental representation?
2. Create a more complete knowledge structure of eating at a fast food restaurant that includes all types of structures in Figure 5.1: taxonomic, spatial, goal-action procedure, and causal.
3. According to the text, there are computerized tutoring systems that help people learn as well as human tutors. What sort of subject matters have representations that are very difficult for computer tutors to simulate, and why?
4. One very abstract concept is “peace.” To what extent can this concept be represented by embodied perception, action, and emotion? What features of peace would be impossible to capture with embodied cognition?
5. Consider a class you are currently taking. Which of the characteristics in Table 5.1 are part of the class activities? For any characteristics that are missing, how could they be incorporated by changing the class activities?

Hot Topic

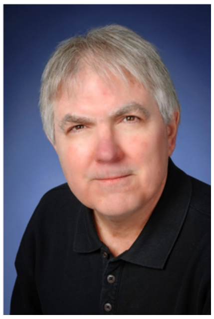

Art Graesser

Our research, along with colleagues in the interdisciplinary Institute for Intelligent Systems, investigates language, discourse and learning. Our primary focus is on the mastery of deep knowledge rather than shallow knowledge in adults. Examples of shallow knowledge are facts, definitions, and routine procedures, whereas deep knowledge involves causal reasoning, justification of claims with evidence, resolution of contradictions, precise quantification of ideas, and problem solving (Graesser, 2015). The workforce in the 21st century has an increased expectation to acquire deep knowledge to the extent that routine tasks are handled by robots and other digital technologies. Unfortunately, the process of deep learning is challenging because the material is difficult, useful strategies are sometimes novel, and some of the accompanying emotions are negative (such as confusion and frustration, D'Mello, Lehman, Pekrun, & Graesser, 2014). Moreover, our current educational systems are typically designed for acquiring shallow knowledge rather than deep knowledge.

Keith T. Shubeck

One approach to acquiring deep knowledge is to develop computerized intelligent tutoring systems that help adults acquire deep knowledge. These systems have pedagogical strategies that are tailored to the knowledge, skills and abilities of individual students. We have developed a system called AutoTutor (Graesser, 2016), where a student learns by having conversations with animated conversational agents (computer-generated avatars). AutoTutor presents difficult questions or problems, often with associated figures and diagrams; the student and AutoTutor have a multiturn conversation to co-construct an answer/solution. AutoTutor has been developed and tested on a number of difficult subject matters, such as computer literacy, physics, electronics, scientific reasoning, and comprehension strategies. These conversational ITS have shown significant learning gains on deep knowledge compared with pretests and control conditions such as reading text. Some versions of AutoTutor implement "trialogues" that involve a conversation between the student and two computer agents, a tutor and a peer (Graesser, Li, & Forsyth, 2014). The two agents can model good social interaction, productive reasoning, and at times argue with each other to show different perspectives and resolutions of conflicts (D'Mello et al., 2014).

We have investigated other approaches to improve deep learning through language and discourse (Graesser, 2015). These include investigating inference generation and mental models during the comprehension of stories, technical text, illustrated texts, hypertext, and hypermedia. We have developed computer systems (available on the internet for free) that scale texts on difficulty (Coh-Metrix, http://cohmetrix.com) and questions on comprehension problems (QUAID, http://quid.cohmetrix.com). We have investigated collaborative problem solving where groups of people in computer-mediated communication tackle problems that individuals cannot solve alone. A curriculum for 21st-century skills is destined to include discourse technologies that facilitate deeper knowledge acquisition.

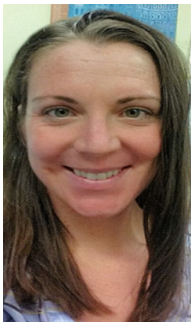

Anne M. Lippert

References

D'Mello, S., Lehman, B., Pekrun, R., & Graesser, A. C. (2014). Confusion can be beneficial for learning. *Learning and Instruction*, *29*, 153–170. doi:10.1016/j.learninstruc.2012.05.003

Graesser, A. C. (2015). Deeper learning with advances in discourse science and technology. *Policy Insights from Behavioral and Brain Sciences*, *2*, 42–50. doi:10.1177/2372732215600888

Graesser, A. C. (2016). Conversations with AutoTutor help students learn. *International Journal of Artificial Intelligence in Education*, *26*, 124–132. doi:10.1007/s40593-015-0086-4

Graesser, A. C., Li, H., & Forsyth, C. (2014). Learning by communicating in natural language with conversational agents. *Current Directions in Psychological Science*, *23*, 374–380. doi:10.1177/0963721414540680

References

Barsalou, L. W., & Wiemer-Hastings, K. (2005). Situating Abstract Concepts. In Pecher D. & Zwaan, R. (Eds.), *Grounding cognition: The role of perception and action in memory, language, and thinking* (pp. 129–163). Cambridge: Cambridge University Press. doi:10.1017/cbo9780511499968.007

Bower, G. H., Black, J. B., & Turner, T. J. (1979). Scripts in memory for text. *Cognitive Psychology*, *11*(2), 177–220. doi:10.1016/0010-0285(79)90009-4

Chi, M. T. H. (2009). Active-Constructive-Interactive: A conceptual framework for differentiating learning activities. *Topics in Cognitive Science*, *1*, 73–105. doi:10.1111/j.1756-8765.2008.01005.x

Clark, J. M. & Paivio, A. (1991). Dual coding theory and education. *Educational Psychology Review*, *3*(3), 149–170. doi:10.1007/BF01320076

Cohen, P. A., Kulik, J. A., & Kulik, C. C. (1982). Educational outcomes of tutoring: A meta-analysis of findings. *American Educational Research Journal*, *19*(2), 237–248. doi:10.3102/00028312019002237

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, *82*(6), 407–428. doi:10.1037/0033-295X.82.6.407

Craik, F. I. M., & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior*, *11*(6), 671–684. doi:10.1016/S0022-5371(72)80001-X

D'Mello, S., Lehman, B., Pekrun, R., & Graesser, A. C. (2014). Confusion can be beneficial for learning. *Learning and Instruction*, *29*, 153–170. doi:10.1016/j.learninstruc.2012.05.003

Dror, I. E. & Harnad, S. (Eds.). (2008). *Cognition distributed: How cognitive technology extends our minds*. Amsterdam: John Benjamins Publishing. doi:10.1075/bct.16

Ekman, P., & Rosenberg, E.L. (Eds.). (2005). *What the face reveals: Basic and applied studies of spontaneous expression using the facial action coding system (FACS)*. Change to Oxford Scholarship Online. doi:10.1093/acprof:oso/9780195179644.001.0001

Elliot, S. (2017). *Computers and the future of skill demand*. Retrieved from the OECD Publishing. doi:10.1787/9789264284395-en

Fiore, S. M., & Wiltshire, T. J. (2016). Technology as teammate: Examining the role of external cognition in support of team cognitive processes. *Frontiers in Psychology: Cognitive Science*, *7*(1531), 1–17 doi:10.3389/fpsyg.2016.01531

Foltz, P. W. (2016). Advances in automated scoring of writing for performance assessment. In Y. Rosen, S. Ferrara, & M. Mosharraf (Eds.), *Handbook of Research on Technology Tools for Real-World Skill Development* (pp. 659–678). Hershey, PA: IGI Global. doi:10.4018/978-1-4666-9441-5.ch025

Glenberg, A. M., Goldberg, A. B., & Zhu, X. (2011). Improving early reading comprehension using embodied CAI. *Instructional Science*, *39*(1), 27–39. doi:10.1007/s11251-009-9096-7

Goldman, S.R., Britt, M. A., Brown, W., Cribb, C., George, M. A., Greenleaf, C., Lee, C. D., & Shanahan, C. (2016). Disciplinary literacies and learning to read for understanding: A conceptual framework for disciplinary literacy. *Educational Psychologist*, *51*(2), 219–246. doi:10.1080/00461520.2016.1168741

Graesser, A. C. (2016). Conversations with AutoTutor help students learn. *International Journal of Artificial Intelligence in Education*, *26*,124–132. doi:10.1007/s40593-015-0086-4

Graesser, A. C., Hu, X., & Sottilare, R. (2018). Intelligent tutoring systems. In F. Fischer, C. E. Hmelo-Silver, S. R. Goldman, and P. Reimann (Eds.), *International handbook of the learning sciences* (pp. 246–255). New York: Routledge. doi:10.4324/9781315617572-24

Hattie, J. A. C., & Donoghue, G. M. (2016). Learning strategies: A synthesis and conceptual model. *Nature Partner Journal: Science of Learning*, *1*, 1–13. doi:10.1038/npjscilearn.2016.13

Hutchins, E. (1995), *Cognition in the wild*, MIT Press, Cambridge, MA.

Jurafsky, D., & Martin, J. (2008). *Speech and language processing* (2nd ed.). New Jersey: Prentice-Hall.

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

Koedinger, K. R., Corbett, A. C., & Perfetti, C. (2012). The Knowledge-Learning-Instruction (KLI) framework: Bridging the science-practice chasm to enhance robust student learning. *Cognitive Science*, *36*(5), 757–798. doi:10.1111/j.1551-6709.2012.01245.x

Landauer, T., McNamara, D. S., Dennis, S., & Kintsch, W. (Eds.). (2007). *Handbook of latent semantic analysis*. Mahwah, NJ: Erlbaum.

Levy, F. & Murnane, R. J. (2006). Why the changing American economy calls for twenty-first century learning: Answers to educators' questions. *New Directions for Youth Development*, *110*, 53–62. doi:10.1002/yd.167

Mayer, R. E. (2009). *Multimedia Learning* (2nd ed.). New York: Cambridge University Press.

McClelland, J. L., Rumelhart, D. E., & the PDP Research Group (1987). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 2: Psychological and Biological Models*, Cambridge, Massachusetts: MIT Press.

McNamara, D. S., Graesser, A. C., McCarthy, P. M., Cai, Z. (2014). *Automated evaluation of text and discourse with Coh-Metrix*. Cambridge, MA: Cambridge University Press. doi:10.1017/cbo9780511894664

Millis, K., Long, D. L., Magliano, J. P., & Wiemer, K. (Eds.). (2019). *Deep comprehension: Multidisciplinary approaches to understanding, enhancing, and measuring comprehension*. New York: Routledge. doi:10.4324/9781315109503

National Academy of Sciences, Engineering, and Medicine (2018). *How People Learn II: Learners, contexts, and cultures*. Washington, D.C.: National Academies Press. https://www.nap.edu/catalog/24783/how-people-learn-ii-learners-contexts-and-cultures

National Research Council (2012). *Education for Life and Work: Developing Transferable Knowledge and Skills in the 21st Century*. Washington, D.C.: The National Academies Press. doi:10.17226/13398

Pickering M.J., & Garrod S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, *27*(2), 169–225. doi:10.1017/S0140525X04000056

Rapp, D. N., & Braasch, J. L. G. (Eds). (2014). *Processing inaccurate information: Theoretical and applied perspectives from cognitive science and the educational sciences*. Cambridge, MA: MIT Press. doi:10.7551/mitpress/9737.001.0001

Rubin, D. (1995). *Memory in oral traditions: The cognitive psychology of epics, ballads, and counting-out rhymes*. Oxford, UK: Oxford University Press.

Stevens, A., & Coupe, P. (1978). Distortions in judged spatial relations. *Cognitive Psychology*, *10*(4), 422–437. doi:10.1016/0010-0285(78)90006-3

van den Broek, P. (2010). Using texts in science education: Cognitive processes and knowledge representation. *Science*, *328*, 453–456. doi:10.1126/science.1182594

VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems and other tutoring systems. *Educational Psychologist*, *46*(4), 197–221. doi:10.1080/00461520.2011.611369

von Davier, A. A., Zhu, M., & Kyllonen, P. C. (Eds.). (2017). *Innovative Assessment of Collaboration*. Cham, Switzerland: Springer. doi:10.1007/978-3-319-33261-1

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

Wiley, J., Goldman, S. R., Graesser, A. C., Sanchez, C. A., Ash, I. K., & Hemmerich, J. A. (2009). Source evaluation, comprehension, and learning in Internet science inquiry tasks. *American Ed-*

ucational Research Journal, *46*(4), 1060–1106. doi:10.3102/0002831209333183

Zacks, J. M., Speer, N. K., & Reynolds, J. R. (2009). Segmentation in reading and film comprehension. *Journal of Experimental Psychology: General*, *138*(2), 307–327. doi:10.1037/a0015305

Zwaan, R. A. (2016). Situation models, mental simulations, and abstract concepts in discourse comprehension. *Psychonomic Bulletin & Review*, *23*(4), 1028–1034. doi:10.3758/s13423-015-0864-x

Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin*, *123*(2), 162–185. doi:10.1037//0033-2909.123.2.162

Glossary

arcs In a knowledge structure, arcs are what connect two different nodes and represent how the nodes are related: is-a, has-a-part, property of, contains, cause, reason. 74

causal networks A knowledge structure consisting of event nodes that are connected by "enables" arcs. 76

embodied cognition The idea that knowledge and mental representations are influenced by experiences of the human body (e.g., emotion, perception, actions). 78

goal-action procedure A hierarchical knowledge structure where nodes represent goals or desired states which are connected by "in order to" arcs. 75

graph The set of nodes that are connected by arcs. 74

knowledge component Describes a mental structure used by learners to understand a topic. Any given topic may consist of many different knowledge components. 72

knowledge structure Relational structure between concepts in a particular topic. Describes how ideas are conceptually related in terms of their proximity with each other. 74

neural network Structure of nodes organized in multiple layers that are interconnected by arcs that either activate or inhibit nodes given arc direction and arc weight. 77

nodes In a knowledge structure, nodes are concepts, states, events, processes, goals or actions of basic ideas that can be expressed by words, phrases, or sentences. 74

script A structure that encompasses the goal-action procedures of all participants in an organized activity. 75

spatial structure A hierarchy of regions that are connected by "is-in" or "contains" arcs. 75

taxonomic structure A hierarchical knowledge structure in which concepts are connected by "is-a" arcs. 74