Comprehending and Learning from Visualizations:

A Developmental Perspective

David H. Uttal and Katherine O'Doherty

Northwestern University

Portions of this work were supported by grant R305H050059 from the Institute of Education Sciences and 0087516 from the National Science Foundation. We thank Linda Liu Hand, Miriam Reiner, and John Gilbert for their comments. Authors' Contact Information: David Uttal, Department of Psychology, 2029 Sheridan Road, Evanston, IL 60208-2710 (duttal@Northwestern.edu; http://psych.northwestern.edu/~duttal).

Comprehending and Learning from Visualizations:

A Developmental Perspective

The use of *visualizations* has become nearly ubiquitous in the practice, teaching, and learning of mathematics, science, and engineering. We define *visualizations* as any type of physical representation designed to make an abstract concept visible. These include but are not limited to concrete items such as photographs, 2-D graphs, diagrams, charts and 3-D models. It is almost impossible to imagine working in complex visual-based sciences such as chemistry or geoscience without the insights that visualizations can afford.

What makes visualizations useful for learning and thinking? Why do they help students learn? Part of the answer is obvious: visualizations help because they make complex information accessible and cognitively tractable. Visualizations highlight the portions of the information that the designer intends for the learner to see and hence support both learning among novices and new discoveries among experts. They allow us to perceive, and to think about, relations among items that would be difficult to comprehend otherwise.

Consider, for example, the visualization shown in Figure 1. This very simple map of the relative locations of a few United States cities makes accessible and tractable what is actually a very complex set of relations. Imagine the difficulty of describing this same set of relations in words (Taylor & Tversky, 1992; Uttal, Fisher, & Taylor, 2006). A great deal of work in a variety of disciplines, including psychology (e.g., Hegarty, Carpenter & Just, 1991; Larkin & Simon, 1987; Novick & Hurley, 2001), computer science (Allwein & Barwise, 1996; Ferguson & Forbus, 2000; Glasgow, Narayanan & Chandrasekaran, 1995) and geography (MacEachren, 1995) has focused on the value of visualizations for depicting scientific phenomena that may otherwise remain opaque or inaccessible.

Put simply, much of the power of visualizations stems from their ability to make us think in visual rather than in abstract, symbolic terms. We do not have to describe in words, for example, the complex relations that allow an enzyme or drug to bond at a specific location on a molecule. With a visualization, we can instead see and think about this complex relationship and location in terms of spatial relations. The visualizations become tools for thinking about the underlying structures and the relations among them.

In this chapter we argue that there is an important prerequisite for learning from visualizations. Before students can benefit fully from the visual-spatial properties of visualizations, they must understand that, and how, visualizations stand for or *represent* particular concepts or complex objects. The fact that visualizations have a visual-spatial nature does not guarantee that the student will comprehend the intended relation between the visualization and what it stands for (the referent).

We begin by developing the argument that visualizations must be conceived as representations. Borrowing from work in developmental psychology and the philosophy of symbols (e.g., Goodman, 1976; Hecht, Schwartz, & Atherton, 2003), we develop the case that grasping the representational relation between a visualization and what it stands for is an interesting and at time difficult challenge. Next, we discuss the development of young children's understanding of simple visualizations and suggest that this development can shed important light on older (i.e., adolescent and adult) students' understanding of more complex visualizations. We then consider the implications of the developmental work for understanding how older students understand and learn from scientific visualizations. Finally, we discuss research questions that follow from our discussions of visualizations and representations.

Understanding Visualizations as Representations.

Our work in this area arose from discussions with colleagues who teach geology and chemistry. They told us that their that students often struggle to master the basic correspondence between visualizations and what they represent. Students' errors sometimes revealed that they did not seem to grasp fully that the visualization was intended to represent something. For example, colleagues have said that their students saw visual representations of proteins only as "red dots or green circles" or complex geology maps as "a bunch of blobs on paper". These anecdotal reports are interesting because they reveal that the students focused only on the visualization itself, not on what it is intended to represent. The students saw only what was displayed on the paper or the computer screen and failed to grasp what the visualization was intended to represent.

We believe errors such as these are indicative of a general challenge in using visualizations or any representations. To the expert user or professor, the intended purpose of the visualization, and its relation to the referent, is obvious. For example, a chemistry professor may look at a visualization of a molecule and think *as if* they were looking at the molecule itself. But to novices, the relations that are so obvious to the expert may be totally opaque.

As an example of what we mean, consider how people understand and use what might be considered a relatively simple visualization, a road map. Many adults may assume that the relation between the map and what it represents is immediately clear or obvious. But on second thought, we realize that this sense of easy understanding belies the many years of development and learning that support our understanding. Upon reflection, we may realize that the map really doesn't look much like the world at all. For example, a road map is two-dimensional and is often drawn on white paper. The world isn't. In addition, the map uses various colors, to represent aspects of the world that are not immediately visible. Red may be used to represent a superhighway, and yellow may be used to indicate that a city has more than 100,000 residents. The road is not red in the world, and the city is not yellow (see Wood, 1991, for a discussion of the non-obvious properties of maps).

The point of this analysis is to show that maps (and all visual representations) are not copies of the world; they are instead representations of some aspect of the world. The relation between a visualization and its referent is seldom obvious to novice learners. The visual nature of visualizations does not obviate the basic prerequisite to understand what the visualization represents. Reading a map first requires that people understand that it represents a particular space in a particular way.

The same is true of more complex scientific visualizations. To the experienced user, and to advanced students, the relation between the visualization and what it is intended to represent is obvious. But again, the feeling of simplicity takes for granted the years of development and learning that supported the understanding of the visualization. For example, in representing the orbit of electrons, a choice is made, of necessity, to distort scale. The electrons would actually be *much* further away than is suggested by the visualization. The author of the visualization thus has decided to sacrifice accuracy of scale in order to represent the path of an orbit, or the number of electrons. In addition, the visualization often includes colors that, of course, do not correspond to actual colors in the represented molecule. For example, representations of molecules may use colors to represent positive or negative charges. This can be useful information, but it is useful only if the learner understands what the colors represent. He or she can not simply look at the visualization and comprehend this information, in part because the visualization is not simply a copy of the molecule.

This analysis raises the question of how we define representations. The answer to this question is rarely, if ever, given by the object itself. One cannot say a priori what is and is not a representation. People create representations, through their intention to have one thing stand for something else (Bloom & Markson, 1998; Deacon, 1997; DeLoache, 2000;

Tomasello, Striano & Rochat, 1999). Anything can be a representation of something else if the intention to use it as such becomes clear. For example, consider two people sitting at a dinner table, with one giving directions to another. It would not be surprising to see the interlocutors using silverware, glasses, or candlesticks to stand for various locations. All that has to happen is for one person to say, "This glass is the Sears Tower, and this fork is the John Hancock building." From that point on, the glass and fork become representations (in this case, of the locations of important landmarks in Chicago), because someone intended for them to be. The spatial relations between the glass and fork *now* become meaningful information, but only because the individuals have met the prerequisite of understanding that the items are intended to be a representation of something else.

We refer to the critical understanding of the relation between a representation and its referent as *representational insight.* It is the process of coming to understand that, and how, a representation stands for something else. Our central thesis is that representational insight is always required when one learns from a representation. This is true even for highly visual representations such as maps or models; to understand or learn from them, we have to know what they are intended to stand for. The fact that visualizations are visual in nature is not sufficient to guarantee representational insight.

As mentioned above, expert users are so accustomed to using visualizations that they may forget how much we had to learn before they could use them. However, when we look at children's struggles to grasp the intended meaning and use of seemingly simple visualizations, we are reminded of the challenges that we faced, and that our students may be facing today. Therefore, in the next section, we suggest that a developmental perspective can help us to understand the important requirement of obtaining representational insight. Research on the development of children's understanding of visualizations has very important implications for understanding people's initial grasp of the crucially important relation between a visualization and its intended referent.

Towards a Developmental Perspective.

In this section we examine the development of children's understanding of simple visualizations. Our focus is on the development of children's understanding of the "stands for" relation between visualizations and what they are intended to represent. We will consider how children come to understand the simplest visualizations, such as photographs, scale models, pictures, etc. In each case, we see that part of the challenge involves learning to understand that, and how, the visualization stands for something.

Grasping at Photographs of Bottles: Infants and Toddlers' Understanding of Photographs. The developmental story begins with very young children, infants in fact. Several interesting studies involving the development of young children's understanding of photographs shed light on the role of representational insight into understanding even the simplest visualization.

One could easily argue that a photograph is the simplest possible visualization in that it looks almost like a copy of the represented object, person, or space. But even in this very simple case we see that gaining representational insight into the relation between photographs and their referents is not a straightforward or simple development. Researchers have examined how and when young children come to understand or appreciate *both* the similarities and differences between photographs and their referents in the world.

In reviewing this work, we need to consider two different sets of questions that are reflected in two different bodies of literature. The first set of questions concerns young children's *perception* of the similarity and differences between photographs and their referents. Do young children recognize objects that are pictured in photographs? And if they do, can they tell the difference between what is shown in the photograph and its real-world referent?

The second set of question concerns children's *use* of photographs as visualizations. When and how do young children use photographs as tools for learning or solving problems?

Before addressing these questions, we need first to discuss briefly how developmental psychologists can ask (and answer) questions about perception and cognition in infants. One frequently used technique is called *habituation.* It relies on the simple fact that events become less interesting the more we see them. When we first see something new, we attend to it, perhaps looking at it for a long time. But if the event is repeated, we quickly get used to it, and it no longer grabs our attention. This basic fact can be used to assess what children know or don't know about their world (e.g., Hespos & Spelke, 2004). Researchers present one stimulus to an infant until she or he *habituates* to the presentation. For example, a researcher might show a photograph or a doll, or play a particular sound. The researcher measures some aspect of the infant's behavior that is thought to indicate interest. One common example is *looking time*, the amount of time an infant's eyes are focused on a particular object or event. When a new event or stimulus is presented, infants tend to look at it, often for several seconds. The researcher would then keep presenting the same stimulus. After a few presentations of the event, most infants become bored with the event, and consequently will look at the additional presentations for a shorter length of time. The researchers will notice that the looking time (a) has decreased substantially, and (b) that the looking time is now fairly stable. Put simply, the infant is now *habituated* to the presentation of the stimulus and looks at the event only for a brief moment.

The basic fact that infants become habituated can be used as a means of determining whether they can tell the difference between two objects, events, or other stimuli. For example, a researcher could habituate children to the presentation of a particular object such as a doll and then switch to a new doll or a realistic photograph of the doll. The critical question is whether the child *dishabituates* to the presentation of the new stimulus. That is, do the children now start to pay attention again, looking at the new stimulus for a relatively long time? If they do, then the researchers can conclude that the children *must* have perceived the difference between the prior and new stimulus.

This basic technique has been applied to the question of whether children understand the similarities and differences between photographs and their referents in the world. Researchers (e.g., DeLoache, Strauss, & Maynard, 1979) have, for example, exposed children to a doll until they are habituated, and then presented them with color or black-and-white photographs of the same doll. Most of the babies failed to dishabituate to either type of photograph, indicating that they recognized the similarity between the photographs and their real-world referents. It is important to note that the failure to dishabituate was not due to babies being unable to discriminate between the photographs and the doll. Additional studies showed that when babies viewed a photograph and its referent side by side simultaneously, they looked longer at the more realistic of the two stimuli—the doll. Taken together, these results indicate that young babies recognize that photographs correspond to live objects but that they are not identical to live objects.

Understanding Photographs as Representations.

One might think that once children can perceive both the similarities and differences between photographs and their referents, then they would be able also to *understand* or use this relation. But this is not necessarily the case. Further explorations of young children's understanding of and interaction with photographs highlight clearly the important distinction between perceiving an object that is intended to be a representation and understanding that intention. Even though young babies can perceive similarities and differences between photographs and their referents, this knowledge is not enough to *understand* and to use the representational relation between photographs and their referents. An ongoing line of research suggests that even when children can perceive the difference between a photograph and its referent, they still may not fully grasp what this relation means. In other words, they do not understand that the photograph is intended to be a representation.

In this work (DeLoache, Pierroutsakos, Uttal, Rosengren & Gottlieb, 1998), we presented photographs, in the form of simple picture books, to 9- and 18-month-olds. For example, children were given a very simple book with one realistic-looking photograph on each page. The photographs showed objects with which infants would typically interact, such as a bottle and a rattle. We placed the picture book in front of the children, turned to the first page, and observed what the children did with the photographed object. After 15 seconds, we turned the page to the next photograph.

We were struck by a very consistent result: Nine-month-olds often attempted to pick up or otherwise grasp the photographed objects. The children's hand motions were directly and specifically aimed at the objects in the photographs. For example, they put their thumb and index finger together, as if they were trying to pick up a small object. Some of the children were quite persistent, attempting to pick up many of the depicted objects. For these reasons we concluded that the infants were not simply trying to pick up the photograph; they were trying to pick up the specific object that was shown in the photograph.

The grasping behavior decreased significantly with age; 18-month-olds did not attempt to grasp the photographed objects very often. Instead, these children pointed to the photographed objects and made noises. Developmental psychologists call this behavior *proto-labeling*; it involves what is probably an attempt to label an object before the child knows the noun. Thus, within a span of about 9 months, children's thinking about photographs develops substantially. Initially, they seem to treat photographed objects as if they were the objects themselves, but with development, they treat the photographs as representations of something in the real world. They are now more focused on the photographs as representations.

Why can babies perceive the difference between photographs and their referents yet still treat photographs as if they were the objects themselves? We believe that the babies note a similarity between the object and the referent. They also see that it doesn't look quite like the typical object does. But not knowing precisely what to make of this difference, they do what they normally do with an object: attempt to pick it up; put it in their mouths, etc. Thus they are attempting to explore the photographic depiction of the object to determine whether it differs from the actual object, and if so, how. However, 18-month-olds have much greater experience in reading picture books with their parents or caretakers. When the picture book reader points to and labels a picture rather than treating the picture as they would the actual object, the children learn about the representational relation between the photograph and what it stands for. They learn that photographs are intended to be treated as representations, not as objects.

This finding is particularly important because it highlights, in a simple and direct way, the central distinction that motivates much of this chapter. The 9-months-olds *perceived* the spatial and perceptual similarities between the photograph of the objects and the real objects. However, they did not *conceive* of the photographs as representations. They had not yet gained representational insight into the relation between the photographed object and the object itself.

Learning from Visualizations.

The research discussed thus far has focused on the development of very young children's understanding of photographs as representations. We next address a different, although obviously related, question: When and how do children *use* visualizations as tools for learning or problem solving? Recognizing the similarity (and difference) between a visualization and its intended referent is a necessary condition for using the visualization as a tool for learning, but this alone is not sufficient. Children need to understand *that,* and *how*, a visualization stands for what it represents.

We review a series of studies on the development of children's *use* of simple visualizations as tools for acquiring information. These studies illustrate very well the challenges of interpreting a visualization as a representation. In addition, they help to provide a theoretical foundation for further explorations of learning from visualizations.

Children's Use of Scale Models. Perhaps the most thoroughly researched topic in this area is the development of young children's understanding of scale models (DeLoache, 1987, 1989, 2000; Uttal, Schreiber, & DeLoache, 1995). This work has shed substantial light on the importance of gaining representational insight. The studies all use a very simple task: Children are asked to use a simple scale model of a room to find a toy that is hidden in the room. The model and search task seem so simple to adults that it is astonishing to see young children having trouble establishing a correspondence between the model and the room. That young children often do experience difficulty again highlights the very important role of representational insight in using visualizations. Even though the model and room look very much alike, this visual-spatial correspondence does not guarantee that children will gain representational insight.

The scale-model task also provides further insight on an issue that is of great interest to science educators (and the focus of this chapter), that is: Learning from visualizations. The children need to use the model to learn where the toy is hidden in the room. Note that this way of learning differs greatly from young children's typical way of learning about the world. Usually, when children look for a missing object (e.g., a favorite toy), they do so on the basis of their prior experiences, looking at places where they remember last seeing the toy or where they typically put it (e.g., Wellman, 1985). But to succeed in our scale-model task, they must do something very different; they must rely on information gained purely from a visualization (the scale model). The development of this skill and its challenges for young children is very illustrative of some of the challenges that older students and adults may encounter when using a new visualization , particularly in the early stages of instruction.

In the original studies (e.g., DeLoache, 1987), the children were 2.5- and 3.0-year-olds. Both the room and the model were rather sparsely decorated, with only a few items of furniture in the room (and the corresponding miniature versions in the model). The task began with a detailed explanation that was designed to help the child grasp the correspondence between the model and the relevance of this correspondence for finding the hidden toy in the room. The experimenter pointed out correspondences between the model and the room. For example, the experimenter said that "Little Snoopy's room was just like Big Snoopy's room." The experimenter also pointed out correspondences between individual items of furniture in the model and in the room and told the child that Little Snoopy and Big Snoopy liked to hide in the same places in their respective rooms.

The child watched as the experimenter hid the toy behind, under, or in one of the pieces of furniture in the model. The experimenter then asked the child to go into the larger room and find the toy. He or she was reminded that the toy was hidden in the same location in the room as the miniature toy was located in the model. Of course, the experimenter was careful not to label the location (e.g., by saying "it's behind the couch") because then the child could solve the problem without needing to think about the relation between the model and the room. The child searched the room until he or she found the toy, but a search was scored as correct only if the child's first search was at the correct location. The procedure was repeated several times, and the experimenter kept track of the number of correct searches.

The results revealed a dramatic developmental change. The 2.5-year-olds performed poorly; their searchers in the room were essentially random, indicating that they did not use the information from the model as a guide to search in the room. In contrast, the 3.0-year-olds performed much better, averaging approximately 75% correct searches, far greater than would be expected by chance. These children were able to use what they saw in the model as a guide for searching in the room.

After the children found the toy in the room, the participants returned to the model for an important check on their memory. The experimenter asked the child to show where the toy was hidden in the model. This ensures that forgetting the location of the toy cannot be the cause of the younger children's problem in finding the toy in the room. If the children are able to point out the location of the toy in the model, then they clearly knew and remembered the model location. Most of the children had no trouble with this memory check; the 2.5-year-olds remembered where the toy was hidden in the model, regardless of whether they could find it in the room.

In summary, then, the data indicate that the young children fail not because they cannot remember where the toy is in the model, but because they do not see a connection between the model and the room that it represents. When they entered the room, they knew the location of the miniature toy in the model. But this information was of no value to them once they entered the room. Thus we see that the presence of a strong perceptual relation between a visualization and what it stands for does not guarantee that children (or adults) will be able to use the visualization. We should not be surprised when older students do not immediately grasp the relation between a more complex visualization and what the teacher or scientist intended for the visualization to represent.

The results discussed thus far establish that 3.0-year-olds are capable of using a visualization to solve a simple search task, but this is not the end of the story. Subsequent research established that 3.0-year-olds' understanding of this relation is quite fragile; ostensibly minor changes had dramatically negative effects on children's use of the model as a representation. For example, giving sparser instructions, in which the experimenter did not explicitly point out the correspondence between the model and the room, led to dramatically lower performance (DeLoache, deMendoza & Anderson, 1999). Likewise, inserting a delay between the time when the child saw the toy in the model and when he or she was asked to search for it in the room also led to substantial declines in performance (Uttal, Schreiber, & DeLoache, 1995). Importantly, the declines observed were more than children getting a bit worse; they were dramatic shifts in performance, from very good to nearly-random searching. Once the task was tweaked so that children had to think a bit more about the relation between the model and the room, the 3.0-year-olds performed like the 2.5-year-olds. What seems to have been lost is the understanding either that the model stood for the room or that this relation was relevant for finding the toy. These findings again highlight the importance of gaining representational insight and demonstrate its fragility; even if children initially do grasp the representational relation between the model and the room, they can easily lose sight of this relation. Similarly, it seems possibly that a novice student might at times lose sight of what the visualization he or she is using is intended to represent and become overwhelmed by the many new colors and shapes.

The Dual Representation Hypotheses. Based on this pattern of results, DeLoache and colleagues formulated the *dual representation hypothesis* to account for children's success and failure in the model-search task and in understanding representations more generally. Central to this account is the notion that all representations have a dual nature; they are intended to stand for something else, but at the same time, they are also objects in their own right. For example, the scale model is intended to be a representation of the room, but it is also an interesting object, regardless of its connection to the represented room. Each visualization (i.e. pictorial chart, graph, 3-D model, etc.) has this dual nature; it is an object in its own right with its own visual properties but it is *also* an intended representation of something else (i.e. a cell, photo-synthesis, the human brain, etc.). As adults, we are so accustomed to thinking about what common visualizations represent that we may forget that they are also objects in their own right. But sometimes we are reminded. For example, when someone uses an unusual or garish font in a PowerPoint presentation, we may then focus on the letters themselves, rather than on the words that the letters represent. Likewise one would not be surprised if an adult initially interpreted the Tokyo subway map shown in Figure 2 as "strands of yarn or spaghetti". The many different colors, combined with a lack of a clear referent, may lead people at least initially to focus on the object itself rather than what it is intended to represent. Similarly, when a teacher first introduces a scientific visualization to his or her class, the students may see only its object properties (e.g., different colors, ribbons, lines, etc.) because they have not learned yet what the teacher intends for the visualization to represent.

One way to think of the dual representation is as a balance scale. On one side of the scale are factors that lead children (or adults) to interpret visualizations as objects in their own right. On the other side are factors that lead people to interpret the visualization as a representation. To use a model (or map, or photograph, or any visualization) as a representation of something else, we must focus on what it is intended to represent not on its properties an object in its own right.

In support of this theory, DeLoache and colleagues have conducted a fascinating series of studies on factors that influence children's use of scale models as representations. For example, in one series of studies (DeLoache, 2000), children were allowed to play with the model before they were later asked to use it as a representation of the room. Many people might assume that allowing children to play with the model beforehand would facilitate their performance, perhaps by helping them become familiar with it. However, note that the dual representation hypothesis predicts the opposite; playing with something that is intended to be a representation may focus children's attention on the object itself, rather than on the intended referent. The results confirmed this counterintuitive prediction: children who played with the model first actually had more trouble using it as a representation of the room.

Additional research also confirmed the opposite prediction: That restricting children's access to the model would make it *easier* for them to use it as a representation of something else. In this research (DeLoache, 2000), the model was placed behind a pane of glass so that children could not touch it or otherwise interact with it. In this case, 2.5-year-olds, who normally fail the task completely performed much better. In terms of the dual representation hypothesis, restricting access helped to *decrease* the children's attention to the properties of the model as an object in its own right and hence allowed the children to focus more on what the model represented.

Perhaps the most convincing work to date in support of the dual-representation hypothesis comes from a series of studies in which the dual nature of the model is made irrelevant. In this line of work (DeLoache, Miller, & Rosengren, 1997), the children were convinced that the model was a shrunken version of the room. A magic trick was used to give the appearance that a room (this time composed of fabric supported by PVC tubing) could be made to shrink or expand. The experiment began with an introduction to a "shrinking machine", which was actually a sham device with many dials that emitted strange sounds. The child was told that the machine was capable of shrinking any object. As a demonstration, the experimenter "shrunk" a toy troll; he or she placed a troll doll on the top of the machine, activated the "machine", and left the room. The child could hear the noises of the machine "working" as he or she stood outside. Unknown to the child, a confederate quietly substituted a miniature version of the troll doll, making it appear as if the doll had shrunk dramatically. When the noises stopped, the experimenter and child then entered the room, and the experimenter pointed out that the troll had shrunk.

The troll-shrinking demonstration was used to motivate the possibility of the room shrinking. After the child was convinced of the "functionality" of the shrinking machine, he or she was then introduced to the standard search task. The child watched as the experimenter hid the toy in the large-sized portable room. The experimenter and child then reactivated the shrinking machine and both left the room. Several confederates worked quietly to substitute a miniature version (i.e. miniature versions of the furniture were substituted for the full-size furniture) of the room. This smaller version of the room was stored behind curtains so that the child did not know that there were two versions of the room. When the child and experimenter returned to the space, it appeared as if the larger room had been shrunk. The critical question was whether the child could find the toy after the shrinking transformation.

The answer to this question was a decisive yes. 2.5-year-olds, who normally fail, performed much better. In interpreting this result, it is important to note that all elements of this task, save one, remained the same. For example, the child still had to use the location of the toy in one space (the larger room) to find the toy in the other space (the miniature room). The only difference is that the shrinking room procedure eliminated the need for the child to think about the model as a representation of the room. In the mind of the child, the small and large rooms were one and the same.

The results of this experiment provide strong evidence in support of the dual representation hypothesis. The children succeeded because the problem of dual representation was eliminated by the "shrinking" procedure. Even though the basic requirements remained the same, the way the child needed to think about those requirements changed. The child no longer had to think about a representational relation between the model and the room. In their minds, the shrunken room and the full-size room were one.

Removing the need to think about the two different components allowed 2.5-year-olds to succeed when they normally would fail.

Learning from Video. A related series of studies involved understanding representations and learning from television and video. What children see on television does not (usually) correspond to something specific or even necessarily real in the world. Many, if not most, children's television shows involve fantasy characters, cartoon, or other fantastical notions. Thus it would not be surprising to find that young children do not associate what they see on a television screen with a particular or specific scene in reality.

This point is demonstrated in a series of studies in which children were asked to use information provided in a brief video clip to find a toy that was hidden in a room (Troseth, 2003; Troseth & DeLoache, 1998). The video showed a room with several hiding locations and was similar in design and layout to the space that was used in previous studies of children's understanding of scale models. In this case, however, the children watched a video that showed the experimenter hiding a toy in the room. After watching this vignette, the child was asked to enter the room and find the toy.

This task proved particularly challenging for 2.0-year-old children. They did not think of the video vignette as a source of information for finding the hidden toy. Further research established again the specific importance of understanding a representational relation between the video scenes and the hiding event in the room. As in the "shrinking room" experiments described above, removing the need to think of the video as a *representation* of the room allowed children to succeed when they otherwise would fail. The researchers (Troseth & DeLoache, 1998) placed a video screen in a custom-fit window pane between the two rooms which led children to believe that they were actually observing the specific hiding event in the room, when in reality they were looking at a video of the hiding event. As in the shrinking room experiment, children who had trouble using the relation between the video and the room now performed well. Once again, removing the need to think about a representational relation led children to succeed when they would otherwise fail, providing clear evidence that the challenge for the young children was thinking about the video as a representation of the room.

Understanding Maps. The development of children's appreciation of geographic maps both illustrates and extends the points that we mentioned above. A variety of methods have been used to investigate the development of children's understanding of maps. Some studies have used methods similar to those that were used in the studies described above. For example, young children have been asked to use a simple map to find a hidden toy (Blades & Spencer, 1987), to navigate a route (e.g., Uttal & Wellman, 1989), or to plan a trip. However, what is perhaps most interesting here is studies of children's understanding of "real", geographic-scale maps, like those that typically adorn the walls of classrooms. When kindergartners or young elementary school children are asked to interpret these maps, they often demonstrate some of the same errors that younger children do when interpreting simple visualizations. For example, when children were asked to interpret a map that showed roads represented in various colors (e.g., red to show major freeways, gray to show two-lane highways, etc.), some said that a red line could not represent a road because the line was too narrow for a car to fit on it. Other children said the line could not represent a road because there are no red roads in the real world (Liben, 1999, 2001; Liben, Kastens, & Stevenson, 2002). Children also made errors of scale, such as correctly identifying a body of water (Lake Michigan, on an aerial photograph of Chicago) as water, but then claiming to see fish in the lake.

These examples are important because they demonstrate a more sophisticated but at times nevertheless incorrect understanding of the relation between visualizations and their referents. The children clearly know that there is some sort of relation between the map and the space, and they readily identify features that look like their referents in the real world. Yet, young children sometimes make mistakes that illustrate how fragile this understanding can be. The children seem to believe that there must be a literal correspondence between the map and the space that it represents. These errors again illustrate that visual correspondence is not enough to promote representational insight. The children do grasp the notion of a visual correspondence, but their strong belief in the *necessity* of visual correspondence actually gets in the way of their comprehension of the map as a representation. These young children may hold a "copy theory" of the relation between maps and the world, believing incorrectly that the map should be a literal copy of the world.

Summary. The studies that we have reviewed thus far have covered a range of visualizations, including photographs, models, video, and maps. Of central importance in all of these studies has been the need to understand, and to exploit, a representational relation between the visualization and what it represents. Initially, young children need to perceive a similarity between a visualization and what it represents, but this alone is not enough if we want them to use a visualization for learning. In fact, in the case of photographs, a high degree of visual similarity can lead babies to treat a photograph as if it were an actual object. Likewise, a belief in the necessity of visual similarity seems to interfere with, rather than to facilitate, elementary school children's understanding of maps. All of the examples that we have reviewed clearly indicate the centrality of understanding the intention to use a visualization to represent something. In the next section we consider the implications of these findings for older students' use of more complex visual representations, such as those used in a typical organic chemistry class.

Implications for the Use of Visualizations in Education.

To what extent are the results and theoretical perspective that we have presented thus far relevant to understanding the use of visualizations in high school or college classrooms?

One could argue that the results we have described here are directly and immediately relevant to high school or university students' learning from visualizations in fields such as chemistry, biology, geoscience, or engineering. On this view, novice learners face the very same challenges that young children face when they first see a new visualization. They may simply fail to recognize that the new visualization represents something else, and thus they may sometimes say they see "blobs" or "dots" when they are asked to explain the visualization. Put simply, the novice learner may fail to obtain representational insight when interpreting a new visualization.

Clearly, students do need to appreciate that visualizations are representations. Instructors are often so accustomed to thinking about the visualization as a representation that they may easily forget that students may not understand the correspondence between the visualization and its referent. What may seem obvious to the instructor may be totally opaque to the student. In such extreme cases, it is possible to imagine that students might literally behave like children; they may fail to grasp the intended representational relation between the representation and its referent.

We think, however, that a second, more nuanced interpretation of the implications of the developmental work is more informative. Adolescent and adult learners differ from young children in a very important way, one that influences greatly whether, and how, they learn from visualizations. At least by middle school, most students have developed what DeLoache (2000, 2004) and others have termed *symbolic sensitivity.* They can bring knowledge and experience gained from many prior learning experiences to new learning experiences. Over time, working with a variety of visualizations and graphics helps students to *expect* that a new visualization stands for something, although they still may have great difficulty bringing this to bear in a new learning situation. Put another way, adolescent and adult learners have acquired some degree of what diSessa and colleagues (e.g., diSessa &

Sherin, 2000) have termed *meta-graphical* awareness, which is a general understanding of the purposes of visual representations and of the value of different kinds of visualizations for solving different types of problems. There is at least the possibility that adults have acquired a more general understanding of the purposes of visualizations and of the match between visualizations and problems. The younger children in the experiments described earlier were, in essence, "blank slates" when it came to learning from the visualizations. In the case of adolescents and adults, this is unlikely to be true.

In addition, there is another important difference between the examples we have discussed from research with children and the challenges that adolescent or adult learners face in the classroom. In most science or mathematics classrooms, visualizations or visual-based learning is important but not sufficient for mastery of the topic. Students must also master more abstract, written symbols. For example, students may need to learn to understand and use equations or written representations of chemical reactions. Visualizations are designed in part to facilitate conceptual understanding that can provide a basis for understanding and learning these more abstract representations. Researchers and educators must therefore consider the relation between the use of visualizations and students understanding of written symbols.

Despite these differences between child and adult learners, the issue of dual representation is still relevant for understanding how adults learn from visualizations. For example, visualization designs that lead students to focus more on the properties of the visualization itself, rather than on what it is intended to represent or teach, can be counterproductive even in adult learners. In the mind of the expert, the intended representation is clear, but in the mind of the novice, the properties of the visualization as an object may be much more salient than the teacher realizes. Students may initially understand that the map or other visualization is a representation, but they may get consumed with the local, physical properties of the object and fail to focus on what the visualization is intended to represent. Therefore, garish or highly attractive visualizations may actually make it harder for students to grasp what the visualization is intended to teach specifically because they will focus more on the properties of the visualization as an object in its own right.

A particularly poignant example of the costs and benefits of using interesting or concrete visualizations comes from the work of Goldstone and Sakamoto (2003). Their specific focus was on transfer of knowledge. It is extremely important that students be able to transfer principles that they acquire from using a visualization to other learning situations that depend on the same general concept. To do so, they must be able to look beyond the properties of a particular visualization to learn about, retain, and ultimately transfer their knowledge to a new domain. Unfortunately, this has often been difficult to accomplish. "Transcending superficial appearances to extract deep principles is as critical to science as it is difficult to achieve." (Goldstone & Sakamoto, 2003, p. 415).

Goldstone and colleagues have used visualizations to teach college students a variety of scientific principles. One example is the notion of competitive specialization, which is the idea that adaptation in a competitive environment can be facilitated by adopting specialized strategies. This principle is relevant to many different scientific domains, including evolutionary biology, perceptual pattern learning, and the economics of business growth. It is therefore an ideal way in which to look at transfer of knowledge and the relevance of different kinds of visualizations for promoting (or inhibiting) transfer.

Goldstone and colleagues systematically manipulated the visualizations they used to investigate the influences of attractive, interesting, or concrete representations on students learning and transfer of knowledge. For example, in one experiment (Goldstone & Sakamoto, 2003), students learned the concept of competitive specialization either from a concrete visualization in which they had to help "ants" select "food" (both depicted pictorially in the visualization) or from a more idealized visualization in which the food and ants were simply represented by small dots and larger blobs. The results indicated that the concrete visualization, which showed pictures of ants and food, helped students learn initially. However, some students who learned from the concrete visualization found it *more* difficult to transfer their knowledge to other domains than did students who learned from the more idealized representation. Of particular interest was that students who initially had trouble learning the concept were much more likely to transfer what they did learn when they used the idealized visualization. These results have been replicated and extended in a series of experiments that indicate that in general, concrete representations may facilitate initial learning but make transfer more difficult than more idealized representations that do not focus students' attention on the particular characteristics of the visualization (Goldstone and Son, 2005).

The results of this research clearly reveal that the dual representation hypothesis is relevant to learning from complex visual representations of the kind that are used in modern science instruction. Even though adults are quite familiar with many visualizations, educational designers still need to balance the attractiveness or concreteness of the visualization with the desire to help students learn from it. The putative assumption that highly interesting or attractive visualizations enhance learning may not be true, either for children or adults. More generally, what we need is a more detailed consideration of the benefits and costs of learning from concrete and idealized visualizations. The research reviewed here clearly indicates that even for very simple visualizations, there is always a tradeoff between making something interesting in its own right and helping people learn from it.

Towards a Research Agenda.

The research and theoretical perspectives that we have reviewed here raise several important questions that could be explored in future research. In this final section we consider the questions that future research could address and suggest methods by which they could be approached.

Applying the Developmental Model. The first, and most general, suggestion is that the developmental model that we have outlined here can be used as a framework even for studying the development of visualization use among college students. Of course, as discussed above, there are important differences between children and adults, but when it comes to understanding new or complex visualizations, there may also be important similarities. Most of what we do know about differences between novice and expert users of visualizations is based on anecdotes or is part of the cultural lore of teaching.

Accordingly, the first research recommendation we have is that we should study adolescents' and adults' learning from new visualizations from a developmental approach. We should investigate what knowledge they bring to bear when first interpreting a new representation and how this knowledge changes as they learn more both about the representation itself and what it stands for. What do students first think about when they encounter a new visualization? Does their prior experience in interpreting other visualizations influence their interpretation of the novel one, and if so, how? Researchers in education and the learning sciences have already begun to answer questions like this (e.g., Lee & Sherin, 2006) through intensive studies of the process of learning through visualizations. But what is missing, for the most part, is how such understanding develops over time, across courses, and whether knowledge gained in one class transfers to other classes. This kind of research is neither easy nor inexpensive to do, but it is very much needed if we are to understand how, when, and why students learn from visualizations.

A second important question concerns the process of abstraction and role of attractive or highly concrete visualizations. Goldstone and colleagues' work is a very important beginning, but we still know very little regarding what type of information should be depicted at various stages in the learning process. For example, we do not know how or when students transform their knowledge or make connections among different forms of representations, nor do we know how or when instructors should switch representations. One can imagine a tradeoff between the misunderstanding that might arise from the frequent use of different visualizations and the difficulties in transfer that might arise from repeated use of the same visualization or other model or representation. Given that students will be exposed to an increasingly large number of visualizations, we need more research on how they relate one visualization to another (See Ainsworth, 2006 for a review of this issue).

A third set of questions concerns the possibility of helping students to develop a more general understanding of visual representations. It may be possible to develop meta-graphical awareness that is at least in part domain general. That is, students could learn to think *about* visualizations in a manner that is not completely tied to a content domain. The work of Novick and Hurley (2001) suggests that students may already understand the appropriate uses of different types of relatively simple visual representations, such as matrices, networks, and hierarchical diagrams. It may be possible, therefore, to create classes or instructional units that emphasize visual-spatial thinking and visualizations. Such courses already exist in some engineering schools (See Sorby, 2001 for discussion of an example), and they seem to be highly successful. Facilitating the development of meta-graphical awareness might help to minimize the chances that students will focus on the concrete properties and hence fail to grasp what the visualization is intended to represent. One important question that will arise in designing such a course or intervention will be where to begin: One could imagine teaching relevant concepts at a variety of ages or grades. The ultimate goal will be to make visual-spatial fluency as important as linguistic fluency.

Conclusion.

Visualizations have transformed both the practice and teaching of science, engineering and mathematics. High-speed, inexpensive computers have made it possible for even beginning students to have access to rich representations of highly complex phenomena. Most textbooks in the natural sciences now come with CDs, DVDs, or links to web sites that allow students to learn from these visualizations without even leaving their homes. Research on the process of learning from visualizations has lagged behind the creation of visualizations, but there are some signs that it is beginning to catch up. One example is this volume, which focuses specifically on the process of learning from visualizations.

Perhaps the most important conclusion that can be drawn from this paper is that visualizations are representations. Indeed, much of the trouble that children, and perhaps adults, have in learning to use visualizations involves recognizing that the person who created the visualization intended it to represent something. Once this basic insight is gained, the spatial properties of the visualization become available and useful for learning. But these properties remain useless until this basic prerequisite is met. The effective design of visualizations therefore must focus on how to facilitate, rather than to obviate, the need to think about them as representations.

References

Ainsworth, S.E. (2006). DeFT: A conceptual framework for learning with multiple representations. *Learning and Instruction, 16*, 183-198.

Allwein, G., & Barwise, J. (1996). *Logical reasoning with diagrams.* Oxford: Oxford University Press.

Blades, M., & Spencer, C. (1987). The use of maps by 4-6-year-old children in a large scale maze. *British Journal of Developmental Psychology, 5,* 19-24.

Bloom, P., & Markson, L. (1998). Intention and Analogy in Children's Naming of Pictorial Representations. *Psychological Science, 9*, 200-204.

Deacon, T. W. (1997). *The symbolic species: The co-evolution of language and the brain.* New York: Norton.

DeLoache, J. S. (1987). Rapid change in the symbolic functioning of very young children. *Science, 238*, 1556-1557.

DeLoache, J. S. (1989). Young children's understanding of the correspondence between a scale model and a larger space. *Cognitive Development, 4,* 121-139.

DeLoache, J. S. (2000). Dual representation and young children's use of scale models. *Child Development*, *71*, 329-338.

DeLoache, J. S. (2004). Becoming symbol-minded. *TRENDS in Cognitive Science, 8,* 66-70.

DeLoache, J. S., deMendoza, O. P., & Anderson, K. (1999). Multiple factors in early symbol use: Instructions, similarity, and age in understanding a symbol-referent relation. *Cognitive Development, 14*, 299-312.

DeLoache, J. S., Miller, K. F., & Rosengren, K. S. (1997). The credible shrinking room: Very young children's performance with symbolic and nonsymbolic relations. *Psychological Science, 8,* 308-313.

DeLoache, J. S., Pierroutsakos, S.L., Uttal, D. H., Rosengren, K. S., & Gottlieb, A. (1998). Grasping the nature of pictures. *Psychological Science, 9,* 205-210.

DeLoache, J.S., Strauss, M.S., & Maynard, J. (1979). Picture perception in infancy. *Infant Behavior and Development, 2*, 77–89.

diSessa, A., & Sherin, B. (2000). Meta-representation: An introduction. *Journal of Mathematical Behavior, 19*, 385-398.

Ferguson, R. W., & Forbus, K. D. (2000). *GeoRep: A flexible tool for spatial representation of line drawings.* Paper presented at the National Conference on Artificial Intelligence, Austin, TX.

Glasgow, J., Narayanan, N. H., & Chandrasekaran, B. (Eds.). (1995). *Diagrammatic reasoning: Cognitive and computational perspectives.* Menlo Park, CA: AAAI Press/The MIT Press.

Goldstone, R., & Sakamoto, Y. (2003). The transfer of abstract principles governing complex adaptive systems. *Cognitive Psychology, 46,* 414-466.

Goldstone, R., & Son, J. Y. (2005). The transfer of scientific principles using concrete and idealized simulations. *Journal of the Learning Sciences, 14,* 69-110.

Goodman, N. (1976). *Languages of Art.* Indianapolis, IN: Hackett Publishing.

Hectht, H., Schwartz, A., & Atherton, A. (2003). Looking *into Pictures: An Interdisciplinary Approach to Pictorial Space.* Cambridge, MA: MIT Press.

Hegarty, M., Carpenter, P. A., & Just, M. A. (1991). Diagrams in the comprehension of scientific texts. In R. Barr, M. L. Kamil, P. Mosenthal, & P. D. Pearson (Eds.), *Handbook of reading research* (Vol. 2, pp. 641–668). New York: Longman.

Hespos, S., & Spelke, E. (2004). Conceptual precursors to language. *Nature, 430,* 453-456.

Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth a 1000 words. *Cognitive Science, 11,* 65-99.

Lee, V. R., & Sherin, B. (2006). Beyond transparency: How students make representations meaningful. *Proceedings of the Seventh International Conference of the Learning Sciences*: 397-403.

Liben, L. S. (1999). Developing an understanding of external spatial relations. In I. E. Sigel (Ed.), *Development of mental representation: Theories and applications* (pp. 297-321). Mahwah, NJ: Erlbaum

Liben, L. S. (2001). Thinking through maps. In M. Gattis (Ed.), Spatial schemas and abstract thought, (pp. 44-77). Cambridge, MA: MIT Press.

Liben, L.S., Kastens, K. A., & Stevenson, L. M. (2002). Real-world knowledge through real-world maps: A developmental guide for navigating the educational terrain. *Developmental Review, 22,* 267-322.

MacEachren, A. M. (1995). *How maps work: Representation, visualization, and design.* New York: Guildford Press.

Novick, L R., & Hurley, S. M. (2001). To matrix, network, hierarchy: That is the question. *Cognitive Psychology, 42*, 158-216.

Sorby, S. A. (2001). A Course in Spatial Visualization and its Impact on the Retention of Women Engineering Students. *Journal of Women and Minorities in Science and Engineering, 7*, 153-172

Taylor, H.A., & Tversky, B. (1992). Spatial metal models derived from survey and route descriptions. *Journal of Memory and Language, 31,* 261-292.

Tomasello, M., Striano, T., & Rochat, P. (1999). Do young children use objects as symbols? *British Journal of Developmental Psychology, 17,* 563-584.

Troseth, G. L. (2003). TV guide: Two-year-old children learn to use video as a source of information. *Developmental Psychology, 39,* 140-150

Troseth, G. L, & DeLoache, J. S. (1998). The medium can obscure the message: Young children's understanding of video. *Child Development, 69*, 950-965.

Uttal, D. H., Fisher, J.A., & Taylor H.A. (2006). Words and maps: Developmental changes in mental models of spatial information acquired from descriptions and depictions. *Developmental Science, 9,* 221-235.

Uttal, D. H., Schreiber, J. C., & DeLoache, J. S. (1995). Waiting to use a symbol: The effects of delay on children's use of models. *Child Development, 66,* 1875-1889.

Uttal, D. H., & Wellman, H. M. (1989). Young children's mental representation of spatial information acquired from maps. *Developmental Psychology, 25,* 128-138.

Wellman, H. M. (1985). *Children's searching: The development of search skill and spatial representation*. Hillsdale, NJ: Erlbaum.

Wood, D. (1992). *The Power of Maps.* New York: Guilford Press

Boston

Seattle

New York

Chicago

Denver

Los Angeles

Houston

Miami

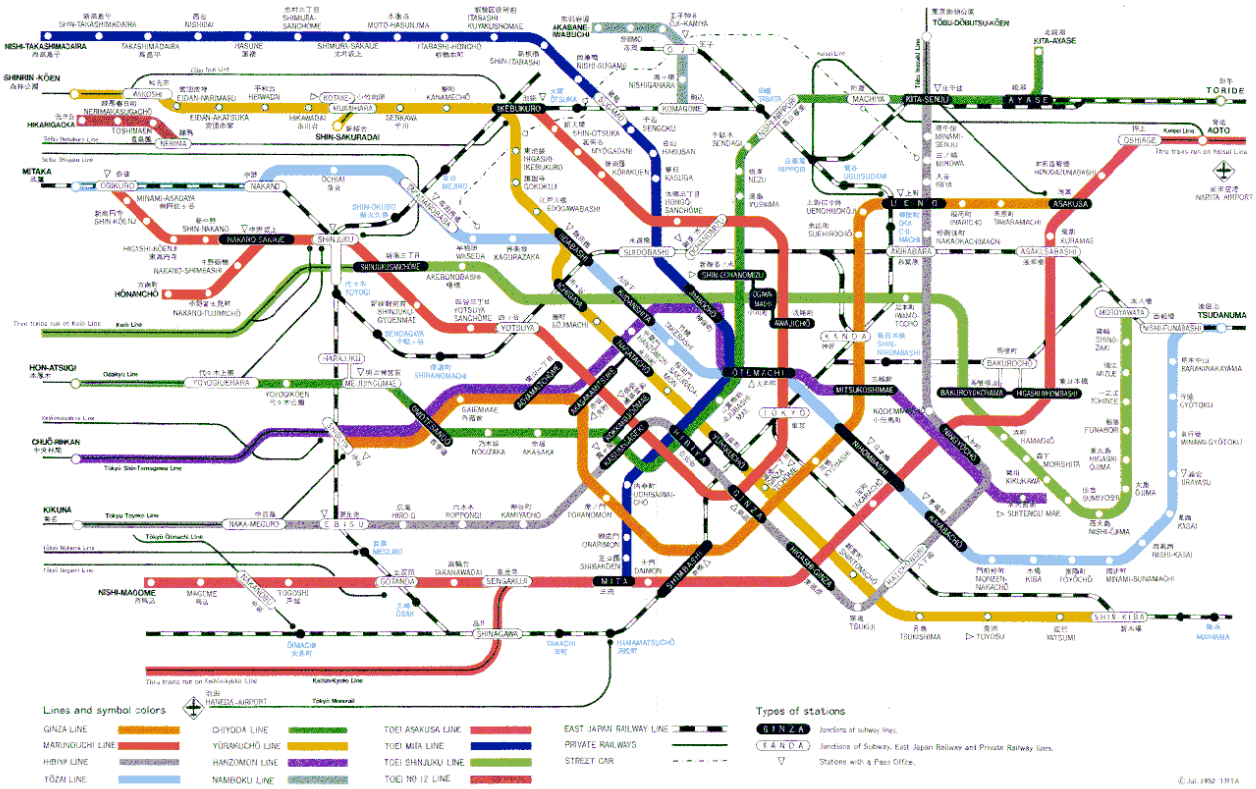

Source: http://web.yl.is.s.u-tokyo.ac.jp/jp/map.gif