RC22742 (W0303-075) March 17, 2003
Computer Science

IBM Research Report

Flexibility of Expressiveness in Concept Maps: Bridging the Gap from State of the Art to State of the Practice

Sherman R. Alpert
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

Research Division
Almaden - Austin - Beijing - Delhi - Haifa - India - T. J. Watson - Tokyo - Zurich

LIMITED DISTRIBUTION NOTICE: This report has been submitted for publication outside of IBM and will probably be copyrighted if accepted for publication. It has been issued as a Research Report for early dissemination of its contents. In view of the transfer of copyright to the outside publisher, its distribution outside of IBM prior to publication should be limited to peer communications and specific requests. After outside publication, requests should be filled only by reprints or legally obtained copies of the article (e.g. , payment of royalties). Copies may be requested from IBM T. J. Watson Research Center ,
P. O. Box 218, Yorktown Heights, NY 10598 USA (email: reports@us.ibm.com). Some reports are available on the internet at http://domino.watson.ibm.com/library/CyberDig.nsf/home .

Flexibility of Expressiveness in Concept Maps: Bridging the Gap from State of the Art to State of the Practice

Sherman R. Alpert
IBM T.J. Watson Research Center
PO Box 704
Yorktown Heights, NY 10598 USA
salpert @ us.ibm.com

Abstract: Over a decade ago, Heeren and Kommers (1991) cited "flexibility of expressiveness" as a major factor in designing concept map tools for learning. In the interim, personal computers have developed substantially in their capabilities and many concept map tool enhancements have appeared. However, the concept map tools in actual use by students have not fully capitalized on such innovations to enhance their flexibility and breadth of expressiveness. This paper reviews the state of the art in concept map tools with the objective of improving those tools actually used in schools.

More than a decade ago, Heeren and Kommers (1991) cited "flexibility of expressiveness" as a major factor in designing concept map tools for learning. In the interim, personal computers have developed substantially in their capabilities. New concept map tools have appeared as well, with additional features not available when Heeren and Kommers performed their study, that offer greater representational options. Many of these capabilities can enhance concept maps' flexibility of expressiveness as well as support concept mapping as a mirror of cognitive knowledge. For example, concept maps applications can be greatly enhanced as learning tools by incorporating such features as rich media and full Web integration. And if concept maps are being used as knowledge representation tools, then concept map tools should incorporate features that provide representations analogous to cognitive representations.

This chapter takes a look at how researchers and developers have attempted to enhance the flexibility, breadth, representational and expressive power of concept maps as a knowledge representation mechanism. The primary tools considered are the Inspiration and Webster concept map tools, although several others will be part of the discussion. Inspiration is an enormously commercially successful concept mapping application and because it is broadly used in schools will represent the state of the practice. Indeed, Inspiration represents the best of the state of the practice with regard to what is actually going on in schools: while Inspiration is very widely used, students in schools use generic graphical drawing editors to construct concept maps. Webster is a concept map application developed by the author (Alpert, in press; Alpert & Grueneberg, 2000, 2001). Webster and other tools will represent the state of the art in concept map research and design.

Interestingly, when I started writing this chapter I began to illustrate the many ways in which the most popular commercial product—and therefore concept mapping practice—has not kept pace with technological enhancements embodied in a number of research concept map tools. While this is still true in many respects, the just-released (circa year-end 2002) version of Inspiration demonstrates convergence in a number of dimensions of the capabilities and philosophies of research and commercial tools, and thereby potentially effecting something approaching convergence of the state of the art with the next generation of the state of the practice. There still exist, however, a number of ways in which the two—art and practice—remain divergent.

Flexibility of Expressiveness and Cognitive Fidelity

Heeren and Kommers (1991) looked at concept map tools that offer differing amounts or degrees of flexibility of expression in a number of dimensions. They found that those concept mapping tools that offer greater flexibility by providing multivariate ways of representing information (and thereby broader representational capabilities) were preferred by users. The authors conjectured that enhanced flexibility may extend to enhanced learning, although they were unable to conclusively demonstrate this in their study. The expressiveness or expressive power of a representation is also a significant concern in the computational knowledge representation community as well: what knowledge can be adequately expressed?

Concept maps are intended to be tools for knowledge representation. Indeed, Jonassen (1992) claims that concept maps are accurate reflections of their authors' cognitive structures. Others support this notion in asserting that concept maps make tacit, internal knowledge explicit (e.g., McAleese, 1992) (although, as should be expected, the notion of concept map as direct mirror of cognitive mechanisms is open to question; Fisher, 1992). Jonassen poses what is therefore a crucial question: "What constraints does the software impose on the product?" (p. 20). In other words, what limitations do concept map tools impose on the knowledge that can be expressed using those tools? If concept maps are tools for representing knowledge and people possess a broad range of cognitive knowledge representations, concept maps ought to offer a concomitantly broad range of representational facilities.

Heeren and Kommers (1991) raised these issues in the context of concept maps that portrayed text-based propositions only—that is, at the time of their experimental investigation, concept map tools provided fairly limited flexibility with regard to representing knowledge. While the tools used in these learner studies provided for simply textual representations of nodes and relationships among them (and some existing tools do not even provide for labeling of links), such tools did not allow users to represent other forms of knowledge, such as imagery and sounds. If we can only express textual propositions in a concept map tool, then we can only portray a subset of our knowledge of a domain. Thus, such tools fail with regard to what the artificial intelligence knowledge representation community refers to as *representational adequacy*, the ability for a knowledge representational formalism or mechanism to represent all the kinds of knowledge that are required for understanding a domain (Rich, 1983).

In the years since Heeren and Kommers' (1991) user studies, personal computers have changed in many ways in terms of the capabilities they provide and the knowledge and information they have ready access to. Computers in schools and homes can now easily display graphics and animated graphics, can play video and audio, and have connections to the world of content available on the World Wide Web. By capitalizing on the capabilities of modern personal computers and the Internet, we can provide for much richer knowledge representation and greater flexibility of expressiveness in computer-based concept map tools.

At the least, concept maps ought to incorporate static imagery simply because this sort of knowledge is possessed by people (e.g., Kosslyn, 1980; Paivio, 1986). Concept maps ought to include temporally dynamic content as well, including animated images, video, and audio. Clearly, people cognitively represent memories for dynamic imagery (e.g., Johnson-Laird, 1983) and auditory information (e.g., Dowling & Harwood, 1986). Adding such capabilities to concept map tools supports enhanced flexibility of expressiveness, and the potential pedagogical benefits of dynamic media are numerous. Concept maps would thereby be able to portray, map authors could then express, and learners could then perceive from maps a much richer set of knowledge and information, with concomitant pedagogical benefits. Concept maps could portray what the entities of a domain look like when in motion, what particular things sound like, how domain elements behave, react, move, and sound in specific contexts or situations, dynamic interactions between people and objects. With static and dynamic imagery as first-class elements, concept maps can portray *essentially visual* elements of a domain per se—for example, in a map about birds, the image of an eagle soaring across the sky and swooping down onto the surface of a lake to catch a fish in its talons—as well as *information visualizations* related to a domain—for example, a map of North America annotated with circles indicating natural eagle refuges.

Another fundamental characteristic of human cognition is the ability, and in fact necessity, to exploit knowledge abstraction. Abstraction inherently implies the ability to represent a concept, action, or object by a single node at one level of detail while possessing the knowledge to expand that single node into an elaborated definition of its own. A single knowledge element at one level of abstraction may comprise a number of lower level elements at a more detailed knowledge level (Anderson, 2000; Rich, 1983). While many concept map tools do provide for the notion of submaps within concept maps, few do so in a facile and intuitive manner. Just as in the case of different forms of knowledge, concept maps should also provide for the easy representation of and simple navigation among multiple abstraction levels. In terms of expressiveness, multiple-layer concept maps also permit users to apply heuristics of visual perspicuity and aesthetics.

A small digression regarding the cognitive fidelity of concept maps: While there are generally accepted theories of cognition and cognitive mechanisms, and psychological explanations and accounts for empirical experimental observations and data, no one can say with absolute certainty how people represent knowledge cognitively. What we can say for sure is that humans do, somehow, represent, remember, and recall knowledge and information. Let us not forget that not very long ago the zeitgeist in experimental psychology was that all behavior was simply the result of stimuli and our animal-like responses to them. Ultimately, then, it would be hubris to say a particular style of concept mapping is more inline with, isomorphic, or analogous to cognitive representations. However, we know that people represent a variety of disparate forms of information cognitively. We have experimental evidence for some sort of associative conceptual memory. We know as well that people, in whatever manner, represent memories for things other than textual and propositional knowledge, such as imagery and sounds. And that people deal with the world by making efficient use of abstraction mechanisms. Thus, even if concept maps do not represent knowledge in ways truly isomorphic to cognitive mechanisms, the mere fact that people do possess memories for information about a domain in a variety of formats and styles implies that knowledge representation tools such as concept maps should also provide mechanisms for representing information in an analogous variety of ways. That is, concept map tools should provide the flexibility for users to express their knowledge in a variety of forms. If indeed concept maps are to be used as knowledge representation tools, they should provide the capability to represent—at the least—the types of knowledge just described. Even if concept maps' representation mechanisms turn out not to truly be analogous to cognitive mechanisms, they objectively will nonetheless provide for greater breadth and flexibility of expressiveness.

These capabilities and options may also help to accommodate individual differences among students using concept maps to demonstrate their own knowledge or acquire new knowledge, thereby better supporting students with differential abilities, learning styles, expressive or learning preferences or needs. For example, hearing impaired learners tend to prefer or require visually oriented learning materials; as a Teacher of the Hearing Impaired has told me, "I spend half my time drawing pictures for my students [to explain textually-represented concepts]" (Bomus, 2001). Perhaps concept maps that incorporate image-based nodes rather than, or in addition to, text-only nodes would better suit the learning needs of such students.

A more recent enhancement to the capabilities of personal computers is accessibility to the information and content available on the Internet. Allowing concept maps to incorporate access to this vast store of knowledge can enhance further the expressiveness of such maps. Concept maps ought to be able to include hyperlinks out to the Web at large.

Concept map tools built by researchers in universities and research labs have attempted to capitalize on the capabilities available in modern computers and move towards providing greater representational coverage and flexibility. Commercial concept map products have generally lagged behind, although very recent developments are encouraging. We now look at specific features and representational capabilities that concept maps may or may not offer.

Nodes and Links

Concept maps consists of nodes representing concepts, objects, ideas, or actions, connected by links that represent relationships between nodes. A node may be represented graphically by a simple geometric object, such as an oval or a rectangle, containing a textual concept name. Inter-node relationships are visually portrayed by directed links between nodes, that is, lines connecting nodes with an arrowhead at one or both ends indicating the directionality of the relationship defined by the link. In some tools, arrowheads, and therefore directionality, are not available to the concept map author. In some, textual labels describing links' relation are not available—links are simple lines indicating there is *some* relationship between the connected nodes. Together, labeled nodes and links define *propositions*, assertions about a domain.

Thus, the simplest constructs of a concept map are nodes and links. Yet even here there is room for flexibility. Heeren and Kommers (1991) found that concept map users preferred to distinguish different types of concepts. Concept map authors may wish to differentiate nodes based on, for example, the nature of the nodes or the relative importance of particular concepts in a domain. Such distinction should also help concept map consumers to better understand the represented domain. For instance, things an object does (e.g. an ocelot can *run fast*; a bird can *fly*) might be distinguished from an object's characteristics or attributes (e.g., an ocelot is a *cat*; an ocelot has *fur* and it is uniquely *spotted*; greenhouse gases are *pollutants*).

It may be desirable to distinguish links as well, perhaps according to relationship categories. For example, the following link relationships might fall into the *causal dependency* category: the *purpose* of an ocelot's spotted coat is camouflage, evolutionary adaptation *led to* ocelots having spotted coats, burning fossil fuels *produces* greenhouse gases, greenhouse gases *cause* global warming, greenhouse gases *alter* climate. On the other hand, *familial and hierarchical* relations, such as an ocelot *is a* cat and greenhouses gases *are a kind of* pollution, might be represented by links visually distinguishable from causal links.

There are a number of possible ways for a concept map tool to allow authors to distinguish nodes. These include providing multiple node shapes—different shapes can be assigned differing semantics—as well as the ability to change a node's fill and outline colors, the ability to modify the font, font size, style (bold, italic, plain), and color of a node's textual label, a node's size, and the z-order of overlapping nodes. These facilities are commonplace in concept map tools at this point in time. Inspiration offers a large library of node shapes and styles and, in addition to the above characteristics, allows users to change a node's background pattern and apply a drop shadow to a node's appearance. Inspiration also allows a node's shape to be morphed: a user may select an existing node and then click on any shape icon in the node toolkit thereby changing the existing node's appearance. Webster allows a different sort of flexibility with regard to node shapes. While the library of available shapes for simple textual nodes is limited, a node may be reshaped by users to suit any idiosyncratic semantic or aesthetic requirement. When a node is selected, a shaper ball appears within the node; the user may drag this shaper to reshape the node in constrained ways. Figure 1 shows examples of this functionality. Thus, the user may create a wide variety of shapes.

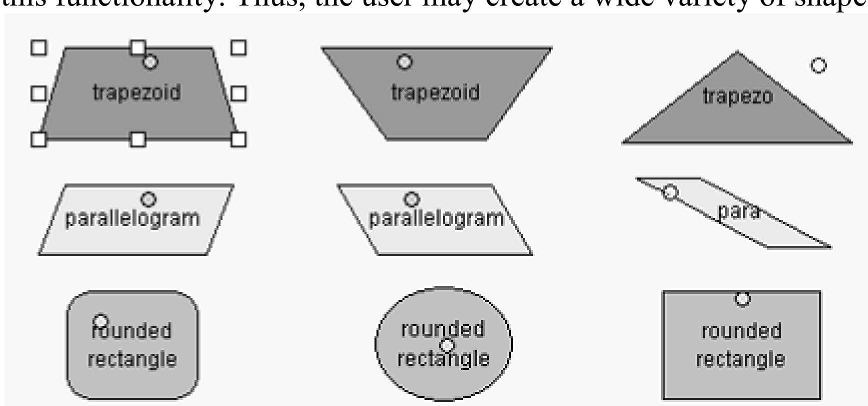

Figure 1. When a node is selected in Webster, its outline handles and "shaper" appear (here the handles are elided for all but the upper left node for the sake of clarity). The shaper—a yellow ball—may be dragged to reshape the node.

Similarly for links: concept maps ought to provide varying link styles so as to allow maps to differentiate semantically different link relationships. The TCU concept map (Lambiotte, Dansereau, Cross, & Reynolds, 1989) offers different link types, such as "barbed" lines with arrowheads along the length of the line. Webster and Inspiration offer links with a varying number of vertices resulting in differing line shapes, to accommodate visual aesthetic preferences and that may also be used to indicate differing semantic relationship types. Links in these tools may also be annotated (that is, distinguished) by color, font, font style, font size. Inspiration takes this idea to its fullest potential by offering a virtually infinite variety of line styles, thicknesses, colors, and shapes—with regard to shape, for example, a user may drag vertices of a link to any location. An advantage of Webster's single vertex line style is that the vertex itself may be dragged to any point on the link thereby repositioning the link's textual label; this is useful again for enhancing flexibility with regard to the visual layout of a concept map's elements. Webster and Inspiration also allow links to be morphed from one style to another: for example, in Webster a user may click on any link type icon (say the two-vertex, right-angle style) in Webster's toolkit and drop it onto an existing link in a map (say, a straight-line link) thereby changing the existing link's appearance.

With regard to links' textual labels, the TCU tool permits the map producer to choose from a small fixed set of "canonical" link names. The TCU team believed there would be cognitive benefits of constraining link types, and therefore their labels, to a fixed set of choices—in particular, they felt a small set would ease the short-term memory burden for map producers and consumers. The set of choices could be changed depending on the domain being represented. This can be useful for map authors in producing links with meaningful domain-specific and domain-independent semantics. Unfortunately, the labels in the TCU tool are abbreviated to a single character and producers and consumers of maps must remember what these abbreviations actually mean, thereby, perhaps, diminishing the potential cognitive benefits of the constrained set. And the very notion of providing only a constrained set of relations is itself disputable. Fisher (1992) concludes that such constraints are not desirable when dealing with a complex knowledge domain. Also, for example, links such as *results in*, *produces*, *enhances*, and *causes* are all causal relations, but if the user is constrained to label all of these simply as *causes* the user and the concept map lose the subtle nuances natural language offers.

Inspiration, on the other hand, provides the flexibility of allowing link labels to be anything the author types. Webster adopts a hybrid approach that offers the benefits of both approaches. When a user wishes to enter the label for a link (when the link is first created or when editing an existing link) a combo-box is displayed; this widget contains a list of commonly used and useful link names as well as a text entry field. The predefined labels include many that should be semantically meaningful and expressive in typical maps, such as "is a," "causes," "is part of," "can," "results from" and so on. The user may select a pre-existing label from the popped-up list or may type any text of the user's choice (there is also a type-ahead feature wherein the user starts to type a label and if a matching label exists in the list the full label appears in the text entry field). Thus the user may avail herself of predefined link labels or create new ones. New, user-created labels are dynamically added to the list so if a user intends to reuse a particular label or labels within a map, they are readily available without retyping. This is useful for domains in which the same idiosyncratic semantic relationship exists between multiple different objects, that is, that require the same domain-specific, user-created, link label in multiple places. In Webster as in the TCU tool, the list of predefined labels may be changed for different domains.

Static Image Nodes

To this point, we have discussed capabilities that completely define the breadth of representational mechanisms for many concept map tools—that is, many concept maps may contain only textually labeled nodes and links. As we have discussed, certainly people possess knowledge other than textual propositions and thus concept map tools need to represent knowledge other than textual nodes and links. Even if we are speaking only of flexibility of expression, again we ought not be constrained to text-only nodes. The first enhancement that has been seen in concept map tools in this regard is the incorporation of image-based nodes, the ability to include any image that may be relevant to any particular domain being represented

Many concept map tools allow the use of *static* images in a map. Learning Tool™ (Kozma, 1992) permits the *association* of static imagery with a textual concept node; it has the provision for a “notecard” to be associated with each concept node, and this notecard may contain text and images. But the image itself is not visible in the network of nodes and links per se. Webster, Inspiration, and other tools such as SemNet™ (Fisher, 1992) incorporate *image nodes*, allowing any image to appear directly within a map. Image nodes in these tools initially appear in a map with a default thumbnail size but may be stretched to show or hide details or otherwise accommodate aesthetic and spatial considerations. Here, image nodes are first class elements of a map, just as textually labeled nodes are.

Webster, and now Inspiration as well, goes a step further in that image nodes may be animated-GIF (GIF89a) images (CompuServe, 1990), an image format commonly used on the Web. Thus we may have animations that “play” in a static location directly within a concept map. For example, the animated bird image in Figure 2 appears to be flying within the map.

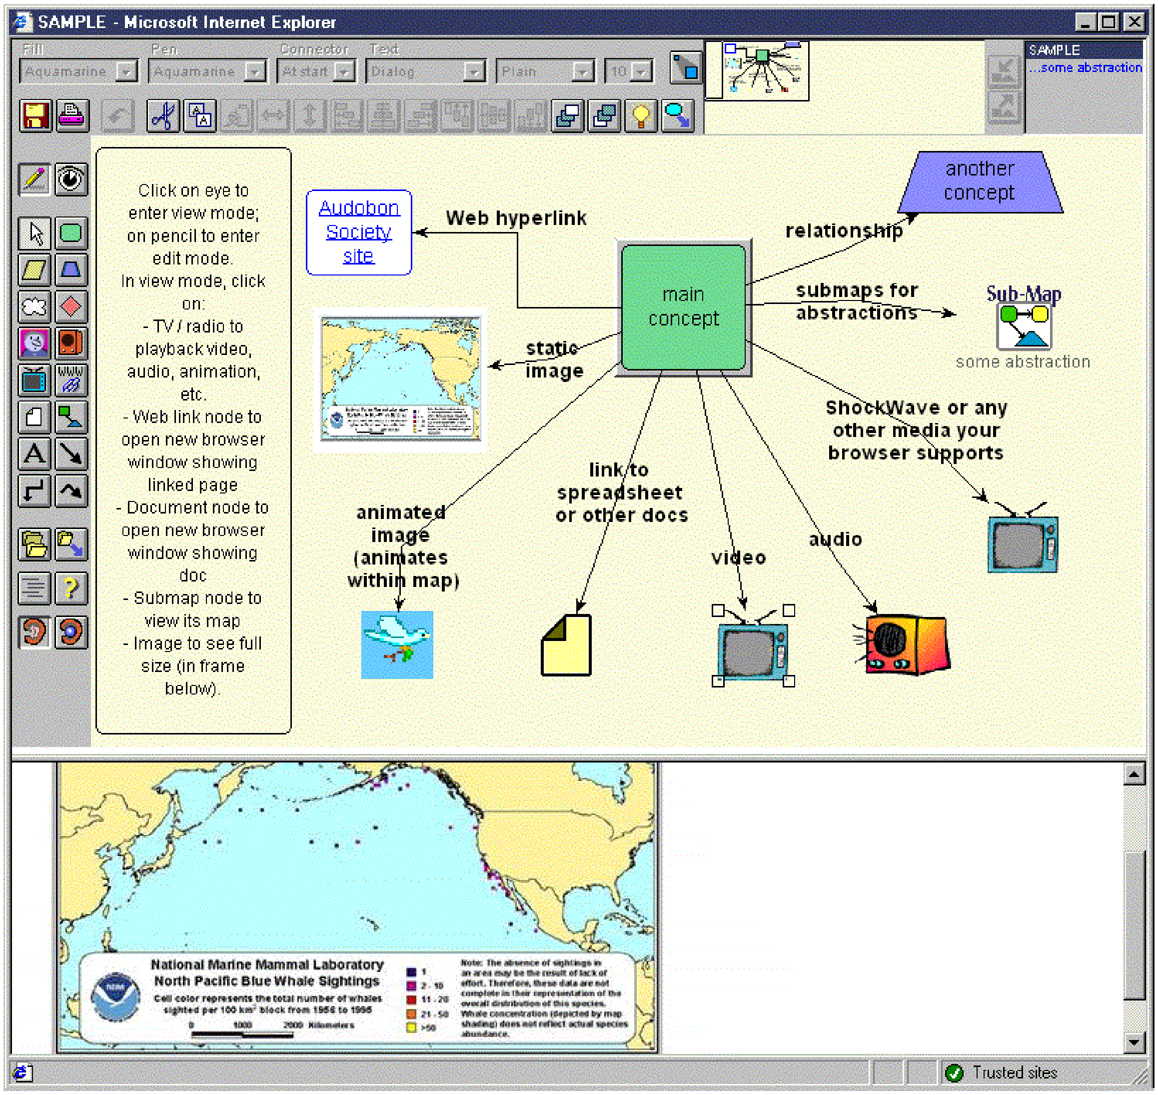

Figure 2. The Webster concept map tool in a Web browser. This concept map shows the types of information that can be incorporated. Image nodes appear in the map as (stretchable) thumbnails; users may view them full size in the bottom frame of the browser in which Webster appears.

Dynamic Media Nodes

The tools that offer image-based nodes offer a step in the right direction beyond text-only representations. But concept map tools ought to extend further. One's knowledge of horses might include not only what a horse looks like, but what one looks like in particular when galloping, trotting, or racing; the sight and thundering sound of a group of wild horses stampeding across a dusty plain; the sights of a bronco bucking wildly attempting to throw a rodeo cowboy off its back; the sound of a horse neighing. In a map concerning whales, a student may wish to represent propositional assertions such as "a whale is a mammal" and "blue whales are the largest animals on Earth," as well as representing and presenting dynamic imagery of whales leaping spectacularly into the air and the sound of whales' unique songs. A history of the Civil Rights Movement in the United States would be diminished if one could not actually hear the emotionality and oratorical style of Martin

Luther King's speech during the 1963 March on Washington. Certainly these elements may be parts of our long-term memories for these respective domains, and we ought to be able to represent them in concept maps to demonstrate our own knowledge of those domains or to use concept maps as instructional resources.

Thus when asking students to represent their knowledge, we need to be cognizant of the fact that some aspects of a student's mental model of a domain may be expressible only through visual or auditory media. Standard PCs are now capable of representing and "playing back" video, audio, and animation. If we are using the computer to represent students' knowledge, we should more fully benefit from what the computational medium offers over more static media. This should include the ability to incorporate dynamic content in concept maps thereby more comprehensively representing the knowledge of a domain.

Webster allows for the representation of dynamic media components as integral constructs of concept maps. In Webster, multimedia objects are represented by distinct nodes. Thus multimedia elements are first-class elements of concept maps with the same status and ease-of-use as textual concept nodes. To incorporate a sound clip, a user includes an audio node (which appears as a radio, albeit an old-fashioned one) in the map; for video information the user includes video nodes (which look like televisions). These nodes are added to a map in the same manner as adding simple textual concept nodes (and all other node types). To play an audio node's sound file, a user simply clicks on a radio node in the map. Video playback is invoked in the same fashion via TV nodes in a map.

Inspiration has now also incorporated the ability to include multimedia in concept maps. In the latest version of this software, media can appear in maps in two ways. First, users may make recordings within the Inspiration environment, presumably for commentary or narration, and associate such recordings with concept nodes. Once done, a small sound icon appears adjacent to the associated concept node. But, we wish, of course, to use any arbitrary sound that may be relevant to a domain being represented. Inspiration now allows a single concept node to hyperlink to *any* file in the computer's file system—when the user clicks on that node the application associated with the file opens on that file (much the same as when a user double-clicks on a file in the system's file browser). Hence, despite the lack of explicit video or audio nodes per se, any node can link to a file whose type is video or audio, and the appropriate media player application will execute when the node is activated by clicking.

Very few other concept map tools incorporate dynamic media elements, and these tools are from the research community. In HandLeR's "topic map" tool (Sharples, 2000), topic nodes may reference the same types of media files as supported in Webster. Unfortunately, all topic nodes in HandLeR are identical ovals with embedded text, connected by unlabeled links. Thus different node *types* are not visually distinguished or apparent. Hence, multimedia elements are not readily manifest. In the IHMC CMap toolkit (IHMC, 2000) users may associate *resources* with a concept node including images, sounds, or video files. Again, in the visible map the user sees only a textually-named concept node and must follow a hyperlink from the node to playback its associated media files. Thus there are no relationship links *between* textual concepts and media nodes; media resources are, in a sense, *part of* the textually-labeled concept nodes. A potential problem with this approach is that labeled links showing the relationships between media nodes and other nodes represent additional knowledge, thereby enhancing the semantics of a map. Indeed, based on studies with learning disabled students, Griffin, Simmons, and Kameenui (1992) strongly recommend that concept maps and other graphical organizers explicitly depict (i.e., label) the relationships among conceptual elements to better support the learning and memorability of the represented domain. (In this specific case, however, the “additional” semantics provided by the link may be minimal when, for example, a link from a concept node to an associated media node is labeled simply *looks like* or *sounds like*.)

On the other hand, making media nodes visually distinguished *can* add semantically meaningful information to the map. In Webster concept maps, nodes provide visual cues regarding their content and type, including the clear indication of the existence of media nodes, which appear as televisions (video nodes) and radios (audio nodes). This notion can be labeled *visual typing*, that is, making each node’s content type evident by its appearance. It appears that The KMap tool (Gaines & Shaw, 1995) adopts a similar approach. Nodes in a KMap concept map appear as rounded rectangles along with differing configurations of lines in the rectangles’ interior indicating each node’s type. In Inspiration, a user may on his own “adopt” the notion of visually typing media nodes: any node may be a hyperlink to a multimedia file and the user may also have an image associated with such a node; users may thus make a node that links to a video file, for example, appear as a television or a film segment or any other image evocative of video.

Visual typing of nodes should support users using a concept map to learn about a new domain; if a user is viewing a map about whales, he can readily find, say, audio elements of a whales’ songs. Webster attempts to make the representation of knowledge elements straightforward and simple, both conceptually and in terms of user interaction, in part by treating different types of knowledge elements in a uniform fashion. Multimedia knowledge is represented as distinct nodes (rather than information attached to, part of, or within a concept node) and these nodes are manipulated in the same manner as other nodes in a map. If there are indeed distinct cognitive representations of textual-propositional and dynamic imagery elements and we wish to reflect knowledge in concept maps, it follows that these knowledge elements of different types should be represented by distinct nodes and distinct node types rather than being bundled into a single node.

From a pedagogical perspective, if, as a popular expression goes, a picture is worth a thousand words, then a scene of pictures in motion may be worth 10,000 words. For example, one can explain with words how to perform a physical skill but a video *showing how to*, say, swing a golf club can be an invaluable tool for learning to perform the skill. I guess one can also say a sound clip is worth *n* words. If studying the domain of blues music, one can explain in words that the blues scale is a modified minor pentatonic scale or try to get more concrete and say that it is derived by flatting the 3rd, 5th, and 7th notes of a major scale, but without also hearing and experiencing an electric Chicago blues band or a song played by a solo country-blues guitarist-singer, one would really know nothing about the blues. We can attempt to explain sensory phenomena in text, but in the words of an old song, the frustration of trying to do so is "like trying to tell a stranger 'bout rock and roll.”[1]

The incorporation of temporally dynamic visual and aural elements in a concept map tool enhances its flexibility of expressiveness. The number and breadth of semantic and aesthetic options are increased for users. It may also help to accommodate individual differences, that is, to better support students with different expressive or learning preferences or needs when using concept maps to demonstrate their own knowledge or

[1] Lyrics from "Do You Believe In Magic?" by John B. Sebastian, performed by the Lovin' Spoonful in 1965, © Kama Sutra/Buddah.

acquire new knowledge. Temporally dynamic content in concept maps will better support those using maps for learning new domains and concepts because these users will be presented with, and be able to perceive for themselves, what objects of these domains look like when in motion and what they sound like. The use of multimedia in concept map tools may also provide for a more engaging user experience, thereby motivating students to spend more time on task, thereby leading to enhanced pedagogical effects.

Multiple Abstraction Layers

Many concept map tools lack adequate visual or structural abstraction mechanisms. In some cases maps appear as a single diagram, that is, nodes and links representing all of the knowledge of a domain are drawn in a single network or layer. In these networks, a weak notion of abstraction is represented only via generality-specificity relationships between nodes in the single layer: a node representing a specific concept (say, *bird*) may have a link, labeled *is a* or *a kind of*, pointing to a more general, or abstract, concept (say, *animal*). The more abstract node might have additional links connected to other concept nodes (e.g., portraying attributes of *animals*, such as *animals have skin* and *animals breathe oxygen*). There might also be relationship links and associated concepts connected to the *bird* node (e.g., a *bird has wings* and a *bird can fly*). This format can extend several levels of abstraction in either direction (e.g., a *goose is a bird*, a *lizard is an animal*, an *animal is a living thing*). Very quickly, such a map can become visually busy and confusing. Further, the map may not reasonably mirror the cognitive representation of this knowledge, which would exploit abstraction mechanisms to represent categories or concepts at differing levels of detail.

Abstraction mechanisms enable the economical and visually lucid expression of knowledge in a concept map. But we must also remember that such benefits are meaningless if we disregard the ease of *using* the representation. One aspect of a representation's usability and effectiveness is the ease of finding information therein (Larkin and Simon, 1987). Thus, the mere *existence* of submap capabilities in concept map software is insignificant if submaps are difficult to discern or navigate. If a learner is unaware that there is knowledge contained in submaps or has difficulty viewing that knowledge or becomes confused due to cognitive overload in viewing the overall representation, this obviously hinders the usefulness of the knowledge contained in submaps!

There are concept map tools that provide submap, or child map, mechanisms to attempt to ameliorate some or all of the above problems; here we'll again focus on Webster and Inspiration. In Webster concept maps, an abstraction may be represented by a submap node and associated submap. These two elements reside at different levels (or, in different layers) of the overall concept map. A submap node represents a single concept—is visually represented by a single node—at one level of a concept map. This single concept can be expanded into another map (a submap) of knowledge elements representing the fuller, more detailed meaning and constituent parts of the single concept. We thereby gain the ability to portray the knowledge as a person might, and in terms of expressiveness further gain the benefit of making the knowledge representation more graphically parsimonious. Now, a concept, say *lizard*, can be represented at one level of abstraction or detail by a single node, without cluttering that map level with all the detailed information about lizards, and yet those details can be available in another layer of the overall concept map. In Webster, submap nodes and their associated maps can be created in a variety of intuitive and flexible ways, including importing an existing external concept map as a submap, or pushing a group of nodes "down" from one level of a map to create another layer (see Alpert, in press, for details)

To reify further the preceding notions in the context of Webster, Figure 3a shows the topmost level of a concept map about whales. This level contains a submap node labeled *Blue whale*, connected via an *is a* link to the main node labeled *Whale*, and the same configuration for the *Humpback* submap node. At the currently visible level of this concept map, *Blue whale* and *Humpback* are individual nodes, singular concepts. A user may "enter" the submaps associated with those individual submap nodes in order to view their constituent knowledge elements. Doing so entails visiting a different abstraction level of the overall concept map. When a submap has been entered, it becomes the visible, active map (as in Figure 3b).

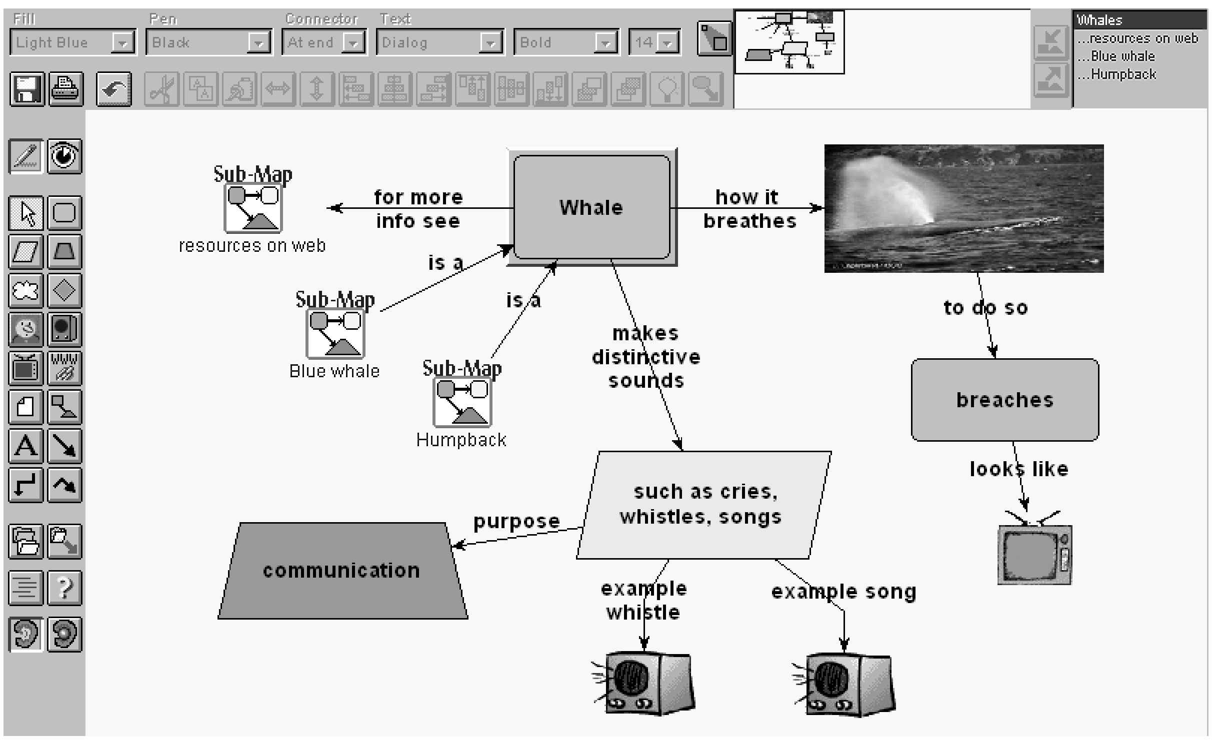

Figure 3a. A Webster concept map about Whales.

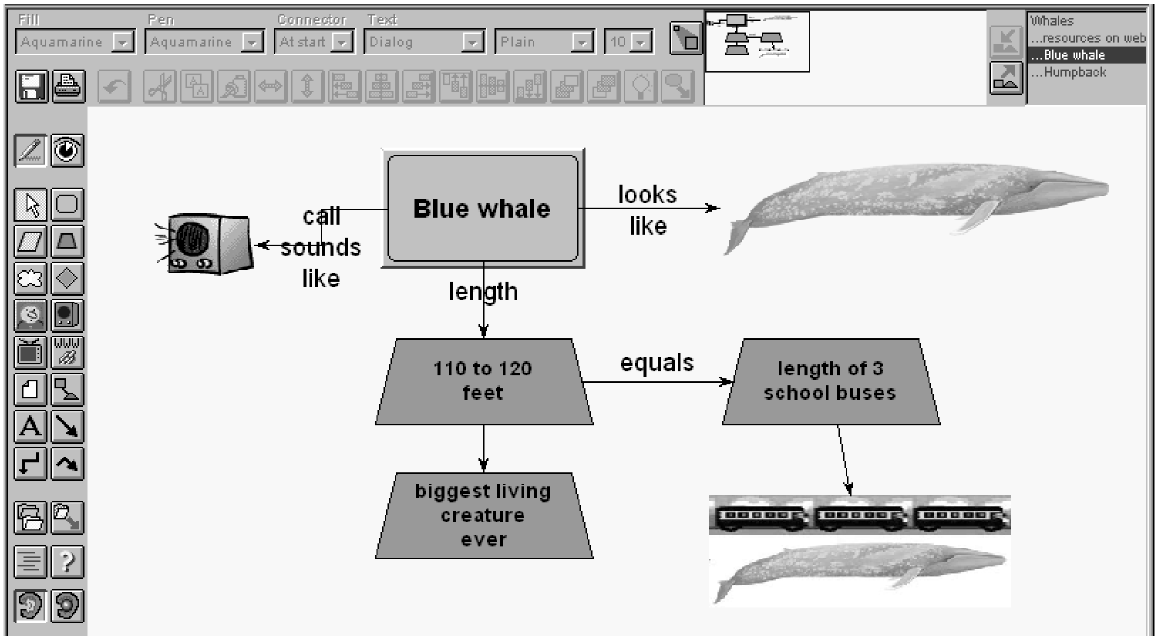

Figure 3b. Continuing the example concept map, portrayed here is the (sub)map associated with the *Blue whale* submap node in 3a.

In older versions of Inspiration, the product also provides for "child maps." Individual nodes in these Inspiration maps may have a child map; however the fact that any node has a child is indicated to users in a rather obscure fashion. When a concept node is selected (but *only* when the node is selected), "handles" appear on the node's corners and sides; if the node has a child map, the handle in the upper right corner of the node is a filled-in square rather than an outline of a square—a not particularly salient indicator. Worse, when a node is not selected, there is no evidence at all that a child node exists. On the other hand, Webster clearly indicates the existence of submap nodes (and therefore submaps) by its visual typing of nodes—a submap node's appearance is clearly distinct from other node types (e.g., see the *Blue whale* submap node in Figure 3a).

In the recently released version of Inspiration, child maps have been removed altogether as a construct. Instead, an individual labeled node may refer (hyperlink) to an entire external map. Clicking on the hyperlink-node causes the linked map to appear. Based on initial usage of this new version, it is not apparent how to navigate *back* from the newly displayed map to the original hyperlink node.

Navigating among layers of a map or among submaps can therefore be seen as problematic. The problem is further complicated by the fact that abstraction mechanisms may result in multiple levels of knowledge at differing degrees of detail. A single layer of a concept map may include any number of abstractions (submap nodes), and the associated submaps may have their own submap nodes, and so on—that is, submaps may be recursively embedded. The result is a "three dimensional" knowledge representational scheme wherein each level of a map contains knowledge elements (nodes and links) arranged in two dimensions, and individual submaps are stacked in the third dimension. The vertical stacking of submaps represents multiple levels of abstraction within the overall concept map. One problem in navigating such a structure is identical to a difficulty that occurs when users interact with hypermedia in general, namely, getting "lost in hyperspace" (Conklin, 1987). This describes a situation wherein a user navigates to a topic (or in our case, to a map at a particular layer in a stack of submaps), then to a linked or associated topic (or submap), and another, and so on, and in the course of doing so becomes confused and disoriented, not able to remember or reconstruct "where" he is with respect to the information traversed. According to Conklin, this disorientation is due to cognitive overload.

To facilitate the use of its abstraction capabilities, Webster provides simple tools for navigation among submaps. One clear approach to ameliorating problems of cognitive overload is to provide scaffolding via external memory mechanisms, in the case of software in the interface of the application. In Webster, the Abstraction Levels Navigator makes the hierarchical relationships among submaps (which reflect the superordinate-subordinate relationships among abstraction levels) explicit in the interface on a constant basis (the Abstraction Levels Navigator appears in the upper right of the user interface; see Figure 4 for a detailed view). This explicit representation provides semantic information that helps users make sense of the three-dimensional structure of the overall map. Further, the Abstraction Levels Navigator always indicates which submap is currently in view: as different submaps are visited, the currently active submap name is highlighted, thereby informing the user "where" he is in the stack of submaps, that is, orienting the user in his three-dimensional navigation of multiple abstraction levels. This approach is isomorphic to recommendations regarding assisting users in hypertext and hypermedia navigation. For example, Shum (1990) asserts that navigation problems can be ameliorated by making the structure of the information space explicit for the user.

Webster thereby attempts to better support the user in identifying the existence of, and navigating among, different abstraction levels in a concept map. Webster's interface makes the existence of all submaps as well as map-submap relationships persistently explicit, makes all submaps immediately accessible, and makes apparent which submap is currently active.

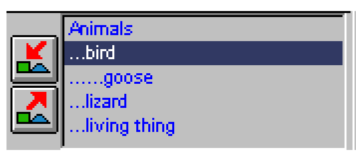

Figure 4. The navigation tools that appear in the upper right of Webster's user interface. The two buttons on the left permit navigating up or down a single abstraction level. The top button lets users "dive into" a submap (it is enabled only when a single submap node is selected); the lower button allows a user to "jump up" one level to the parent map of the currently active submap (it is enabled only while viewing a submap). The Abstraction Levels Navigator (right) makes abstraction relationships explicit via indentation—*goose* is a submap of the *bird* map, which in turn is a submap of the top-level *Animals* map. From another perspective, the *Animals* map contains three submap nodes labeled *bird*, *lizard*, and *living thing*, each of which has an associated submap. Users may also navigate directly to any submap (abstraction level) by clicking on the submap name. At all times, the currently visible submap is highlighted to orient the user.

Mapping and Collaborating on the Web

The Web offers an unprecedented opportunity to educational technologists. No longer are we required to deal with the logistics of getting software applications onto individual computers so they can be used by students. Instead, Webster and other concept map tools (e.g., Gaines & Shaw, 1995; Flores-Mendez, 1997) have been architected to be accessible anywhere via the Web and a standard Web browser. Figure 2 shows Webster running in a Web browser. Some mapping tools have been implemented as plug-ins

Web-based concept map tools (or any educational application) offer the ability to reach greater numbers of students. Students are no longer constrained with respect to where or when they can access the tools. The Internet allows us to easily have a single centralized database so that all student maps reside in a single location on a server. Further, having applications—in Webster's case a downloadable client applet and a server-side servlet to access the centralized database—reside in a single location results in students always using the current version of the software. Bug fixes and tutor enhancements are immediately available to everyone, rather than again having to deal with the logistics of distributing and installing software updates.

This allows incremental work by an individual student on an individual concept map to begin at school and progress at home, library, or community center, collaboratively with friends, anywhere they have access to the Internet, at times other than only when school is open and its computer resource room is available. Another dimension of flexibility is thereby addressed, namely flexibility of when and where concept mapping occurs. Teacher assessment of student maps may also occur anywhere at any time. (And hopefully, the fact that many libraries now possess publicly accessible, Internet-enabled computers may help to mitigate the "digital divide" and provide something approaching egalitarian access.)

Webster's approach to Web accessibility is a portable, highly interactive Java applet that executes in popularly used browsers and interacts with servlets running on the Webster server for database access and other functionality. Another approach for Web-based concept map creation and editing is to implement a browser plug-in that executes a proprietary mapping application (e.g., Kremer & Gaines, 1996). The minor difficulty with this approach is that users must have already downloaded and installed the plug-in code to edit maps. Several concept map tools that do not permit concept map creation or editing on the Web do offer the ability to export a concept map in a Web-enabled form. Inspiration, for example, allows one to export maps as graphic images or as HTML pages. Most of the interactive features and elements in the original concept map appear in the HTML translation—for example, the hyperlinks defined in the map do exist in the corresponding Web page, although image animations and audio narrations do not. Nonetheless, although map creation cannot be performed online, the resulting concept maps can be published and then accessed by map consumers.

A centralized database of maps opens the door for formal collaboration in map creation. For example, in Webster, students can import external maps created by others into their own maps. Thus, a team of students, perhaps in disparate locations, can work together to create a complex concept map that contains multiple components created by the team's individuals. For example, one student might create a *Blue whale* concept map. At another point in time, a second student working on a *Whales* map can incorporate the first student's *Blue whale* map as a submap. Similarly, an Inspiration map created by student A can include a node that hyperlinks to a second concept map previously created by student B, although, because Inspiration is not Web-enabled, to do so would involve logistics of collocating the two map files on a single machine or local area network. One can imagine further a concept map tool that permits *synchronous* collaborative concept map editing, wherein multiple users in different locations can work together on a single, shared concept map in which changes by individuals are seen by all (a point also made by Kremer and Gaines).

Collaborative map creation increases the flexibility with which concept map may be created, and being able do so over the Internet, independent of location and file-sharing considerations, increases this flexibility further. Indeed, collaborative map building addresses an entirely new dimension of flexibility of expressiveness that probably never occurred to Heeren and Kommers when they coined the term. Here, rather than an individual's expression of his or her knowledge, multiple users contribute components of the overall knowledge of a domain. Collaboration in map creation should also strengthen the very qualities of concept mapping that constitute its pedagogical rationale, such as deeper reflection on one's own and now others' knowledge. Modifying a map in collaborative fashion exercises users' skills of critical analysis and argumentation regarding one's own and others' updates. Collaborative map creation may also be more engaging (Cicognani, 2000), and engaged analysis of the knowledge, especially of others, represented in the group's map can highlight gaps in one's own knowledge and bring about learning. There is indeed empirical evidence of positive pedagogical effects of peer interaction, evaluation, and cooperation in online collaborative educational settings (see, e.g., Hiltz, 1998). With regard to concept mapping per se, there is also evidence supporting the pedagogical usefulness of collaborative construction and editing (Stoyonova & Kommers, 2001; Gaines & Shaw, 1995). To facilitate collaborative concept mapping, the tools students use should be accessible and executable over the Internet with a shared centralized database of maps. Further, Stoyonova and Kommers found the strongest learning effects for synchronous collaboration and therefore perhaps the best method would be via an online, synchronous multi-user concept map tool, none of which exist at this time.

In addition to making concept mapping and concept maps accessible from the Web, concept map tools should also provide links out from maps to the vast store of information on the Web. This again enhances the expressiveness of a concept map and the expressive possibilities open to map authors, allowing them to refer map consumers to Web-based information and materials external to the concept map per se. This also encourages a "new" style of research, specifically Web-mining to find relevant sites and online information and materials to incorporate into one's maps as additional resources for learning. A study by the Center for Applied Special Technology (Follansbee et al., 1997) demonstrated that such online research can result in enhanced learning, critical thinking, and presentation effectiveness, including improved ability to "see the big picture," to view knowledge from multiple perspectives.

Webster and Inspiration (and other concept map tools) provide Web hyperlink nodes, in which a Web address (URL) can be specified. Both tools allow the node's visible text to be something other than the URL itself (see Figure 2). When a user clicks on a Web-link node, the associated Web page is displayed in a separate browser window. Thus, Web-based resources can be "part of" a concept map.

Other Node Types

Similarly, concept maps ought to be able to reference external materials other than Web sites. For example, Webster is soon to be used by Chemistry students at an Ivy League university. The plan is to use concept maps in an attempt to link students' laboratory experiments and results to their knowledge obtained in the classroom and from textbooks. Student will wish to incorporate lab data and other lab-related information into such concept maps; their lab data and lab reports reside in files associated with standalone spreadsheet and text editor applications. Students should be able to integrate external documents of this nature into concept maps.

Both Inspiration and Webster once again provide such a capability. As described earlier, Inspiration now allows a single node to hyperlink to any file, and when a user clicks on that node the application associated with the file opens. Webster provides the Web-based version of the same functionality. Many applications provide browser plug-ins for accessing application-specific documents on the Web—for example, Microsoft Excel Viewer for spreadsheets, Microsoft Word Viewer and Adobe Acrobat Reader for text documents, Microsoft PowerPoint Viewer and Lotus Freelance Graphics Plug-In for slide presentations. There also exist third party plug-ins (e.g., Quick View Plus) that allow viewing of files associated with a broad range of applications. To include documents associated with external applications in concept maps, Webster users may create document nodes (see Figure 2). When a user clicks on a document node, the document is displayed in a separate browser window. The map creator must upload the documents to the Webster server, but this also implies that, unlike in standalone applications like Inspiration, those documents that are incorporated into a concept map may be viewed from anywhere on the Web.

Many other sorts of Web-based media can be accessed and viewed via Webster's "video" nodes. In addition to traditional video, "television" nodes in Webster may have associated with them any computational media that the user's Web browser is capable of playing—that is, any media type for which the browser has a plug-in. Concept maps may therefore incorporate three-dimensional virtual walkthroughs of, say, the Egyptian Pyramids, as well as Flash and Shockwave animations, VRML virtual reality scenes, and so on.

Outline Translations

There is abundant experimental evidence that multiple representations are an aid to problem solving. Alternative representations also support developing deeper understanding of a domain, and this may be particularly true when multiple representations of the same information are linked to one another in a learning environment (Kaput, 1989). A rational alternative representation for the information contained in a concept map is a tradition outline. Many students are comfortable with seeing their thoughts arranged in outline form. This may be especially true for students organizing ideas in preparation for prose composition. Outlines present concepts in a line-by-line numbered format with subordinate or associated terms in physically subordinate locations in the outline—subordinate terms appear indented below its immediately higher level concept—and this physical superconcept-subconcept scheme may repeat itself. Inspiration and Webster are among the concept map tools that offer automatic translation of concept maps into outlines.

Inspiration's outline facility is, in one sense, tightly integrated with the corresponding concept map. The creation of a new subtopic entry in the outline automatically results in a newly created node in the concept map; similarly, relocating an entry in the outline results in relocating (detaching from one node and reattaching to another) the corresponding link in the concept map. On the other hand, Inspiration's outline translations of concept maps are rather problematic in a number of other ways. Based on Kaput's (1989) assertion regarding the pedagogical effective of multiple representations, as well as common sense regarding a user's needs and expectations, outline representations *should* contain the equivalent semantic content as the corresponding concept map. Even though Inspiration concept maps may incorporate static and animated-in-place images in addition to textual nodes, the image nodes are absent in its outline translation of a map. Similarly, though Inspiration nodes may link to a video or audio file, its outline translation of a map with such nodes includes only the textual labels of those nodes. Further, the textual labels of all links in a concept map are also elided in the outline! The knowledge elements in child maps are absent as well in the outline translation of a concept map. The only content present in a concept map to appear in the outline "translation" is the names of the nodes in the single level of the overall concept map that was visible when the "outline" button was pressed. (These node names are appropriately indented depending on the relationship links that appear in the map). These deficiencies may defeat users' goals in using the alternative outline representation.

In Webster, multimedia concept maps become multimedia outlines. The static and animated images, videos, and audios that are part of a concept map appear in the associated outline representation, as shown in Figure 5. Webster outlines incorporate all of the knowledge, information, media, and relationship names contained in the corresponding concept map. This implies that all of the information contained in all concept map abstractions (i.e., submaps) appears in a *single* outline at the appropriate indentation levels—thus, as we expect the user desires, the outline presents the organization of all thoughts and concepts in the overall map in one place. With regard to static and animated images, they are handled much as they are in concept maps: images (including in-place animated-GIFs) are embedded and visible directly within outline items. Images are displayed in the outline with a default thumbnail size; to view the image in its original size, a user clicks on the thumbnail in the outline. Radio and TV nodes also appear in the outline and act as hyperlinks that reference specific audio or video files; hence, users again click on embedded radio and television images to "play" them.

Webster outlines also incorporate the textual labels that appear on relationship links in the concept map. This is crucial semantic information in the concept map and comprises associations without which propositions are no longer represented. For example, an outline entry might appear subordinate to the "Whale" entry with the text "is a: mammal" thus representing the propositional knowledge "A whale is a mammal." Other link-inclusive entries may include, for instance, the text "looks like:" followed by a clickable TV image, or "example song:" plus a radio-image hyperlink.

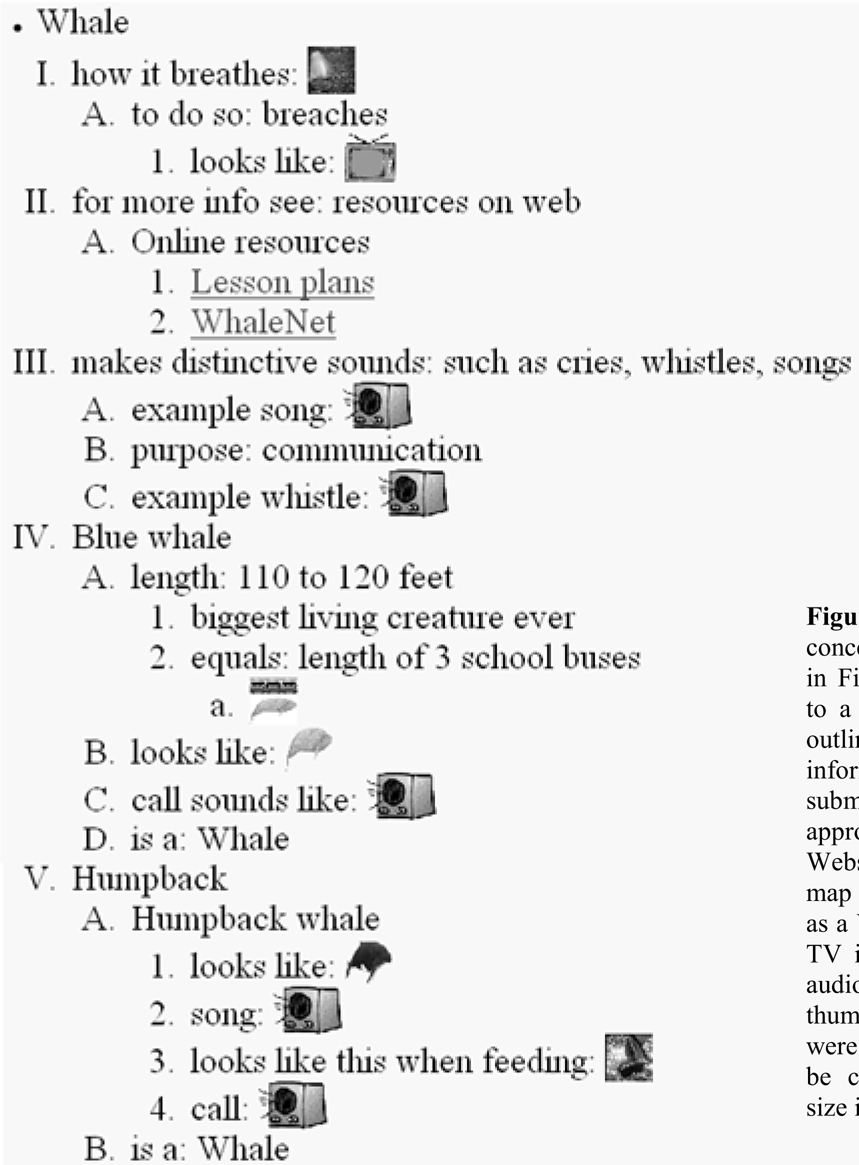

Figure 5. The Whales concept map partially shown in Figures 4 and 5 translated to a multimedia outline. The outline includes the information contained in all submap abstractions, at the appropriate indentation levels. Webster translates the concept map to HTML and displays it as a Web page. The radio and TV images are hyperlinks to audio and video files; the thumbnails of the images that were in the concept map may be clicked to view the full size images.

Conclusion

Many concept map tools have offered or do offer enhanced capability, expressiveness, and flexibility beyond what is provided by the commercial software used in schools and the current state of the practice in concept mapping. Fully and intelligently integrating the capabilities presented here into concept mapping software provide the opportunity to:

- allow students to represent their knowledge more comprehensively;
- provide richer expressive power and representational choices, that is, enhanced flexibility of expressiveness;
- offer the illustrative and pedagogical advantages of dynamic visual imagery and audio for students learning new concepts and domains;
- better accommodate individuals with differential learning needs or preferences;
- allow concept map authors to apply heuristics of visual aesthetics and clarity via multiple map layers and abstractions;
- perhaps provide for a more engaging student experience;
- perhaps, allow students to represent their knowledge in ways analogous to their own cognitive representations;
- better capitalize on functionality available in modern personal computers; and
- bring the state of the practice of educational concept map usage in line with the state of the art functionality.

Figure 6 portrays capabilities discussed in this chapter and their high-level rationale. Commercial concept map tools should be updated with these capabilities to more fully realize the potential pedagogical benefits of concept mapping (but see next paragraph). By offering broader knowledge representation mechanisms and broader expressiveness, more semantic and aesthetic options will also be available for users. It is heartening to see that *many* ideas and capabilities described in academic journals and conferences and implemented in research tools are indeed being adopted by the most widely used commercial concept mapping tool, Inspiration. We are now seeing something approaching convergence of the state of the art and practice in a number of dimensions. Nonetheless, Inspiration is less than adequate in its approach to multi-level concept maps, which if properly implemented could offer the benefits of abstraction both conceptually and with regard to visual parsimony, in its lack of collaborative facilities, and its anemic outline translations of maps.

The academic concept map community should set a research agenda to experimentally determine whether such enhancements are in fact preferred by learners and whether they indeed influence learning. As mentioned earlier, Webster is soon to be used by university Chemistry students; from this experience we will be able to gather evidence regarding the usefulness, usability, likability, and pedagogical implications of a number of features, such as full Web integration, multiple abstraction levels, and links to data in spreadsheets and other external documents. But there are other issues. For example, while there is evidence that memory for visual imagery is more robust than that for purely textual information (Shephard, 1967) and that information encoded both visually and verbally is more memorable than when encoded in either format alone (Paivio, 1986), can we therefore justifiably expect positive learning effects from the incorporation of static and dynamic imagery in concept maps? This is an empirical question. The concept map research community should seek to determine experimentally whether rich media and other enhancements discussed here affect concept mapping in positive pedagogical and affective ways. That is, we need to determine more clearly whether further convergence of the state of the art and state of the practice is in fact beneficial. Of course, if so, we must continue to move forward toward full convergence.

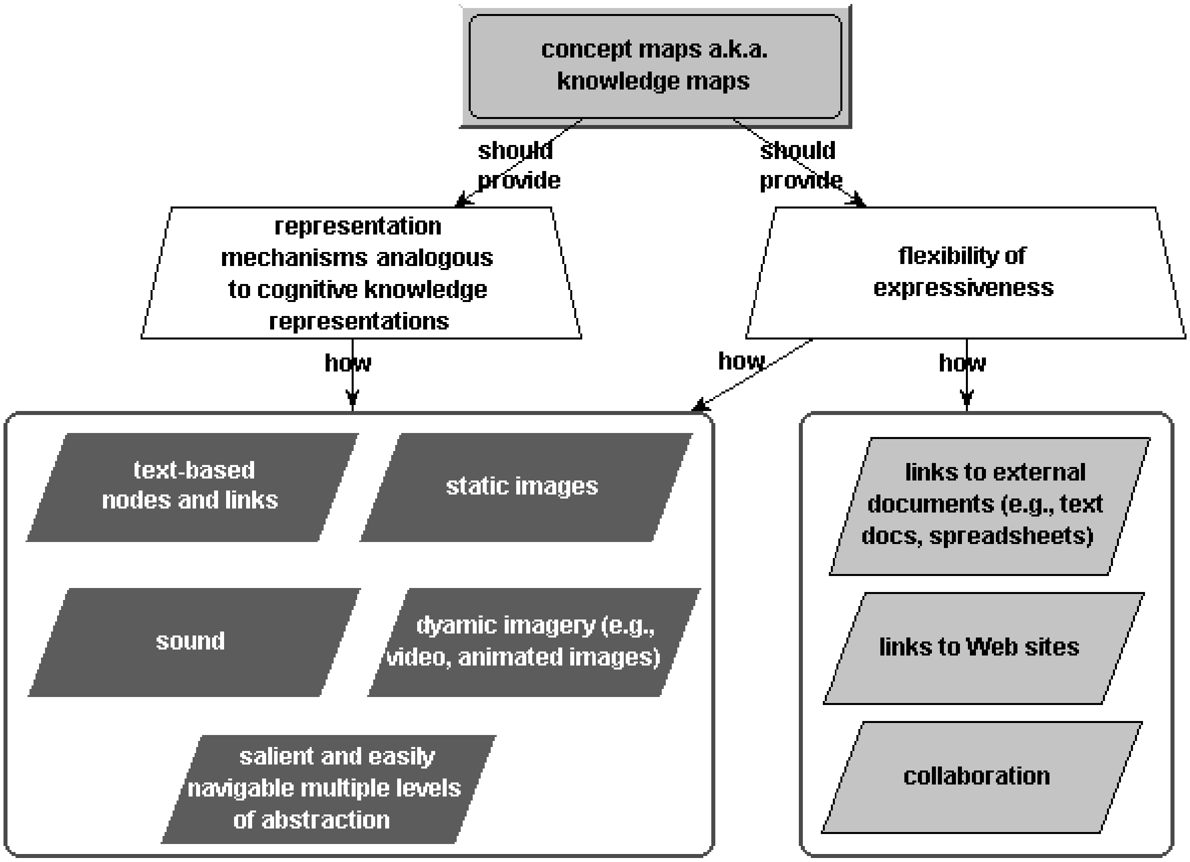

Figure 6. Types of representations and functionality available in state of the art concept map tools.

References

Alpert, S.R. (in press). Abstraction in Concept Map and Coupled Outline Knowledge Representations. *Journal of Interactive Learning Research.*

Alpert, S.R. & Grueneberg, K. (2001). Multimedia in Concept Maps: A Design Rationale and Web-Based Application. In *Proceedings of Ed-Media 2001: World Conference on Educational Multimedia, Hypermedia and Telecommunications*.

Alpert, S.R. & Grueneberg, K. (2000). Concept Mapping with Multimedia on the Web. *Journal of Educational Multimedia and Hypermedia*, *9*(4), 313-330.

Bomus, A. (2001). Personal communication.

Cicognani, A. (2000). Concept mapping as a collaborative tool for enhanced online learning. *Educational Technology & Society*, *3*(3), 150-158.

Follansbee, S., Hughes, R., Pisha, B., & Stahl, S. (1997). The Role of Online Communications in Schools: A National Study. *ERS Spectrum*, *15* (1), 15-26. Also http://www.cast.org/udl/OnlineRep.RTF.

CompuServe (1990). CompuServe Corporation GIF89a specification, http://www.w3.org/Graphics/GIF/spec-gif89a.txt.

Conklin, J. (1987). Hypertext: an introduction and survey. IEEE Computer, 20(7), 17-41

Dowling, W.L. & Harwood, D.L. (1986). *Music Cognition*. Orlando, FL: Academic Press.

Fisher, K.M. (1992). SemNet: A tool for personal knowledge construction. In P.A.M. Kommers, D.H. Jonassen, and J.T. Mayes (Eds.), *Cognitive Technologies for Learning* (pp. 63-75). Berlin: Springer-Verlag.

Flores-Méndez, R.A. (1997). Java concept maps for the learning web. In *Proceedings of ED-MEDIA'97*. Also http://www.cpsc.ucalgary.ca/~robertof/publications/edmedia97.

Gaines, B.R. & Shaw, M.L.G. (1995). WebMap: Concept mapping on the Web. *World Wide Web Journal*, *1*(1) 171-183. Also: http://ksi.cpsc.ucalgary.ca/articles/WWW/WWW4WM/.

Gaines, B.R. & Shaw, M.L.G. (1995). Collaboration through Concept Maps. In *Proceedings of CSCL'95, Computer Supported Cooperative Learning conference*. Also: http://ksi.cpsc.ucalgary.ca/articles/CSCL95CM/.

Griffin, C.C., Simmons, D.C., & Kameenui, E.J. (1992). Investigating the effectiveness of graphic organizer instruction on the comprehension and recall of science content by students with learning disabilities. *Reading, Writing, and Learning Disabilities*, *7*(4), 355-376.

Heeren, E., & Kommers, P.A.M. (1992). Flexibility of expressiveness: A critical factor in the design of concept mapping tools for learning. In P.A.M. Kommers, D.H. Jonassen, and J.T. Mayes (Eds.), *Cognitive Technologies for Learning* (pp. 85-101). Berlin: Springer-Verlag.

Hiltz, S.R. (1983). Collaborative Learning in Asynchronous Learning Networks: Building Learning Communities, *Invited talk, Web98 Conference*. Also http://eies.njit.edu/~hiltz/collaborative_learning_in_asynch.htm

IHMC (2000). The Institute for Human & Machine Cognition (IHMC) Concept Map Software: A Knowledge Construction Toolkit. http://cmap.coginst.uwf.edu/.

Johnson-Laird, P.N. (1983). *Mental Models*. Cambridge, UK: Cambridge University Press.

Kaput, J. J. (1989) Linking Representations in the Symbol Systems of Algebra, in S. Wagner & C. Kieran (Eds.) *Research Issues in the Learning and Teaching of Algebr*a. Hillsdale, NJ: LEA.

Kosslyn, S.M. (1980). *Image and Mind*. Cambridge, MA: Harvard University Press.

Kozma, R.B. (1992). Constructing knowledge with Learning Tool. In P.A.M. Kommers, D.H. Jonassen, and J.T. Mayes (Eds.), *Cognitive Technologies for Learning* (pp. 23-32). Berlin: Springer-Verlag.

Kremer, R. & Gaines, B.R. (1996). Embedded Interactive Concept Maps in Web Documents. *Proceedings of WebNet'96: World Conference of the Web Society*. Also: http://pages.cpsc.ucalgary.ca/~kremer/webnet96/webnet_kremer.html.

Lambiotte, J.G., Dansereau, D.F., Cross, D.R., & Reynolds, S.B. (1989). Multirelational Semantic Maps. *Educational Psychology Review, 1*(4), 331-367.

Larkin, J. H. and Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, *11*, 65-100.

McAleese, R. (1992). Cognitive Tools: The Experience of CASP, NoteCards, SemNet. In P.A.M. Kommers, D.H. Jonassen, and J.T. Mayes (Eds.), *Cognitive Technologies for Learning* (pp. 77-83). Berlin: Springer-Verlag.

Paivio, A. (1986). *Mental Representations: A Dual Coding Approach*. New York: Oxford University Press.

Rich, E. (1983). *Artificial Intelligence*. NY: McGraw-Hill.

Sharples, M. (2000). *Disruptive Devices: Personal Technologies and Education*. Inaugural Lecture of Kodak/Royal Academy of Engineering Research Chair in Educational Technology, University of Birmingham. http://www.eee.bham.ac.uk/handler/documents/inaugural.pdf.

Shephard, R.N. (1967). Recognition memory for words, sentences, and pictures. *Journal of Verbal Learning and Verbal Behaviour*, *6*, 156-163.

Shum S.B. (1990). Real and Virtual Spaces: Mapping from spatial cognition to Hypertext, *Hypermedia*, *2*(2), 133-158.

Stoyanova, N. & Kommers, P. (2001). Learning Effectiveness of Concept Mapping in a Computer-Supported Collaborative Problem Solving Design. In *Proceedings of Euro-CSCL 2001, European Conference on Computer-Supported Collaborative Learning* (pp. 561- 568). Also http://www.mmi.unimaas.nl/euro-cscl/Papers/153.doc.

Acknowledgments. Webster uses the Java-based JHotDraw graphical editor framework. Great thanks to Erich Gamma for providing the JHotDraw source code.
Inspiration® is a registered trademark of Inspiration Software, Inc., http://www.inspiration.com/. Microsoft® is a registered trademark of the Microsoft Corporation. Adobe® Acrobat® Reader® are registered trademarks of

Adobe Systems Incorporated. Quick View Plus® is a registered trademark of Jasc Software, Inc. Shockwave™ Player and Flash™ Player, are trademarks of Macromedia, Inc. Java® is a registered trademark of Sun Microsystems, Inc.