Self-Regulated Learning with Graphical Overviews: When Spatial Information Detracts from Learning

Sarah Davies (sarah.davies@utah.edu)
Kirsten R. Butcher (kirsten.butcher@utah.edu)
Corey Stevens (coreybstevens@gmail.com)
Department of Educational Psychology, 1705 Campus Center Drive, MBH 327
Salt Lake City, UT 84112 USA

Abstract

Graphical overviews have been studied as a method to improve hypertext learning and digital search. Although previous studies have found learning benefits to graphical overviews of single hypertext, it is unclear if these benefits extend to online learning across multiple (independent) documents. Previous research also has found that graphical overviews facilitate domain focus during online search, but it has not been established whether these benefits are derived from the spatial organization of the graphic or its textual content. This research examined the impact of using graphical overviews organized either spatially (i.e., network view) or textually (i.e., outline view) during self-regulated online learning. Assessments focused on deep understanding of science concepts and the relationships between them. Results indicated that the outline view promoted deeper understanding of science concepts and fewer errors about the relationships between them. Implications are discussed for the design and implementation of instructional materials to support self-regulated learning.

Keywords: self-regulated learning; graphical representations; online learning; conceptual browsing; comprehension

Introduction

As individual learning tasks increasingly are performed in online environments (Graham & Metaxas, 2003), there is a strong need to understand how the format of different materials impacts successful self-regulated learning (Pintrich, 2000; Winne, 2001). Self-regulated learning refers to learning situations in which students themselves must organize and manage the learning task (Azevedo & Cromley, 2004); it can be contrasted with learning in structured environments such as intelligent tutoring systems, where the computer system typically chooses the problems and decides when the student has reached mastery and is ready to move on to new materials (Anderson et al., 1995).

When students work with online learning materials – for example, hypertext documents – the learning task is inherently self-regulated by virtue of non-linear links that allow the learner to choose a unique path through the digital content. Research has found that students have great difficulty in self-regulating their learning with hypermedia, often utilizing ineffective strategies during self-regulated learning tasks (Azevedo et al., 2008). Other research has demonstrated the potential of organizational materials to facilitate more effective self-directed learning in online environments. For example, graphical overviews have been found to facilitate learning when presented before students work with a hypertext document (Salmerón et al., 2009). However, it is unclear if graphical overviews will have similar facilitative effects in online environments with limited coherence between independent online resources (rather than within a single hypertext document).

There is some evidence that a graphical interface can facilitate learning with varied, independent online resources. Research studying the use of a graphically-organized interface for online browsing showed that it facilitated processing of domain information in a digital library environment when compared to a keyword search interface (Butcher, Bhushan, & Sumner, 2006). However, it remains unclear if results were driven by the spatial formatting of the graphical interface or its conceptual (textual) content.

This research investigates the effects of a graphical overview (presented as either a text-based outline view or a spatially-organized network view) on students' self-regulated learning with online digital resources drawn from an educational digital library.

Self-Regulated Learning with Hypermedia

When students are asked to self-regulate their learning from hypermedia, they often struggle to organize and process information in ways that support deep understanding (Azevedo et al., 2008). Although successful self-regulated learners engage in strategies such as planning and prior knowledge activation (Azevedo, Guthrie, & Seibert, 2004), students engaged in self-regulated learning with hypermedia frequently choose to prioritize their reading based upon personal interest or text location (Salmerón, Kintsch, & Cañas, 2006). Not surprisingly, this failure to attend to conceptual relationships and coherence in the domain can lead students to miss important semantic connections between ideas and to form a more shallow understanding of hypermedia content (Salmerón et al., 2006).

Students may need significant help – especially in activating prior knowledge, organizing knowledge, and processing conceptual relationships – in order to learn effectively with online content. One way to offer this support is to provide the student with useful organizational materials that can be used to guide study and learning. Graphical overviews, which illustrate high-level ideas and the relations between them for a given text or topic, provide one form of organizational materials that has been shown to support learning among students with low prior knowledge (Salmerón et al., 2009).

Graphical Overviews and Hypertext Learning

Salmerón and colleagues have examined the impact of graphical overviews as the method of navigation through a hypertext (Salmerón et al., 2009; Salmerón & Garcia, 2012). These graphical organizations provided students with freedom to choose navigational paths through the hypertext but organized the content that could be viewed across the hypertext using a conceptual overview of the content. Salmerón and Garcia (2012) found that providing young (sixth grade) learners with a graphical overview of a hypertext document improved their knowledge integration during a comprehension task. These results complement earlier findings which showed that providing students with a graphical overview before hypertext study led to increases in comprehension for undergraduate learners (Salmerón, et al., 2009). Salmerón and colleagues have proposed two potential explanations for the observed benefits of graphical overviews: first, graphical overviews may facilitate learning by providing an organizational framework to support online study; second, graphical overviews may facilitate active processing of difficult texts by providing a text macrostructure that frees up additional resources for comprehension processes.

If graphical overviews facilitate learning by providing learners with an organizational framework for domain knowledge, studying their effects within a single hypertext may underestimate their potential benefits. Whereas a single hypertext likely has an overall coherence and topical focus, self-regulated learning in more authentic online environments requires working across independent digital resources that may not be easily integrated. Thus, it is important to consider whether graphical overviews may facilitate learning when students work with multiple online resources (i.e., independent web pages and sites).

Graphical Overviews and Digital Search

There is some evidence that graphical overviews change learners' processing when engaged in learning tasks that require work with multiple online resources. Butcher, Bhushan, and Sumner (2006) studied the impact of graphical overviews on students' search and evaluation processes as they attempted to locate useful online resources in an educational digital library. Students used either a graphical representation (a domain overview in the form of a node-link diagram) or a keyword interface to search for relevant digital content. Results showed that using the graphical representation as a search interface increased the depth of domain-relevant processing. Whereas students who navigated digital resources using a keyword interface tended to focus on superficial features of the resources, students navigating the resources with the graphical interface focused on analyzing domain concepts. Changes in the depth of students' processing of digital resources does not provide direct evidence of deeper learning with these resources; however, novice learners engaged in educational search tasks likely are engaged in "search to learn" processes which include iterative rounds of cognitive processing and interpretation (Marchionini, 2006). Recent research (Butcher et al., 2011) has confirmed the impact of graphical overviews on digital search and evaluation: when graphical representations were used as the basis for preservice teachers' navigation of resources in an educational digital library, students were more likely to identify educationally-useful online content and to focus on domain-level content when evaluating a web page or site.

Format and Content of Graphical Overviews

Although Butcher and colleagues (Butcher et al., 2006; Butcher et al., 2011) have found clear evidence that graphical representations can impact the processes that students use during online search and the overall success of online searches during educational tasks, it remains unclear whether these observed benefits were derived from the spatial format of the graphic (i.e., the spatial organization of the graphical overviews) or its (textual) domain content. Because keyword interfaces may require significant cognitive effort to generate relevant search terms (Marchionini & White, 2007), it is possible that the benefits of graphical overviews for self-regulated, online learning tasks may be derived from reallocation of cognitive effort from keyword generation to concept analysis. If this were the case, we would expect that removing spatial organization could facilitate even greater benefits by removing processing difficulty associated with examining and understanding spatial information.

If it is largely the textual content of graphical overviews that facilitates learning, more complex spatial formats actually may hurt novice learners. Graphical overviews in the form of a network map (see Figure 1) may depict interrelationships that are too complex for novice learners to understand. Novice learners may be better served by formats that emphasize organizational information in a hierarchical (i.e., linear) manner (see Figure 2). In a comparison of learning from linear and non-linear conceptual overviews, Amadieu and colleagues (2009) found that domain novices reported increased disorientation when learning from a network conceptual overview that depicted important relationships. In contrast, learners reported less disorientation and achieved better recall when learning with a hierarchical conceptual overview. Still, if it is true that graphical overviews promote learning by providing a conceptual framework for domain content, we would expect that a hierarchical graphical overview that removes spatial information would cease to be effective.

The current research extends prior research by examining two forms of graphical overviews during an online learning task: a spatially-organized *network* view vs. a textually-organized (linear) *outline* view. The use of these two conditions facilitates a direct comparison of whether the spatial format or the domain content of the graphical overviews has the greatest impact on learning outcomes. In addition, this research examines impact within a more authentic online environment, using the graphical overview to facilitate learning across a variety of independent online

resources. Because the network view is designed to demonstrate key conceptual relationships between multiple learning goals, we hypothesized that this graphical overview would facilitate greatest understanding of domain relationships.

Factual Knowledge vs. Deeper Understanding

When considering learning outcomes, it is important to recognize that comprehension research has established that different levels of knowledge can be formed during learning (Kintsch, 1998). In this work, we draw upon a well-known, established model of comprehension – Construction-Integration (CI) – that distinguishes between three levels of knowledge representation: the surface level, the textbase, and the situation model (Kintsch, 1994). A surface level representation is formed by encoding the specific details of a text (e.g., exact words and sentences). A textbase representation consists of the semantic meaning of a text; thus, a textbase representation drives recall of basic ideas derived from learning materials. The most flexible and durable knowledge representation is the situation model, which is formed when the learner integrates to-be-learned content with prior knowledge. A well-developed situation model drives inference, application, and transfer; as such, students who develop the situation model can be considered to *understand* materials rather than simply remember them.

The outcome assessments in this research target knowledge at the textbase and situation model levels. As described below, textbase assessments focus on factual knowledge learned during study and recalled during testing. Situation model assessments focus on students' application of learned knowledge, through explanation of concepts and relationships. Errors in student explanations, which may result from superficial reasoning about perceived relationships, also are examined.

Method

Participants

Twenty-six undergraduate students (8 males, 18 females, *M* age = 23) at a large public university in the western United States participated in this study in partial fulfillment of a class research requirement. One participant was excluded because his major was geology.

Design

This study utilized a two-condition, between-subjects experimental design. Participants were randomly assigned to one of the two experimental conditions upon arrival to the study.

Materials

Graphical Overviews The graphical overviews in this study were drawn from the Science Literacy Maps published on the National Science Digital Library (NSDL) website. NSDL is a digital library which seeks to provide access to up-to-date, high-quality, online resources in varied formats that will support education and learning in science, technology, engineering, and mathematics (Zia, 2000). The NSDL Science Literacy Maps are derived from strand maps developed by the American Association for the Advancement of Science (AAAS, 2001); these maps take the form of node-link diagrams. Nodes contain text that describe key learning goals in a topic area. The spatial organization of the nodes and the links between them demonstrate how student knowledge (as evidenced by the learning goals) should progress over time in a given domain.

In the NSDL, the Science Literacy Maps serve as a conceptual browsing interface (Zia, 2000); that is, the maps serve as a graphical search interface. To retrieve relevant digital resources using a conceptual search interface, users select a specific learning goal from the graphical overview (i.e., the Science Literacy Map). Clicking a learning goal brings up a small window that lists the NSDL-catalogued resources relevant to the conceptual information contained in the learning goal; much like a commercial search interface, each listed result provides users a title, a linked URL, and a short description of the resource.

Network Graphical Overview. The network view of the search interface utilizes the standard form of the Science Literacy Maps as found on NSDL.org. Learning goals are represented as nodes and are connected to one another with arrow links (see Figure 1); links between nodes indicate conceptual relationships between the learning goals. The overall spatial organization of the network indicates a more global knowledge organization, showing how learning goals develop over time, across grade levels and subtopics in the domain (see Figure 1).

Outline Graphical Overview. The outline view of the search interface contains the same node content as the network view. That is, all nodes contain the same text describing the same learning goals. However, in this view, the learning goal nodes are listed vertically rather than spatially. Learning goals in the outline view still are grouped by grade level (see Figure 2), but there are no links indicating conceptual relationships and spatial organization has been removed. As in the network view, clicking a learning goal in the outline view will bring up a window showing relevant resources catalogued in the digital library (see Figure 2). The learning goals in the outline view retrieved the same digital resources as in the network view (i.e., both interfaces searched over the same collection of digital resources and used the same algorithms to retrieve content relevant to each learning goal).

Reference Versions of Network and Outline Views

Before students used the graphical overview as a search interface to find online digital resources, they were given ten minutes to familiarize themselves with a non-interactive version of the graphic. The non-interactive forms of the graphical overviews utilized the same formatting and content as the interactive (search interface) versions of the graphical overviews as described above (see Figures 1 & 2).
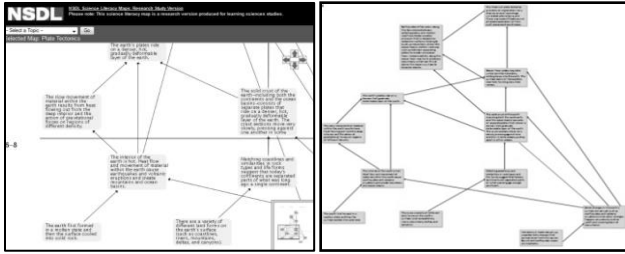
Figure 1: The network conceptual search interface is on the left. On the right is its associated non-interactive reference.
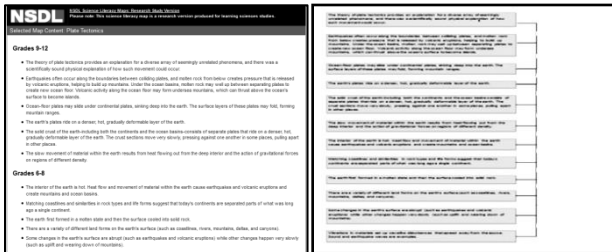

Figure 2: The outline conceptual search interface is on the left. On the right is its associated non-interactive reference.

Learning Assessments Learning assessments were administered at the beginning and end of each session. Questions tested participants' factual knowledge of plate tectonics, as well as their understanding of important plate tectonics concepts and relationships between them

Factual Knowledge. Factual knowledge items were designed to capture participants' textbase-level knowledge of plate tectonics. Factual items consisted of generative as well as non-generative (multiple choice and true/false) questions. Generative questions provided participants with images, such as a cross-section of the Earth, and asked them to generate labels for specific components or processes. Participants were asked to generate 13 diagram labels; correct labels received one point and partially-correct labels received half a point, for a total of 13 points. Non-generative questions tested students on their general knowledge (e.g., the number of Earth's tectonic plates). The non-generative factual assessment consisted of 33 items; participants received one point per correct item, for a total of 33 points.

Conceptual Understanding. Conceptual understanding items were designed to elicit participant explanations about key plate tectonics processes, thereby reflecting participants' situation models. These items asked students to interpret a diagram and explain the plate tectonics processes pictured. Conceptual understanding items were scored using a rubric that categorized explanations from most shallow to most deep, with a maximum of 5 points available per item. See Table 1 for examples of shallow, moderate, and deep answers. There were four conceptual understanding items, for a total of 20 points possible.

Table 1: Conceptual Explanation Examples

Shallow	It is showing the movement and direction in which Earth is moving caused by heat.
Moderate	The arrows are drawn in a circular pattern because that is how the convection heat current travels beneath the surface.
Deep	The rock in the mantle is heated up and due to its then lighter density rises to the surface where it is cooled because it is further away from the core and starts to become more dense and sinks. This process is repeated over and over again and is called convection.

Relationship Explanations. These items were designed to assess the depth with which students understood conceptual relationships between the learning goals. Relationship explanation items provided students with two distinct learning goals from the graphical overview and asked them to explain the relationship between the learning goals. This assessment presented students with 3 pairs of learning goals at pretest and 6 pairs at posttest. Relationship explanation items were scored as shallow or deep (see Table 2 for examples). Because novice learners often fail to identify and understand important relationships during learning, and because the conditions differed in the explicit portrayal of these relationships, the accuracy of relationship explanations was also examined. Explanations containing incorrect reasoning or mechanisms were marked as containing errors.

Table 2: Relationship Explanation Examples

Shallow	They both talk about the movement of the earth and what is causing the earth to move.
Deep	Because of heat flow and gravity, we see a pattern of movement within the earth's mantle (convection). The plates ride on the mantle, so this movement translates into the plates interacting with each other.

Procedure

To begin the study session, participants completed a brief survey which gathered demographic information. Next, the pretests were administered to assess participants' prior knowledge of plate tectonics. The learning task included 10-minute study of the reference version of the graphical overview (as appropriate to randomly assigned conditions), followed by forty minutes of learning with online digital resources as facilitated by the (condition-appropriate) graphical overview acting as the search interface. During online study, the reference version of the graphical overview was displayed on a second monitor so that participants could refer to it when reading/examining a digital resource. Following the learning task, posttest assessments were administered.

Analysis

As a check of random assignment, factual knowledge at pretest was analyzed using a MANOVA. Posttest learning assessment components also were analyzed using a repeated measures analysis of variance (RMANOVA) and a MANCOVA (see below). Alpha level was set at $p = .05$ for all analyses.

Results

Prior Knowledge of Plate Tectonics

A MANOVA for performance on both types of factual knowledge items at pretest did not show an overall condition difference ($F_{(2, 22)} = 2.50$, $p = .11$); however, univariate tests indicated that the two conditions did differ significantly on pretest diagram labels ($F_{(1, 23)} = 5.19$, $p = .03$). At pretest, the network overview condition correctly labeled a higher percentage of diagrams ($M = .29$, $SD = .19$) than the outline overview condition ($M = .15$, $SD = .10$). To control for the variance in learning due to prior knowledge, pretest performance in diagram labels was used as a covariate in a MANCOVA for posttest performance.

Posttest Performance on Learning Assessments

Factual Knowledge A RMANOVA was used to examine pre- and posttest performance on the non-generative factual knowledge items. Overall, participants showed a slight but significant learning gain from pre- to posttest ($M_{diff} = .05$; $F_{(1,23)} = 5.30$, $p = .03$) but there was no significant effect of condition ($F < 1$).

Conceptual Understanding & Relationship Explanations
A MANCOVA was used to examine posttest performance on measures of deep comprehension. There was a significant main effect of condition ($F_{(3, 20)} = 4.32$, $p = .02$). Univariate tests showed a main effect of graphical overview condition on conceptual understanding (see Table 3).

Students in the outline graphical overview condition produced conceptual explanations that evidenced deeper understanding of plate tectonics concepts ($M = .38$, $SD = .19$) than the network graphical overview condition ($M = .34$, $SD = .24$; $F_{(1, 22)} = 9.42$, $p < .01$). There also was a significant

Table 3: M and (SD) for Assessments of Learning

Assessment Scores (%)	Network	Outline
Factual Knowledge		
Non-generative (pretest)	.56(.09)	.53(.11)
Non-generative (posttest)	.62(.13)	.57(.06)
Conceptual Understanding*	.34(.24)	.38(.19)
Relationship Explanations		
% Deep Relationships	.17(.25)	.24(.25)
% Conceptual Errors*	.27(.17)	.14(.16)

* $p < .05$

condition difference in the percentage of errors when explaining relationships between plate tectonics concepts ($F_{(1, 22)} = 8.12$, $p < .01$). Students in the outline condition generated a smaller percentage of errors ($M = .14$, $SD = .16$) than students in the network condition ($M = .27$, $SD = .17$).

A non-significant but note-worthy trend was found in the percentage of deep explanations of relationships between concepts ($F_{(1, 22)} = 3.34$, $p = .08$). The outline condition produced a higher percentage of deep relationship explanations ($M = .24$, $SD = .25$) than the network condition ($M = .17$, $SD = .25$).

Discussion

After learning from multiple resources online, students in both conditions evidenced a similar increase in factual (text-base level) understanding of plate tectonics concepts. Overall, this is consistent with previous research finding that providing a graphical overview before hypertext study supports textbase comprehension (Salmerón et al., 2009). However, the current results also demonstrate that a spatially-organized graphical representation does not facilitate textbase learning more than a linearly-organized representation. Thus, it may be the textual content of the graphical organizer that facilitates macrostructure processing and leads to learning gains.

Although spatial format does not vary learning outcomes when considering factual (textbase-level) knowledge, it does impact the depth of understanding for important concepts and relationships between them. However, the pattern of results was opposite of hypothesized findings. Current results show that an outline graphical overview provided a learning advantage over a network (spatially-organized) overview: students learning with the outline view produced more deep explanations of science concepts and evidenced fewer erroneous ideas about inter-conceptual relationships.

This is a surprising result, since only the network view visually depicted the conceptual relationships among the learning goals. Indeed, previous studies have hypothesized that a schematic representation of relationships between concepts may provide novice learners with a framework for assimilating knowledge (Salmerón et al., 2009; Butcher et al., 2011). In this study, the spatial depiction of domain relationships compromised deep understanding. Concepts depicted in a network organization resulted in more errors when students explained conceptual relationships; students working with the network view also demonstrated less evidence of deep thinking about concepts. It may be that the graphic illustration of relationships actually precluded students from thinking deeply about the nature of those relationships. By explicitly depicting the conceptual relationships between learning goals, the network view may have caused students to generate fewer of their own inferences or predictions during learning. Alternatively, the network representation of content may have been too complex for novice learners. Previous research has found that students report feeling more disoriented with a network organization than with a more linear representation of

information (Amadieu et al., 2009). Because the network view did not specify the nature of potentially-complex relationships, students may have resorted to more shallow strategies of reasoning, integrating concepts based on superficial, easily-perceivable common features such as shared keywords.

When searching for information online, students typically learn from varied sources (Marchionini, 2006). Creating a deep, flexible understanding of the situation under investigation requires that self-regulated learners be able to synthesize multiple sources of information and integrate their learning with prior knowledge (Butcher & Kintsch, 2012; Perfetti, Rouet, & Britt, 1999). By demonstrating potential drawbacks to network-based graphical organizers during online learning, this study contributes an important initial finding to the literature on how to externally support self-regulated learning with multiple online resources. However, more research is needed to understand the specific relationship between the format of graphical overviews and their impact on learning outcomes.

Acknowledgments

References

AAAS (2001). *Atlas of Science Literacy, Volume 1*. Washington, DC, Project 2061, Association for the Advancement of Science (AAAS) and the National Science Teachers Association.

Amadieu, F., Van Gog, T., Paas, F., Tricot, A., & Mariné, C. (2009). Effects of prior knowledge and concept-map structure on disorientation, cognitive load, and learning. *Learning and Instruction*, *19*, 376-386.

Anderson, J. R., Corbett, A. T., Koedinger, K. R., & Pelletier, R. (1995). Cognitive tutors: Lessons learned. *Journal of the Learning Sciences, 4*(2), 167-207.

Azevedo, R., & Cromley, J. G. (2004). Does training on self-regulated learning facilitate students' learning with hypermedia? *Journal of Educational Psychology, 96*, 523–535.

Azevedo, R., Guthrie, J. T., & Seibert, D. (2004). The role of self-regulated learning in fostering students' conceptual understanding of complex systems with hypermedia. *Journal of Educational Computing Research, 30,* 87–111.

Azevedo, R., Moos, D. C., Greene, J. A., Winters, F. I., & Cromley, J. G. (2008). Why is externally-facilitated regulated learning more effective than self-regulated learning with hypermedia? *Educational Technology Research & Development*, *56*, 45–72.

Butcher, K.R., Bhushan, S., & Sumner, T. (2006). Multimedia displays for conceptual search processes: Information seeking with strand maps. *ACM Multimedia Systems Journal, 11*, 236-248.

Butcher, K.R., Davies, S., Crockett, A., Dewald, A., & Zheng, R. (2011). Do graphical search interfaces support effective search for and evaluation of digital library resources? *Proceedings of the Eleventh Annual International ACM/IEEE-CS Joint Conference on Digital Libraries* (pp. 315-324). New York, Association for Computing Machinery.

Butcher, K. R., & Kintsch, W. (2012). Text comprehension and discourse processing. In A.F. Healy & R.W. Proctor (Eds.), *Handbook of Psychology, Volume Four: Experimental Psychology* (*2nd Ed.*). Hoboken, NJ: Wiley.

Graham, L., & Metaxas, P. T. (2003). Of course it's true; I saw it on the internet! *Communications of the ACM, 46*(5), 71-75.

Kintsch, W. (1994). Text comprehension, memory, and learning. *American Psychologist, 49*(4), 294-303.

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

Marchionini, G. (2006). Exploratory search: From finding to understanding. *Communications of the ACM, 49*(4), 41-46.

Marchionini, G., & White, R. (2007). Find what you need, understand what you find. *International Journal of Human-Computer Interaction*, *23*(3), 205-237.

Perfetti, C. A., Rouet, J.-F., & Britt, M. A. (1999). Towards a theory of documents representation. In H. van Oostendorp & S. R. Goldman (Eds.), *The construction of mental representations during reading*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Pintrich, P. R. (2000). The role of goal orientation in self-regulated learning. In M. Boekaerts, P. R. Pintrich, & M. Zeidner (Eds). *Handbook of Self-Regulation*. San Diego: Academic Press.

Salmerón, L., Baccino, T., Cañas, J. J., Madrid, R. I., & Fajardo, I. (2009). Do graphical overviews facilitate or hinder comprehension in hypertext? *Computers & Education, 53*, 1308-1319.

Salmerón, L., & Garcia, V. (2012). Children's reading of printed text and hypertext with navigation overviews: The role of comprehension, sustained attention, and visuo-spatial abilities. *Journal of Educational Computing Research, 47*(1), 33-50.

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

Winne, P. H. (2001). Self-regulated learning viewed from models of information processing. In B. Zimmerman & D. Schunk (Eds.), *Self-regulated learning and academic achievement: Theoretical perspectives (2nd Ed).* Mahwah, NJ: Erlbaum.

Zia, L. L. (2000). The NSF National Science, Mathematics, Engineering, and Technology Education Digital Library (NSDL) Program: A progress report. *D-Lib Magazine*, 6: Available <http://www.dlib.org/dlib/october00/zia/10zia.html#ref08>.