

Understanding Knowledge Models: Modeling Assessment of Concept Importance in Concept Maps

David Leake, Ana Maguitman and Thomas Reichherzer

Computer Science Department, Indiana University
Lindley Hall 215, Bloomington, IN 47405, USA
{leake, anmaguit, treichhe}@cs.indiana.edu

Abstract

Concept mapping is widely used in educational and other settings to aid knowledge construction, sharing, and comparison; concept maps are also used as a vehicle for assessing understanding. To aid the concept mapping process, projects at Indiana University and the Institute for Human & Machine Cognition (IHMC) are developing “intelligent suggesters” to support users as they build concept maps, by presenting them with relevant information from existing knowledge models and the Internet. This depends on identifying important concepts in the concept map under construction. This paper presents and evaluates models of the influence of concept map layout and structure on the selection of concepts expected to be relevant to the topic of concept maps. It presents and assesses a set of potentially-relevant structural factors and evaluates how these factors combine to affect human judgments of concept importance. Twenty subjects were asked to judge the relative importance of concepts in concept maps selected to highlight particular characteristics, and three models were compared to their judgments. Analysis of the results shows that subjects were significantly influenced by concept map topology, but little influenced by other aspects of concept map layout. The results suggest that layout-independent models of concept maps can provide a suitable representation for guiding retrieval of topic-relevant information to support concept map construction, provided that the representation reflects topologically-based influences. The results are applied in the design of the suggesters’ similarity assessment procedures for retrieving relevant concept maps.

Introduction

Concept mapping [Novak and Gowin, 1984] has been widely used to elucidate humans’ knowledge and to facilitate knowledge elicitation, construction, and comparison and sharing. In concept mapping, users construct a two-dimensional, visually-based representation of concepts and their relationships. The concept map representation encodes propositions describing two or more concepts and their relationships, in simplified natural language sentences. In educational settings, concept mapping exercises have been used to encourage students to actively construct an understanding of concepts and relationships within domains of interest. To facilitate concept map construction and sharing, the Institute for Human and Machine Cognition (IHMC) has developed CmapTools, publicly-available tools to support generation and modification of concept maps in an electronic form (<http://cmap.ihmc.us/>). CmapTools enable interconnecting and annotating maps with material such as other concept maps, images, diagrams, and video clips, providing rich, browsable knowledge models available for navigation and collaboration across geographically-distant

sites. The CmapTools software has been downloaded by users in approximately 150 countries, and has been used in major educational initiatives, such as the Quorum project [Cañas et al., 1995], which involved more than one thousand schools in Latin America. It has also been used for modeling and sharing the knowledge of human experts, for example, for modeling NASA experts’ knowledge of Mars (<http://cmex-www.arc.nasa.gov/>).

CmapTools provides a convenient framework for knowledge construction, but users may have difficulty finding relevant resources, remembering specific aspects of a domain to include, or locating relevant concept maps to compare. To alleviate this problem, projects are under way at Indiana University and the IHMC to develop “intelligent suggesters” to support users by retrieving resources such as prior concept maps and multi-media materials [Leake et al., 2003]. Figure 1 shows a screenshot of a Mars knowledge model under construction, with suggestions of propositions, resources, and topics to consider. The suggesters’ effectiveness depends on their ability to retrieve topic-relevant information, which in turn depends on modeling users’ own judgments as they examine concept maps. Thus modeling users’ judgments of the importance of concepts to a map’s topic has practical value—for suggester software to support concept mapping—and scientific value, for better understanding what influences human understanding of the knowledge that concept maps convey.

The assessment of concept importance may depend on the concepts they include (based on their labels in the concept map), on the concept map topology, or on layout differences between isomorphic maps. Especially for users unfamiliar with a domain, we would expect topology and layout to play an important role in their assessment of the topic of a concept map. However, to our knowledge, no previous studies have investigated whether/how the topology and layout of a concept map actually influence judgments of its topic. To hypothesize candidate topological and layout factors that might influence decisions of which concepts are most topic-relevant, we considered general structure and layout guidelines for building “good” concept maps in the concept mapping literature, as well as methods for identifying important nodes from the structure of hyperlinked environments. These were used to develop candidate models for the influence of structural features on identifying the concepts most important to the topic of a concept map. We then performed experiments in which twenty paid subjects judged the relative importance of concepts in concept maps selected to investigate particular structural influences. We used this data to set pa-

Models for Analyzing Concept Maps

We developed four candidate models of the influence of structural and layout characteristics on expectations for the importance of particular concepts to the topic of concept maps. In the models, concepts are represented as nodes in the concept map graph. The baseline model treats map topology and layout as unimportant. The three remaining models use the topology of the concept map to compute a weight predicting each concept’s importance in describing the topic of the map.

To determine which factors to include in the models, we first considered factors from the concept mapping literature. Novak proposed that meaningful learning is facilitated when new concepts or concept meanings are subsumed under broader, more inclusive concepts, which suggests that concept maps should have a hierarchical structure. All of the non-baseline models can reflect such a structure, with weightings reflecting that important concepts are at the top of the map, and less important at the bottom. However, the models are parameterized so that the actual contribution of hierarchical structure—if any—can be determined empirically. We also considered the applicability of topological analysis methods from other domains, in particular, Kleinberg’s algorithm [Kleinberg, 1999] for topological analysis of graphs, used to identify important nodes in a hyperlinked environment. Kleinberg’s work characterized nodes on the World Wide Web as “hubs” and “authorities” based on their interconnections. When applied to concept maps, we expected hub and authority concepts to be especially important to determining the topic of concept maps.

Connectivity Root-Distance Model (CRD)

The connectivity root-distance model is based on two observations. First, concepts that participate in more than one proposition, as indicated by their connectivity—the number of incoming and outgoing connections—may be more important in defining a map’s content than concepts with lower connectivity. Second, Novak argues that concept maps are best constructed if a “focus question” or a single root concept guides the selection of concepts and their hierarchical organization in the map. The root concept, typically located at the top of a map, tends to be the most general and inclusive concept and to specify the map’s topic. This suggests that concept importance may increase with proximity to the root concept.

The CRD model determines proximity by counting the number of direct links between the map’s root concept and a given concept. For example, in figure 2, the concept “masses of ice” has a connectivity of four (two outgoing and two incoming links) and a distance of one to the root concept “glaciers”. If concept k in a map has o outgoing and i incoming connections to other concepts and is d steps distant from the root concept of the map, then the weight assigned to k by the CRD model is

$$W(k) = (\alpha \cdot o(k) + \beta \cdot i(k)) \cdot (1/(d(k) + 1))^{1/\delta}$$

The model parameters α , β , and δ determine influence of the incoming connections, outgoing connections, and distance to the root concept. The formula implies that the higher a concept’s connectivity and the shorter its distance to the root con-

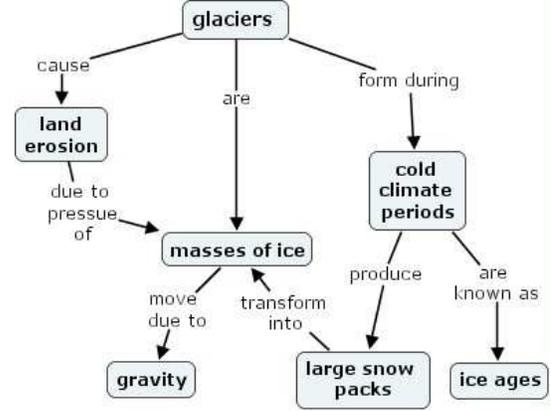


Figure 2: A simple concept map about glaciers.

cept, the larger its weight and therefore relevance in the topic of the map.

Hub Authority and Root-Distance Model (HARD)

The Hub Authority and Root-Distance Model also explores the importance of the root node and the hierarchical organization of concepts in maps. However, while CRD performs a local analysis, only taking immediate neighbors into account, HARD performs a global analysis on the influences of the concepts on each other. Its analysis centers on three different types of concepts that may be found in a concept map:

- *Authorities* are concepts that have multiple incoming connections from hub nodes.
- *Hubs* are concepts that have multiple outgoing connections to authority nodes.
- *Upper* nodes include the root concept and concepts closest to the root concept.

To determine a node’s role as a hub or authority, we adapted Kleinberg’s algorithm for analyzing hyperlinked graphs to concept maps. Our algorithm, described in detail in [Cañas et al., 2001], associates each concept with three weights between 0 and 1, each reflecting the concept’s role as a hub, authority, or upper node. A given concept may simultaneously have properties of all three, but in Figure 2, “glaciers” is primarily a hub concept, due to the number of outgoing connections, and “masses of ice” is primarily an authority, due to its mostly incoming connections. Among the three concepts with outgoing links to the concept “masses of ice”, “glaciers” is the one with the greatest influence in making “masses of ice” an authority node, because of the comparative strength of “glaciers” as a hub.

In the HARD model, the three weights of a selected concept k are combined into a single weight as follows:

$$W(k) = (\alpha \cdot h(k) + \beta \cdot a(k) + \gamma \cdot u(k))$$

In the above formula a , h , and u are the corresponding authority, hub, and upper node weights of a concept in a map and α , β , and γ are the model parameters. As above, the parameters reflect the influences of the different roles that a concept may play.

Path Counter Model (PC)

The Path Counter Model, like the CRD model, reflects the expectation that concepts participating in more propositions will tend to be more important to the topic of a map. However, instead of considering only a concept node's immediate connectivity the PC model considers indirect relationships as well. It counts all possible paths, starting from the root concept, that contain the concept in question and either (1) end on a concept with no outgoing connections, or (2) end on a concept that has already been visited in a path. We note that if a concept has high connectivity (which allows for many paths to form in the map), then the number of paths crossing a concept also increases for concepts indirectly linked to the high-connectivity concept. For example, the PC value for the concept "gravity" in figure 2 is three, because there are three paths extending from the root concept to "gravity," due to "masses of ice" which is well connected in the map. Formally, to determine the weight $W(k)$ of a concept k in a map, assume that n is the number of paths crossing k . Then the weight is computed as $W(k) = n$. Unlike the previous two models, this model considers only a single influence on concept weight, and consequently requires no parameters.

Experiments and Results

We conducted a human-subjects experiment to study the influences of the hypothesized factors on human judgments of concept importance, and the overall fit of the four models' predictions to human judgments, with the parameter settings that best fit the CRD and HARD models to the subject data.

Method

Twenty paid subjects, all students admitted to Indiana University, were recruited by postings on electronic message boards and bulletin boards for a one-hour experiment conducted on the Web. The experiment was divided into a training phase, to familiarize participants with the study and to provide background information on concept maps, and a test phase. In the training phase, participants were given a brief description of concept maps and their applications, and then asked to write a short summary of two concept maps from different domains. In the test phase, subjects answered 56 questions about a total of 12 small concept maps (fewer than 15 concepts each). The maps were designed with controlled differences in their topological structure and layout, to investigate the presence or absence of influences from particular types of changes (e.g., changing position of a node without affecting topology). Each question presented a concept map and two concepts selected from that map. Participants were asked to examine a map and to answer which of the two concepts best described the map's topic, or whether both described it equally well.

To allow participants to first practice decision making on regular concept maps, the first 2 of the 12 concept maps used regular words in the concepts. To prevent domain knowledge from influencing participants' decisions, concept labels were replaced with artificial terms in the remaining 10 maps, and only responses concerning the latter 10 test maps were used in evaluating the models. The use of artificial terms as labels, the topological and layout changes between the concept maps, and randomization of the order of options to answer a

question were all done to ensure that the participants made their choice independently of the concept maps they have already examined.

The concept maps in the experiment were designed to test specific hypotheses about the topological and layout factors that may influence subjects' evaluation of relevance of concepts to a concept map's topic. Because domain knowledge is absent, evaluations had to rely entirely on topology and layout.

Results

To test whether subjects' judgments of the importance of two concepts changed significantly from one map to another, we used a χ^2 test of independence when comparing the subjects' selections from two different maps. Table 1 summarizes the statistical results.

Distance to root concept: To test the influence of distance to the root concept, subjects evaluated two concept maps in which the distance from a test concept to the root concept was changed from 2 to 1, by inserting an intermediate node. In a series of questions, subjects were asked to compare importances of the test concept, which was moved in the map's hierarchy, to the root concept and neighboring concepts of the moved concept. The results show that the root concept was considered most important compared to the other concepts, and that the importance of the test concept increased as it moved up the hierarchy. The differences in the selection of the moved concept over its neighboring concepts between the two concept maps were statistically significant.

Connectivity of a concept: To test the influence of connectivity, we used two concept maps which differed by increasing a test concept's connectivity—the number of incoming and outgoing connections to neighboring concepts—from 1 in the first map to 6 in the second. Subjects were asked to compare importances of the test concept to the root concept and the neighboring concepts of the modified concept. When the test concept's connectivity was increased, participants favored it over neighboring concepts and sometimes even over the root concept. All differences were statistically significant except for the preference over the root concept.

Layout of a map: To test whether a difference in layout affects subject's selections, two concept maps were constructed with identical topology but substantially different layout. The layout changes primarily involved horizontal organization, but in one instance a single concept was moved from the center right to the bottom left position. The questions asked for both layouts compared the concept that changed its position to its neighboring concepts. The statistical evaluation revealed that the layout changes had no significant affect on the concept ratings.

Direct and indirect influences of hub and authority nodes in a map: To test the effects of direct and indirect influences, a total of four concept maps were constructed with strong hub and authority concepts connected to other concepts in the map. The results showed that hub and authority concepts have an influence on the selection of concepts, and that authorities play a stronger role than hubs. However, the indirect influence of either a hub or authority concept on other

Influence	Significant	χ^2 Test of Independence
distance to root concept	yes	$(1, N = 40) = 17.04, p < 0.05$
concept connectivity	yes	$(1, N = 40) = 19.37, p < 0.05$
map layout	no	$(1, N = 40) = 0.23, p > 0.05$
direct, hub concept	yes	$(1, N = 40) = 7.74, p < 0.05$
direct, authority concept	yes	$(1, N = 40) = 15.82, p < 0.05$
indirect, hub concept	no	$(1, N = 40) = 3.73, p > 0.05$
indirect, authority concept	no	$(1, N = 40) = 3.73, p > 0.05$

Table 1: Statistical evaluation of influences on concept importance.

Model	Parameters for Best Fit			RMSE	Cumul. Error
	α	β	γ / δ		
CRD	0.930	4.959	3.603	0.072	27.5%
HARD	0	2.235	1.764	0.1487	32.8%
PC	N/A	N/A	N/A	0.170	27.8%
Baseline	N/A	N/A	N/A	0.564	66.8%

Table 2: Summary of model parameters and RMSE.

concepts (when a hub or authority is indirectly connected to a test concept) did not significantly affect concept importance.

Fitting the Models to the Data

A hill-climbing algorithm was used to determine the parameter settings for the CRD and the HARD models which gave the best fit between the models and user data. Table 2 summarizes the chosen parameter values, the root-mean-square error (RMSE) of user and model data, and the cumulative error. The cumulative error is the percentage of the total questions (44 questions per subject, involving the 10 test concept maps) for which the models determine different responses from the subjects. To determine a model’s preference between two concepts in a concept map, we compared the model’s importance values for the two nodes. The model was considered to treat the concepts as equally relevant when their relevance values were within a fixed threshold of each other, for a threshold distance determined by hill-climbing. The last row of the table shows the RMSE and the cumulative error for a baseline model. In this model each concept in a map is rated equally important by assigning it a weight of 1.

The results show that the CRD model provides the best fit to the user data, followed by HARD and PC. All models except the baseline agree with more than 67% percent of the decisions reached by the participants, who were in a few cases strongly divided in their vote for the best topic-describing concepts. For the remaining 33%, in most cases the models’ predictions match the decisions of some subjects. Only once for the CRD model, twice for the HARD model, and four times for the PC model were model and user predictions entirely disjoint. Overall, CRD, HARD, and PC perform better than the baseline model.

Further analysis of the best-fit parameters for the CRD and HARD models supports the importance of authority nodes (nodes with incoming connections). For the CRD model, nodes with incoming connections (nodes that play the role of an authority) are more relevant than nodes with outgoing connections (nodes that play the role of a hub) because their β

is greater than α . With the best-fit parameters for the HARD model, hub nodes are not considered relevant when computing the weight of a node. However, we note that hub nodes still play an important role when computing the level of authority of other nodes in the map.

Discussion

The experiments studied how topology and layout affect assessments of the importance of concepts within concept maps. They compared four candidate models which, using only analysis of a map’s topology, compute a weight for each concept in a map. The computed weights provide an estimate of the importance of each concept as a descriptor of the topic of the map, according to subjects’ judgments of topic importance.

The studies highlighted the importance of topological information; to our knowledge, this is the first study to show this effect. They also suggested that specific layout does not have a significant effect. This is important for being able to recognize similarity across concept maps developed by different individuals, despite superficial differences that might affect user judgments. It is also interesting to note that despite the importance of topology, local information alone was sufficient to account for the observed results. The CRD model, which considers distance from the root node and local connectivity, outperformed the more sophisticated HARD model, which takes indirect influences into account as well.

The current experiment studied small concept maps, and considered only the topological and layout factors of the maps, rather than their content. We are conducting additional studies to explore the role of content in assessments of concept importance. However, preliminary results suggest that structure plays a surprisingly strong role, with structural information alone often sufficient to make high-quality predictions.

Application in the Suggesters

The experimental results are reflected in the design of the CmapTools suggesters, two of which are shown in use in the lower center of Figure 1. The first suggester uses the calculated importance values to weight keywords from concept labels in a concept map, in order to retrieve similar prior concept maps for comparison and to suggest propositions from those maps. This approach to supporting concept map generation is inspired by case-based reasoning [Kolodner, 1993]; concept maps constructed by different users are considered as case-bases of their concept-mapping activity, with each concept map considered to be a separate case. When a user wants

to “extend” a concept—to add a new connected concept—the system draws upon prior concept maps that include the original concept, as examples of how that concept was extended in similar past contexts. The second suggester uses the similarity weighting to weight keywords for Web search, to derive topics for the user to consider when starting a new concept map to broaden the knowledge model. These and other implemented suggesters are described in detail in [Leake et al., 2003].

Conclusion

This paper explores factors affecting human judgments of concept importance in determining the topic of concept maps. Modeling such judgments helps elucidate the knowledge captured in concept maps and aids the development of intelligent support systems to provide relevant material during concept mapping. Our experiments assessed the influence of specific factors and examined the ability of four different models to reflect human assessments of concept importance.

Among the three models, the CRD model, which considers connectivity and distance to the root concept, provided the best match to human data: Its predictions were consistent with the average predictions made by the participants for forty-three out of forty-four questions. The results highlight the importance of local topology and suggest that human topic decisions are robust to layout differences, which is encouraging for the generality of concept mapping for knowledge sharing and the development of support tools to retrieve similar concept maps and topic-relevant information. We are performing followup studies to examine the role of domain content and the fit between the predictions of these models and the concept maps developed by domain experts for sample domains.

Principles suggested by the results have been applied to “intelligent suggesters” to aid the human knowledge modeling process, and the implemented systems appear to give good results in practice. We consider the type of evaluation presented here as important step for guiding the design of such tools, and are now designing experiments to more formally test the relevance of the suggester systems’ recommendations during the concept map construction process.

Acknowledgments

This research is supported by NASA under award No NCC 2-1216. We thank Alberto Cañas and the IHMC CmapTools development team for their many contributions to the project, Tei Laine for her comments on a draft of this paper, and John Kruschke for assistance on experimental design and analysis.

References

- [Aidman and Egan, 1998] Aidman, E. and Egan, G. (1998). Academic assessment through computerized concept mapping: validating a method of implicit map reconstruction. *Int. Journal of Instructional Media*, 25(3):277–294.
- [Ausubel, 1963] Ausubel, D. (1963). *The Psychology of Meaningful Verbal Learning*. Grune and Stratton, NY, NY.
- [Cañas et al., 1995] Cañas, A., Ford, K., Brennan, J., Reichherzer, T., and Hayes, P. (1995). Knowledge construction and sharing in quorum. In *World Conf. on Artificial Intelligence in Education, AIED’95*, pages 218–225. AACE.
- [Cañas et al., 2001] Cañas, A., Leake, D., and Maguitman, A. (2001). Combining concept mapping with CBR: Experience-based support for knowledge modeling. In *Proceedings of the Fourteenth Int. Florida Artificial Intelligence Research Society Conf.*, pages 286–290. AAAI Press.
- [Collins and Quillian, 1969] Collins, A. M. and Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of verbal learning and verbal behavior*, 8:240–248.
- [Esposito, 1990] Esposito, C. (1990). A graph-theoretic approach to concept clustering. In *Pathfinder associative networks: studies in knowledge organization*, pages 89–99. Ablex Publishing Corp.
- [Ganter and Wille, 1999] Ganter, B. and Wille, R. (1999). *Formal Concept Analysis: Mathematical Foundations*. Springer, Berlin - Heidelberg - New York.
- [Gärdenfors, 2000] Gärdenfors, P. (2000). *Conceptual Spaces: The Geometry of Thought*. MIT Press.
- [Goldsmith and Davenport, 1990] Goldsmith, T. E. and Davenport, D. M. (1990). Assessing structural similarity of graphs. In *Pathfinder associative networks: studies in knowledge organization*, pages 75–87. Ablex.
- [Kleinberg, 1999] Kleinberg, J. (1999). Authoritative sources in a hyperlinked environment. *J. of the ACM*, 46(5):604–632.
- [Kolodner, 1993] Kolodner, J. (1993). *Case-Based Reasoning*. Morgan Kaufmann, San Mateo, CA.
- [Leake et al., 2003] Leake, D., Maguitman, A., Reichherzer, T., Cañas, A. C., Carvalho, M., Arguedas, M., Brenes, S., and Eskridge, T. (2003). Aiding knowledge capture by searching for extensions of knowledge models. In *Proceedings of the Second Int. Conf. on Knowledge Capture (K-CAP)*, New York. ACM Press, pages 44–53.
- [Michael, 1994] Michael, R. S. (1994). *The Validity of Concept Maps for Assessing Cognitive Structure*. PhD thesis, School of Education, Indiana University.
- [Novak and Gowin, 1984] Novak, J. and Gowin, D. (1984). *Learning How to Learn*. Cambridge University Press, NY.
- [Quillian, 1968] Quillian, M. R. (1968). Semantic memory. In Minsky, M., editor, *Semantic Information Processing*, pages 216–270. MIT Press.
- [Sowa, 1984] Sowa, J. F. (1984). *Conceptual Structures: Information Processing in Mind and Machine*. Addison-Wesley.
- [Tulving, 1972] Tulving, E. (1972). Episodic and semantic memory. In *Organization of Memory*. Academic Press.
- [West et al., 2002] West, D., Park, J., Pomeroy, J., and Sandoval, J. (2002). Concept mapping assessment in medical education: a comparison of two scoring systems. *Medical Education*, 36(9):820–826.