Concept Maps: Theory, Methodology, Technology
Proc. of the First Int. Conference on Concept Mapping
A. J. Cañas, J. D. Novak, F. M. González, Eds.
Pamplona, Spain 2004

"GOOGLING" FROM A CONCEPT MAP: TOWARDS AUTOMATIC CONCEPT-MAP-BASED QUERY FORMATION

David Leake, Ana Maguitman, Thomas Reichherzer, Indiana University, U.S.A.
Alberto J. Cañas, Marco Carvalho, Marco Arguedas, Tom Eskridge, Inst. for Human & Machine Cognition, U.S.A.
Email: {leake, anmaguit, treichhe}@cs.indiana.edu; {acanas, mcarvalho, marguedas, teskridge}@ihmc.us

Abstract. Electronic concept mapping tools provide a flexible vehicle for constructing concept maps, linking concept maps to other concept maps and related resources, and distributing concept maps to others. As electronic concept maps are constructed, it is often helpful for users to consult additional resources, in order to jog their memories or to locate resources to link to the map under construction. The World Wide Web provides a rich range of resources for these tasks—if the right resources can be found. This paper presents ongoing research on how to automatically generate Web queries from concept maps under construction, in order to proactively suggest related information to aid concept mapping. First, it examines how concept map structure and content can be exploited to automatically *select terms to include in initial queries*, based on studies of (1) how concept map structure influences human judgments of concept importance, and (2) the relative value of including information from concept labels and linking phrases. Second, it examines how a concept map can be used to *refine future queries* by reinforcing the weights of terms that have proven to be good discriminators for the topic of the concept map. The described methods are being applied to developing "intelligent suggesters" to support the concept mapping process.

1 Introduction

Concept mapping (Novak & Gowin, 1984) is widely used to enable individuals at many different levels—from elementary school students to scientists—to construct new knowledge, externalize knowledge in a human-usable form, and share and compare that knowledge to advance human learning and understanding. To facilitate this process, the Institute for Human and Machine Cognition (IHMC) has developed CmapTools (Cañas *et al*. 2004a), a suite of tools to support generating and modifying concept maps in electronic form, interconnecting those maps, and annotating them with additional material such as images, diagrams, and video clips to develop rich browsable knowledge models. CmapTools provides a convenient framework for building and sharing multimedia concept maps and organizing them into linked knowledge models. A recent research focus is to augment CmapTools with support methods to aid the users' knowledge construction process. Developing tools to guide Web search during concept map construction is an important part of this effort: The Web provides an enormous body of information, but finding the right information may be difficult. Carvalho, Hewett and Cañas (2001) presented methods for exploiting the propositional and hierarchical nature of concept maps to improve filtering and ranking of results from search engine queries generated by users while browsing or constructing concept maps. Those methods have been successfully applied in a CmapTools search aid which has proven popular with users: It enables users to initiate Web searches for information that is related to a concept within the context of a concept map. This paper presents continuing research aimed at developing proactive "intelligent suggesters" to automatically provide information during concept map generation and access (e.g., Leake *et al*. 2003). In particular, it describes methods for automatically generating queries from concept maps, in order to provide users with related information as they build knowledge models. The methods identify important aspects of a concept map under construction and exploit this information (1) to generate queries to Web search engines, (2) to automatically refine those queries, based on returned results and the current concept map, and (3) to filter returned results. The methods exploit the structure of concept maps, combining structural information with simple term-based techniques in a robust framework that can be applied without requiring additional domain knowledge or full natural language processing. Initial results from this effort have been incorporated into CmapTools (Cañas *et al*., 2004b).

The paper begins with an overview of the CmapTools software and its built-in methods for enhancing Web searches. It then describes recent research results on developing extended methods for automatic query formation based on concept maps. Section 3 studies factors affecting how information from a concept map should be used for initial query generation. Section 4 describes our approach to automatically refining queries by exploiting concept map information, and Section 5 describes how the pieces are combined into an end-to-end system for automatic, concept-map based query generation. This work provides a toolbox of approaches, which, taken together, provide an end-to-end approach to automatic querying based on concept maps.

2 CmapTools and the CmapTools Search Enhancer

CmapTools is a suite of tools for generating and sharing concept maps in electronic form. CmapTools supports generating and modifying concept maps, as well as adding navigational links from concepts to other concept maps and multi-media material such as images, diagrams, and video clips, enabling the construction of rich knowledge models. The tools facilitate the storage and access of concept maps on multiple servers, providing the network services required to support knowledge sharing across geographically-distant sites. The CmapTools network tracks additions of new knowledge models as they become available, and users can search for information to refine and extend their concept maps, or to link their concept maps to relevant shared concept maps and multi-media material. In addition, integrated collaboration tools empower groups of users to construct concept maps simultaneously or to engage in discussions of the information in the concept map by posting notes or by opening discussion forums. Figure 1 shows a knowledge model displayed by CmapTools. The tools are freely available for nonprofit use, and have been downloaded by users in over 130 countries. Concept mapping with CmapTools has been used successfully to facilitate learning in educational settings, for capture and management of expert knowledge, and for just-in-time training. Full information is available from http://cmap.ihmc.us.

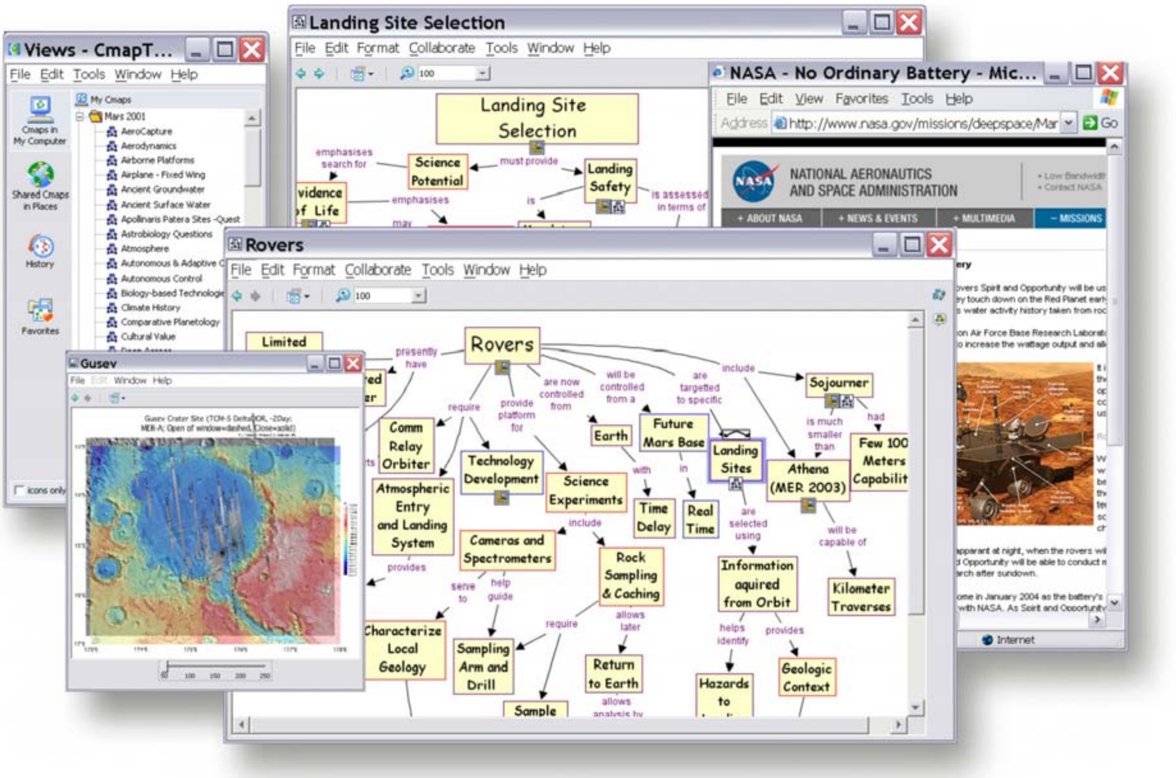

Figure 1: Knowledge model on the Mars domain

During concept mapping, access to appropriate Web-based resources can play a valuable role in aiding concept map construction. A current effort in the CmapTools project develops “intelligent suggesters” to aid in tasks such as *knowledge expression*, e.g., choosing terms to use in a concept map; *knowledge connection*, e.g., determining propositions to include or images, video clips, or other resources with which to annotate the nodes in a concept map; *knowledge communication*, e.g., presenting relevant concept maps to users developing related knowledge models; k*nowledge comparison*, to aid users in focusing on differences between their own and others’ conceptualizations; and *topic selection*, to determine which additional concept maps to construct.

CmapTools provides a search tool, which takes queries based on a concept map—either under construction, or being browsed—and searches the Web (and/or CmapServers) for information related to the map. The user can easily and concisely specify the context of the search in a concept map, and that context is used for the automatic construction of queries. The Web-search algorithm allows the user to select a concept and ask the system to search for Web information that is relevant to the concept within the context of the concept map. The process consists of: (a) Analyzing the concept map to prepare a relevant query to use in searching the Web, (b) Retrieving relevant documents from the Web, (c) Ranking the retrieved Web pages according to relevance, and

(d) Presenting the results to the user. To generate the query, key concepts are selected from the map to reflect the domain of the map. These include the selected concept itself, the root of the concept map, and the map's authority nodes, nodes with the greatest number of incoming links from other nodes (a detailed explanation on selecting key concepts from a concept map is presented in Section 4 below). Based on the query, the tool retrieves Web pages to build a collection of documents. We have developed a meta-search engine, based primarily on Google (Brin & Page, 1998), to retrieve an initial set of documents from the Internet. Once retrieved, these documents are added to a cache for ranking. The ranking is based on comparing distance matrices calculated from the concept map and from each of the candidate documents. The purpose of the distance matrices is to favor documents which may contain propositions similar to the ones found in the map and those containing concepts that appear close to the root of the map. The top ranked documents are presented to the user. In a human subjects experiment, subjects judged the algorithm's rankings superior to those of traditional search engines, highlighting the value of exploiting contextual information to rank search results (Carvalho, Hewett and Cañas 2001). In practice, the tool has proven popular with users.

3 Using Concept Maps to Guide Automatic Query Generation and Refinement

Search engines provide access to a vast and ever-growing repository of information on the World Wide Web. However, finding relevant information remains challenging, because of the need to formulate queries to select on-point resources from the enormous range of possibilities. Users often generate very short queries (in a study of over a million queries to the Excite search engine, Spink *et al*. (2001) found that 60% of queries were one or two terms long), making query ambiguity a serious problem. Users may be inexperienced in selecting suitable keywords, may not know enough about the domain to select good query terms, or may simply overlook useful keywords. Some research addresses these problems by automatically augmenting user queries based on the task context, e.g., to make related suggestions as users read or write a document (e.g., Budzik *et al*., 2001). Our hypothesis is that the rich structure of concept maps can be exploited to automatically select effective keywords for Web search queries. Our research focuses on how to automatically generate queries for information related to a concept map under construction, and then to automatically refine those queries by analyzing the returned results in context of the original concept map. Our research strategy is to examine particular facets of the process, and the factors that affect them, in order to fill in pieces of the puzzle of how to exploit the nature of concept maps to guide search.

Specifically, we are studying three questions: (1) Can concept map structure be used to identify important concepts in a concept map?, (2) Are labels of those concepts sufficient to retrieve information related to the map?, and (3) Can query terms be refined automatically, based on analysis of their results and the maps from which they were generated? Because the goal of the retrieval process is to present *useful* suggestions, the ideal method for evaluating result quality would be an end-to-end evaluation, in which subjects directly assessed the usefulness of system suggestions. However, to guide the incremental development of the methods, it is crucial to be able to assess incremental steps for which human-subjects evaluations would be impractical. Consequently, the evaluations reported here use modified versions of information retrieval metrics, designed to test properties which we hypothesize to correlate with usefulness in practice. Guided by the results of these initial experiments, we have implemented a prototype system, and we are currently developing human subjects experiments to directly evaluate the usefulness of the full system. Results of a human-subjects study of a "concept suggester" implementing initial results from this research effort are reported in Cañas *et al*.(2004b).

4 Using Concept Map Topology to Weight Concept Terms

Web search engines limit query size (e.g., Google's query size limit is 10 terms). Thus a challenge for automatic query generation is to generate short queries focusing on the most important terms in a concept map. To achieve good recall, some of the terms should be selected for their quality as descriptors of the topic of the map. To improve precision, others should be selected for their ability to discriminate between documents. Ideally, topic descriptors will reflect the information that users will consider most relevant. This is hard to assure, however, because user importance judgments may depend on many factors, such as the structure of the map, the specific content of its concept labels and linking phrases, or the user's focus of attention when adding concepts to a map under construction. To develop robust methods that can be applied without knowledge engineering or full natural language processing, we have studied whether concept map structure can be used to predict concept importance, independent of the content questions which would require "understanding" the concept map.

To study how concept map structure influences human predictions of concept importance, we developed three candidate models of how structure affects concept importance (Leake *et al*. 2004). Some of the factors considered by our models are inspired by general guidelines for constructing “good” concept maps, taken from the concept mapping literature (Novak & Gowin, 1988), e.g., reflecting the importance of concept maps’ hierarchical structure by weighting upper and lower concepts differently. Others are inspired by methods for analyzing the topology of hyperlinked network structures (Kleinberg, 1999), e.g., that nodes may be characterized based on outgoing and incoming connections as “hubs” or “authorities,” where, e.g., authority nodes tend to have many incoming links. This view suggests the hypothesis that hub and authority nodes may play a significant role in describing the map’s content.

Our models consider several structural influences on concept importance including (1) distance of a concept to a root concept, measured in terms of the number of links on the shortest path to the root concept in the concept map graph, (2) connectivity of a concept measured in terms of the number of incoming and outgoing connections, and (3) the concept’s global connectivity to the root concept measured by a “path frequency” (PF) measure. In addition, models that consider multiple influences have parameters to weight the different structural influences. The three models, CRD, HARD, and PF, are summarized in Table 1. For CRD, the model parameters α, β, and δ adjust the effect of the number of incoming connections (i), the number of outgoing connections (o) and the distance to the root concept (d). For HARD, the model parameters α, β, and γ adjust the effect of authority (a), hub (h) and upper nodes (u) (nodes appearing near the top of the concept map). The values for authority, hub, and upper nodes correspond to the concept’s role as an authority, hub, and upper node, while the upper weight reflects proximity of a concept to the root concept (for full details, see (Leake *et al*. 2003)). PF counts the number of distinct paths that reach a given concept, traversing the paths of the concept map graph starting from the root concept (n), and requires no parameters.

Connectivity Root Distance (CRD)	W(C) = $\left(\alpha \cdot o + \beta \cdot i\right) \cdot \left(1/(d+1)\right)^{1/\delta}$, $\alpha, \beta \geq 0, \delta \geq 1$
Hub Authority Root Distance (HARD)	W(C) = $\left(\alpha \cdot a + \beta \cdot h + \gamma \cdot u\right)$, $\alpha, \beta, \gamma \geq 0$
Path Frequency (PF)	W (C) = n.

Table 1: Models for assessing concept importance.

To our knowledge, no previous studies have explored the role of such factors in human judgments of concept importance. We conducted a set of experiments involving 20 subjects admitted to Indiana University, using concept maps specifically designed to investigate structural influences by varying factors such as the degree of connectivity or the distance to the root node of certain concepts in the map. The maps’ concept and link labels were replaced with random letter combinations, in order to observe structural influences independently of map content. Subjects were presented with pairs of concepts, and asked to select the more important one, or to indicate that both were equally important. We then fitted the models to the subjects’ preferences by adjusting the models’ parameters using a hill-climbing algorithm. The results revealed two trends in structural effects on concept importance: (1) both authority nodes and nodes with incoming connections are considered more important than hub nodes or nodes with outgoing connections, and (2) nodes close to the root node are considered more important than nodes more distant from the root node (Leake *et al*. 2004). These results enable us to model structural influences on concept importance when selecting terms from a concept map to include in a search engine query, by summing the weights of the concepts in which the term occurs and using terms with the highest weights as query terms.

5 The Sufficiency of Concept Labels to Guide Retrieval

In the weighting method described in the previous section, links between concepts influence the choice of concepts, but linking phrases themselves are not reflected in the set of search terms. Thus an important question is whether concept labels, combined with information derived from concept map structure, are sufficient for query generation. Informal results suggest that the prevalence of generic linking phrases (e.g., “has”, “includes”, “consists of”) may make linking phrases less informative, decreasing their usefulness for an automatic system that cannot determine their specific meanings. However, to our knowledge, no controlled studies have examined this hypothesis. To investigate the effects on retrieval of concept-label-based queries compared to queries also reflecting linking phrases, we performed experiments comparing two types of queries using terms extracted from the concept and link labels in a concept map: (1) queries using terms selected from one or several concept keywords, or (2) queries using terms from a predefined ratio of concept labels and linking phrases. The queries were submitted to Google, and the terms appearing in the Web pages and other text documents returned by Google were then compared with the terms in the concept map from which the keywords were extracted. Terms were compared using two metrics, an adaptation of the Jaccard coefficient (a commonly-used association coefficient), and a coverage measure which we defined for this task. Table 2 below shows the measures.

S(Q,M,D), the modified Jaccard coefficient, measures the similarity between a document *D* and a concept map *M* as the proportion of terms in a retrieved result that are in the source map, but are not in the query *Q*. If the set of search results for a given query is empty, the measure for that query is considered to be 0. *C(Q,M,D)*, our coverage measure, is the ratio of shared keywords in *D* and *M* to the number of keywords in *M*, not counting the keywords in *Q*. The measures' values range from 0 (no similarity or coverage) to 1 (identical or full coverage).

$$S(Q,M,D) = \frac{|(D \cap M) - Q|}{|(D \cup M) - Q|} \qquad C(Q,M,D) = \frac{|(D \cap M) - Q|}{|M - Q|}$$

Table 2: Comparison measures for concept maps and Web documents.

For the experiments, we randomly selected ten concept maps from the Mars 2001 knowledge model (Briggs *et al*., 2004). In each experiment, a set of queries was computed, differing in the selection and number of concepts and linking phrases from which the keywords were drawn. Table 3 summarizes the average results of comparisons between concept maps and documents across the ten concept maps and the set of queries submitted to Google. The results suggest that concepts are critical in the search for information and search terms extracted from more than one concept yield better matching results. In contrast, linking phrases play a subordinate role in the search for information. Compared to queries from concepts, queries constructed only from linking phrases returned documents exhibiting little similarity to the concept maps from which they came. While some linking phrases aid in finding information somewhat similar to a concept map, they were generally far less effective, and often insufficient, for finding suitable information compared to concepts. Thus, we hypothesize that concepts are generally more valuable to improve precision of Web queries than linking phrases. This provides support for the ability of concept-based queries, such as those based on terms selected by our previous models, to retrieve resources with good coverage.

single concept		*link-link-link*		*link-concept-link*		*concept-concept-concept*	
Jaccard	**Coverage**	**Jaccard**	**Coverage**	**Jaccard**	**Coverage**	**Jaccard**	**Coverage**
0.02669	0.12618	0.01493	0.11380	0.02445	0.18492	0.02915	0.25636

Table 3: Results from the Google concept-linking phrase experiments.

6 Learning to Refine Query Terms

The previous sections describe methods for selecting and weighting concept map terms to include in a query. However, the first terms generated for a Web search may not provide the definitive results. For human-generated queries, users frequently decide, based on initial results, to refine their queries or pursue new directions (in Spink *et al*.'s study, 52% of users' sessions involved multiple queries).

When seeking for material on a topic, it is natural to form queries using the most descriptive terms of the topic. Terms are *good descriptors* of the topic of a concept map if they answer the question "What is this concept map about?" As discussed in Section 3, topological factors are useful for finding good descriptors of the topic of a concept map. We conducted an experimental study, using the Mars 2001 knowledge model, to investigate the hypothesis that terms that have higher weighting values according to our topological models tend to be good query terms. This evaluation showed that although these terms are useful, queries composed of descriptors alone are not sufficient to assure high precision. To address this problem, we developed a new approach based on the notion of *topic discriminators*.

Terms are good topic discriminators if they answer the question "Which are the best query terms to access similar documents?" A term is a good discriminator for a topic if most documents that contain that term are topically related. Thus finding good topic discriminators requires finding terms that tend to occur only in the context of the given topic. We propose that both topic descriptors and discriminators play an important role in concept map retrieval: That including topic descriptors in queries is important for recall, while including topic discriminators is important for precision. We have developed a method to exploit this within an automatic query generation system, automatically refining query terms by dynamically extracting good topic discriminators from search results and combining them with good topic descriptors identified by means of topological analysis.

Identifying Topic Discriminators for Concept Maps: Given a collection of m documents and n terms we can build a m×n matrix H, such that H[i,j]=k, where k is the number of occurrences of term t_j in document d_i. We define discriminating power of a term in a document as a function $\delta:\{t_0,\ldots,t_{n-1}\} \times \{d_0,\ldots,d_{m-1}\} \rightarrow [0,1]$:

$$\delta(t_i, d_j) = \frac{\operatorname{sgn}(H^T[i, j])}{\sqrt{\sum_{k=0}^{m-1} \operatorname{sgn}(H^T[i, k])}}$$

Analogously, we define descriptive power of a term in a document as a function $\lambda:\{d_0,\ldots,d_{m-1}\}\times\{t_0,\ldots,t_{n-1}\} \rightarrow [0,1]$:

$$\lambda(d_i, t_j) = \frac{H[i, j]}{\sqrt{\sum_{k=0}^{n-1} (H[i,k])^2}}$$

We can also define the descriptive power of a term *t* in a concept map as W(*t*), where W is the weight of the highest weighted concept in the map containing *t* according to some model from Section 3.

These simple notions of document descriptors and discriminators share some insight with standard IR proposals (Jones 1972; Salton & Buckley, 1988). Another recurrent notion in IR is document similarity. Let $\sigma(d_i,d_j)$ stand for the similarity measure between documents d_i and d_j. This measure can be computed in terms of term descriptive power as follows:

$$\sigma(d_i, d_j) = \sum_{k=0}^{n-1} (\lambda(d_i, t_k) \cdot \lambda(d_j, t_k))$$

We are interested in identifying good topic discriminators to form queries that will result in high precision. Function δ allows discovering terms that are good discriminators of a *document*, as opposed to good discriminators of the *topic* of a document. Because our goal is to refine queries to best reflect the topic of a concept map, we propose a topic-dependant definition of topic discriminators based on the notion of similarity between documents. We define the discriminating power of a term in the topic of a document (or concept map) as a function $\Delta:\{t_0,\ldots,t_{n-1}\}\times\{d_0,\ldots,d_{m-1}\}\rightarrow [0,1]$ calculated as follows:

$$\Delta(t_i, d_j) = \sum_{k=0,\quad k\neq j}^{m-1} (\sigma(d_k, d_j) \cdot \partial(t_i, d_k)^2)$$

Thus the discriminating power of term t_i in the topic of document d_j is an average of the similarity of d_j to other documents discriminated by t_i. Therefore, a term's discriminating power on the topic of a document is computed using the definitions of document similarity and term discriminating power on a document.

We claim that concept map descriptors can be recognized locally, by looking at the topology of a concept map. However, in order to find good discriminators of the topic of a concept map, as we mentioned earlier, the topology of the map is not sufficient. However, a relatively cheap *distillation* process has proven to be helpful to extract good topic discriminators. This distillation process consists of submitting a few queries to a search engine and using the information readily available from the search results to compute a Δ value for each found term. We argue that terms with high Δ value are good for attaining high precision. In the next section, we evaluate our methods for refining queries based on the notion of topic descriptors and discriminators.

An Experimental Test of Query Refinement: In order to test our proposal for learning to refine queries, we again used the Mars 2001 knowledge model. For each concept map, a baseline static method (described below) and our method for topic distillation were applied to select query terms. To evaluate the performance of these methods we computed the mean similarity of the returned results to the source concept map. We measure similarity as the proportion of novel terms (terms not in the query) in a retrieved document that are also part of the source map, using the function *S(Q,M,D)* described previously.

In our experiment, we use *Inverse Map Frequency* (IMF) as the baseline static feature weighting method. IMF is an adaptation of the standard IDF (*Inverse Document Frequency*) weighting scheme to measure the overall rarity of a term in a knowledge model. Thus, each term t in a concept map label was weighted as $IMF(t)=\log((1+|KM|)/|KM_t|)$, where *KM* represents the number of concept maps in the knowledge model (KM = 118 for the ``Mars 2001" model) and $|KM_t|$ stands for the number of concept maps containing term t. The IMF method was used to sort the terms occurring in a concept map. We then incrementally generated queries of increasing sizes, starting from a query of size 1 consisting of the most highly weighted term and incrementally adding terms in order of decreasing weight.

In our evaluation, we constructed a query for each concept in a concept map as a Boolean combination of terms occurring in the given concept plus the terms occurring in the map's root concept. For each query, we used the top 30 Google search results (if fewer than 30 results were returned, all were used). The search results were divided into 3 sets of equal size. In a three-stage evaluation, we used one of the sets for distillation and the other two for testing, rotating the roles of the sets at each stage. For each stage, the distillation data was used to learn an approximation of the discriminating power of each concept and the testing data was used for performance analysis. Only the "snippets" returned by Google were used for distillation, and the full documents were used for testing. (The *snippet* returned by Google is a text excerpt from the page summarizing the context in which the search terms occur.) The distillation phase in the query-refinement method consisted in computing the Δ value for the terms involved in a query to determine which terms are the best discriminators. For a fair comparison of the performance of the query-refinement method against IMF, we assured that queries were the same length by setting the size of the IMF queries to the number of terms occurring in the query resulting from the query-refinement method.

Figure 2 shows a scatter plot and table of results comparing the query-refinement method to the IMF method. We can see from the scatter plot comparing mean similarity between returned results and the source map that the query-refinement method outperforms IMF (64% vs. 36%) In the comparison table, we present the mean similarity confidence interval resulting from the query-refinement method, and we compare it against the mean similarity confidence interval resulting from applying the IMF method with query size adjusted as we explained above. This comparison table shows that the query-refinement method results in significant improvements over IMF.

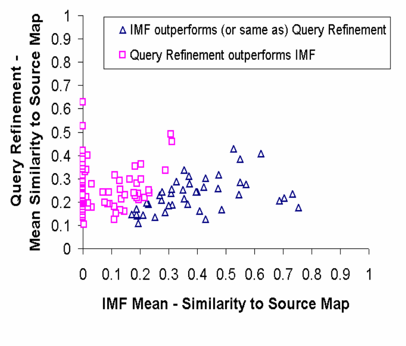

	N	Mean	Standard Deviation	Standard Error	95% Confidence Interval
Query Refinement	118	0.2498	0.0903	0.0083	**(0.2335, 0.2661)**
IMF	118	0.1880	0.1955	0.0180	**(0.1527, 0.2232)**

Figure 2: Performance comparison of the query refinement method against IMF.

7 Combining the Approaches

The previous steps provide methods for automatically generating and refining concept-map-based queries. The following algorithm summarizes how the entire set of methods is combined in order to "Google from a concept map":

ALGORITHM:
INPUT:
C: a concept map.
T_σ: threshold for filtering irrelevant results.
q: number of queries per round.
BEGIN
//topological analysis
Apply topological analysis to assign a weight *W(c)* to each concept *c* in the concept map *C*.
//distillation
Submit *q* queries using the concepts *c* in *C* with highest *W(c).*
Compute Δ for each term in *C*, using only "readily available" results.
//search
Combine the concepts with highest *W(c)* value and the terms *t* with highest Δ*(t,C)* value to form *q* additional queries and submit the formed queries to a search engine.
//using concept map to filter results
Only display those documents *d* such that $\sigma(d,C) > T_\sigma$ where the content of *d* is approximated for this judgment using "readily available" information from search results.
END

The techniques discussed in the previous sections, and the results of the evaluations described there, are reflected in the design of the algorithm. Because it is possible to estimate concept importance based on concept map topology, and because concept labels tend to provide a good description of the content of a concept map for the purposes of retrieving similar information, the query formation process is based on concept labels: For the first round of queries, only labels of concepts with high descriptive power are used as query terms, where descriptive power is assessed by using the topological analysis models. For the *distillation* and *search* phases of the algorithm, the query refinement techniques discussed in section 6 are used to identify good discriminators of the topic of the map and to form queries that combine terms with high descriptive and discriminating power. As a final step, the returned results can be filtered according to the map, to improve the ranking of results; this is implemented in our system but cannot be described here due to space limitations.

8 Conclusion

This paper has described ongoing research on exploiting the information in concept maps to automatically generate and refine queries to Web search engines, to aid concept mapping. The component algorithms have been implemented in robust prototypes, and have been evaluated individually with promising results. The results provide initial support for the hypothesis that information extracted from concept maps can provide an effective starting point for automatically-generated search queries, and for the efficacy of the specific methods described. We are now designing experiments for an end-to-end human-subjects evaluation of the system as a whole, to directly assess subjects' judgments of the relevance and usefulness of the information provided.

9 Acknowledgements

This work was supported in part by NASA under awards No NCC 2-1216 and 2-1297. We thank Google for the use of their API.

References

Briggs, G., Shamma, D., Cañas, A. J., Carff, R., Scargle, J., & Novak, J. D. (2004). Concept Maps Applied to Mars Exploration Public Outreach. In A. J. Cañas, J. D. Novak & F. González (Eds.), Concept Maps: Theory, Methodology, Technology, *Proceedings of the First International Conference on Concept Mapping.* Pamplona, Spain: Universidad Pública de Navarra.

Brin, S. & Page, L. (1998) The Anatomy of a Large-Scale Hypertextual Web Search Engine. *Computer Networks and ISDN Systems*. 30(1):107-117.

Budzik, J., Hammond, K., and Birnbaum, L. (2001). Information Access in Context. *Knowledge-Based Systems,* 14(1-2): 37-53.

Cañas, A. J., Hill, G., Carff, R., Suri, N., Lott, J., Eskridge, T., Gómez, G., Arroyo, M., & Carvajal, R. (2004a). CmapTools: A Knowledge Modeling and Sharing Environment. In A. J. Cañas, J. D. Novak & F. M. González (Eds.), Concept Maps: Theory, Methodology, Technology, *Proceedings of the First International Conference on Concept Mapping.* Pamplona, Spain: Universidad Pública de Navarra.

Cañas, A. J., Carvalho, M., Arguedas, M., Leake, D. B., Maguitman, A., Reichherzer, T, (2004b). Mining The Web to Suggest Concepts During Concept Map Construction, In A. J. Cañas, J. D. Novak & F. González (Eds.), Concept Maps: Theory, Methodology, Technology, *Proceedings of the First International Conference on Concept Mapping.* Pamplona, Spain: Universidad Pública de Navarra.

Carvalho, M., Hewett, R., & Cañas A. (2001). Enhancing Web searches from concept map-based knowledge models. In *Proceedings of the SCI Conference*. Orlando, Florida: AAAI Press.

Jones, K. S. (1972). A Statistical Interpretation of Term Specificity and its Application in Retrieval. *Journal of Documentation,* 28:11-21.

Kleinberg, J. (1999) Authoritative sources in a hyperlinked environment. *Journal of the ACM*, 46(5):604-632.

Leake, D., Maguitman, A., Reichherzer, T., Cañas, A., Carvalho, M., Arguedas, M., Brenes, S., & Eskridge, T. (2003). Aiding knowledge capture by searching for extensions of knowledge models. *In Proceedings of KCAP-2003*. St. Augustine, Florida: ACM Press.

Leake, D., Maguitman, A., Reichherzer, T. (2004). Understanding Knowledge Models: Modeling Assessment of Concept Importance in Concept Maps. *In Proceedings of CogSci 2004*. Mahwah, NJ: Erlbaum. In press.

Novak, J. D., & Gowin, D. B. (1984). *Learning How to Learn.* New York: Cambridge University Press.

Salton G., & Buckley C. (1988). Term weighting approaches in automatic text retrieval. *Information Processing and Management*, 24(5): 513-523.

Spink, A., Wolfram, D., Jansen, B.J., Saracevic, T. (2001). Searching the Web: The Public and their Queries. *Journal of the American Society for Information Sciences and Technology* 52(3): 226-234.