Suggesting Novel but Related Topics: Towards Context-Based Support for Knowledge Model Extension

Ana Maguitman, David Leake, and Thomas Reichherzer
Computer Science Department, Indiana University
Lindley Hall 215, Bloomington, IN 47405, USA
{anmaguit, leake, treichhe}@cs.indiana.edu

ABSTRACT

Much intelligent user interfaces research addresses the problem of providing information relevant to a current user topic. However, little work addresses the complementary question of helping the user identify potential topics to explore next. In knowledge acquisition, this question is crucial to deciding how to extend previously-captured knowledge. This paper examines requirements for effective topic suggestion and presents a domain-independent topic-generation algorithm designed to generate candidate topics that are novel but related to the current context. The algorithm iteratively performs a cycle of topic formation, Web search for connected material, and context-based filtering. An experimental study shows that this approach significantly outperforms a baseline at developing new topics similar to those chosen by an expert for a hand-coded knowledge model.

Categories and Subject Descriptors

H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval—*Query formulation, Search process*; H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*Graphical User Interfaces (GUI)*; I.2.6 [**Artificial Intelligence**]: Learning—*Knowledge Acquisition.*

General Terms

Algorithms, Experimentation.

Keywords

Human-Centered Knowledge Acquisition Tools, Concept Mapping, Context, Automatic Topic Search.

1. INTRODUCTION

Concept mapping is widely used in education to help students organize, extend, and share their knowledge. More recently, concept mapping has also been explored as a vehicle for capturing and sharing expert knowledge. Because the concept mapping process is simple and relatively unconstrained, concept-map-based interfaces for knowledge modeling are appealing as a way to empower experts to play an active role in the knowledge capture process (e.g., [3]). This paper addresses the problem of how to enhance such interfaces to help users select topics to include.

When experts generate concept maps, their task serves both *knowledge acquisition*—capturing their pre-existing knowledge—and *knowledge construction*—furthering their development of new knowledge about a domain. Electronic concept mapping tools (e.g., CmapTools [5]) provide an interface for generating concept maps and linking them to other concept maps in order to form concept-map-based *knowledge models*. The basic CmapTools interface supports "pencil and paper" operations on concept maps, such as drawing networks of concepts and links, augmented with capabilities for annotating maps with multimedia resources and browsing networks of concept maps. We are developing methods to augment this interface with "intelligent suggesters" to automatically provide context-relevant suggestions of information to support the expert's concept-mapping process [15]. Many of these suggesters focus on how to fill in the current concept map with related information, proposing concepts and relationships to include in a concept map and resources with which to annotate the nodes in a concept map already under construction. The focus of this paper is on a suggester addressing a larger question: given a concept map under construction, reflecting a certain topic, what *other* topics might be suitable to include in the knowledge model. Thus the goal is to propose context-relevant ways to *extend* an existing knowledge model.

Numerous intelligent interfaces examine the user's task context or information access behavior, to suggest related resources (e.g., [24, 4]). The topic-suggestion task contrasts with this task in at least two ways. First, its goal is not to suggest individual *resources*, but rather, to make suggestions at the higher level of *topics*. A single topic may be partially reflected by a number of different Web pages, with none focusing solely on the topic or summarizing the topic in its entirety. Second, its goal is not to suggest the topics *most related* to the current concept map, but rather to suggest new topics that are *related but novel*.

Developing topic suggestion interfaces depends on addressing three central questions: (1) how to characterize topics, (2) how to evaluate the appropriateness of new candidate topics, and (3) given an initial context, how to generate the new topics themselves. After providing a brief overview of concept mapping and the CmapTools system, this paper addresses those questions, focusing especially on how they apply to suggestions for concept mapping. It then presents EXTENDER, an implemented system which mines search engines to generate topic suggestions for presentation by the CmapTools interface. It closes with an evaluation demonstrating encouraging results for the approach's ability to generate novel, related, and cohesive topic suggestions.

IUI'05, January 9–12, 2005, San Diego, California, USA.

2. CONCEPT MAPPING

Concept maps [19, 20] are collections of simplified natural language sentences displayed as a two-dimensional, visually-based representation of concepts and their relationships. Concept maps depict concepts as labeled nodes and inter-concept relations as labeled links. In educational settings, teachers assign students to draw concept maps as a way to encourage them to organize their knowledge, and to make their understanding explicit for knowledge assessment and sharing. Concept mapping has been used for a wide range of age groups; even elementary school students can generate concept maps successfully. The naturalness of the concept mapping process makes it promising as a method for direct knowledge capture by experts themselves, and the conciseness and structure of concept maps assists understanding the captured information. Note that concept mapping produces an "informal" representation; its goal is to facilitate knowledge capture for human examination and sharing, rather than for automated reasoning.

3. CMAPTOOLS

Electronic concept mapping tools generate concept maps in an electronic form that is browsable—concept nodes can be used to organize resources, including other maps—and sharable, both for individual examination and for collaborative development. CmapTools, developed by the Institute for Human and Machine Cognition (IHMC), is a suite of publicly-available software tools for knowledge acquisition, construction, and sharing based on concept maps. The CmapTools system has been widely used as a vehicle for knowledge capture and sharing, both in educational and commercial settings. The system and numerous sample uses are discussed in [5]. Figure 1 shows the CmapTools' display of portions of two concept maps about the topic of Missions to the Moon developed by a NASA expert as part of an extensive knowledge modeling effort using the system [3].

The basic CmapTools interface empowers experts to construct knowledge models of their domains without the need for a knowledge engineer's intervention, or to actively participate in knowledge elicitation if a knowledge engineer leads the process. However, observations of user behavior building knowledge models revealed a number of opportunities for augmenting the CmapTools interface with intelligent support. When experts and ordinary users build concept maps, they often stop for significant amounts of time, wondering how to extend their models, in some cases searching the Web looking for new material and ideas to enhance their in-progress representations. This search activity can be done more effectively if mechanisms for information access and delivery are included as part of the knowledge modeling tools. Consequently, an effort is under way to augment CmapTools with intelligent aids that start from a concept map under construction, and propose information to aid the user's knowledge capture and knowledge construction [15]. The focus of this paper, the EXTENDER system, uses information automatically extracted from the current knowledge model to guide mining the Web to identify and suggest novel but relevant topics.

4. THE EXTENDER SYSTEM

EXTENDER's goal is to aid experts building knowledge models by "jogging the user's memory," providing suggestions for new topics to cover. EXTENDER provides its topic suggestions as small sets of terms, meant to convey the sense of a topic (e.g., a label of the form *lunar, moon, prospector...* is used to describe the topic *Lunar Prospector Mission to the Moon*.) These terms are presented in a suggestions panel visible at the upper right of Figure 1. To avoid distracting the user, the suggestion panel becomes visible only when the user decides to open it. When the panel is closed, an unobtrusive icon shows that new suggestions are available.

Figure 2 outlines EXTENDER's processing cycle. The system starts from a concept map and iteratively searches the Web for novel information. EXTENDER's interface allows the user to highlight a concept or set of concepts from the starting concept map in order to bias the system's search towards topics related to the highlighted concepts. Alternatively, the search can be initiated from the full map, without introducing any additional bias.

At each iteration, the collected material is represented internally by document-term matrices, clustering is applied to identify topics in the collection, and unimportant material is discarded. This process is repeated, with the stopping criterion depending on topic convergence and a user-selected limit on iterations. In our preliminary tests, three iterations is usually sufficient to generate a rich variety of artificial topics.

Once EXTENDER completes its iterations, it presents the generated topics as suggestions to the user. In addition, it presents the Web pages that gave rise to those topics, grouped by topic, to facilitate access to topic-relevant information. EXTENDER's preferences panel, shown at the bottom of the suggestion window in Figure 1, allows the user to adjust the suggestions' level of focus on the chosen concepts and the range of topics generated. The interface enables users to easily import a generated topic into an in-progress concept map as a set of concepts, from which the user can start the mapping process. The concept map at the bottom left of the figure contains some concepts that the user selected from a topic suggested by EXTENDER.

5. DESIDERATA FOR TOPIC SUGGESTION

EXTENDER's task is an instance of a more general one: suggesting novel topics related to a user's focus. This task could apply in many contexts. For example, a topic suggestion interface could be useful to a researcher, to propose related but distinct areas to consider for connections and synergies or to help assure that relevant areas have been considered. Consequently, one of to goals of research on EXTENDER has been to develop a task- and domain-independent approach to topic suggestion.

In general, given a corpus **C** of *items* I (e.g., documents or concept maps), and a universe **U** of possible terms, we consider a *topic* T of an item I to be a nonempty subset of the terms contained in **U**, satisfying a set of domain-relevant criteria for topic quality. Note that the mapping from an item I to topics may be one-to-many.

A topic can be defined as a set of cohesive terms (as we define formally in [17]), but it can also be defined implicitly, as a set of items (e.g., Web pages) that share a common theme. A topic generation system starts with an item I, one or more topics associated with I, and possibly additional resources, and generates a set of new topics **N**. The performance of a topic suggester can be judged according to the following metrics:

- **Local quality**. Each generated topic must be of high quality according to the criteria for the domain. Such criteria might include measures for conciseness (that the topic is summarized in a few terms, for easy user comprehension), term coherence (that each topic description is constituted of tightly related terms and documents), etc.
- **Global Coherence:** The topics in **N** must be *relevant* to I.
- **Coverage:** **N** must contain a sufficient subset of the topics considered to be relevant.

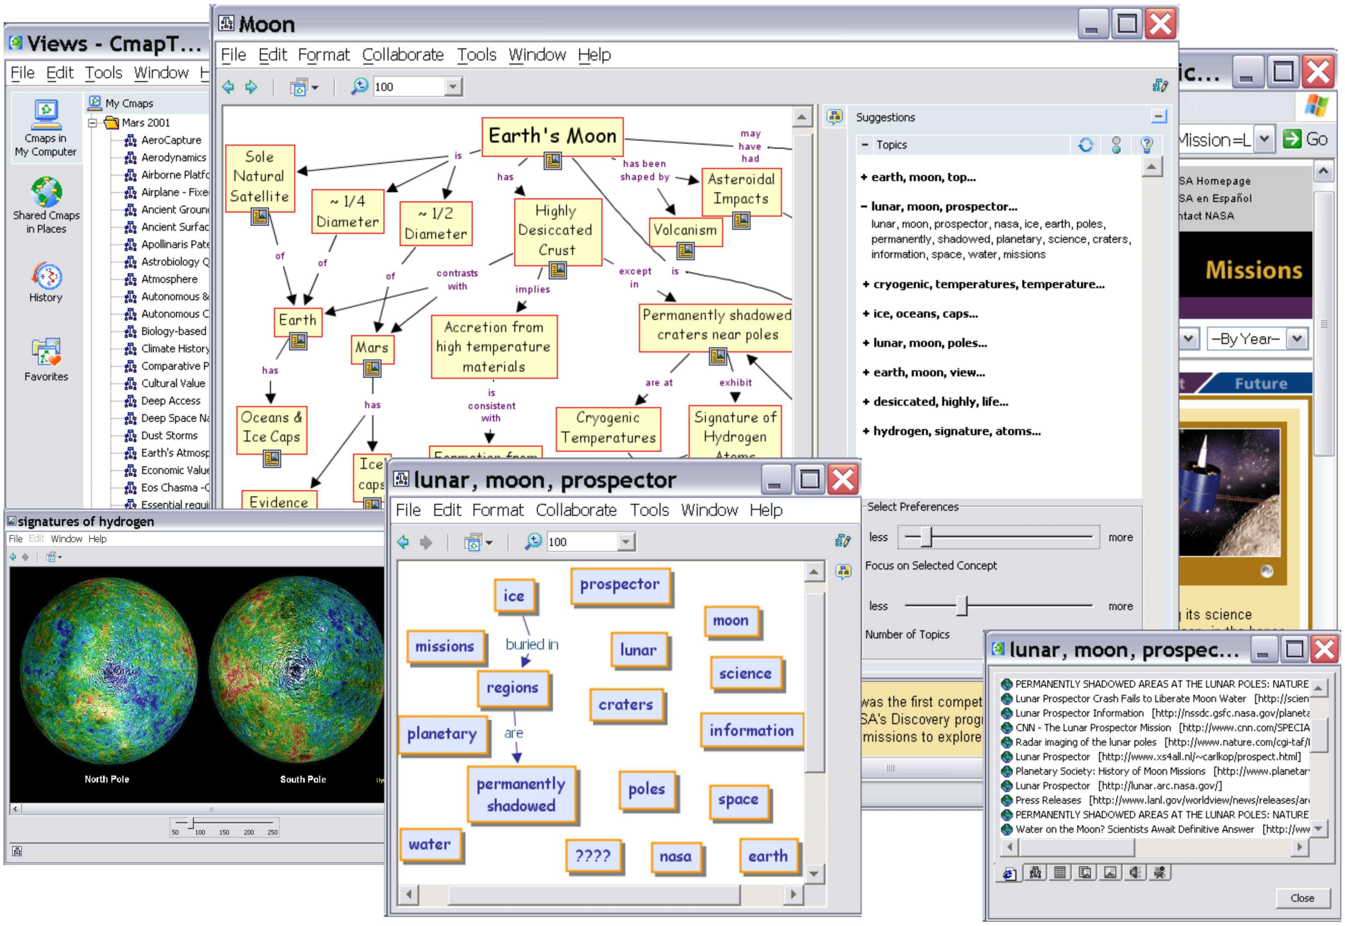

Figure 1: Portion of a Knowledge Model with EXTENDER suggesting new topics.

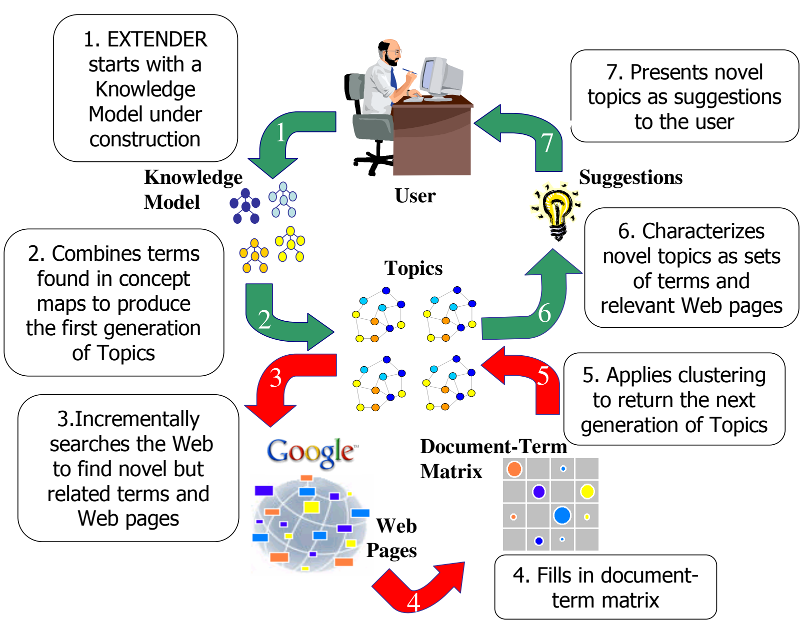

Figure 2: EXTENDER's Cycle.

- **Novelty**. **N** must include a sufficient subset of topics not included in the originating corpus **C**.
- **Diversity**. Within **N**, topics must be sufficiently diverse from each other for additional topics to be useful.

The following sections describe specific applications of these criteria to topic generation and performance evaluation in EXTENDER.

6. HOW EXTENDER ACHIEVES THE DESIDERATA

The topics generated by EXTENDER must be relevant to the user's "in progress" knowledge model but must also go beyond that model. Consequently, achieving an appropriate balance of relevance and novelty, and achieving diversity while preserving global coherence, are crucial for EXTENDER.

6.1 Preserving Global Coherence

For EXTENDER's task, topics are globally coherent if they relate to the user's initial concept map, which may provide rich information to exploit as a context. Unfortunately, because search engines restrict queries to a small number of terms (e.g., the 10-term limit for Google), a single query can only reflect limited information. Consequently, EXTENDER uses a multi-step approach to focus its topics by collecting and filtering terms over multiple retrievals. During its cycle, EXTENDER maintains the relationship between candidate topic terms and the initial concept map in three ways:

- **Term-weight reinforcement.** Terms collected during EXTENDER's retrievals are associated with weights summarizing the terms' goodness as query terms. Terms may be weighted highly as *topic descriptors* or *topic discriminators*. A term is a good topic descriptor if it occurs often in documents similar to the initial concept map, while a term is a good topic discriminator if it occurs primarily in similar documents, where similarity is computed by the cosine similarity metric between the document and concept map's term-vector representations. For reasons of space, we do not define descriptors or discriminators more formally here, but we provide formalizations in [17], which proposes and evaluates methods for the dynamic extraction of topic descriptors and discriminators in the context of concept maps.

  At the start of its cycle, EXTENDER calculates terms' descriptive power directly from the topology of the user's concept map as described in [6, 14]. If the user has selected focus concepts to bias the topic search, the weights of the terms in the selected concepts' labels are adjusted by a constant weighting factor greater than one. For subsequent iterations, weights are adjusted according to the dynamic extraction procedures of [17], to reinforce the weights of terms that have proven to be good descriptors or discriminators for the topic represented by the search context.

- **Context-based filtering.** For a document's terms to be considered candidates for inclusion as part of a new topic, the document has to survive a selection process that requires a minimum similarity between the document and the search context. Novel terms that are not good descriptors or discriminators of the topic reflected by the search context are also discarded.

- **Query refinement.** The first query terms generated for a Web search may not provide the definitive results. However, initial search results can help to automatically refine subsequent queries. Terms that occur often in documents with high term-wise similarity to the initial concept map may help to achieve good recall when used in a query. On the other hand, terms that tend to occur only in similar documents are useful for achieving high precision. EXTENDER computes term-wise similarity between retrieved documents and the initial concept map, adding terms that occur often and primarily in documents similar to the search context.

In EXTENDER's final stage, when the system presents the final set of topics to the user, the terms with highest descriptive value are used to produce labels for the suggested topics.

6.2 Generating Cohesive Topics

Local coherence reflects the degree to which each generated topic is composed of tightly related terms. In our context, cohesiveness is measured by the ability of the topic to prompt retrieval of documents that are similar to each other. For generating cohesive candidate topics, EXTENDER uses only short text excerpts (the text "snippets" provided by Google, which are readily available from the search results) to represent documents. The need to group a collection of short text excerpts from highly related documents contrasts with common clustering scenarios. When full access to documents' text is available, document clustering appears preferable over term clustering, to give the clustering algorithm greater discerning power to identify topics. However, we have observed that when documents are represented by a small number of terms (as is the case for the text excerpts collected by EXTENDER), and the collection under analysis consists of material that shares a common general theme (which is a consequence of EXTENDER's attempt to preserve global coherence), terms may be as informative as documents for identifying topics within the collection.

With a few exceptions (e.g., [9]) most existing clustering algorithms apply single-purpose clustering—they cluster documents and terms separately. EXTENDER applies a medoid-based co-clustering algorithm, to cluster documents and terms simultaneously. For each document, EXTENDER finds the terms that best characterize the document's topic. In a subsequent step, it uses the selected terms to identify the documents that best specify the topic of those terms. This process is repeated until either (1) two consecutive iterations produce the same set of terms and documents, or (2) the same result is detected for two non-consecutive iterations. Because after each iteration the sizes of the sets containing selected terms and documents decrease or remain the same, the selection process is guaranteed to terminate in a finite number of iterations. The sets of selected terms and documents are used as term-medoids and document-medoids to define a set of cohesive topics. This clustering algorithm tailored for EXTENDER addresses the goal of achieving local coherence for this domain.

6.3 Coverage, Novelty and Diversity

Coverage reflects the ability of a topic-generation strategy to generate the full set of the existing relevant topics, even if those topics are novel and go beyond previously captured information. In our scenario, generating new topics from Web searches, it is not possible to formally assess coverage.

Diversity and global coherence are conflicting goals. However, a reasonable topic-generation strategy should be able to produce topics with a suitable balance. EXTENDER uses a "curiosity mechanism" to diversify during initial processing stages and to focus towards the end. The application of EXTENDER's curiosity mechanism is in the spirit of searching and learning techniques (e.g., simulated annealing and reinforcement learning) in which a temperature factor is used to favor exploration at the beginning and exploitation during the final stages.

Throughout the system's iterations, while attempting to extend a given topic, new-found terms are collected. Because the number of collected terms grows rapidly, novel terms to retain are selected based on a curiosity mechanism. For each term, the system tracks both the goodness of the term in describing and discriminating the current topic. It retains those that surpass a threshold for the survival of descriptors/discriminators, where the threshold is a function of the number of iterations. Another curiosity threshold is used by EXTENDER to filter irrelevant documents according to the search context. During the initial steps EXTENDER collects documents on diverse topics, which must preserve the global theme of the originating concept map. After each iteration is completed, the current topic gives rise to a new set of descendant topics. As the system moves its focus through the new set of topics, the search context is updated and the curiosity threshold required for term retention is increased. Because the threshold increases with the number of iterations, novel terms and documents are seldom collected during the final stages. Consequently, the final stages are an exploitation phase that primarily reinforces the weights associated with particular material that have been already added to the collection.

7. THE EXTENDER ALGORITHM

The previous techniques form the core of EXTENDER's topic extension algorithm. However, because retrieving and processing large numbers of Web pages is costly, EXTENDER first applies a less expensive *distillation phase*, in which a series of queries is submitted to a search engine and only the information that is readily available from the search results (e.g. title, "snippet" of text, url, Open Directory Project summary) is used to identify promising topic descriptors and discriminators. After this preliminary step, the best topic descriptors and discriminators are used as query terms in a *search phase* to search for additional material on the Web. The new set of search results is filtered according to the search context and then clustered to produce the next generation of topics. Finally, each of the topics is refined, keeping only those documents and terms that are good topic representatives. The topics resulting from different branches usually have significant overlap of coverage. To ensure diversity, after an iteration is completed EXTENDER merges similar topics. This is done by applying a simple single-linkage clustering procedure. A parameter r defines the similarity threshold between two topics. If the similarity between two topics is at least r, then the two topics are merged. Table 1 provides a high-level description of this algorithm.

8. EVALUATION

Topic selection is hard to assess in a controlled way because the usefulness of topic suggestions is highly subjective. To perform an objective test, we used the Mars 2001 knowledge model [3], an expert-generated set of concept maps, as our "gold standard" for an automatic evaluation of EXTENDER's topics. This knowledge model on Mars exploration was created by NASA experts and contains 118 concept maps, presenting an extensive coverage of topics in the field.

In our tests, the top-level concept map from the knowledge

PROCEDURE EXTEND-TOPIC
INPUT:
 M: source concept map.
 s: total number of iterations.
 q_d: number of queries submitted for distillation.
 q_s: number of queries submitted for search.
OUTPUT:
 A set of topics related to T.
BEGIN
 Topics[0]= { M }.
 for (i=0; i $<$ s; i++)
 do
  Topics[i+1]=$\emptyset$.
  for each Topic T $\in$ Topics[i]
  do
   N = NEXT-GENERATION-OF-TOPICS(T, i).
   Topics[i+1]= Topics[i+1] $\cup$ **N**.
  Merge similar topics in Topics.
 return Topics.
END
PROCEDURE NEXT-GENERATION-OF-TOPICS
INPUT:
 T: topic to extend.
 i: present iteration.
OUTPUT:
 N: A new set of topics.
BEGIN
 //distillation
  Use those terms with highest descriptive value
   to form q_d queries and submit the queries
   to a search engine.
  Use search result's "readily available information"
   to compute descriptive and discriminating power
   for each term.
 //search
  Combine the terms with highest descriptive and
   discriminating value to form q_s queries.
  Submit the queries to a search engine and collect the
   returned document excerpts.
 //filtering
  Use curiosity mechanism to filter returned documents
   according to the map.
  Use curiosity mechanism to filter terms according to
   their descriptive and discriminating value.
 //clustering
  Cluster remaining data to generate cohesive topics.
 //clean-up
  For each topic only keep terms that are good
   descriptors or discriminators.
  For each topic only keep documents that are similar
   to the medoid of the topic.
  Collect resulting topics into set **N**.
 return N.
END

Table 1: Pseudocode of the Topic Extension Algorithm

model was used as the starting point (corresponding to the human user's map under construction). EXTENDER's topic extension algorithm was used to produce a collection of topics, without access to any of the other maps in the knowledge model. As a baseline method for comparison, we implemented a simple algorithm which constructs queries using all the concepts from the same concept map EXTENDER used as a starting point, submits them as queries to the Google Web API, and clusters the results to generate topics.

We expected EXTENDER's mechanism to provide results with superior global coherence, novelty, and coverage for equal numbers of Web queries. An evaluation based on coherence and coverage requires an operational definition of *topic relevance*. Here, we consider the expert-generated Mars 2001 topics as *target topics*, with the relevance of a system-generated topic measured by the accuracy with which a system-generated topic replicates an expert-generated topic. Note that the accuracy measure also provides an indication of topic quality, because its results depend on the similarity between EXTENDER's topics and the expert-generated set, which we expect to be of good quality for the domain.

The measures of accuracy, coherence and coverage are formalized in the next section.

8.1 Criterion Functions for Evaluating Topic Generation

To measure global coherence assume that $R = \{r_1, \ldots, r_m\}$ is a target set of relevant topics and $A = \{a_1, \ldots, a_n\}$ is a set of topics generated by the topic-generation strategy under evaluation. Because novelty is one of our desiderata for topic generation, we want to favor strategies that produce relevant topics with a high number of novel terms. Consider the set o, containing the terms of the originating topic, i.e., the knowledge model that is used as a starting point to search for topics. We propose a modified similarity measure reflecting the proportion of *novel terms* (terms not in the starting knowledge model) in a system-generated topic a_i that are also part of an r_j from the set of relevant topics:

$$\textbf{Similarity}^N(a_i, r_j, o) = \frac{|(a_i \cap r_j) - o|}{|(a_i \cup r_j) - o|}.$$

Then, we can define the *accuracy* of topic a_i in R as follows:

$$\textbf{Accuracy}^N(a_i, o, R) = \max_{r_j \in R} \textbf{Similarity}^N(a_i, r_j, o).$$

The *accuracy* function measures the precision with which a given topic replicates some topic in a given set of topics, disregarding those terms that are in the starting knowledge model. We use the *accuracy* function to define *global coherence* as follows:

$$\textbf{Global\_Coherence}^N(o, A, R) = \frac{\sum_{a_i \in A} \textbf{Accuracy}^N(a_i, o, R)}{|A|}.$$

The *global coherence* function measures the fraction of relevant topics generated by the algorithm being evaluated, weighted by the algorithm's level of accuracy replicating the relevant topics.

Global coherence is a generalization of the IR notion of precision, and as such, it has its limitations. This criterion function can be maximized if the system generates a single topic identical to some relevant topic, which clearly does not guarantee acceptable topic generation performance. Hence, a *coverage* factor must be introduced to favor topic-generation strategies that cover many topics of a target set of relevant topics. To address this issue, we define a coverage function as a generalization of the standard IR notion of recall:

$$\textbf{Coverage}^N(o, A, R) = \frac{\sum_{r_i \in R} \textbf{Accuracy}^N(r_i, o, A)}{|R|}.$$

8.2 Parameter Settings

EXTENDER's methods depend on parameters such as the number of generations (iterations), the number of queries submitted from the source concept map and from each generated topic, the maximum number of topic descendants for each topic, the starting and stopping thresholds for the curiosity mechanisms, and the similarity threshold for merging topics. This results in a large parameter space. In practice, however, pragmatic concerns for the interface, such as the desire for rapid response and low memory use, suggest constraining some parameters. Accordingly, our tests limited the number of generations to 3, the number of queries from each new topic to 20, and the number of topic descendants at each stage to 8.

8.3 Experimental Results

Our evaluation involved 48 trials, with different settings for EXTENDER's parameters. When comparing the performance of EXTENDER against the baseline, we set the number of queries for the baseline to the total number of queries submitted by EXTENDER. For each trial, EXTENDER and the baseline method used the same similarity threshold and method for merging topics.

Figures 3(a) and 3(b) compare the performance of EXTENDER's topic generation algorithm to the baseline method in terms of global coherence/novelty and coverage. A particular parameter setting corresponds to a trial and is represented by a point. The point's horizontal coordinate corresponds to the performance of EXTENDER for that case, while the vertical coordinate corresponds to the performance of the baseline method. In Tables 2 and 3 we present the mean confidence interval resulting from computing the performance criterion functions for EXTENDER and the baseline. These comparison tables show that EXTENDER results in statistically significant improvements over the baseline method.

	N	MEAN	STDEV	SE	95% C.I.
EXTENDER	48	0.267	0.05	0.007	**(0.253, 0.281)**
Baseline	48	0.101	0.085	0.012	**(0.077, 0.125)**

Table 2: Confidence intervals for the average global coherence of EXTENDER and baseline (considering novel material only.)

	N	MEAN	STDEV	SE	95% C.I.
EXTENDER	48	0.116	0.059	0.008	**(0.099, 0.132)**
Baseline	48	0.019	0.009	0.001	**(0.017, 0.022)**

Table 3: Confidence intervals for the average coverage of EXTENDER and baseline (considering novel material only.)

When we analyzed the relationship between parameter settings and EXTENDER's results we noticed that different parameter settings favor different aspects of EXTENDER's performance. These results shed light on the selection of appropriate thresholds for the curiosity mechanism parameters and for the number of iterations, helping us to improve the design of both EXTENDER's algorithm and EXTENDER's interface. For example, higher thresholds for the curiosity mechanism favor global coherence while lower thresholds favor coverage. Therefore, the interface enables the user to adjust these parameters, to choose to focus on topics more or less similar to the user's current topic.

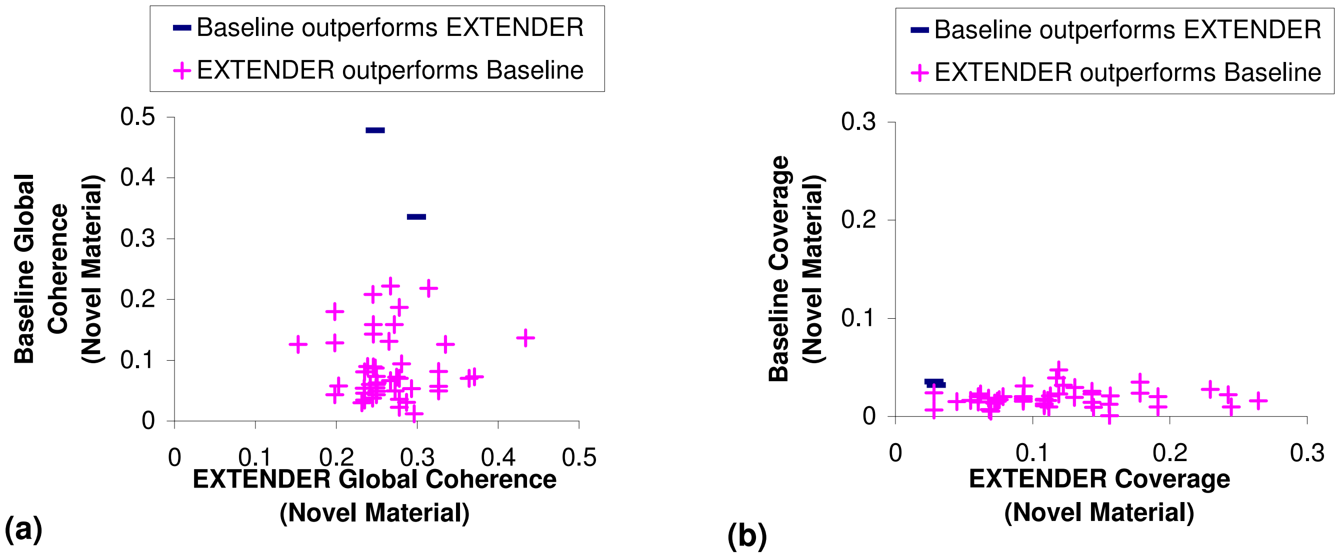

Figure 3: EXTENDER **vs. Baseline: (a) Global Coherence, (b)Coverage**

9. RELATED WORK

Frameworks aimed at capturing knowledge have centered mostly on the construction of standardized representations. The knowledge modeling community has long been concerned with devising ontologies as formal specifications that machines can read and process [12]. Recently, with the growing attention to the development of a Semantic Web [2], research on language design for developing ontologies has increased. Ontology construction is a tedious process; therefore systems have been built to expedite the design of ontologies and to facilitate sharing and integration of different frameworks. Examples of systems that facilitate collaborative development of ontologies include the ONTOLINGUA server [11]. RIBOWEB [1], COMMUNITY WEB PORTALS [25], ONTOSHARE [8], and the PROTÉGÉ family [21]. These tools provide a graphical environment for ontology-development and knowledge acquisition. However, the goal of these tools is to facilitate the construction of standardized representations, while the goal of EXTENDER is to provide human-centered support for knowledge extension.

Numerous Web agents have been developed to facilitate access to resources on the Web. Some of these agents, such as SoftBots [10] operate on top of Internet tools and services, with the purpose of abstracting away the technology underlying the accessed resources. Web crawlers [23] exploit the graph structure of the Web to follow hyperlinks, discover resources, and map them into searchable index structures. Some Web crawlers are exhaustive, and perform an extensive exploration of the resources available online, independently of a pre-defined set of topics. Other Web crawlers are topical or focused [7, 18], in which case the mining process is guided not only by following existing links but also by considering content to focus on pages relevant to a specific theme. EXTENDER contrasts in relying entirely on a search engine to mine the Web—and in not being aimed at discovering specific pages. Instead, it attempts to dynamically generate short topic descriptions to jog the user's memory during knowledge modeling.

Several suggester systems exploit user interaction with computer applications to determine the user's current task and contextualize information needs. This gives rise to context-aware suggester systems [4, 24, 16]. As opposed to these systems, EXTENDER's goal is not to suggest the *most related* material, but rather to suggest topics that *go beyond* previously captured information.

Our research on topic extraction also shares insights and motivations with proposals aimed at clustering search results (e.g., [13, 27]) and refining queries (e.g., [26, 22]). However, in contrast to our approach, these systems provide browsing interfaces requiring explicit user intervention. In addition, their goal is to help users to focus on specific information and to remove alternatives, rather than to discover novel but related material.

10. CONCLUSION

The aim of topic generation is to aid the user in pursuing *new directions* relevant to his or her work. Topic generation provides a new research area for intelligent user interfaces, complementary to the considerable work on interfaces to provide documents relevant to the user's current task

One area in which such suggestions promise to be useful is supporting knowledge capture, by identifying new topics related to a current knowledge model. This paper identifies general desiderata for topic generation and presents a domain-independent topic-generation algorithm developed for supporting concept-map-based knowledge modeling. The process reflects the knowledge modeling context through an iterative process of topic generation, Web search, and context-based filtering. An experimental study shows that this approach significantly outperforms a baseline at recovering topics close to those of an expert's hand-coded knowledge model.

11. ACKNOWLEDGMENTS

We gratefully acknowledge the contributions of our collaborators Alberto Cañas and the IHMC CmapTools team. This research is supported in part by NASA under award No NCC 2-1216. We thank Google for the use of their Web APIs service.

12. REFERENCES

[1] R. Altman, M. Bada, X. Chai, M. Carillo, R. Chen, and N. Abernethy. RiboWeb: An ontology-based system for collaborative molecular biology. *IEEE Intelligent Systems*, 14(5):68–76, 1999.

[2] T. Berners-Lee, J. Hendler, J., and O. Lassila. The Semantic Web. *Scientific American*, May, 2001.

[3] G. Briggs, D. Shamma, Cañas, R. Carff, J. Scargle, and J. D. Novak. Concept maps applied to Mars exploration public outreach. In A. J. Cañas, J. D. Novak, and F. González, editors, *Concept Maps: Theory, Methodology, Technology. Proceedings of the First International Conference on Concept Mapping*, Volume 1, pages 109–116, U. of Navarra, 2004.

[4] J. Budzik, K. J. Hammond, and L. Birnbaum. Information access in context. *Knowledge Based Systems*, 14(1–2):37–53, 2001.

[5] A. Cañas, G. Hill, R. Craff, N. Suri, J. Lott, T. Eskridge, G. Gómez, M. Arroyo, and R. Carvajal. CmapTools: A knowledge modeling and sharing environment. In A. J. Cañas, J. D. Novak, and F. González, editors, *Concept Maps: Theory, Methodology, Technology. Proceedings of the First International Conference on Concept Mapping*, Volume 1, pages 125–133, U. of Navarra, 2004.

[6] A. Cañas, D. Leake, and A. Maguitman. Combining concept mapping with CBR: Experience-based support for knowledge modeling. In *Proceedings of the Fourteenth International Florida Artificial Intelligence Research Society Conference*, pages 286–290. AAAI Press, 2001.

[7] S. Chakrabarti, M. van den Berg, and B. Dom. Focused crawling: a new approach to topic-specific Web resource discovery. *Computer Networks (Amsterdam, Netherlands: 1999)*, 31(11–16):1623–1640, 1999.

[8] J. Davies, A. Duke, and Y. Sure. OntoShare: a knowledge management environment for virtual communities of practice. In *Proceedings of the international conference on Knowledge capture*, pages 20–27. ACM Press, 2003.

[9] I. S. Dhillon. Co-clustering documents and words using bipartite spectral graph partitioning. *SIGKDD-01*, pages 269–274. ACM Press, 2001.

[10] O. Etzioni and D. Weld. A Softbot-based interface to the Internet. *CACM*, 37(7):72–76, 1994.

[11] A. Farquhar, R. Fikes, and J. Rice. The Ontolingua server: A tool for collaborative ontology construction. *International Journal of Human-Computer Studies*, 46(6):707–727, 1997.

[12] T. R. Gruber. Towards Principles for the Design of Ontologies Used for Knowledge Sharing. In Guarino, N. and Poli, R., editors, *Formal Ontology in Conceptual Analysis and Knowledge Representation*, Deventer, The Netherlands. Kluwer Academic Publishers, 1993.

[13] M. A. Hearst and J. O. Pedersen. Reexamining the cluster hypothesis: Scatter/gather on retrieval results. *SIGIR-96*, pages 76–84, Zürich, CH, ACM Press, 1996.

[14] D. Leake, A. Maguitman, and T. Reichherzer. Understanding knowledge models: Modeling assessment of concept importance in concept maps. *CogSci-04*, pages 785–800, 2004.

[15] D. Leake, A. Maguitman, T. Reichherzer, A. Cañas, M. Carvalho, M. Arguedas, S. Brenes, and T. Eskridge. Aiding knowledge capture by searching for extensions of knowledge models. *KCAP-03*, pages 44–53. ACM Press, 2003.

[16] D. B. Leake, T. Bauer, A. Maguitman, and D. C. Wilson. Capture, storage and reuse of lessons about information resources: Supporting task-based information search. In *Proceedings of the AAAI-00 Workshop on Intelligent Lessons Learned Systems. Menlo Park*, pages 33–37. AAAI Press, 2000.

[17] A. Maguitman, D. Leake, T. Reichherzer, and F. Menczer. Dynamic extraction of topic descriptors and discriminators: Towards automatic context-based topic search. *CIKM-04*, pages 463–472, ACM Press, 2004

[18] F. Menczer, G. Pant, and P. Srinivasan. Topical Web crawlers: Evaluating adaptive algorithms. *ACM TOIT (to appear)*, 2004.

[19] J. Novak. *A Theory of Education*. Ithaca, Illinois, Cornell University Press, 1977.

[20] J. Novak and D. B. Gowin. *Learning How to Learn*. Cambridge University Press, 1984.

[21] N. Noy, R. Fergerson, and M. Musen. The knowledge model of Protégé-2000: Combining interoperability and flexibility. In *Proceedings of EKAW*, 2000.

[22] S. Oyama, T. Kokubo, T. Ishida, T. Yamada, and Y. Kitamura. Keyword spices: A new method for building domain-specific Web search engines. *IJCAI-01*, pages 1457–1466, 2001.

[23] G. Pant, P. Srinivasan, and F. Menczer. Crawling the Web. In M. Levene and A. Poulovassilis, editors, *Web Dynamics: Adapting to Change in Content, Size, Topology and Use*. Springer-Verlag, 2004.

[24] B. Rhodes and T. Starner. The remembrance agent: A continuously running automated information retrieval system. *PAAM-96*, pages 487–495, London, UK, 1996.

[25] S. Staab, J. Angele, S. Decker, M. Erdmann, A. Hotho, A. Maedche, H. Schnurr, R. Studer, Y. and Sure. AI for the Web—ontology-based community Web portals. In *AAAI-2000*, Menlo Park, USA. MIT Press, 2000.

[26] B. Vélez, R. Weiss, M. A. Sheldon, and D. K. Gifford. Fast and effective query refinement. *SIGIR-97*. Philadelphia, PA, pages 6–15. ACM Press, 1997.

[27] O. Zamir and O. Etzioni. Grouper: a dynamic clustering interface to Web search results. *Computer Networks (Amsterdam, Netherlands: 1999)*, 31(11–16):1361–1374, 1999.