Concept Mapping: Connecting Educators
Proc. of the Third Int. Conference on Concept Mapping
Tallinn, Estonia & Helsinki, Finland 2008

AUTOMATIC CLASSIFICATION OF CONCEPT MAPS BASED ON A TOPOLOGICAL TAXONOMY AND ITS APPLICATION TO STUDYING FEATURES OF HUMAN-BUILT MAPS

Alejandro Valerio, David B. Leake, Indiana University, U.S.A
Alberto J. Cañas, Institute for Human and Machine Cognition (IHMC), U.S.A.
Email: (avalerio, leake)@cs.indiana.edu

Abstract. The flexibility of concept mapping enables users to construct a wide range of maps for a given domain. This variation raises the question of what constitute “good” concept maps. This paper contributes towards answering this question by presenting the development and evaluation of a tool for automatic classification of concept maps based on a topological taxonomy. In addition to showing successful demonstrations of the tool’s ability to distinguish novice and expert maps, the paper shows the usefulness of the tool for understanding the types of features that occur in human-generated concept maps.

1 Introduction

Concept maps are an explicit graphical representation of a human’s understanding of a domain of knowledge. Concept maps represent this understanding by means of a two-dimensional network in which nodes correspond to concepts and links correspond to concept relationships. In a concept map, *concepts* are the labels used to refer to objects or events and *linking phrases* (the text on the links) are usually verbs (Novak & Gowin, 1984). Given that each person’s understanding of a domain is different, even if people construct concept maps on the same topic, the maps constructed by every individual are different, reflecting their personal knowledge structures. Concept maps differ from other graphical knowledge representation schemes in the freedom map builders have when selecting concepts and linking phrases, which are limited only by the constraint that concepts linked by a linking phrase should form a propositional sentence, a claim that “makes sense” when read separately from the map. This freedom results in further variation among the maps constructed on a same topic. The resulting flexibility is commonly regarded as an advantage of concept mapping for use in many fields, such as education, in which the map reflects what the student knows, and for representing the knowledge of experts, in which the map is used to represent the idiosyncrasies of each expert.

The freedom for concept map builders results in a variety of concept maps. This variety has at times considered by some to be a “problem.” Computer scientists, for example, often complain about the lack of formality in concept maps, which makes it difficult for computer programs to “process” concept maps, where “processing” could be referring to performing inference on the knowledge expressed on the map, making some kind of decision based on the map, or trying to automatically rate or compare maps. To facilitate such processing, many researchers and practitioners restrict the list of concepts and/or linking phrases from which the map builder can select, resulting in varying degrees of “formality” in the resulting types of maps. Accepting that the flexibility provided by concept mapping in the selection of concepts and linking phrases makes it difficult to perform fully automatic analysis of concept maps, our research efforts have been aimed at the construction of tools that aide the user in all facets of interaction with electronic concept maps, including tools for concept map construction (Cañas, Hill et al., 2004), searching the Web using concept maps (Leake et al., 2004), suggesting concepts during map building (Cañas, Carvalho et al., 2004), categorizing documents based on concept maps (Valerio, Leake, & Cañas, 2007), among others.

A common issue within the concept mapping community is how to “assess” concept maps. Despite the variety of concept maps that arise from the differences among map builders, some maps can be considered “better” than others, based on a variety of criteria. One concept map could show a “deeper understanding” of a topic than another, perhaps reflecting that the first was constructed by an expert and the second by a novice. One map could be considered by some to be “easier to read” than another map on the same topic. Among the features that can be used to assess the quality of a concept map, we can distinguish between topological features (e.g. hierarchical structure, linking phrases, number links into and out of concepts, etc.), and semantic features (are the propositions correct? how expressive are the linking phrases? is the focus question answered by the concept map?). Of course, the semantic content is always more important than the topological structure, but just as in the English language a well structured sentence is considered to be better written than a badly structured sentence, even if their content is the same, a “well structured” concept map is considered better than a badly structured map, even if their contents are “equivalent.” There seems to be some agreement among concept mapping “experts” regarding some of the characteristics of a “well structured” concept map, even they only get a “glance” at them (Carvajal, Cañas, Carballeda, & Hurtado, 2006).

In this paper we report on an effort to automatically classify concept maps based on topological features. Clearly a classification based on topological features is not the automatic assessment rubric that teachers often look for to automatically assess their students' concept maps; we are certainly not attempting to assign an "absolute score" for a map. This work makes three primary contributions: (1) the development and validation of a tool for automatic classification of maps based on a topological taxonomy, (2) development and validation of models of well-constructed concept maps that can be used to facilitate automatic construction of concept maps (Valerio et al., 2007), and (3) the refinement of previous models used for automatic concept mapping support, e.g. the concept Suggester (Cañas, Carvalho et al., 2004). The tool can integrate a broader environment to study the evolution of individual student concept maps in time, to determine how these maps change as the student progresses in both understanding the subject matter and developing skills for constructing concept maps.

2 Topological taxonomy classification and feature annotations

Recent research studies the topological properties of concept maps. (Novak & Cañas, 2006) describe guidelines for concept map construction and identify a list of features that well constructed maps commonly have. Further study resulted in the development of a topological taxonomy for concept maps (Cañas, Novak et al., 2006) based on these features. The taxonomy is intended as a tool to estimate the quality of human-made concept maps based on their structural complexity. It has been successfully used as an evaluation tool for "Proyecto Conéctate al Conocimiento" in Panama (Tarté, 2006) as part of a larger scheme to measure the impact of concept mapping in education. In the context of automatic support of concept mapping tasks, other structural features were identified: High-value concepts in maps, used for searching the Web (Leake et al., 2004), and concept descriptors and discriminators used to describe the topic of a map (Maguitman, 2005), and to associate document libraries to concepts (Reichherzer & Leake, 2006).

Based on the topological taxonomy and other features used in earlier research, we developed a set of algorithms to automatically classify electronic concept maps by their topological category and make annotations of structural features. As described in Section 5, the tools can be used to efficiently annotate large collections of human-built concept maps and assist on quantitative analyses of their structure. If accurate, these tools can also be used to support other concept mapping tasks, for example aiding users by suggesting structural changes to a map under construction.

This section describes the feature annotation algorithms in detail. Most algorithms are based on combinations of previously published work, with some notation changed to have a uniform description. The algorithms make two kinds of annotations: map level and node level annotations. The annotations could be absolute counts (e.g., the total number of concepts) or scores (e.g., hub and authority scores of nodes in the concept map graph). All score annotations are normalized in the range [0, 1].

2.1 Topological taxonomy of maps

The topological taxonomy (Cañas, Novak et al., 2006; Miller, 2008) classifies concept maps into seven levels of increasing structural complexity. In the taxonomy, a map is represented by five features known to be good descriptors of concept map structure: the existence of hierarchical structure, size of concept labels, presence of linking phrases, number of branching points, and number of cross links. Values for these features determine the level of complexity.

To determine the topological class of a map, our automatic topological classifier executes the following procedure: In an iterative process which starts from Level 0, if the map belongs to Level N, check if it belongs to Level $N+1$ until the topmost Level 6 is reached. If during the sequential checking the map fails one of the conditions of Level N, then classify the map as Level $N-1$. Table 1 describes the conditions evaluated by the classifier at each level of the taxonomy. As described in later sections, the tool could be used to analyze large sets of maps. Because some of these maps might be under construction or too small for analysis, consequently not satisfying conditions for any level, for completeness the table introduces a Level -1 to label to such maps.

2.2 Root concept

Structurally, a concept map is a directed graph of concept and linking phrase nodes. The root concept of a map is defined as "the most general and most inclusive concept" (Novak & Cañas, 2006) and is considered the head of the concept map hierarchy, from which all other nodes can be reached though some path. The root is commonly found in the topmost vertical position.

Previous work on automatic concept map processing describes different methods to discover the root concept. In Cañas *et al.* (2001) the root is the concept with the highest upper node score; a measure to characterize concepts that appear in the top of a tree hierarchy. For most concept maps, the upper node score (described later on) can be used as an accurate estimator for the root concept, but due to its recursive definition it assigns the root a very low value when a concept in the lower levels of the tree has a cross link to the root. Another definition of root concept is used in Reichherzer & Leake (2006), where the root is the node with the highest value of $\#outgoing\_links - \#incomng\_links$, however it assigns the root using only limited information about the structure of the map. These definitions are independent of two features that can help find the root more accurately: the "outreach" of the concept (the number of concepts that can be reached though some path starting on the node) and the relative vertical position of the concept.

Level #	**Conditions by level as indicated in Cañas, Novak et al. (2006)**	**Conditions evaluated by to classify concept map *M*, which has concept *c* and linking phrase *l* the as nodes of *M*.**
Level -1	No conditions	(default)
Level 0	At least 4 connected concepts Mostly long concept labels Empty linking-phrases 0 to 1 branching points	$\lvert\{c : concept\}\rvert \geq 4$ (default) (default) (default)
Level 1	More concepts than long concept labels Half or more missing linking-phrases 0 to 1 branching points	$\lvert\{c : concept \mid labelSize(c) < 12\}\rvert \geq \lvert\{c : concept\}\rvert / 2$ (default) (default)
Level 2	More concepts than long concept labels Less than half missing linking-phrases At least 2 branching points	(checked at Level 1) $\lvert\{l : linkingPhrase \mid labelSize(l) > 0\}\rvert \geq \lvert\{l : linkingPhrase\}\rvert / 2$ $\lvert branchingPoints(M)\rvert \geq 2$
Level 3	No long concept labels No linking-phrases missing At least 3 branching points Less than 3 hierarchy levels	$\forall c : concept, labelSize(c) < 12$ $\forall l : linkingPhrase, labelSize(l) > 0$ $\lvert branchingPoints(M)\rvert \geq 3$ (default)
Level 4	No long concept labels No linking-phrases missing At least 5 branching points At least 3 hierarchy levels	(checked at Level 3) (checked at Level 3) $\lvert branchingPoints(M)\rvert \geq 5$ $\exists c : concept, depth(c, M) > 3$
Level 5	No long concept labels No linking-phrases missing At least 5 branching points At least 3 hierarchy levels At least 1 cross-link	(checked at Level 3) (checked at Level 3) (checked at Level 4) (checked at Level 4) $\lvert crossLinks(M)\rvert \geq 1$
Level 6	No long concept labels No linking-phrases missing At least 7 branching points At least 3 hierarchy levels At least 3 cross-links	(checked at Level 3) (checked at Level 3) $\lvert branchingPoints(M)\rvert \geq 7$ (checked at Level 4) $\lvert crossLinks(M)\rvert \geq 3$
Conditions labeled with (default) are met if the map fails a condition of a following level. Conditions labeled with (checked at Level N) are revised at a previous level.		

Table 1: Required conditions for the levels of the topology classifier.

In the approach of developed in this paper, the root concept score is defined as a linear combination of the four features mentioned above. The root concept of a map M is:

$$root(M) = \arg\max_{c \in M} \left(w_U \cdot upperNodeScore(c) + w_O \cdot outgoingScore(c) + w_R \cdot outreachScore(c) + w_P \cdot verticalPositionScore(c) \right)$$

Currently, the values of w_U, w_O, w_R, w_P are 0.1, 0.16, 0.37, 0.37 respectively, which are the best-fit parameters for a small training set of concept maps manually marked with the root concept.

2.3 Cross-links and Cycles in the concept map graph

Safayeni *et al*. (2005) studied the presence of cross-links and cycles in concept maps as indicators of a rich structure and examined their importance in educational settings. According to Novak & Cañas (2006), a cross link is "a link between concepts in different segments or domains of the concept map." For the purposes of our algorithm, a link between two concepts is a cross link if they belong to two different branches of the concept map hierarchy, defined as follows. Concepts *A* and *B* are in two different branches if the shortest path from *A* to the root node $\langle Root,...,A\rangle$ does not share more than *K* nodes with the path of B $\langle Root,...,B\rangle$; in other words, the closest common ancestor of *A* and *B* is at least at distance *K*. This guarantees enough separation between the two concepts for which a cross link is found. In the algorithm implementation *K* was set to 5.

A cycle in a concept map is as defined for directed graphs: it is a path of size more than 1 that starts and ends with the same node. The definition is simple, but an algorithm to find all cycles in an arbitrary concept map can be computationally expensive for large graphs. Fortunately, concept map graphs are relatively small.

2.4 Hub, Authority, and Upper concepts

The definition of hub, authority and upper concepts was introduced in Cañas, Leake, & Maguitman (2001) in the context of automatic support for human concept mapping tasks. These serve as features to help find high value concepts in a concept map graph, which can be used in interactive concept suggesters (Cañas, Carvalho *et al.*, 2004) or automatic identification of terms that describe the topic of a concept map (Maguitman, 2005).

We are interested in calculating the scores for concept nodes only, so we take the full concept map graph $G_M = (V_M, E_M)$ and define a new graph G_C – the concept-to-concept graph of map *M* – as:

$G_C = (V_C, E_C)$, $V_C = \{c : concept, c \in V_M\}$, and $E_C = \{(c_1, c_2) \mid l : linkingPhrase, l \in V_M, (c_1, l) \in E_M, (l, c_2) \in E_M\}$.

The hub and authority scores for concepts in the graph G_C are calculated by a mutually recursive definition, originally described in Kleinberg (1999):

$hubScore(c_1) = \sum_{(c_1,c_2)\in V_C} authorityScore(c_2)$ and $authorityScore(c_1) = \sum_{(c_2,c_1)\in V_C} hubScore(c_1)$

maintaining the constant $\sum_{c\in E_C} hubScore(c)^2 = \sum_{c\in E_C} authorityScore(c)^2 = 1$

The upper node score for the concepts in graph G_C is calculated by a recursive definition as well:

$upperScore(c_1) = 1\ if\ \neg\exists(c_2, c_1)\in V_C\ and\ \sum_{(c_2,c_1)\in V_C} upperScore(c_2)\ otherwise.$

2.5 Lower level features

The features described previously are defined in terms of simpler lower level characteristics defined as follows:

- **Node label size:** The size of a node label is the number of terms separated by whitespace characters. Some maps have concepts that are connected directly without a linking phrase. In this case, an artificial linking phrase is created with an empty label.
- **Concept depth value:** The depth of a concept in a map is the length of the shortest path from the root concept to the node. The depth of the root concept is 0.
- **Branching points in a map:** A branching point is a concept or linking phrase node with more than one outgoing link.
- **Outgoing reach concept score:** The outgoing reach score of a concept *c* is based on the number of concepts that can be reached by walking the map graph exhaustively starting from *c*. The scores of all concepts in a map are normalized in the range [0,1].
- **Outgoing concept score:** Measures the local output connectivity of a concept. The outgoing score of a concept is the normalized value of $\#outgoing\_links - \#incomng\_links$.
- **Vertical position concept score:** Assigns a linear score based on the relative vertical position of the concept in the concept map layout; giving the topmost concept a score of 1 and the bottommost position a 0.
- **Number of concept map resources:** Electronic concept maps can be annotated with multimedia resources and with topically related concept maps. These attached resources could be: concept maps, documents, images, or videos. A simple count is taken for each type.

3 Validation of the automatic topological taxonomy classifier

The topological taxonomy classifier was used to segment a collection of human made concept maps, to study the characteristics of maps on each topology level. In order to assess the classifier's performance, an independent categorization of the maps was needed. Consequently, the algorithm was validated using a preclassified set of concept maps as testing set. The set contained maps of two major categories: (1) expert concept maps: 180 maps built by concept mapping experts; these include STORM-LK (Hoffman, Coffey, Ford, & Carnot, 2001), Mars 2001 (Briggs et al., 2004), and the concept maps in the CmapTools website[1]; and (2) novice concept maps: 108 maps selected manually from a public repository showing different construction deficiencies as described in Novak & Cañas (2006).

Table 2 presents the result for the classification by the topology classifier. Ideally, all expert concept maps would be classified in the higher levels of the hierarchy and all novice maps in the lower levels. The results show that 93% of the expert maps were classified in Levels 4, 5 and 6, while 88% of novice maps were classified in levels 1, 2 and 3, confirming the accuracy of the classifier. The maps on the "expert" set were not manually inspected to verify that all belong to the highest levels of the taxonomy, therefore some noise may be present in this set which can be consistent with relatively few misclassifications.

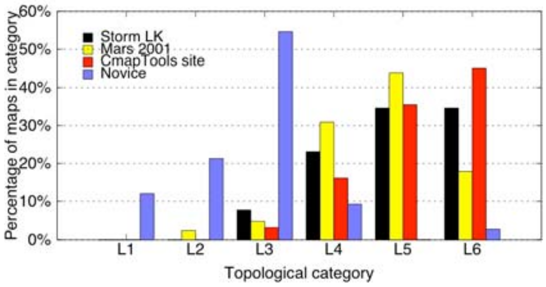

	Storm LK		Mars 2001		CmapTools site		Novice	
	number	%	number	%	number	%	number	%
Level-1	0	0%	0	0%	0	0%	0	0%
Level 0	0	0%	0	0%	0	0%	0	0%
Level 1	0	0%	0	0%	0	0%	13	12%
Level 2	0	0%	3	2%	0	0%	23	21%
Level 3	2	8%	6	5%	1	3%	59	55%
Level 4	6	23%	38	31%	5	16%	10	9%
Level 5	9	35%	54	44%	11	35%	0	0%
Level 6	9	35%	22	18%	14	45%	3	3%

Table 2: Performance of taxonomy classifier on the test set.

The topological taxonomy definition described in Cañas, Novak *et al.* (2006) considers mostly structural characteristics of concept maps to decide their levels. It is unclear if structural features alone are sufficient to achieve accurate classification. Content features could also be explored as additional information about the map to improve the classification performance. Some of these features have been successfully used for other concept mapping tasks, for example to find good descriptor and discriminator terms for the topic of a map (Maguitman, 2005). A semantic rubric has been developed as part of the Conéctate project, which could be evaluated for incorporation into this tool (Miller & Cañas, 2008).

4 Selection, cleanup, and processing of human built concept maps

The second major goal of this work is to analyze the composition of human-made concept maps through an empirical study. Three main sources of concept maps were selected: (1) the 180 expert knowledge maps used during the validation procedure described in Section 3, (2) 8018 maps from the public CmapServer[2] (Cañas, Hill et al., 2004) *IHMC Public Cmaps (3)*'s *Users* directory, and (3) 13,287 maps from the public *Conéctate Público (Panama)* public CmapServer, which mostly correspond to the "Who Am I?" maps of the project (Sánchez et al., 2008). The maps were crawled from the CmapServers using KEA (Cañas, Hill et al., 2006) and downloaded into a local repository to simplify their processing.

Before processing, each map was cleaned to remove concepts that were not connected to the main map. These floating concepts are frequently used to record the focus question of the map and are not considered part of the concept map itself. Occasionally, the author of the map can introduce noise to the concept map graph during map generation, especially by accidentally reversing the direction a link between concepts. These errors are difficult to detect and introduce noise that could affect the performance of the algorithms.

After the clean up step, each map is processed individually and all annotations described in Section 2 are made. Finally, the maps are split by levels using the topological taxonomy classifier and grouped by level to gather the aggregated information.

---

[1] These knowledge models can be accessed at: http://cmapskm.ihmc.us/servlet/SBReadResourceServlet?rid=1103742096094_178662418_7950&partName=htmltext, http://cmex.ihmc.us, and http://cmap.ihmc.us.

[2] A "public" CmapServer is one on which anyone can store concept maps using the CmapTools client, and therefore there is no 'control' or filter on the quality or type of maps stored. The two public CmapServers used in this study can be accessed at: http://cmapspublic3.ihmc.us, and http://cmapspublico.conectate.edu.pa.

5 Characteristics of human built concept maps

This section presents the results of processing maps aggregated by levels of the topological taxonomy. Table 3 presents the distribution of concept maps into topology categories grouped by source. As described in Section 4, the “Expert” category includes all maps from “STORM-LK”, “Mars 2001”, and “CmapTools documentation maps”. We observe that most expert maps are classified in the high level of the taxonomy. Also, the maps from the “Conéctate” project have substantially better classifications than the maps from *IHMC Public Cmaps (3)/Users* server and this result could reflect the impact of the concept mapping workshops in Conéctate.

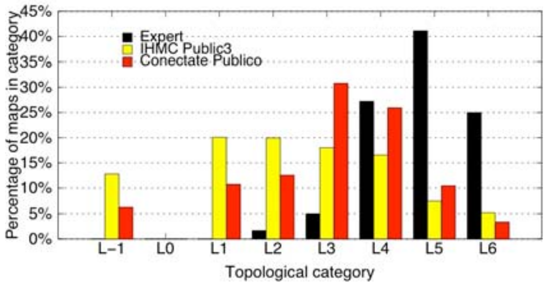

	Expert		IHMC Public3		Conéctate Público	
	number	%	number	%	number	%
Level-1	0	0%	1026	13%	823	6%
Level 0	0	0%	0	0%	0	0%
Level 1	0	0%	1609	20%	1424	11%
Level 2	3	2%	1600	20%	1668	13%
Level 3	9	5%	1444	18%	4093	31%
Level 4	49	27%	1327	17%	3447	26%
Level 5	74	41%	594	7%	1388	10%
Level 6	45	25%	418	5%	444	3%

Table 3: Distribution of maps in taxonomy levels grouped by source

The following tables present variables from maps correlated with their topological taxonomy level. These were obtained by the annotations described in Section 2. Each variable is represented by its average (AVG) normalized by the size of the map and the root mean square error (RMSE) is reported as the standard deviation. All maps from the three sources were combined to get information about the aggregated data.

Table 4 presents the number of cycles and crosslinks by category, where the Levels 5 and 6 show a higher number of these features. The variance of the number of cycles at Level 6 is surprisingly high possibly indicating the existence of a small number of maps with very high connectivity.

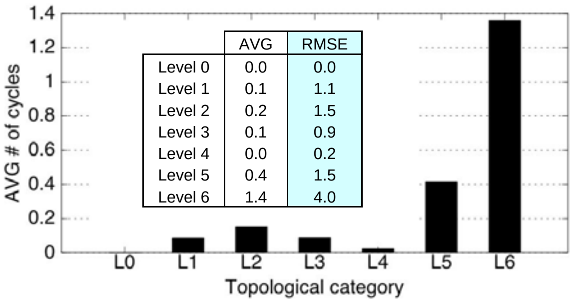

	AVG	RMSE
Level 0	0.0	0.0
Level 1	0.1	1.1
Level 2	0.2	1.5
Level 3	0.1	0.9
Level 4	0.0	0.2
Level 5	0.4	1.5
Level 6	1.4	4.0

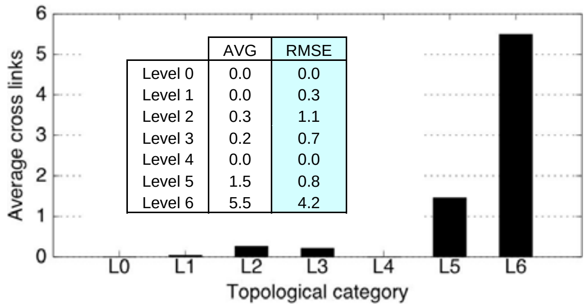

	AVG	RMSE
Level 0	0.0	0.0
Level 1	0.0	0.3
Level 2	0.3	1.1
Level 3	0.2	0.7
Level 4	0.0	0.0
Level 5	1.5	0.8
Level 6	5.5	4.2

Table 4: Average number of cycles and crosslinks per category.

Table 5 shows the average vertical position of the root node in the concept map layout. It shows that regardless of the level it is consistently positioned in the topmost part of the map.

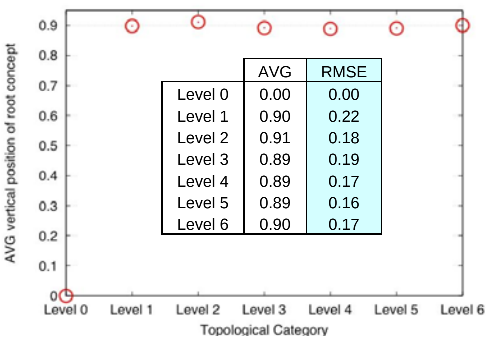

	AVG	RMSE
Level 0	0.00	0.00
Level 1	0.90	0.22
Level 2	0.91	0.18
Level 3	0.89	0.19
Level 4	0.89	0.17
Level 5	0.89	0.16
Level 6	0.90	0.17

Table 5: Average vertical position of the root node in a map layout.

Table 6 shows the distribution of hub and authority concepts. The number of authority nodes is consistently higher than the hub nodes, reflecting the hierarchical structure of concept maps. Also, the proportion of authorities to hubs decreases on Levels 5 and 6, due to cycles and crosslinks from the lower nodes of the map to higher nodes.

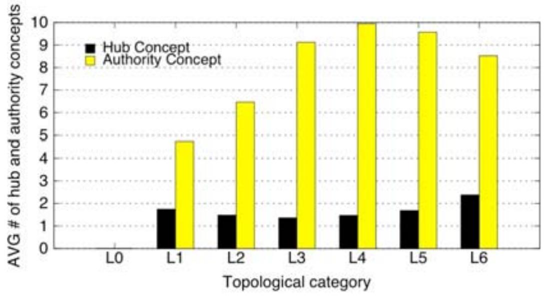

	# hub concepts		# authority concepts	
	AVG	RMSE	AVG	RMSE
Level 0	0.0	0.0	0.0	0.0
Level 1	1.7	2.3	4.7	3.6
Level 2	1.5	1.3	6.5	4.0
Level 3	1.4	1.2	9.1	5.5
Level 4	1.5	1.3	9.9	5.6
Level 5	1.7	1.4	9.6	5.7
Level 6	2.4	2.0	8.5	5.4

Table 6: Distribution of hub and authority nodes by category.

Table 7 shows the distribution of resources per category, which illustrate the increasing number of attached resources as the structure of the map improves. This indicates that authors of more structurally refined maps also refine their maps with more resources and suggest that these maps may be part of a more complex knowledge model.

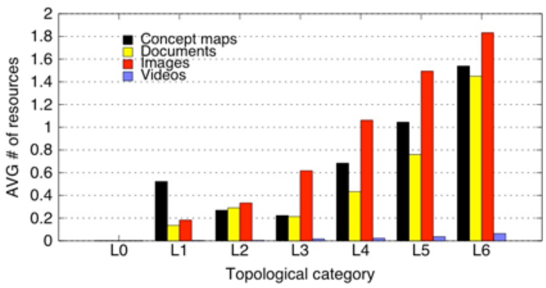

	Concept maps		Documents		Images		Videos	
	AVG	RMSE	AVG	RMSE	AVG	RMSE	AVG	RMSE
Level 0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Level 1	0.5	8.2	0.1	1.2	0.2	1.2	0.0	0.1
Level 2	0.3	1.3	0.3	2.5	0.3	1.6	0.0	0.1
Level 3	0.2	1.2	0.2	1.6	0.6	2.0	0.0	0.3
Level 4	0.7	3.5	0.4	2.7	1.1	3.1	0.0	0.3
Level 5	1.0	3.3	0.8	3.4	1.5	3.4	0.0	0.5
Level 6	1.5	4.9	1.4	6.1	1.8	5.0	0.1	0.6

Table 7: Number of attached resources by category

6 Summary and Future work

In this paper, we present the construction and validation of an automatic classifier of concept maps based on its topology along with a suite of annotators to help analyze concept map topological features. The classifier and annotators were applied to a collection of human built concept maps to identify the structural composition of real-life concept maps.

Further work needs to be done to refine the existing methods to go beyond structural information, bringing in semantic considerations to extend existing models of well constructed concept maps. Also, we foresee the use of the implemented classifier as the basis for a tool to study the evolution of individual concept maps in time by comparing a sequence of snapshots in time through periodic crawling.

In addition, starting from a collection of concept maps annotated with information on concept map quality, the annotators could be used as the basis for a bottom-up analysis of which features are most important for assessing concept map quality, potentially providing more refined assessment methods.

References

Briggs, G., Shamma, D. A., 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* (Vol. I, pp. 109-116). Pamplona, Spain: Universidad Pública de Navarra.

Cañas, A. J., Carvalho, M., Arguedas, M., Leake, D. B., Maguitman, A., & Reichherzer, T. (2004). Mining the Web to Suggest Concepts during Concept Map Construction. In A. J. Cañas, J. D. Novak & F. M. González

(Eds.), *Concept Maps: Theory, Methodology, Technology. Proceedings of the 1st International Conference on Concept Mapping.* (Vol. I, pp. 135-142). Pamplona, Spain: Universidad Pública de Navarra.

Cañas, A. J., Hill, G., Bunch, L., Carff, R., Eskridge, T., & Pérez, C. (2006). KEA: A Knowledge Exchange Architecture Based on Web Services, Concept Maps and CmapTools. In A. J. Cañas & J. D. Novak (Eds.), *Concept Maps: Theory, Methodology, Technology. Proceedings of the Second International Conference on Concept Mapping* (Vol. 1, pp. 304-310). San Jose, Costa Rica: Universidad de Costa Rica.

Cañas, A. J., Hill, G., Carff, R., Suri, N., Lott, J., Eskridge, T., et al. (2004). 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* (Vol. I, pp. 125-133). Pamplona, Spain: Universidad Pública de Navarra.

Cañas, A. J., Leake, D. B., & Maguitman, A. (2001). *Combining Concept Mapping with CBR: Towards Experience-Based Support for Knowledge Modeling*.

Cañas, A. J., Novak, J. D., Miller, N. L., Collado, C. M., Rodríguez, M., Concepción, M., et al. (2006). Confiabilidad de una Taxonomía Topológica para Mapas Conceptuales. In A. J. Cañas & J. D. Novak (Eds.), *Concept Maps: Theory, Methodology, Technology. Proceedings of the Second International Conference on Concept Mapping* (Vol. 1, pp. 153-161). San Jose, Costa Rica: Universidad de Costa Rica.

Carvajal, R., Cañas, A. J., Carballeda, M., & Hurtado, J. (2006). Assessing Concept Maps: First Impressions Count. In A. J. Cañas & J. D. Novak (Eds.), *Concept Maps; Theory, Methodology, Technology. Proceeding of the Second International Conference on Concept Mapping* (Vol. 1, pp. 28-31). San José, Costa Rica: Universidad de Costa Rica.

Hoffman, R. R., Coffey, J. W., Ford, K. M., & Carnot, M. J. (2001). *STORM-LK: A Human-Centered Knowledge Model For Weather Forecasting.* Paper presented at the Proceedings of the 45th Annual Meeting of the Human Factors and Ergonomics Society.

Kleinberg, J. M. (1999). Authoritative Sources in a Hyperlinked Environment. *Journal of the ACM*, *46*(5), 604-632.

Leake, D. B., Maguitman, A., Reichherzer, T., Cañas, A. J., Carvalho, M., Arguedas, M., et al. (2004). Googling from a Concept Map: Towards Automatic Concept-Map-Based Query Formation. 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* (Vol. I, pp. 409-416). Pamplona, Spain: Universidad Pública de Navarra.

Maguitman, A. (2005). *Intelligent Support for Knowledge Capture and Construction.* Unpublished Ph.D. Thesis, Indiana University, Bloomington.

Miller, N. L. (2008). *An Exploration of Computer-mediated Skill Acquisition in Concept mapping by Panamanian in-service Public Elementary Schoolteachers.* Unpublished Ph.D. Thesis (submitted), Universitat Oberta de Catalunya, Barcelona, Spain.

Miller, N. L., & Cañas, A. J. (2008). A Semantic Scoring Rubric for Concept Maps: Design and Reliability. In A. J. Cañas, P. Reiska, M. Åhlberg & J. D. Novak (Eds.), *Concept Mapping: Connecting Educators. Proceedings 3rd Int. Conference on Concept Mapping*. Tallinn, Estonia & Helsinki, Finland: University of Tallinn.

Novak, J. D., & Cañas, A. J. (2006). *The Theory Underlying Concept Maps and How to Construct Them* (Technical Report No. IHMC CmapTools 2006-01). Pensacola, FL: Institute for Human and Machine Cognition.

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

Reichherzer, T., & Leake, D. (2006). Towards Automatic Support for Augmenting Concept Maps with Documents. In A. J. Cañas & J. D. Novak (Eds.), *Concept Maps: Theory, Methodology, Technology. Proceedings of the Second International Conference on Concept Mapping* (Vol. 1). San José, Costa Rica: Universidad de Costa Rica.

Safayeni, F., Derbentseva, N., & Cañas, A. J. (2005). A Theoretical Note on Concept Maps and the Need for Cyclic Concept Maps. *Journal of Research in Science Teaching, 42*(7), 741-766.

Sánchez, E., Bennett, C., Vergara, C., Garrido, R., & Cañas, A. J. (2008). Who Am I? Building a Sense of Pride and Belonging in a Collaborative Network. In A. J. Cañas, P. Reiska, M. Åhlberg & J. D. Novak (Eds.), *Concept Mapping: Connecting Educators. Proceedings of the Third International Conference on Concept Mapping*. Tallinn, Estonia & Helsinki, Finland: University of Tallinn.

Valerio, A., Leake, D., & Cañas, A. J. (2007, October 2007). *Automatically Associating Documents with Concept Map Knowledge Models.* Paper presented at the XXXIII Conferencia Latinoamericana de Informática (CLEI 2007), San José, Costa Rica.