Concept Maps: Theory, Methodology, Technology
Proc. of the Fifth Int. Conference on Concept Mapping
Valletta, Malta 2012

ASSESSING DIVERSITY AND SIMILARITY OF CONCEPTUAL UNDERSTANDING VIA SEMI-AUTOMATED SEMANTIC ANALYSIS OF CONCEPT MAPS

Natalia Derbentseva, Peter Kwantes, David R. Mandel, Defence R&D Canada – Toronto, Canada

Abstract. The development of common conceptual understanding is one of the essential steps in professionalization of a field of practice. To support this process, we collected Concept Maps (CMaps) from practitioners in the field answering the same focus question. Our goal was to expose variability in the understanding of the concept under investigation and to identify common themes, which could be later taken as the basis for the development of a common definition. Collected maps were open-ended, they consisted of varying numbers of concepts and propositions, had different structures, and had little overlap of concepts and concept connections. In order to extract common themes from this set of maps, we used Latent Semantic Analysis to compare CMaps' semantic elements (e.g., concepts, propositions and maps) to each other. Based on this semi-automated analysis we were able to extract a set of concepts and propositions that are semantically representative of the collection.

1 Introduction

Having a shared understanding of common terms and core concepts is an important precondition for communication all professions and disciplines. Clear, consensual definitions reduce the ambiguity and vagueness of terms, thus minimizing the chances for misunderstanding. The development of a common lexicon is one of the essential steps toward a field's professionalization, and this process begins with the assessment of the existing diversity of conceptual meanings. One community of practice where professionalization efforts have only recently begun is intelligence analysis. Intelligence analysis has a long history of practice; however it has been practiced more as a "craft" with largely undefined standards, terminology, and best practices (Bruce & George, 2008). Consequently, many terms lack widely shared definitions. In order to assist the intelligence community in developing a common lexicon, and building on our earlier work applying CMapping in the domain of intelligence analysis (e.g., Derbentseva & Mandel, 2011), we collected definitions of one central concept—analytic integrity—from a number of intelligence practitioners. We used the elicited definitions to assess the variability in the conceptual understanding that exists among the practitioners.

The goal of our study was to expose variability in the understanding of the concept and to identify common themes, which could be later taken as the basis for the development of a common definition. The task of assessing variability in conceptual understanding is not a trivial one and requires integration of methods for data collection and analysis. We used a graphical knowledge representation technique called Concept Mapping (CMapping, Novak, 1998; Novak & Cañas, 2006; Novak & Gowin, 1984) to capture the conceptual understanding of practitioners (see Figure 1 for an example CMap collected during the study), and then focused our analysis on the two main aspects of the CMaps – their structure and their content. The collected CMaps were analyzed and compared along several dimensions. Although there is a great deal of research on scoring and evaluating CMaps in the educational setting, many of these approaches could not be applied to open-ended CMaps in the domain where there is no criterion list of correct propositions or expert-drawn CMaps for comparison.

We chose to use CMapping for data collection because of its unique properties that fit well with the goals of the project. CMapping yields a concise representation of the mapper's subjective understanding of the topic that the map describes. The process of CMapping "…encourages – perhaps even forces – the mapper to reach for crystal clarity about what he or she wishes to express" (Crandall, Klein, & Hoffman, 2006, p.54). Such clarification results in a representation that focuses on essential concepts and their relationships while minimizing word redundancy associated with written texts. A strength of CMapping is that it encourages a mapper to clearly externalize the conceptual relationships that he or she wishes to express, which allows map readers to identify divergent or missing relationships more easily. In addition, CMapping has been successfully applied to various aspects of knowledge elicitation, codification and management (e.g., Coffey et al., 2003; Crandall, et al., 2006; Harter & Moon, 2011; Hoffman & Lintern, 2006; Moon, Hoffman, & Ziebell, 2009). Harter and Moon (2011), in particular, applied CMapping as an intermediate step in the development of a common lexicon in the field of security analysis and risk management. The original definitions in their study, however, were collected in text form and were later converted into CMaps by expert CMappers (rather than expert practitioners).

Analysis of the content of open-ended CMaps without a set of criterion propositions or an expert map for comparison is a laborious task that requires human coding and interpretation. Natural language processing algorithms (NLP) have been applied to automating CMap construction tasks to identify and suggest map elements for inclusion into the map (e.g., Kowata, Cury, & Boeres, 2010; Valerio & Leake, 2006). However, NLP algorithms can also aid in analyzing and comparing content of CMaps on a semantic level. In particular, NLP algorithms can be used to assess semantic similarity of concepts and propositions within and between maps, as well as the semantic similarity among a set of maps.

Latent Semantic Analysis (LSA) is one of the statistical models of language processing that is widely used in information retrieval and study of human memory (Landauer & Dumais, 1997; Landauer, Foltz, & Laham, 1998). Analyzing a large collection of unstructured documents, LSA constructs semantic representations for words based on a statistical analysis of the terms' occurrences across documents. The frequencies are expressed as a matrix, the dimensionality of which is reduced via singular value decomposition. LSA allows one to compare the semantic similarity of words, phrases, and texts, by taking the cosine of the angle formed by their vector representations. Cosine values ranges between -1 and 1, with values closer to 1 indicating higher degree of semantic similarity.

In this study, we used LSA trained on a custom corpus of documents from the intelligence domain. We analyzed the semantic similarity of main contextual elements of the CMaps – concepts and propositions – within each map and between the maps, and we compared the collected maps (as collections of all their propositions) to each other. The following section reports the results of our investigation.

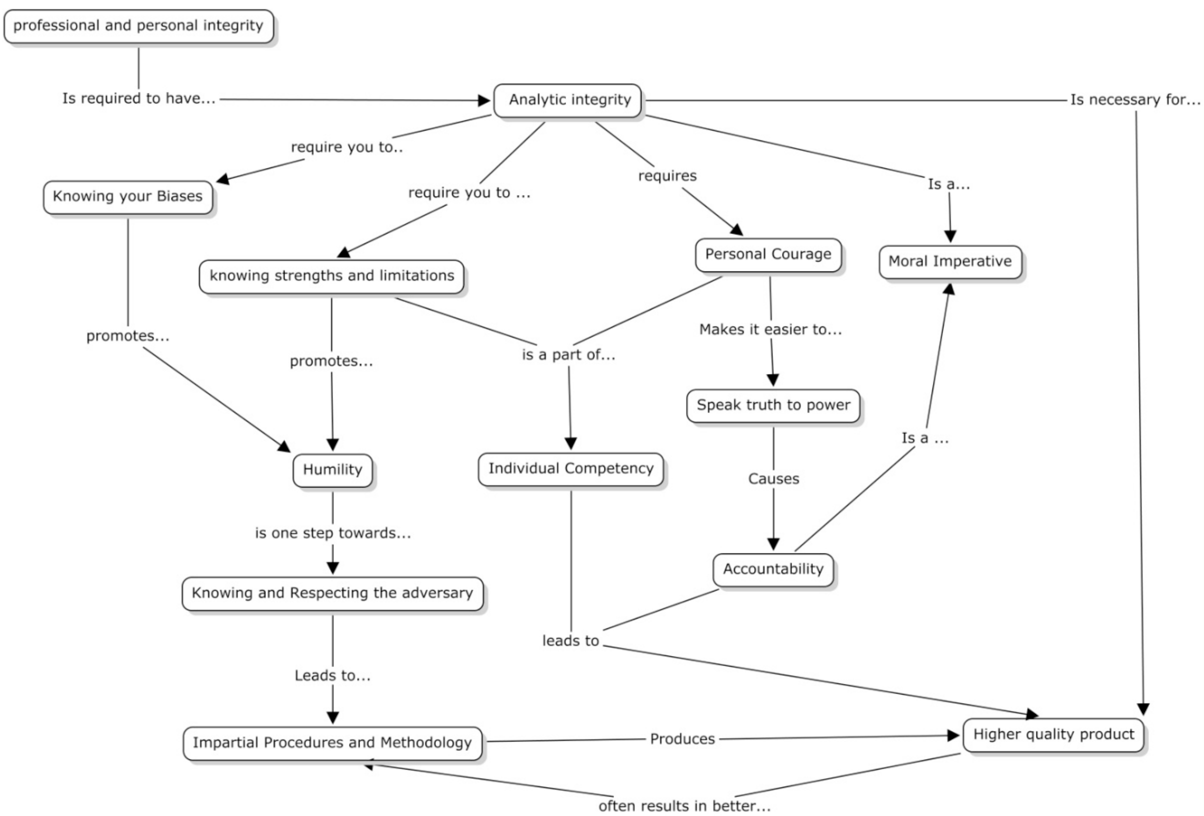

Figure 1. An example CMap constructed by a group of intelligence practitioners to answer the focus question "*What is Analytic Integrity?*"

2 Method

We used CMapping to gather various conceptualizations of *analytic integrity* from practitioners. Instead of collecting textual definitions of the term and then converting them into CMaps, as was done by Harter and Moon (2011), we asked practitioners to construct CMap representations of the term themselves (see Figure 1 for an example of a map collected during the study). This approach served a dual purpose. On the one hand, it introduced practitioners to CMapping as a knowledge representation technique and allowed them to experience its functionality firsthand. We postulated further that such experience is more likely to facilitate adoption of the technique in the future. On the other hand, engaging practitioners in the CMap construction process encouraged them to refine and clarify their ideas on the topic and provide us with a concise representation. This approach removed the necessity to interpret and infer the intended meaning from potentially ambiguous textual accounts; and allowed us to capture the conceptual relationships that practitioners deemed important to represent.

2.1 Data collection procedure

Defence R&D Canada's Human Research Ethics Committee reviewed and approved the study. Overall, we conducted four study sessions, during which we collected definitions of *analytic integrity* from 52 civilian and military practitioners who volunteered to participate in the study. Participants had varying level of experience in the field of intelligence, ranging from novices with little practical experience to senior practitioners with decades of experience in the field. All participants were novice CMappers and the study was their first introduction to the Novakian CMapping technique. Each of the four study sessions had between 8-18 participants. A study session began with a 45-minute training session on CMapping and was followed by a small group collaborative CMap construction task. CMap training was designed for the purposes of this study and included the following components:

- Introduction, a brief history of the development of CMapping and its current applications;
- Main properties of CMaps, including definitions and examples of concepts, linking phrases, propositions, the propositional coherence principle, the overall organization of a CMap;
- The process of CMap construction as outlined in the IHMC flowchart (IHMC, 2006) including exercises that gave participants practice with each of the steps in the CMap construction process;
- Basic introduction to the CmapTools software (IHMC, 2008) that allowed participants to comfortably construct basic CMaps.

After the training, participants formed groups of 4-5 members to collaboratively construct a CMap answering the question, "*What is analytic integrity*?" The map construction task was open-ended with no restrictions imposed. Moreover, no structural constraints were specified and no list of potential concepts or linking phrases was given. Each group worked in separate quarters, equipped with a laptop running CmapTools software (IHMC, 2008) with a new map window opened displaying the focus question. Groups were given about an hour to complete the task. We used a collaborative approach to constructing CMaps for two main reasons. First, the goal of the study was not only to collect conceptual representations from participants, but also to promote the discussion and exchange of ideas on the topic. Group map construction is well suited for this purpose. Participants commented that they enjoyed the collaborative approach, finding it to be intellectually stimulating. Second, participants were novice CMappers and the collaborative exercise allowed them to assist each other in map construction.

2.2 Measures

As noted earlier, the aim of the study was to assess the variability in understanding of the term *analytic integrity* among the intelligence practitioners and to identify similarity in their conceptualizations that could serve as a basis for the development of a common definition. We used several measures to evaluate and compare the maps, including several structural properties of the maps, a semantic analysis of their content, and a subjective overall map evaluation. There was no expert CMap on the topic or a set of correct propositions that could be used for comparison.

2.2.1 Descriptive structural map assessment

The structure of collected maps was analyzed along a number of dimensions. The CmapAnalysis tool (Cañas, Bunch, & Reiska, 2010) was used to calculate some of the measures, such as *average number of propositions per concept* and *map taxonomy score*. Some other measures (e.g., *map density* and *normalized concept centrality*) were calculated based on the network analysis methods (e.g., Knoke & Yang, 2008; Scott, 1991). In the latter case, each map was treated as an undirected binary graph.

In particular, the *map density* measure was calculated as a proportion of links present in an undirected map-network relative to all possible links in the map, and it is an overall measure of a map's interconnectivity. The *average normalized concept degree* measure was computed to show the average proportion of other concepts in a map to which a concept is directly connected. In other words, it is a normalized measure of concept degree centrality when a map is treated as an undirected binary network. Both *map density* and normalized *concept degree* measures are normalized, thus allowing for a direct comparison among maps.

The structural measures were correlated with semantic and subjective measures (both are described below) to determine which measures are closely related.

2.2.2 Semantic measures

The content of the CMaps was analyzed in two ways:

- by identifying identical concepts and propositions among the maps, and
- by assessing semantic similarity of concepts and propositions within and between maps using LSA.

In the former case, identical concepts and propositions among the maps were identified by direct word matching. In the latter case, LSA was performed using Dennis and Stone's (2011) LSA implementation in the SEMMOD package (v 1.5). LSA provides potentially different results depending on which collection of documents (i.e., corpus) it is trained. For example, the term "tree" will have a different meaning (and will co-occur with different kinds of words) in the context of biology that it will in computer science. Therefore, to achieve the best LSA performance, we reasoned that the model should be trained on a collection of documents from the same subject domain to which it is subsequently applied.

A custom training corpus of documents on the topic of "*intelligence analysis*" was created for LSA. The corpus was created using Stone and Dennis's (2012) tool for custom corpus generation from a Wikipedia archive. A set of keywords from the intelligence analysis field was used as a query to find relevant documents, such as "intelligence analysis," "intelligence assessment," "strategic intelligence," "military intelligence," etc. The training corpus included 9,973 documents with 78,896 terms. Several algorithms in the Python programming language (for the Linux environment) were developed to apply LSA to the CMaps and their elements. To facilitate machine processing of CMaps, they were converted to CXL format.

The following LSA analyses were conducted:

1. All pairwise comparisons of concepts in a map: This analysis provides an indication of the *semantic diversity* of concepts in a map. *Concept semantic diversity* was measured as a range between the highest and lowest pairwise concept LSA cosines;
2. Comparison of each concept to its map document. This measure provides an indication of which concepts are most similar to the map as a whole, which can be seen as *semantic concept centrality.* A map document represents a set of all propositions in a map. The contribution of more frequently-occurring concepts to the map document vector was controlled by converting the frequency to its logarithm.
3. All pairwise comparisons of propositions in a map: Similar to concept pairwise comparisons, it provides an indication of *propositional semantic diversity* of a map;
4. Comparison of each proposition to its map document (a map document represents the set of all propositions in that map). This measure provides an indication of which propositions are most similar to the map as a whole; i.e., *semantic proposition centrality*;
5. Each concept from each CMap was also compared to the aggregated map document comprised of all CMaps collected. This measure provides an indication of which concepts are most similar (or central) to the entire collection of maps;
6. Similarly, each proposition was compared to the aggregated map document, indicating which propositions were the most representative of the map collection.

Semantic diversity of a map represents a range of observed cosine values for all pair-wise comparisons of either concepts (the measure of *concept semantic diversity*) or propositions (the measure of *proposition semantic diversity*) in a map. Maps that have some highly similar concepts and some concepts with low similarity will have a greater range of cosine values than maps with all highly similar concepts or all concepts with low similarity. Semantic diversity is an indication of the (moderate degree of) variability of ideas expressed in a map. Based on our measure, both too much variability (i.e., little similarity among the ideas) and too little variability (i.e., all the ideas are relatively similar) in a map will result in a low semantic diversity score.

2.2.3 Subjective rating

Two human raters scored each CMap on the basis of how well a map answered the focus question ("*What is analytic integrity?*") on a scale from 0 ("does not answer the question at all") to 10 ("answers the question very well"), with a good inter-rater consistency (*ICC* = 0.81, $p < .01$). Both raters had experience with CMapping and were familiar with the field of intelligence analysis. The average score from the two raters for each map is reported as its relevance to the focus question score. The subjective ratings represent qualitative evaluations from a reader's perspective and were used compared to other structural and semantic measures.

3 Results

3.1 Overall CMap analysis

We collected 12 group CMaps on the topic of *analytic integrity* (see Figure 1 for an example CMap). The average number of concepts per map was 11.8 (*SD* = 3.2) and the average number of propositions per map was 15.4 (*SD* = 5.5). Table 1 provides a descriptive analysis of each of the maps along several dimensions including the maps' structural properties, taxonomy score, subjective rating, and semantic diversity.

As is clear in Table 1, maps vary considerably along the different measures. Interestingly, the map that was judged to provide the best answer to the focus question by raters, Map 7, received a fairly low taxonomy score, 2. The average map density was 0.13 (SD = 0.04), which indicates that on average, maps included about 13% of all possible connections. Although 13% seems to be a fairly low proportion, it is worth noting that high map density is not necessarily desirable in a CMap, because a large number of propositions could potentially diminish a map's readability and, therefore, its usefulness. The optimum map density level needs to be determined.

#	Level of Experience	Structural measures					Map Taxonomy score	FQ subjective rating	Semantic diversity	
		# Concepts	#of propos.	Map Density	Av. # of propos per concept	Normalized concept degree			Concept	Proposition
1	M	10	14	0.16	1.40	27.8	4	6.00	0.371	0.856
2	M	10	15	0.17	1.50	33.3	2	7.00	0.509	0.739
3	M	13	19	0.12	1.46	23.3	4	8.00	0.845	0.680
4	M	9	10	0.14	1.11	27.5	3	5.25	0.550	0.552
5	M	9	16	0.22	1.78	40.0	4	6.00	0.486	0.795
6	M	13	21	0.13	1.62	26.7	4	7.25	0.427	0.716
7	L	15	26	0.12	1.73	23.2	2	9.25	0.731	0.937
8	L	15	20	0.10	1.33	17.9	5	7.00	0.917	0.806
9	L	10	12	0.13	1.20	26.7	5	6.25	0.682	0.804
10	H	11	13	0.12	1.18	21.7	2	6.00	0.317	0.778
11	H	19	14	0.04	0.74	7.8	5	5.75	0.827	0.744
12	H	8	5	0.09	0.63	15.7	2	4.00	0.501	0.146
Average		11.8	15.4	0.13	1.30	24.3	3.5	6.50	0.600	0.710
StDev		3.2	5.5	0.04	0.36	8.3	1.2	1.40	0.200	0.200

Table1 1. Overall descriptive analysis of the collected CMaps on a number of structural and semantic dimensions

Silva, Romano and Correia (2010) raised concerns regarding the quality of CMaps constructed by novice CMappers and the potential "naïve use of this technique". One of the measures on which novice maps significantly differed from expert maps in the study by Silva et al. was the average number of propositions per concept (the reported average was 1.0 proposition per concept). The observed average number of propositions per concept in the maps we collected was 1.3 (SD = 0.3), and was higher than the average reported for the novice maps in Silva et al. study (t = 2.96, p < .05). This is a positive indication that our CMap training addressed at least some of the concerns experienced by novices.

Most of the five structural measures – *number of concepts*, *number of propositions*, *map density*, *average number of propositions per concept*, and *normalized concept degree* – were significantly correlated. There was no correlation between the *number of concepts* and the *normalized concept degree* measure and the *average propositions per concept*. Also, there was no correlation between the *map density* measure and the *number of propositions*.

Surprisingly, there was no significant correlation between the *map taxonomy score* and any of the other measures, including structural, subjective and semantic.

As expected, there was a positive correlation between *concept semantic diversity* and the *number of concepts* (r[10] = .66, p < .05) and between the *proposition semantic diversity* and the *number of propositions* (r[10] = .72, p < .01). While concept semantic diversity did not correlate with any other measures, there was a positive correlation between *proposition semantic diversity* and the *map subjective rating* (r[10] = .69, p < .05).

This relationship could indicate that map readers look for a moderate range of ideas in a map to provide a rounded answer to their focus question. This relationship requires further investigation.

3.2 Semantic analysis of CMaps

There were total of 142 concepts and 185 propositions in the 12 maps collected. After collapsing identical concepts across maps, the number of unique concepts remaining was 89. Table 2 lists 13 concepts that were shared by three or more maps. Additionally, there were 13 concepts shared by two maps. In the 12 CMaps, only a single proposition matched in two maps in its entirety—namely, “Analytic Integrity requires Honesty”. Four additional pairs of concepts were linked in two maps each, but with different linking phrases. These are indicated in the fifth column of Table 2, *Linked to the same concepts in two maps*.

	Concept	# of maps	% of maps	Linked to the same concepts in 2 maps	Average structural centrality	Semantic similarity to the aggregated map document (as percentage of the highest value)
1	Analytic integrity	8	67	Accountability, Credibility, Honesty	4.5	58 (72)
2	Honesty	6	50	Analytic integrity*	2.5	12 (15)
3	Credibility	5	42	Analytic integrity*, Quality, Sources	2.2	23 (29)
4	Objective	5	42	None	2.6	28 (35)
5	Accountability	4	33	Analytic integrity*	2.5	27 (33)
6	Accuracy	3	25	None	1.7	39 (48)
7	Bias	3	25	None	1.7	34 (42)
8	Evidence	3	25	None	2.7	29 (35)
9	Logical	3	25	None	1.7	36 (44)
10	Responsibility	3	25	None	3.3	28 (35)
11	Sources	3	25	None	4.0	39 (48)
12	Being thorough	3	25	None	2.3	24 (30)
13	Experience	3	25	None	3.3	36 (44)

* Represents the same links as indicated for the concept *Analytic Integrity* in the first row

Table 2. Concept and concept link overlaps in the CMap set

Table 2 also reports the average structural centrality of the concepts that appear in more than one map as well as their semantic similarity to the aggregated map document (i.e., a document that consists of all 185 propositions from the 12 maps), and the normalized measure relative to the highest observed semantic similarity measure is included in the parentheses.

This analysis reveals that there is little direct overlap in the concepts and their relationships that different groups used to express their understanding of *analytic integrity*. However, it does not necessarily mean that the maps are strikingly different. In fact, we conducted all pairwise LSA comparisons of the maps, and the average cosine value was 0.78 (SD = 0.08), which indicates a considerable semantic overlap in the maps despite the lack of the directly matching concepts and propositions.

The next step in our analysis was to identify concepts and propositions that are representative of the collection of the maps. This was achieved by comparing each of the concepts and each of the propositions to the aggregated map collection. The top 10% of the concepts having the highest semantic similarity with all the maps are listed in Table 3. These concepts have reasonable face validity. Note that there is not much overlap between the most frequently occurring concepts (see Table 2) and the list of most semantically similar concepts. The top 10% of propositions with highest cosine values are presented in a CMap in Figure 2.

Based on the semantic analysis, the map elements in Table 3 and Figure 2 are most representative of the overall collection of the maps generated by professionals. These elements could be taken as the basis for the development of a common definition by the community practitioners. Note that Figure 2 shows only those propositions that exist in the map collection, and thus is segmented. The map segments can easily be connected by adding relationships.

4 Conclusion

Analyzing CMaps’ content and identifying commonalities in a set of maps with very little verbal overlap can be a daunting task. NLP algorithms could facilitate this process. With LSA we were able to extract a set of concepts and propositions that, according to our analysis, are representative of the collection of CMaps about analytic integrity that we collected from practitioners. An advantage of this approach is that the selection of most representative elements can be fully automated. Although this process will not produce a finished map or a concept definition, it can serve as a valuable starting point and input stage for later human assessment.

#	Concept	Cosine	Norm.	#	Concept	Cosine	Norm.
1	Broad subject knowledge	0.809	1.00	9	Policies and procedures	0.604	0.75
2	Quality sources	0.778	0.96	10	Logical thinking	0.594	0.73
3	Critical thinking	0.692	0.86	11	Analytic integrity	0.584	0.72
4	Impartial procedures and methodology	0.673	0.83	12	Personal ethics	0.577	0.71
5	Non-bias	0.639	0.79	13	Higher quality product	0.576	0.71
6	Quality control	0.635	0.78	14	Analysis	0.574	0.71
7	Professional and personal integrity	0.633	0.78	15	Comprehensive examination	0.566	0.70
8	Individual competency	0.632	0.78	16	Speak truth to power	0.562	0.69

Table 3. List of the top 10% of concepts based on their semantic similarity to the aggregated map collection

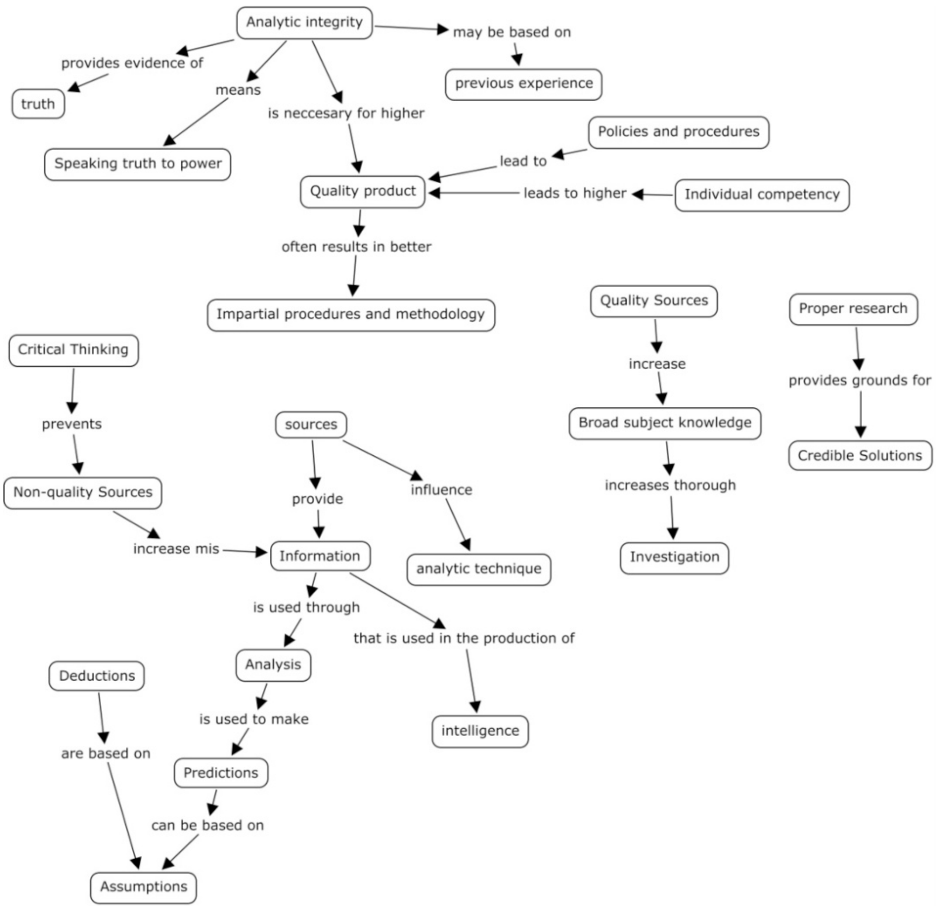

Figure 2. A CMap representation of the top 10% of propositions based on their semantic similarity to the aggregated map collection

Semantic analysis, such as LSA, can also facilitate CMap structural analysis. Most of the structural analyses of CMaps ignore semantic information carried by the linking phrases, treating all propositions as equal connections (and often ignoring their directionality). The requirement of labeling all concept connections is one of the distinguishing features of CMapping. Linking phrases carry information and differentiate among propositions. Structural analysis of CMaps requires methods that are able to take this information into account. One potential solution to this problem is to differentiate among propositions in a map based on their semantic centrality (as measured by a cosine between a proposition vector and a map vector). Weights could be assigned to links in a map network based on the propositions' semantic centrality, which would allow transforming a graph of a map from binary into a valued network. The integration of semantic and structural analyses of CMaps is an area that deserves further investigation.

References

Bruce, J. B., & George, R. Z. (2008). Intelligence Analysis - The Emergence of a Discipline. In R. Z. George & J. B. Bruce (Eds.), Analyzing Intelligence: Origins, Obstacles, and Innovations (pp. 1-15). Washington, DC: Georgetown University Press.

Cañas, A. J., Bunch, L., & Reiska, P. (2010). CmapAnalysis: An extensible Concept Map analysis tool. In J. Sánchez, A. J. Cañas, J. D. Novak (Eds.), *Concept Maps: Making Learning Meaningful, Proc. of the Fourth Int. Conference on Concept Mapping*, Viña del Mar, Chile: Universidad de Chile.

Coffey, J. W., Cañas, A. J., Hill, G., Carff, R., Reichherzer, T., & Suri, N. (2003). Knowledge modeling and the creation of El-Tech: a performance support and training system for electronic technicians. Expert Systems with Applications, 25(4), 483-492.

Crandall, B., Klein, G., & Hoffman, R. (2006). Working Minds: A Practitioner's Guide to Cognitive Task Analysis. Cambridge, MA: MIT Press.

Dennis, S., & Stone, B. (2011). SEMMOD: Semantic models package (Version 1.5): Ohio State University.

Derbentseva, N., & Mandel, D. R. (2011). Using Concept Maps to improve the practice and organization of intelligence in Canada. In B. Moon, R. Hoffman, J. D. Novak & A. J. Cañas (Eds.), Applied Concept Mapping: Capturing, Analyzing, and Organizing Knowledge (pp. 109-130). New York, NY: Taylor & Francis Group: CRC Press.

Harter, A. G., & Moon, B. (2011). Common Lexicon Initiative: A Concept Mapping Approach to Semiautomated Definition Integration. In B. Moon, R. Hoffman, J. D. Novak & A. J. Cañas (Eds.), Applied Concept Mapping: Capturing, Analyzing, and Organizing Knowledge (pp. 131-150). New York: Taylor & Francis Group: CRC Press.

Hoffman, R., & Lintern, G. (2006). Eliciting and Representing the Knowledge of Experts. In K. A. Ericsson, N. Charness, J. Feltovich & R. Hoffman (Eds.), Cambridge Handbook of Expertise and Expert Performance (pp. 203-222). New York: Cambridge University Press.

IHMC. (2006). Steps in building a Concept Map. Retrieved 15 November, 2010, 2010, from http://cmapskm.ihmc.us/servlet/SBReadResourceServlet?rid=1064009710027_279131382_27088&partName=htmltext

IHMC. (2008). CmapTools software (Version 4.18). Pensacola, Florida.

Knoke, D., & Yang, S. (2008). Social Network Analysis. Thousand Oaks, CA: Sage Publications.

Kowata, J. H., Cury, D., & Boeres, M. C. S. (2010). Concept Maps core elements candidates recognition from text. In J. Sánchez, A. J. Cañas, J. D. Novak (Eds.), *Concept Maps: Making Learning Meaningful, Proc. of the Fourth Int. Conference on Concept Mapping*, Viña del Mar, Chile: Universidad de Chile.

Landauer, T. K., & Dumais, S. (1997). A Solution to Plato's Problem: The Latent Semantic Analysis Theory of Acquisition, Induction, and Representation of Knowledge. Prychological Review, 104(2), 211-240.

Landauer, T. K., Foltz, P. W., & Laham, D. (1998). Introduction to Latent Semantic Analysis. Discourse Processes, 25(2-3), 259-284.

Moon, B., Hoffman, R., & Ziebell, D. (2009). How did you do that? *Electric Perspectives*, 20-29.

Novak, J. D. (1998). Learning, creating, and using knowledge: Concept Maps(R) as facilitative tools in schools and corporations. Mahwah, NJ: Erlbaum.

Novak, J. D., & Cañas, A. J. (2006). The Origins of the Concept Mapping Tool and the Continuing Evolution of the Tool. Information Visualization Journal, 5(3), 175-184.

Novak, J. D., & Gowin, D. B. (1984). Learning how to learn. New York: Cambridge University Press.

Scott, J. (1991). Social network analysis: a handbook. Thousand Oaks, Calif.: Sage.

Silva Jr., S. N., Romano, J. G., & Correia, P. R. M. (2010). Structural Analysis of Concept Maps to Evaluate the Students' Proficiency as Mappers. In J. Sánchez, A. J. Cañas, J. D. Novak (Eds.), *Concept Maps: Making Learning Meaningful, Proc. of the 4th Int. Conf on Concept Mapping*, Viña del Mar, Chile: Univ. de Chile.

Stone, B., & Dennis, S. (2012). Wikipedia Subcorpora Tool (Version 1.0): Ohio State University.

Valerio, A., & Leake, D. (2006). Jump-Starting Concept Map Construction with Knowledge Extracted from Documents. In A. J. Cañas & J. D. Novak (Eds.), *Concept Maps: Theory, Methodology, Technology. Proceedings of the 2nd Int. Conf. on Concept Mapping*. San José, Costa Rica: Universidad de Costa Rica.