AMERICAN INTELLIGENCE JOURNAL

THE MAGAZINE FOR INTELLIGENCE PROFESSIONALS

New Paradigms in Intelligence Analysis

NMIA

Vol. 32, No. 1, 2015

Propositional Diagrams for Intelligence Sensemaking:

Examples and Case Studies

by Dr. Robert R. Hoffman, Tom Eskridge, Simon Henderson, Jonathan Jenkins, and Brian Moon

INTRODUCTION

Newspaper reports of intelligence-related activities conducted by the U.S. Intelligence Community have occasionally referred to large "Pens and Post-Its" wall charts that were created, for example, to understand and represent adversarial networks or the structures of NGOs. Similar practices have been reported in other nation's intelligence services, such as in the UK. Another rationale for such diagramming is that collaboration requires the externalization of understanding and supporting conversations, to achieve a shared understanding in which concepts and their meanings are made precisely clear. Team creation and analysis of meaningful diagrams encourage, even force, people to achieve consensus and clarity. When the information that is being shared is critical understanding or intent, upon which lives may depend, there is a clear imperative that both sender and receiver do everything in their power to ensure that shared information leads to shared, and accurate, understanding.

Recent guidance on analytical methodology has included recommendations regarding the applications of such meaningful diagrams (Heuer and Pherson, 2011), but recommendations have not been accompanied by realistic, detailed examples showing how to make good diagrams and how they can be used to best results. This article presents guidelines on what makes for a good meaningful diagram, and expresses how meaningful diagrams can support the process of intelligence sensemaking.

MEANINGFUL DIAGRAMS

Research on diagrammatic reasoning, from fields spanning geography, statistics, and instructional design, has investigated the value of maps, schematic diagrams, and many other forms of diagram. It has converged on a set of conclusions concerning the value of diagramming, and offers an explanation for why diagramming has value (see Mandl & Levin, 1989; Vekirl, 2002). Diagrams can "externalize" cognition, guide reasoning, reduce cognitive demands, support working memory, present information "at a glance," and shift some of the burden of text processing over to the visual perception system. In a team context, diagrams can support dialogue, help uncover hidden assumptions, facilitate the development of shared understanding, and act as a tool for supporting the communication of meaning and intent.

Concept Maps are meaningful diagrams composed of labeled nodes (concepts) and relational links. The original form of meaningful diagram called Concept Mapping was invented in the 1970s by Joseph Novak of Cornell University, who was interested in capturing the knowledge of school children (Novak and Gowin, 1984). Since then, an extensive background and substantive research foundation has validated Concept Mapping for a variety of applications spanning primary education to professional brainstorming (Moon, Hoffman, Cañas, & Novak, 2011). Concept Mapping encourages critical thinking (Mintzes, Wandersee & Novak, 2000), and results in measurable gains in knowledge. Building good Concept Maps leads to longer retention of knowledge and greater ability to apply knowledge in novel settings (Cañas et al., 2003; Mintzes et al., 2000; Novak, 1991, 1998). Of particular interest in this article is the use of Concept Maps to express expert knowledge and to capture the complex concepts and relations involved in analytical problems (Crandall, Klein and Hoffman, 2006).[1]

An example Concept Map is presented in Figure 1. This diagram represents an attempt to explain a contradiction about refugee status—why the Bhiari refugees do and do not qualify as refugees under international law (Faranza, 2008). The icons beneath some of the nodes hyperlink to text pieces and URLs that present supporting evidence. The full set of diagrams that captured the analysis of the Bihari situation consisted of 15 Concept Maps, including this one.

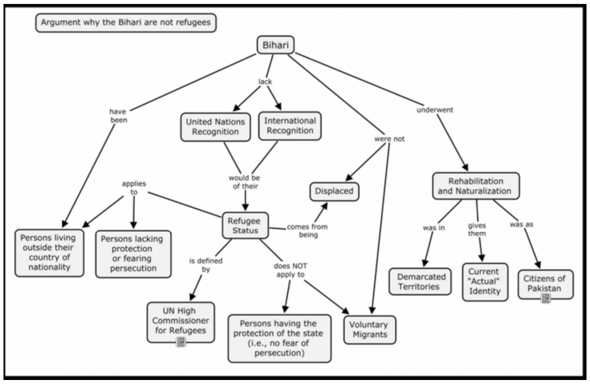

Figure 1. A Concept Map used in the process of sensemaking about the Bihari refugee situation.

Concept Maps include concepts and relationships among concepts indicated by lines linking two concepts. Linking phrases specify the relationships. The most important concepts and the context are generally toward the top and the more detailed concepts and relations generally toward the bottom. The diagram morphology is generally hierarchical (although the cross-links mean that the diagram is not a hierarchy in a graph-theoretic sense). In a well-formed Concept Map, each node-link-node triple can be read as a stand-alone proposition, a feature that renders Concepts Maps searchable on the basis of more than just a "bag of words."

It should be noted that Concept Maps are not "argument maps." Concept Maps do not restrict the propositions in terms of semantics of argument structure. Therefore, for instance, Concept Maps can represent causal relations or temporal structures as well as arguments.

With the support of the "CmapTools" freeware, Concept Mapping is being used around the world, at all levels of education (Cañas, 1999; Cañas et al., 2003) and in many locations as part of the core infrastructure in schools and entire school systems (Ford, Coffey, Cañas, Andrews & Turner, 1996). Concept Mapping has come to be used in numerous government and business applications as well (Moon et al., 2011). Concept Mapping is being used to create and edit ontologies for intelligent decision aids and use on the semantic Web (Eskridge and Hoffman, 2013).

ADVANTAGES FOR ANALYTICAL WORK

Consider the practical advantages of such diagrams over typical analytical worksheets or data matrices. While spreadsheets or synchronization matrices make an analyst record and analyze certain kinds of information in certain ways, they are usually not a useful tool for conveying meaning to others—i.e., what the "big picture" is, or the "so what?"

Concept Mapping supports a number of cognitive processes that are crucial to critical thinking and fluid intelligence (Hoffman et al., 2011): Assimilation (changing current knowledge as a result of the discovery of new knowledge), differentiation (distinguishing sub-concepts and their relations), superordination (seeing how previously unrelated concepts are in fact related), subsumption (seeing how previously unrelated concepts actually fall under a higher-order concept), and reconciliation (achieving coherence and consistency). Concept Maps made by domain experts tend to show high levels of agreement (see Gordon et al., 1993; Hoffman, Coffey & Ford, 2000). Reviews of the literature and detailed discussion of methods for making good Concept Maps can be found in Cañas et al. (2004), Crandall, Klein & Hoffman (2006), and Moon et al. (2011).

It is important to think of Concept Mapping as a process, versus the qualities of finished Concept Maps. Technically stated, when creating a Concept Map, the Mapper uses spatiality (i.e., different areas of the diagram space) as a tool to de-convolute meanings. As nodes and partially-linked sets of nodes are grabbed and moved around in the diagram space, the Mapper considers various relationships and ideas to be expressed. The Mapper struggles to add in cross-links while avoiding the creation of a "spaghetti graph" having too many overlapping cross-links. Clusters of nodes will be parked somewhere, and that region of the Concept Map space becomes, in effect, a memory aid.

In a study conducted with the support of DARPA's "Rapid Knowledge Formation" project, Concept Maps were made by domain experts but were subsequently "tidied up" overnight by computer scientists. Upon next seeing their Concept Maps so tidied up, the experts were upset because things "weren't where they were supposed to be" (Hayes, personal communication, 2003). The Mappers had been using diagram spatiality as a tool.

Node-link-node triples essentially make Concept Maps a surface notation for propositional logic. The expression of meanings in terms of propositions is central to the construction of effective Concept Maps and the meaningful capture of knowledge. CmapTools includes capabilities for extracting and sorting propositions, automated suggestions of related concepts through searches on the Web, automatic layout tools, recording and playback of the stages involved in diagram construction, validation of map coherence, and automatic fixing of broken links.

Additional advantages of Concept Mapping for analytical work stem from the capabilities of the CmapTools freeware. For example, the ability to hyperlink digital "resources" such as text documents, images, video clips, and website addresses is a significant capability. Hyperlinks to resources are indicated by the small icons underneath concept nodes (see Figure 1). Indeed, a Concept Map explaining the situation under analysis and integrating the resources can be the analyst's report. This represents a potential major gain in efficiency, perhaps eliminating the need to turn an analysis into a slide presentation.

CmapTools has a number of capabilities that aid networking. This includes the ability to search on the occurrence of a concept to see if the same concept also appears in other people's Concept Maps. If it does, then there exists the ability to "pull in" another Concept Map to incorporate within one's own, or to point to the other Concept Map with a link attached to the concept of interest. In this manner, knowledge literally becomes networked and the capability for knowledge reuse is created. CmapTools contains a range of functions for data import and export that can assist with better understanding of the data (for example, concept propositions and hierarchies can be exported as text).

Example #1: Thinking in Terms of Propositions

We use this case study to illustrate the process of creating a Concept Map-based analysis.

> **"Thai protesters build barricades and toy with talks." BBC News Website, April 21, 2010**
>
> *Thai anti-government protesters have built formidable barricades of tyres and sharpened bamboo canes in Bangkok as tensions build in the capital. But tentative hints of possible new talks between protesters and the government have emerged, as parliament met for the first time in two weeks. Troops remain behind lines nearby in an increasingly militarised standoff. The red-shirts are demanding that Prime Minister Abhisit Vejjajiva step down and parliament is dissolved. However, analysts say both sides might feel the need for talks as the prospect of another bloody crackdown looms. A failed attempt to clear protesters on 10 April left 25 people dead. [http://news.bbc.co.uk/1/hi/world/asia-pacific/8633893.stm]*

The first step is to extract the concepts of interest (Figure 2). (In CmapTools, clicking on the diagram space creates a concept node, ready for typing the concept label.)

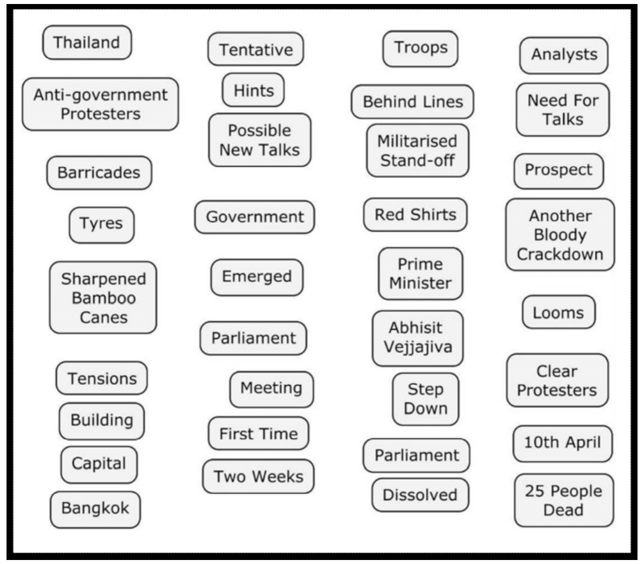

Figure 2. Concepts in the Thai Protest article.

The next step involves grouping related concepts spatially. This helps direct thinking about the relationships that exist among concepts that are related to each other.

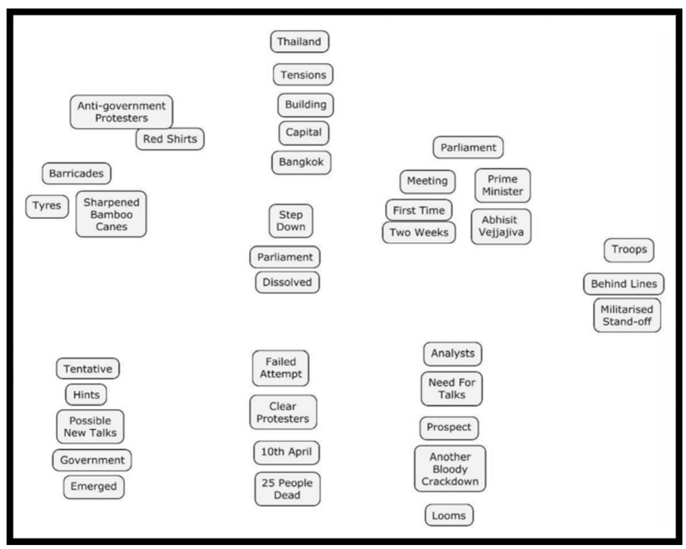

Figure 3. Concepts are grouped.

The next step is to link the concepts and label the relationships (links are easily created in CmapTools by point and drag, and then the typing of the linking phrase). The Concept Map should seek to exhibit "propositional coherence." This means that every node-link triple should make sense when read alone, for example,

[Anti-Government Protesters] <are known as> [Red Shirts]

Thinking in terms of propositions is a skill that takes some practice. Learners often fall back on the inferences and concepts that are tacit in ordinary syntax. For example, the phrase *parliament met for the first time in two weeks* has the concept of "meeting" expressed as a verb, and when going from text to propositions people sometimes mistakenly make all the verbs into linking relations. Those who are new to Concept Mapping might craft the phrase *protesters have built formidable barricades of tires* as

[protesters] <have built> [formidable barricades] <of> [tires]

when propositionally the source text reads as:

[protesters] <have built> [barricades]
[barricades] <are> [formidable barricades]
[formidable barricades] <are made of> [tires]

Note also that the expression:

[formidable barricades] <of> [tires]

does not read as a stand-alone proposition.

Another crucial activity is the creation of cross-links; Concept Maps are not "pure" hierarchies, but instead accommodate the complexity and interconnectedness of ideas and events. A Concept Map representing all of the propositions in the Thai protest text is presented in Figure 4. The reader is invited to look for things in Figure 4 that might be improved, our point being that there is not necessarily one single "best" or "right" way to decompose and represent open text.

Example #2: Hypothesis Exploration

As CMappers review and reorganize their thinking, the Concept Map undergoes various transformations, revisions, additions, and deletions. The Concept Map can also support the representation of causality, temporality, uncertainty, and inference—all features that are critical in analysis. An assertion (in contrast with a "fact") is a statement or declaration, often expressed without supporting evidence or accompanying reasoning. A suggested graphical method for capturing assertions is to code assertions in the linking phrase using both color and symbology since redundant encoding is easier to process. This is illustrated in Figure 5. In this case, we know that it has been asserted that Tribe X has been accused of political corruption, but we do not know where this assertion has come from, or whether it is true.

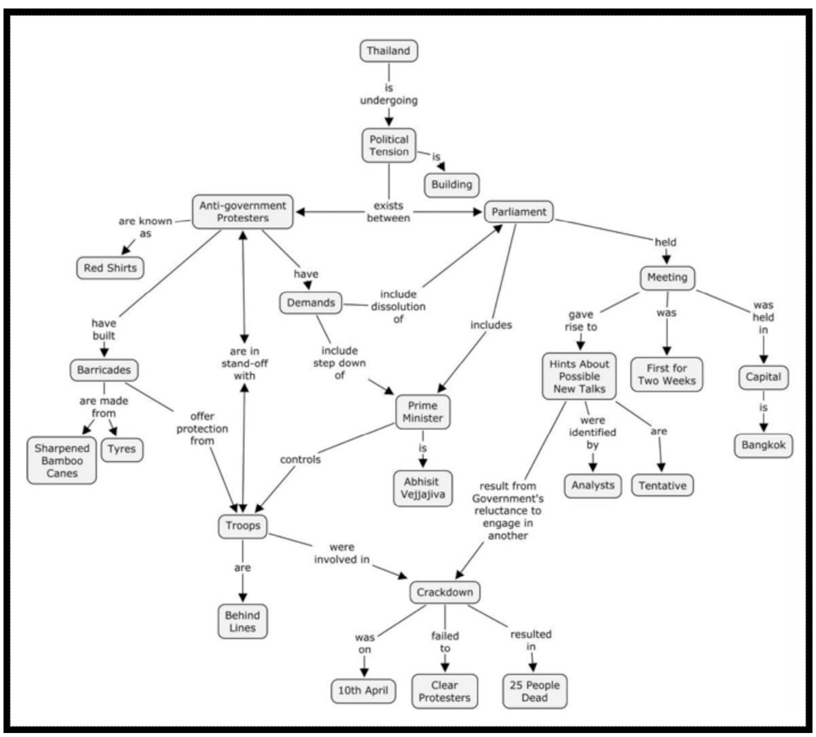

Figure 4. A completed Concept Map based on the "Thai protests" article.

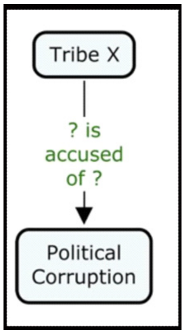

Figure 5. Assertions can be represented using color coding and together with front and back question marks.

In Figure 6, an inference is made that, if there is political corruption present in Tribe X, then there are officials within the tribe who are themselves corrupt. Note that Figure 6 also indicates that two linked concepts are inferred from the premise of political corruption. The inference is not just that Tribe X has officials but that those officials are corrupt. Thus, the proposition is "nested," indicating that it is a proposition that is inferred. As with the representation of assertions, the use of the double question mark to label the linking phrase in an inference enables it to be output collectively as ordered triples for intelligent search.

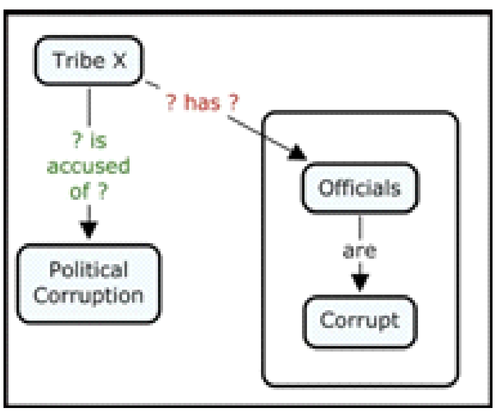

Figure 6. Inferences can be represented using color coding along with front- and back-end question marks

Example #3: Inferencing

These techniques will now be exemplified by analysis of the following text. Although this is comprised of a simple sentence, it contains a great deal of information and can be used to speculate about a wide range of issues. The analysis involves not just extracting the key information contained in the text, but converting the information into a form that is propositionally coherent and inferring additional concepts and relationships.

BBC News Website, July 13, 2010

An Iranian nuclear scientist at the centre of an abduction row between the United States and Iran is free to leave, the US State Department says.
[http://www.bbc.co.uk/news/10617656?utm_source=twitterfeed&utm_medium=twitter]

In Figure 7, a number of relationships are inferred and diagrammed explicitly. For example, while the nuclear scientist is said to be at the center of the abduction dispute, it is not explicitly stated that it was he who was abducted. However reasonable it might be to infer this, the relationship is expressed and color-coded as an inference. Notice also the creation of (and notation for) an inferred concept (? Released ?), and also the use of intersecting nested nodes to capture relationships among sets of concepts. This analysis shows clearly how even simple assertions can contain many implications and entailments, which get "hidden" by the syntactic conventions of ordinary language. It also shows how one proposition can refer to, or comment on, another proposition.

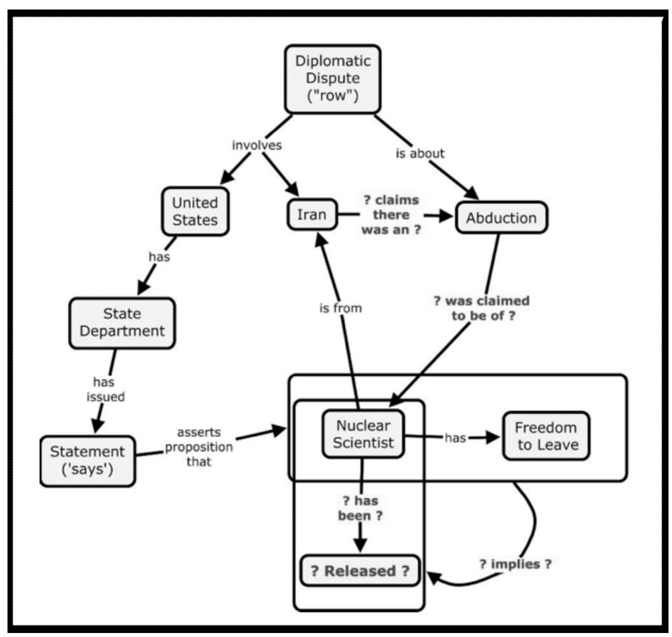

Figure 7. A Propositional-Inferential representation.

Taking the analysis a step further can involve adding annotations, comments, and questions to make the analyst's thinking process an explicit part of the representation. In Figure 8 a range of questions has emerged, which are recorded as their own supplementary Concept Map.

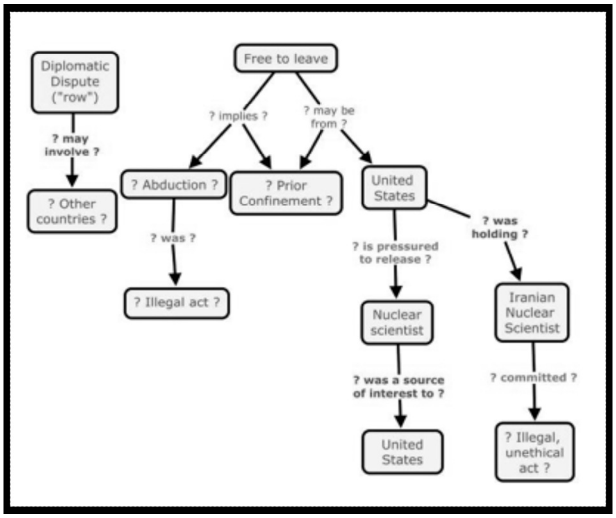

Figure 8. Questions that emerge from the inferences.

COMMUNICATING THE BIG PICTURE

When mapping a large domain or more extensive set of material, a single Concept Map can become unmanageable for the user to comprehend, display, and manipulate. To facilitate the construction of large representations, the CmapTools allows the user to create and hyperlink collections of Concept Maps, enabling the navigation from one Concept Map to another. Another capability is that different authors (for example, from different functional desks or technical disciplines) can support the generation of integrated collections of Cmaps, each from its own perspective, and show explicitly how their map relates to those produced by others. Cmap Tools provides the ability collaboratively and synchronously or asynchronously to construct a joint Concept Map (see Cañas, Suri, Sánchez, Gallo, & Brenes, 2003).

Hyperlinks can connect Concept Maps to other Concept Maps; a set of Concept Maps hyperlinked together is regarded as a "Knowledge Model." The hyperlinking permits navigation among the Concept Maps and serves as a navigational tool that prevents "getting lost in hyperspace." Within the context of a Knowledge Model, an overarching Concept Map can be created to communicate the "big picture." Individual concepts within this Cmap then link to detailed Cmaps that expand the concept into a series of lower-level components. Also part of the big picture, any digital resource can be hyperlinked into a Cmap, bringing in the supporting evidence: imagery, reference documents, video, and websites. Knowledge Models can serve as living repositories of expert knowledge to support knowledge sharing as well as knowledge preservation. This too represents a significant capability for preserving and sharing organizational expertise (Ford et al., 1996). In capturing the expert knowledge within an organization, practitioners can always add to and modify the Concept Maps in the existing pool.

One such Knowledge Model is called STORM (System To Organize Representations in Meteorology) (Hoffman et al., 2001, 2006). It consists of two dozen Concepts Maps created by forecasters at the Naval Training Oceanographic and Meteorology Facility at Pensacola Naval Air Station, FL. Although the project involved creating many dozens of Concept Maps about all aspects of weather and weather forecasting, the knowledge model focuses on weather of particular concern to naval aviation in the Gulf Coast region (e.g., turbulence, fog, thunderstorms, and hurricanes). It covers forecasting processes, such as the use of the radar. STORM Cmaps can be viewed at [http://cmapspublic.ihmc.us/rid=1147120059423_996189320_18181/ROCK-TA%20Navigator.cmap]. Another, and larger, Knowledge Model called ROCK (Representation of Conceptual Knowledge) was created for the U.S. Army and focuses on intelligence preparation of the battlefield (Eccles et al., 2003). It consists of about 200 Concept Maps and has hyperlinked topographic maps, aerial photos, photointerpretation keys, and information about trafficability of such terrain features as types of dunes. Cmaps can be viewed at [http://cmapskm.ihmc.us/rid=1103739939432_102411597_6499/STORM-LK].

An example STORM Concept Map is shown in Figure 9 and an example ROCK Cmap (with a screen shot of some open resources) is shown in Figure 10. Note that Cmaps in Knowledge Models can have many hyperlinks. The STORM hyperlinks include digital videos of expert discussion of weather forecasting procedures and links to real-time weather and radar data. The ROCK hyperlinks stitch the many Cmaps together and link to text pieces including aerial photos and text about trafficability. The ROCK Cmap also has, at its left side, a "Cmap piece" that shows the place of the particular Cmap within the larger Knowledge Model. Using this piece, and the hyperlinks within it, the user can always tell where he/she is in the Knowledge Model, and how he/she can get from anywhere to anywhere in the model, in only one or two mouse clicks. More information about the construction of knowledge models can be found in Cañas, Hill, & Lott (2003); Crandall, Klein, & Hoffman (2006); and Moon et al. (2011).

CAUSAL REASONING

Concept Maps can express causality and temporality. Of particular concern in analysis is "indeterminate causation" where the goal is to anticipate individual

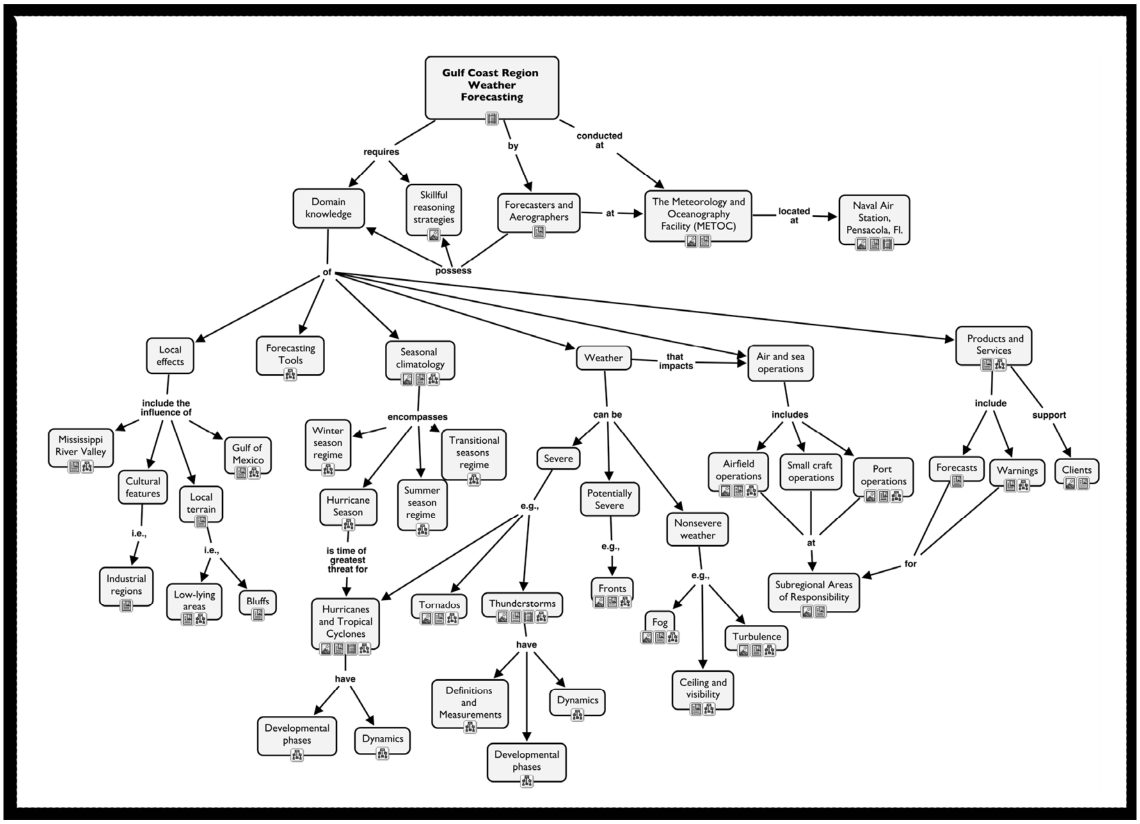

Figure 9. A Concept Map from the STORM Knowledge Model about Navy weather forecasting.

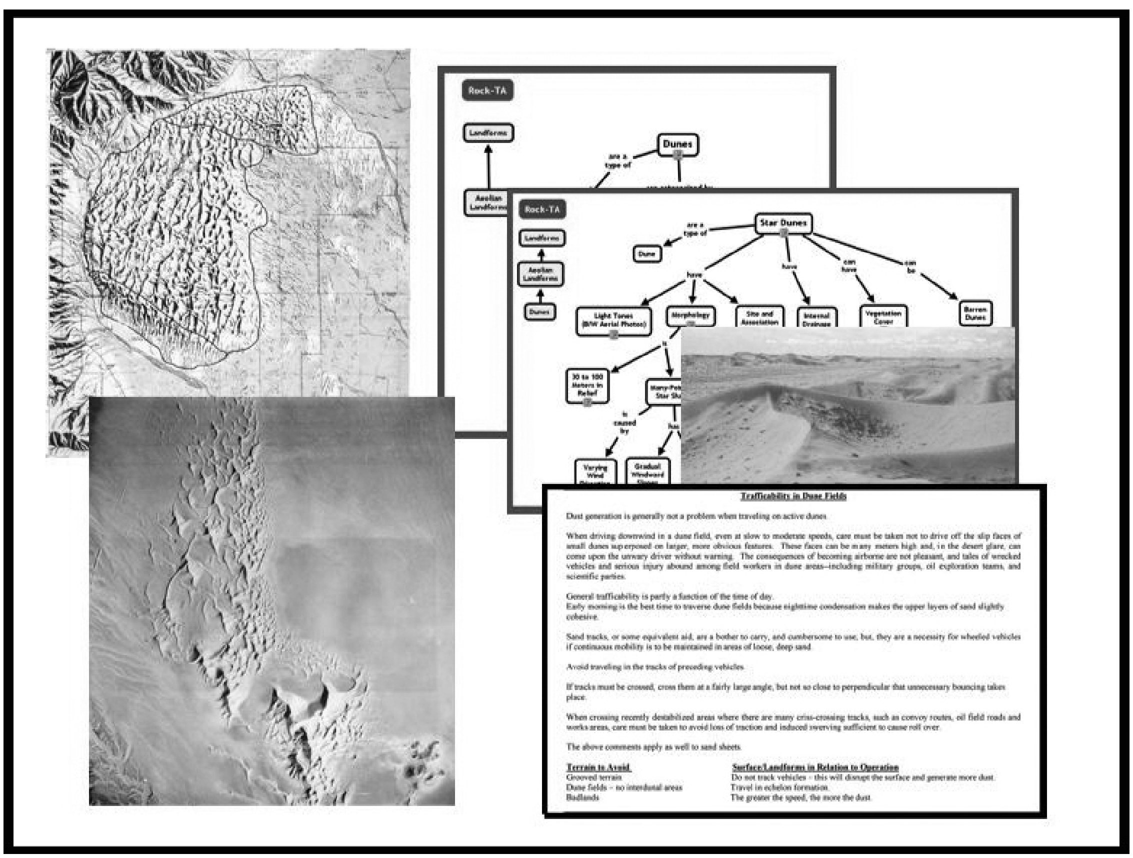

Figure 10. A Concept Map from the ROCK Knowledge Model for intelligence preparation of the battlefield.

or aggregate human activity (Moore and Hoffman, 2011). Causal analysis helps explain how the current situation "came to be," and on this basis supports analysis of where the situation is likely to go (and the impact that different actions might have on shaping that path). Research on the causal reasoning strategies used by practitioners across a range of domains, including economic, political, and military, revealed a dozen patterns (Hoffman and Klein, 2009). A few of these are:

> *The Abstraction.* This is a generalization over evidence (events or conditions). This causal attribution takes several causes, sometimes including counterfactuals, and synthesizes these into a single-cause explanation.
>
> *The Domino.* This is a chain or sequence of causes and effects culminating in the primary effect or phenomenon that is to be explained.
>
> *The Swarm.* This is when a number of independent causes converge to bring about some effect.
>
> *The Spiral.* Events X and Y were both causal of Event Z, but Event X increased (or decreased) the power of Event X.
>
> *The Clockwork.* This is when one or more causes have effects that influence other causes, culminating in the primary effect or phenomenon that is to be explained.
>
> *The Onion.* This is when the analyst wonders about what caused the effect that is used to explain the primary effect or phenomenon that is to be explained.
>
> *The Snark Hunt.* The Snark is a mythical animal for which one can search, but which can never be found because it does not really exist. The Snark Hunt is when the explainer is seeking some particular kind of cause when in fact the to-be-explained effect has some cause that is hidden or might be unknown. The Snark Hunt can be considered a form of counterfactual or disconfirmational reasoning.

All of these themes, and more, can be expressed as a Concept Map, or as a simple "Concept Map Piece" that can be embedded in a larger Concept Map. Awareness of the different causal structures can support critical thinking, that is, the search for alternative causes or causal structures. An example "clockwork" is from economics: bank deregulation permitted mortgaging that entailed relaxed lending criteria; these resulted in risky loans that were used to leverage mortgaging. In other words, the key causal factors interacted. An example "chain," also from economics, would be: Low interest rates caused people to purchase homes they could not really afford, which caused the housing "bubble," which in turn caused the economic decline.

We present three somewhat richer examples in the following figures. Figure 11 presents both a generalized form for the "abstraction"—a template if you will—and a specific example. Figures 12, 13, and 14 illustrate the "onion," "spiral," and "snark hunt," respectively.

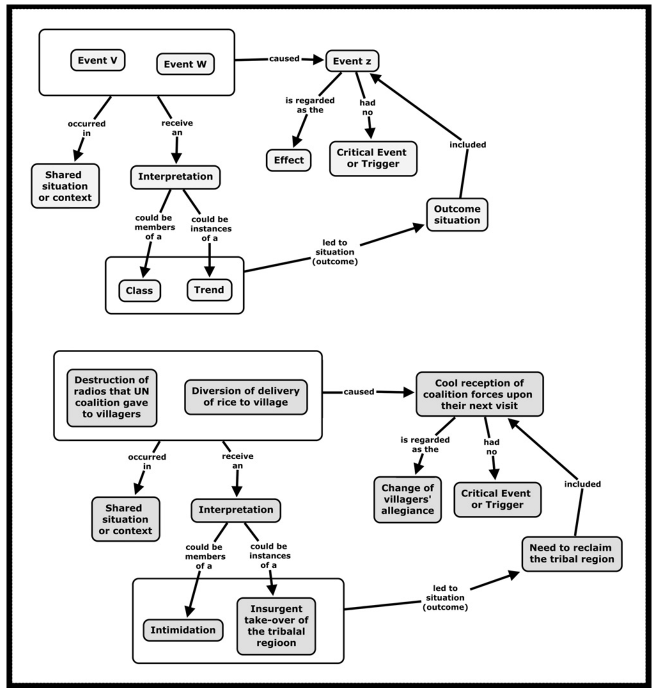

Figure 11. A generic form for the Abstraction causal structure, along with an example.

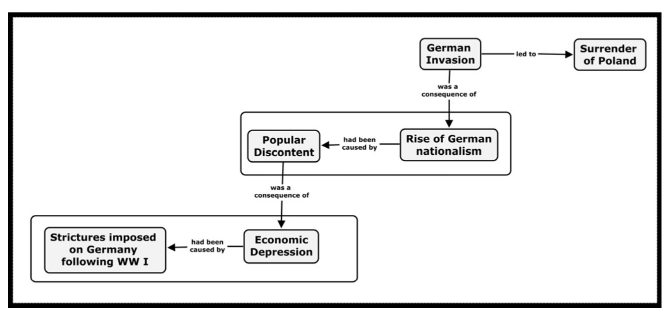

Figure 12. An example of an Onion causal structure, focused on the primary phenomenon of the surrender of Poland to Germany.

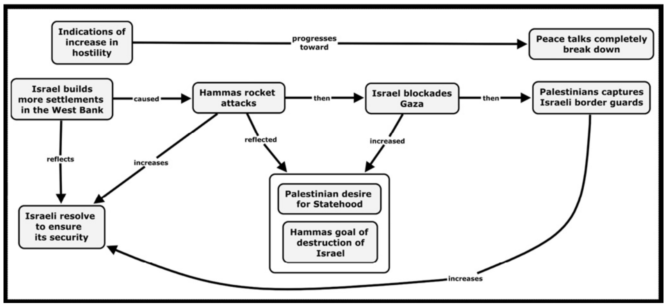

Figure 13. An example Spiral causal structure involving the Israeli-Palestine relation.

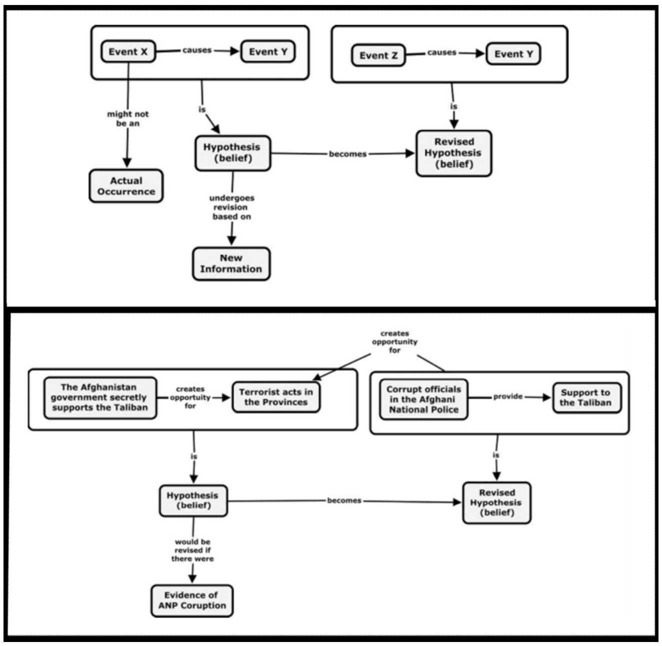

Figure 14. A template plus an example for the Snark Hunt causal structure.

Our final case study brings many of these ideas together, and also hints at the richness and complexity that can be involved in conducting analytical work using meaningful diagrams.

Case Study: The Klathu Scenario

We illustrate the Cmap-based analysis method using the "Klathu Scenario" developed by David Moore (Moore, 2010). While hypothetical, it is a detailed, rich, and realistic scenario about a regional conflict.[2] The text is 13 pages, covering history and background, current situation, assets, recent events, and culminating events as reported by open sources and intelligence sources. The scenario includes maps and listings of evidence, followed by an invitation to the reader to apply such methods as Analysis of Competing Hypotheses. The challenge is to decide what is most likely to happen in the regional conflict, annotated by conflicting evidence. The realism of the scenario is highlighted by the fact that there is no single, clear best answer, and analyses of the scenario by experienced analysts not only do not always agree, but result in hypotheses and findings that were not anticipated by the scenario's creator (Moore, personal communication).

An attempt was made to represent the scenario text exhaustively in propositionally-coherent Concept Maps. The first result was 14 Concept Maps, each representing one of the major paragraphs of the scenario (four Cmaps for each of the four main paragraphs in the "Background," three Cmaps for scenario sections that each focused on one of the hypothetical nations, a Cmap for a paragraph about "mysterious events" occurring on an island in the region, a Cmap covering the discussion of the regional religions and historical religious conflicts, a Cmap about recent news reports, a Cmap about events at a particular shrine, and a Cmap about regional wars). As these were created, separate diagrams were made expressing the hypotheses and speculations that occurred during the analysis (in the manner of Figure 8, above, but not unlike the process involved in the Analysis of Competing Hypotheses). Next, versions of all of the Cmaps were created that expressed assertions and hypotheses (in the manner of Figures 5-7, above). All this took time—the better part of three days of full-time effort.

The Concept Map about one of the nations is presented in Figure 15 (next page). This diagram is representative of the degree of detail one would expect in a useful Concept Map in analytical work. Our heuristic is that for the clearest presentation of meaning a good Cmap should have no more than about 40 concepts (or about 45 propositions). As a Cmap gets larger than this, it is appropriate to break it up into meaningfully appropriate smaller Cmaps that are then hyperlinked together.

One of the Concept Maps about the analyst's hypotheses is presented in Figure 16. A majority of other analysts who had attempted the scenario (up to the time that the Cmap-based analysis was conducted) had attempted conclusions about whether or when a war would break out (Moore, personal communication). The conclusion for this Cmap-based analysis was that a war had already broken out.[3]

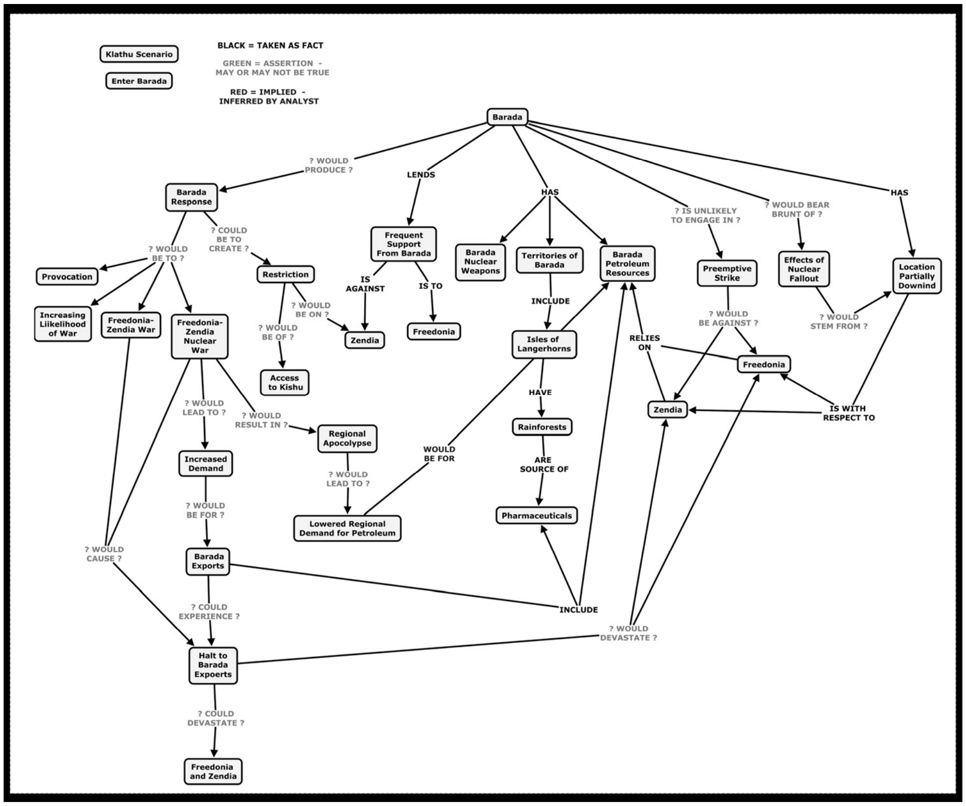

Figure 15. A Concept Map expressing the evidence, assertions, and inferences about one of the nations in the Klathu Scenario.

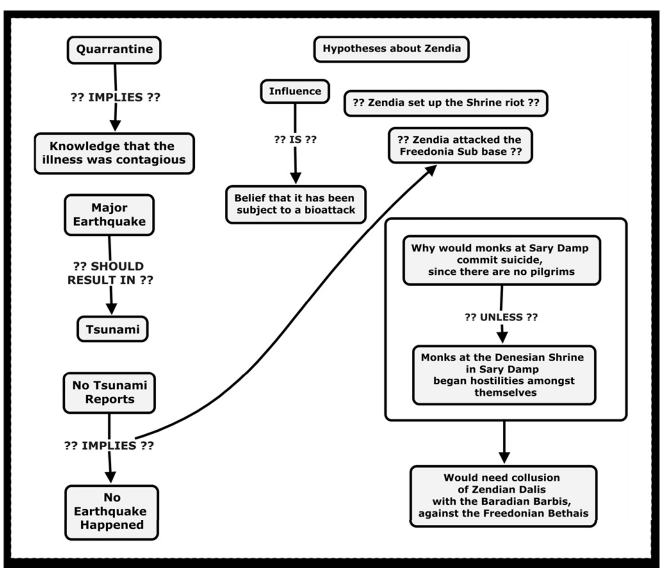

Figure 16. A Concept Map used to express some of the questions and hypotheses formulated by the analyst.

It is not our contention that Cmap-based analysis will uniquely support the creation of high-value analytical products or results. We only suggest, as do Heuer and Pherson (2010), that analysis using meaningful diagrams has its appropriate and valuable uses. While we have focused primarily on Concept Mapping we note also that meaningful diagrams of other kinds can be useful in analytical work.

REPRESENTING ANALYTICAL INTENT

The Analyst's Intent diagram is designed to help practitioners set out, at the beginning of problem analysis, their intended strategy and "line of attack." This can serve a number of purposes. It can enable analysts to return to their foundation during time of information overload or distraction, in order that they can help keep themselves on track. It can help other analysts understand how the problem is being tackled (for example, if the problem was handed over to somebody else midway through). It can also serve as a training device to help communicate tough cases and exemplars to less experienced analysts. By filling in the blanks in the template for Analyst's Intent, the analysts address a range of considerations and captures their resultant thought process. When drafted at the beginning of an analytical activity, the Intent Diagram helps frame and contextualize the work process, but it can be iteratively refined as the analytical work progresses. It can be used to frame reports on, and other products associated with, the activity.

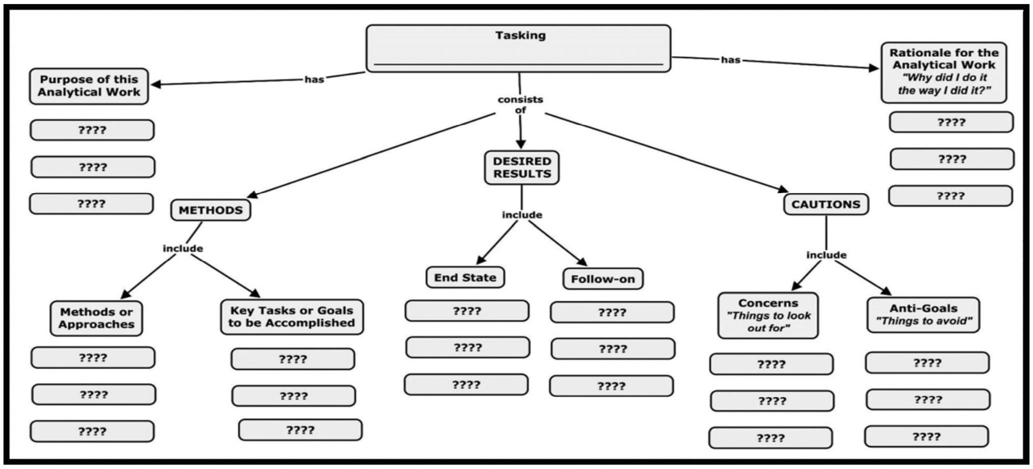

Figure 17. A template for a meaningful diagram to express "Analyst's Intent."

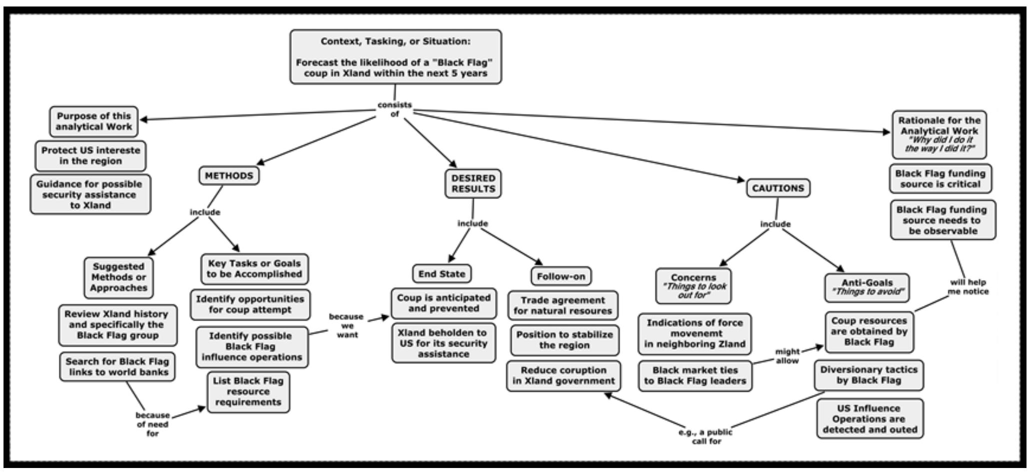

Figure 18. An instance of a completed "Analyst's Intent" diagram.

A template Analyst's Intent diagram is presented in Figure 17 (note that this is not a Concept Map though it is easily created using Cmap Tools). In using this template, the nodes are filled by phrases or short sentences and then meaningful relationships among nodes are added, as deemed useful, using linking phrases to express interrelationships, contingencies, or dependencies. An example of a completed Analyst's Intent diagram is presented in Figure 18 (descriptions, templates, and examples of other causal explanation structures can be provided upon request).

CONCLUSION

The methodology of analysis includes techniques of diagramming. Uses of meaningful diagrams within the Intelligence Community include: eliciting knowledge from experts' analysis, designing new technology by domain experts (bridging the gap between analysis and systems engineering requirements and needs statements), revealing expert-novice differences, acquiring software-assisted knowledge, brainstorming, knowledge sharing (getting data points and information for others who view Cmaps), contrasting alternative perspectives, training, identifying knowledge gaps, creating new knowledge (for example, turning tacit knowledge into an organizational resource), representing team knowledge, constructing and representing shared understanding, structuring conceptual queries, expressing and comparing methodologies, structuring linguistic definitions, designing competency questions, representing networks and organizations, and decomposing analytical problems.

Thus, we have seen a great diversity of diagrams, posted on walls and workboards/workplaces such as command posts and analyst cubicles. Diagrams range from the well-composed and formatted, to nearly useless "spaghetti graphs." There has been little discussion of what makes for a good diagram, and why. Thus, in this article we have presented some principles or heuristics for the creation of *meaningful* diagrams used in problem decomposition, based on findings from research on diagrammatic reasoning in psychology and other disciplines. We have illustrated a variety of forms of meaningful diagrams that have been and are being used in analytical work, accompanied by templates and examples (for a discussion of the use of meaningful diagrams in structuring conceptual queries and forming ontologies, see Eskridge and Hoffman, 2012).

One of the most potentially valuable and effective uses of meaningful diagrams may stem from the fact that the diagram supports the analytical activity and at the same time can serve as a key part of the analyst's report. It is widely known that report preparation can be a huge drain on an analyst's time, and that reporting is a major bottleneck. For an example of using Concept Maps to summarize material, see the synoptic diagrams for the chapters in Hoffman and Militello (2006).

[Authors' Note: Portions of this material were created with the support of an SBIR to Perigean Technologies, Inc. (Contract No. W31P4Q-08-C-0229) in the "Rapid and Accurate Idea Transfer SBIR Program" sponsored by the Defense Advanced Research Projects Agency, Washington, DC. Additional material in this article has been adapted from Henderson, S., Hoffman, R.R., and Jenkins, J. (2011). "A Guide to Influence Mapping: Analysing and Representing Influence Activity over Time." Report 11.00106/1.0, QinetiQ, Farnborough, UK.]

Notes

[1]Other diagramming schemes have been created. Some of these borrow Novak's ideas. Most are limited, for example, by not using labelled links or not having a principled layout or morphology.
[2]Readers may benefit by attempting to work the Klathu Scenario on their own, using their preferred method, before reading about the Cmap-based analysis.
[3]While our figures offer some hints as to how this conclusion was reached, we do not present the full details, by way of inviting readers to conduct their own analysis.

References

Ausubel, D.P. (1960). The use of advance organizers in the learning and retention of meaningful verbal material. *Journal of Educational Psychology* 51, 267-272.
Ausubel, D.P. (1968). *Educational Psychology: A Cognitive View.* New York: Holt, Rinehart and Winston.
Ausubel, D.P., Novak, J.D., & Hanesian, H. (1978). *Educational Psychology: A Cognitive View* (2nd ed.). New York: Holt, Rinehart and Winston.
Bauer, M.I. & Johnson-Laird, P. N. (1993). How diagrams can improve reasoning. *Psychological Science* 4, 372-378.
Cañas, A.J., Coffey, J.W., Carnot, M.J., Feltovich, P., Hoffman, R., Feltovich, J., & Novak, J.D. (2003). *A Summary of Literature Pertaining to the Use of Concept Mapping Techniques and Technologies for Education and Performance Support.* Report to the Chief of Naval Education and Training, prepared by the Institute for Human and Machine Cognition, Pensacola, FL.
Cañas, A.J., Hill, G., Carff, R., Suri, N., Lott, J., Eskridge, T., Gomez, G., Arroyo, M., & Carvajal, R. (2004). CmapTools: A knowledge modeling and sharing environment. In A.J. Cañas, J.D. Novak & F.M. Gonzalez (eds.). *Concept Maps: Theory, Methodology, Technology, Proceedings of the 1st International Conference on Concept Mapping*, Pamplona, Spain: Universidad Publica de Navarra, pp. 125–133.
Cañas, A.J., Leake, D., & Wilson, D.C. (1999). *Managing, Mapping and Manipulating Conceptual Knowledge: Exploring the Synergies of Knowledge Management and Case-Based Reasoning.* AAAI Workshop Technical Report WS-99-10. Menlo Park, CA: AAAI Press.
Cañas, A., Hill, G., & Lott, J. (2003). *Support for Constructing Knowledge Models in CmapTools. Technical Report IHMC CmapTools 93-02.* Available at: http://cmap.ihmc.us/Publications/WhitePapers/Support for Constructing Knowledge Models in CmapTools.pdf Last accessed 20/12/2010.
Cañas, A., Novak, J.D., & González, F.M. (eds.) (2004). *Proceedings of the First International Conference on Concept Mapping: Concept Maps: Theory, Methodology, Technology*. Navarra, Spain: Dirección de Publicaciones de la Universidad Pública de Navarra.
Cañas, A., Suri, N., Sánchez, C., Gallo, J., & Brenes, S. (2003). *Synchronous Collaboration in CmapTools. Technical Report IHMC CmapTools.* Available at: http://cmap.ihmc.us/Publications/WhitePapers/ Synchronous Collaboration in Cmap Tools.pdf. Last accessed 20/12/2010.
Crandall, B., Klein, G., & Hoffman R.R. (2006). *Working Minds: A Practitioner's Guide to Cognitive Task Analysis.* Cambridge, MA: MIT Press.
Eccles, D.W., Feltovich, P.J., and Hoffman, R.R. (2003, April). Modeling Terrain Analysis Expertise to Support Command and Control in the Field. In *Proceedings of the 2003 Collaborative Technology Alliances Symposium.* Sponsored by the Army Research Laboratory and held at the University of Maryland, College Park, MD.
Edwards, J. & Fraser, K. (1983). Concept Maps as reflectors of conceptual understanding. *Research in Science Education* 13, 19-26.
Farzana, K.F. (2008). The neglected stateless Bihari community in Bangladesh: Victims of political and diplomatic onslaught. *Journal of Humanities & Social Sciences* 2(1), 1-19.
Ford, K.M. & Bradshaw, J.M. (1993) (eds.). *Knowledge Acquisition as Modeling*. New York: John Wiley & Sons.
Ford, K.M., Coffey, J.W., Cañas, A.J., Turner, C.W., & Andrews, E.J. (1996). Diagnosis and explanation by a Nuclear Cardiology Expert System. *International Journal of Expert Systems* 9(4), 499-506.
Ford, K.M., Coffey, J.W., Cañas, A., Andrews, E.J., & Turner, C.W. (1996). Diagnosis and explanation by a nuclear cardiology expert system. *International Journal of Expert Systems* 9, 499-506.
Glasgow, J., Narayanan, N.H., & Chandrasekaran, B. (eds.) (1995). *Diagrammatic reasoning: Cognitive and*

computational perspectives. Cambridge, MA: MIT Press.
Gordon, S.E. & Gill, R.T. (1997). Cognitive task analysis. In C. Zsambok & G. Klein (eds.), *Naturalistic Decision Making*, Mahwah, NJ: Erlbaum, pp. 131-140.
Gordon, S.E., Schmierer, K.A., & Gill, R.T. (1993). Conceptual graph analysis: Knowledge acquisition for instructional system design. *Human Factors* 35, 459-481.
Hackett-Fisher, D. (1970). *Historians' Fallacies: Towards a Logic of Historical Thought*. New York: Harper Torchbooks, pp. 164-186.
Henderson, S., Hoffman, R., & Haines, P. (2010*). Influence Support Concepts: Simple Tools for Complex Problems*. QinetiQ Report 10/00788/1.0, QinetiQ, Farnborough, UK.
Henderson, S., Hoffman, R.R., and Jenkins, J. (2011). "A Guide to Influence Mapping: Analysing and Representing Influence Activity Over Time." QinetiQ Report 11.00106/1.0, QinetiQ, Farnborough, UK.
Heuer, R.J. & Pherson, R.H. (2010). *Structured analytic techniques for intelligence analysis*. Washington, DC: CQ Press.
Hoffman, R.R. (ed.) (2007). *Expertise out of context: Proceedings of the Sixth International Conference on Naturalistic Decision Making*. Boca Raton, FL: Taylor and Francis.
Hoffman, R.R., Coffey, J.W., & Ford, K.M. (2000). *A Case Study in the Research Paradigm of Human-Centered Computing: Local Expertise in Weather Forecasting*. Report to the National Technology Alliance on the Contract, "Human-Centered System Prototype." Institute for Human and Machine Cognition, Pensacola, FL.
Hoffman, R.R., Coffey, J.W., Ford, K.M., and Carnot, M.J. (2001, October). "STORM-LK: A human-centered knowledge model for weather forecasting." In *Proceedings of the 45th Annual Meeting of the Human Factors and Ergonomics Society* (p. 752). Santa Monica, CA: HFES.
Hoffman, R.R., Coffey, J.W., Ford, K.M., and Novak, J.D. (2006). A method for eliciting, preserving, and sharing the knowledge of forecasters. *Weather and Forecasting* 21, 416-428.
Hoffman, R.R., Henderson, S., Moon, B., Moore, D.T., and Litman, J.A. (2011). Reasoning difficulty in analytical activity. *Theoretical Issues in Ergonomic Science* 12, 225-240.
Hoffman, R.R. and Markman, A.B. (eds.) (2001). *The interpretation of remote sensing imagery*: *The human factor*. Boca Raton, FL: Lewis Publishers.
Klein, G. & Hoffman, R. (2009, June). Causal Reasoning: Initial report of a naturalistic study of causal inferences. *Proceedings of NDM9, the 9th International Conference on Naturalistic Decision Making*, London, UK.
Mandl, H. & Levin, J.R. (eds.) (1989). *Knowledge acquisition from text and pictures*. New York: Elsevier.
Markham, K., Mintzes, J., & Jones, G. (1994). The concept map as a research and evaluation tool: Further evidence of validity. *Journal of Research in Science Teaching* 31(1), 91-101.
Mintzes, J.J., Wandersee, J.H., & Novak, J.D. (2000). *Assessing Science Understanding: A Human Constructivist View*. San Diego, CA: Academic Press.
Moon, B., Hoffman, R., Novak, J., & Cañas, A. (2011). *Applied concept mapping: Capturing, analyzing, and organizing knowledge*. New York: CRC Press.
Moore, D.T. (2010). What is going on in Klathu? NSA APPROVAL NUMBER: PP-08-135.
Moore, D.T. and Hoffman, R.R. (2011). "A practice of understanding." Chapter 5 In D.T. Moore, *Sensemaking: A structure for an intelligence revolution*. Washington, DC: National Defence Intelligence College Press.
Novak, J.D. (1991). Clarify with concept maps: A tool for students and teachers alike. *The Science Teacher* 58, 45-49.
Novak, J.D. (1998). *Learning, Creating, and Using Knowledge: Concept maps® as Facilitative Tools in Schools and Corporations*. Mahwah, NJ: Lawrence Erlbaum Associates.
Novak, J.D. & Cañas, A. (2008). *The Theory Underlying Concept Maps and How to Construct and Use Them*. Technical Report IHMC CmapTools 2006-01 Rev 2008-01. Available at http://cmap.ihmc.us/Publications/ResearchPapers/TheoryCmaps/TheoryUnderlyingConceptMaps.htm Last accessed 3/14/2010.
Novak, J.D. & Gowin, D.B. (1984). *Learning How to Learn*. New York: Cambridge University Press.
Vekirl, I. (2002). What is the value of graphical displays in learning? *Educational Psychology Review* 14, 261-298.

Dr. Robert R. Hoffman is a recognized world leader in cognitive systems engineering and human-centered computing. He is a Fellow of the Association for Psychological Science, a Fellow of the Human Factors and Ergonomics Society, and a Fulbright Scholar. Hoffman has been recognized in disciplines including psychology, remote sensing, weather forecasting, and artificial intelligence for his research on human factors in remote sensing, the psychology of expertise, the methodology of cognitive task analysis, and the design of complex cognitive work systems. His current research focuses on methodological issues in the analysis of complex systems and performance measurement for complex cognitive work systems.

Simon M. Henderson is a Lecturer in Information Design at Cranfield University, based at the Defence Academy of the United Kingdom, where he has been since 2011. Earlier he was a researcher at QinetiQ (and its forerunner organizations) focusing on the psychological basis of sensemaking, decision-making, and teamwork. He began his career in the areas of knowledge elicitation and system dynamics modeling, developing exploratory models of future command information systems. He then joined a group studying the impact of technology on command and intelligence teamwork and decision- making. His contributions to methodology include the development of a range of data collection and knowledge elicitation tools and techniques; methods for the experimental and field-based assessment of teamwork, sensemaking, and naturalistic

decision-making; methods for the enhancement of team debriefing and self-learning approaches; and methods for the design and evaluation of team support technologies. Work over the past five years has focused on studying the psychological basis of collective problem solving in complex settings, with particular emphasis on the development of knowledge mapping and representational systems. This work has been widely published. Simon has been the UK representative on a number of NATO and other international panels on decision making and teamwork.

Jonathan A. Jenkins is a Director at i to i research. He leads the organization in its application of cutting-edge human science to address theoretical and practical challenges across the domains of influence, intelligence, and counterterrorism. His passion and expertise lie in the practical science of behavior change within groups, organizations, and populations, with his work focusing on practical methods, techniques, and approaches. Having originally studied psychology, Jon has spent the last ten years focused on the areas of behavioral research, behavior change tools and techniques, and measurement for influence campaigns. He is the author of guidance for change practitioners and speaker on behavior change and influence to practitioners across the defense and security domain, and is a visiting speaker at the NATO Information Operations Senior Officers Course.

Brian Moon is Chief Technology Officer for Perigean Technologies LLC. A leading practitioner and researcher in the field of expertise management, he applies diagrammatic and naturalistic interviewing techniques across a spectrum of knowledge management problems, from the elicitation and capture of expert knowledge to the development of decision support and training systems that enable its effective and efficient transfer. He has been widely recognized for his innovative applications of Concept Mapping, which are highlighted in the edited volume, Applied Concept Mapping: Capturing, Analyzing, and Organizing Knowledge. His clients have included several private corporations, the Department of Defense, the New York Power Authority, the Federal Bureau of Investigation, and Sandia National Laboratories. Brian is also an adjunct faculty member for Kent State University's Information Architecture and Knowledge Management, an online MS program.