

FORMALIZING THE INFORMAL: A CONFLUENCE OF CONCEPT MAPPING AND THE SEMANTIC WEB

Thomas Eskridge, Pat Hayes, Robert Hoffman, Florida Institute for Human and Machine Cognition (IHMC), USA

Margaret Warren, Carma Productions, USA

Email: teskridge,phayes,rhoffman@ihmc.us, www.ihmc.us

Email: mm@carmapro.com, http://www.margaretwarren.us

Abstract. Concept Maps are seen as ways to conceptualize domains that can be comprehended within the human attention span, and are often based on terminologies arising in technical, scientific or engineering domains. We report on a set of related projects that go beyond this tradition in two ways: by connecting CmapTools to the new, machine-processable standard notations for formal Web meta-data (RDF and OWL), and by applying the resulting tools to the construction of Web markup for works of art. This domain provides a semantically-rich case study because descriptions of artwork inherently contain a mix of artifacts, some of which can easily be formalized (creator, title, creation date, etc.) and some of which express meanings about the work that are more difficult to formalize. Formalization of knowledge about works of art could benefit web developers by giving greater organization to these types of resources on the internet. It will also help foster the development of ontologies in other, similarly rich and challenging domains. A CmapTools extension called CmapTools Ontology Editor (COE) translates RDF and OWL to and from Concept Maps, providing a bridge between the worlds of human-oriented Concept Mapping and machine-oriented formal inference. Because of the strengths in CmapTools to provide a specific method for capturing meanings and concepts, we believe that the COE extension can provide the means for formalizing this conceptually challenging domain that involves an artist's interpretation of the intentions and motivations behind their artwork.

1 Introduction

The Semantic Web is often viewed as the “next generation” of the current World Wide Web (Berners-Lee, Hendler et al. 2001), which aims to create a simple way to express and store interchangeable data on the internet.

Because the volume of web resources (web pages, images, music, video, etc.) continues to increase exponentially, there is a recognized need for advanced strategies that will enable resources to be interpreted more logically by search engines and other intelligent agents. As such, there are many individuals and organizations beginning to explore the integration of semantic web strategies into their future web development activities.

The most widely accepted technique for employing semantic web technologies involves the use of a formalized language (OWL/RDF¹) to ‘markup’ web resources with detailed descriptions of their content that can be more easily ‘ingested’ by computers. The resulting machine-readable markup is often referred to as an ontology. Ontologies must be created using a tightly controlled and elaborate syntax, and are required to satisfy many logical constraints, some of which may not be intuitively clear to human readers. Because of this, determining the content and structure of an ontology for a given domain requires an expert-level understanding of that content and its purpose.

CmapTools (Cañas, Hill et al. 2004) has been used to capture knowledge of human experts for use in training (Hoffman, Coffey et al. 2006), institutional memory preservation (Coffey, Eskridge et al. 2004), and in-the-field use (Coffey, Cañas et al. 2003). The method for eliciting this knowledge usually involves a moderator who questions and guides the expert and a mapper who documents the interview in one or more concept maps (Hoffman, Shadbolt et al. 1995; McNeese, Zaff et al. 1995). The result of this process, called a *knowledge model*, has proven to be easily navigated and understood by human users. However, because of their semantic expressivity, Concept Map-based knowledge models cannot easily be interpreted and made actionable by computers.

The work reported here is part of an effort to build a bridge between these two worlds (human comprehension and computational analysis), by casting OWL/RDF semantic web ontologies in the form of Cmap knowledge models, thereby making available to users the human-oriented clarity and conceptual organization provided by Cmaps, while retaining the ability to produce exact, formal ontologies for semantic-web markup.

¹ See <http://www.w3.org/RDF/> and <http://www.w3.org/2004/OWL/>

We have developed an extension to CmapTools to facilitate this, called the CmapTools Ontology Editor, or COE² (Hayes, Eskridge et al. 2005). Human users of COE see and build knowledge models, and the software translates these to and from OWL/RDF ontologies when possible. COE can import any OWL/RDF formal ontology and display it as a Cmap, using a variety of notational conventions to render the idiosyncratic and formal content of the ontology in a reasonably human-readable format. It then allows users to edit and modify the Cmap in ways that conform as far as possible to the knowledge-model extraction method, and will automatically export the represented content into formal OWL/RDF for use as machine-interpretable semantic web markup. COE is operational at the beta-release stage and is currently being tested by various users in US government agencies. When completed, it will be made available as part of the basic functionality of the CmapTools software.

Warren is a practicing artist, photographer and technologist who maintains several websites containing a large number and variety of images using current web technologies. One goal of the Conceptual Art Mapping Project (CAMP) is to use COE to mark-up these image websites using OWL/RDF to describe both the images themselves and aspects of the ideas and processes that were involved in their creation and subsequent interpretation.

The formal specification of artistic work is a domain that is currently without a standardized ontology (van Ossenbruggen, Troncy et al. 2006). It is also a domain which might derive a great benefit from a more complete semantic web solution. Many artists and art organizations desire to organize and easily access their content, create presentations, sell to image banks or make available to editors, researchers or anyone interested in finding image content on the internet.

This paper reports on ongoing research into the techniques and methods for the conceptualization of relevant domain information using CmapTools, and the conversion of that information into machine interpretable, formalized ontologies using COE. The results are a confluence of informal knowledge specification and formal ontology design. Our ultimate goal is that an OWL/RDF formalized diagram of such informal and abstract knowledge will enable machine-readable markup for artists' work on the internet that is much more complete than traditional image annotations, and will hopefully result in a general purpose ontology for image makers.

2 Ontologies

Formal domain ontologies already exist. Some have the goal of providing an overarching framework for connecting disparate knowledge (Niles and Pease 2001), while others represent organized bodies of specialist scientific or technical knowledge. For example, the SWEET ontologies being developed at NASA JPL³ represent concepts from the atmospheric and oceanographic sciences. New ontologies can build on existing ones, borrowing concepts and relations from any number of ontologies where convenient, and adding new concepts and relations where necessary. In this way, a common vocabulary for representing and annotating data can evolve. Although a large formalized ontology can be used for many purposes, the first benefit is often the ability to search more 'intelligently' through large corpora containing information from many sources.

Although common vocabularies for annotating images already exist, none of them address the artist's intent. Digital cameras store annotated technical data with each image at the time the image is captured, such as camera model, exposure settings, flash setting, date & time and other information (Committee 2002). Many image database and editing software tools contain annotation methods that draw from a number of vocabularies. IPTC⁴, Adobe XMP⁵, Dublin Core⁶, Visual Resources Association⁷ can be used to tag the image with metadata describing properties such as creator, title, creation date, techniques and descriptions. All of these methods are based around keywords or 'tagging', which are enjoying a surge of popularity on image websites such as Flickr⁸.

² <http://cmap.ihmc.us/coe/>

³ <http://sweet.jpl.nasa.gov/>

⁴ <http://www.controlledvocabulary.com/imagedatabases/iptc-naa.html>

⁵ <http://www.adobe.com/products/xmp/index.html>

⁶ <http://dublincore.org/>

⁷ <http://www.vraweb.org/vracore3.htm>

⁸ <http://www.flickr.com>

Even with the use of keywords and tagging, it should be noted that Google-type image search engines only look for the occurrence of a search term on a page containing the image or in the name of the image itself. Even annotation-based markup for images independent of page content would clearly help focus the work of such search agents. Search engines also have difficulties simply within the determination of keyword connections. For example: the words “chimp” “under” “large” “tree” may be used to tag an image, but it would not be clear that “large” refers to “tree” or to “chimp” unless the keywords were combined, but then it would be hard to determine that “large tree” is a type of “tree” (Schreiber, Dubbeldam et al. 2001).

There is also a distinction to be made regarding what type of work the image resource is; the work of a visual artist demands much more detail in its descriptions than work that is simply documentary in nature (images of real estate property, Uncle Bob in front of the Eiffel Tower, etc.) Much of an artist’s intent may involve incidental aspects of the work in question rather than the obvious subject of the work. For example, the artistic impact of an image may be the mood set by the lighting conditions rather than the sailboat depicted in the image; more likely, it will be a combination of a variety of aspects of both the picture and the scene, and will involve analogies, metaphors, affective connotations, and other ‘informal’ connections between ideas.

With CAMP, we are specifically attempting to create markup that includes such things as: the artist’s intent in creating the work; thematic elements in the work; the artist’s interpretation of the work and elements of the work, and any influences upon their choice of medium or the style of the work.

3 Method

We began by creating unconstrained Concept Maps about a few of the artist’s photographic images. As the domain expert, the artist, while gaining familiarity with CmapTools, used COE to represent factual information about each of the images, motivations behind the production of the works, and ideas involving artistic interpretation.

Early on it became clear that several kinds of information were being described in these maps, including descriptions of pictorial aspects of the image itself, descriptions of the techniques and materials used, of the physical scene being imaged, the circumstances of the picture-taking event (in the case of work produced through a photographic process) and of relationships between all of these - for example, that the time of the shoot influenced the quality of the light – and finally, more general interpretative information used in part to explain the titles of the pieces and how the interpretation relates to all the other elements. For the purpose of this paper, we selected one representational image from our initial test cases - a photograph titled “Confluence” - to illustrate our findings up to this point. The conceptual description has gone through many iterations. Some of these iterations were done using

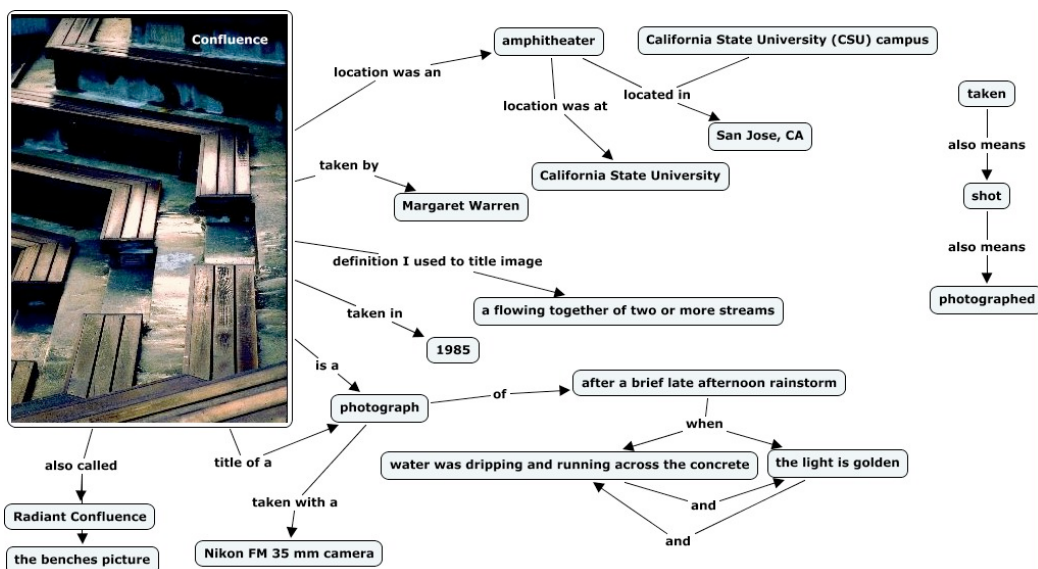


Figure 1. Initial “Confluence” Concept Map.

the PRESERVE interview sessions (Coffey and Hoffman 2003), while others were based on individual reflection and refinement on the part of the artist.

In the initial Concept Map, the artist was deliberately approaching the Concept Map from the point of view of describing the image itself, thereby making it the root node: it is a collection of statements made directly about the image. This initial Concept Map captures several types of information about the image: where and when it was taken, what kind of camera was used, other names for the image, what the image is of (“afternoon rainstorm,” “golden light,” “water dripping”), and how the title was derived from the image (“a flowing together of two streams”).

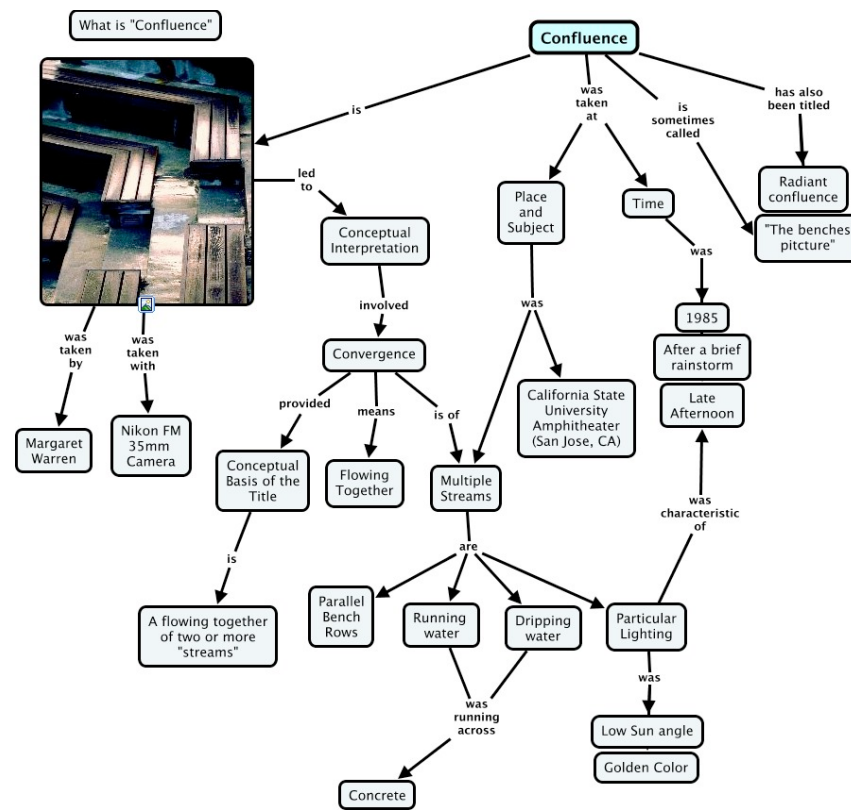


Figure 2. Second “Confluence” Concept Map.

The second Cmap was created through an interview session and uses the artist’s title of the work as a root node as it also notably introduces the concepts “Conceptual Interpretation” and “Convergence” (see Figure 2.) This process of abstraction and explication continued with the development of the Concept Map shown in Figure 3, which abandoned the inclusion of the image altogether. This Cmap is organized with a series of major classifying nodes: “Temporal Elements,” “Scenic Elements,” “Pictorial Elements” and “Interpretive Elements,” whose goal is to impose some order upon the various kinds of assertions represented by the ‘interpretive’ material of the previous Concept Map which reflects generalizations about the image-making process itself. We also introduced some semi-technical terminology to express the relationships involved. The ‘educe’ relation emerged from discussions at this point. The term “educe” (meaning “to draw, bring out; elicit”) was suggested by the artist as a single word to replace a variety of formulations in an earlier version of the Concept Map, including ‘give rise to,’ ‘produce,’ ‘lead to,’ etc. Each of these had been used to talk about the relationship between a *physical* aspect of a depicted scene which was in some way responsible for a resulting *pictorial* element or feature. These pictorial elements were, to an experienced professional photographer, characteristic signs. For example, the angle of the sun in the late afternoon (a ‘scenic element’) educes a golden colored light (a ‘pictorial element’).

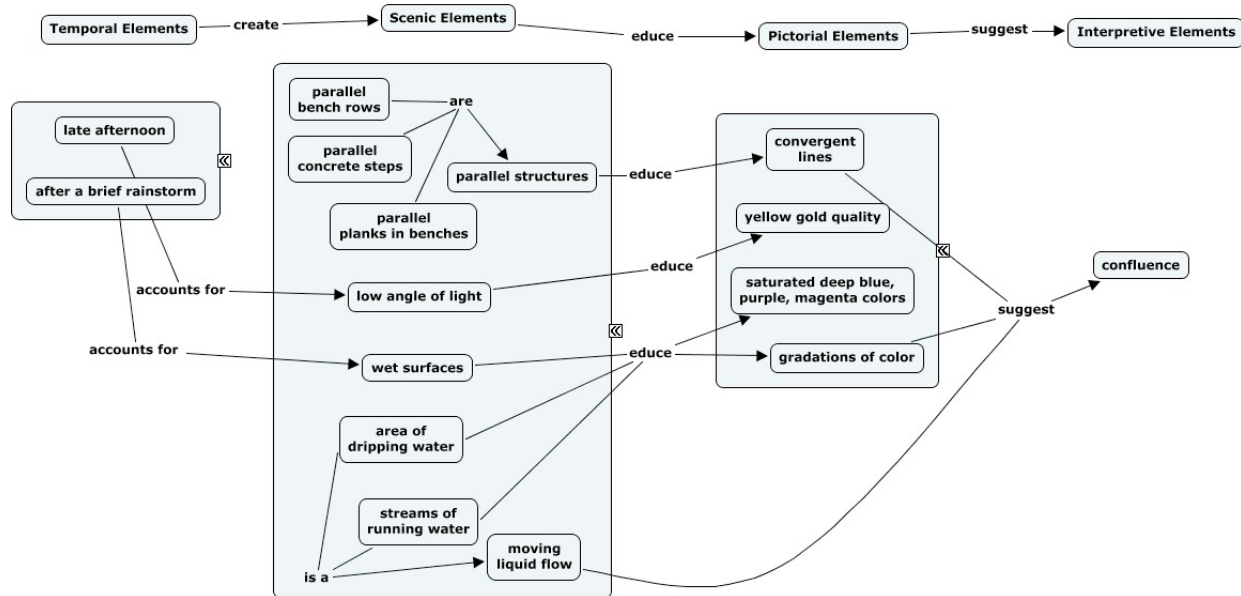


Figure 3. Third “Confluence” iteration, removing image entirely.

Why use a relatively obscure word to express an idea which can be said in a variety of ways already? The answer lies in a tension which runs through this project, between the *general* and the *particular*. One of our goals is to produce useful semantic markup for websites containing artistic images; semantic markup is of utility only to the extent that it uses concepts that are also used in other such markup. A purely private conceptual framework is of no use on the Web, where the ultimate aim is always to create and foster communication between disparate and distant web agents. We are therefore aiming to invent a ‘standard’ vocabulary which might be potentially used by a variety of Web markups. Such standardized vocabularies are critical to the success of Semantic Web applications. On the other hand, we wish to allow the artists to use their own terminology when describing their work, with relatively idiosyncratic and even private terms to express their intuitions. The expert knowledge used by artists may be informed by their academic background and training, but the creative process itself is less well articulated and less informed by any background theory than by the knowledge underlying expertise in scientific and engineering domains. Since there is no accepted terminology to express how the artist conceptualizes their artistic process, it is necessary to use metaphors, approximations, etc. Our approach is to deal at first with the particular, with the hope of being able eventually to abstract general concepts from it and incorporate these into a useful middle-level ‘artistic interpretation’ ontology that will enable many artists to annotate their work using a common vocabulary, and enable computers and people and software agents to collaboratively make use of them. Rather than attempt to *impose* any kind of aesthetic theory at first, we hope to *discover* some common structure or conceptual themes from a large number of detailed analyses of particular images. We believe that this notion that we have labeled ‘educate’ is an example of a central concept in image interpretation. Ironically, providing it with a special ‘formal’ label is one way to encourage its recognizable use in markup, essentially by providing a single name for a relatively fine-grained concept.

It is interesting also to examine Figure 3 from a Concept Mapping point of view, rather than from an ontology development point of view. In this figure, we have introduced four new concepts concerning the elements of an artistic work and have grouped different aspects of our concept map description of “Confluence” under their corresponding elemental type. Novak’s (1998) explanation of Concept Mapping in terms of Ausubel’s assimilation learning theory could just as neatly apply to our work here developing ontologies, which we find encouraging. In particular, introducing superordinate concepts, progressively differentiating existing knowledge, and integrative reconciliation are all theoretical principles that are clearly evident in the successive iterations of the Concept Maps. Being able to transfer the theoretical foundation of Concept Mapping to building ontologies is an added benefit of the COE approach to ontology development.

4 Moving from informal to formal

The last stage of our process is converting the Concept Map into a form that can be exported into OWL/RDF. This is trivial to do poorly, but difficult to do well. A trivial translation simply renders the Concept Map into RDF so that each node-arc-node combination becomes an RDF triple. However, the result is ‘shallow’ since most of the significant meaning is buried in the node and arc labels, and hence is invisible to a semantic web reasoner. A deeper analysis is required in order to more adequately represent the meaning of phrases used as node and arc labels. For example, our Concept Map contains a concept “*yellow gold quality*” which refers to the quality of the light in the picture. Clearly, this has something to do with color. We could naïvely transcribe this label into an RDF ‘label’ which must be a Web Universal Resource Identifier, and so might become http://purl.oclc.org/CAMP/concepts/colors/yellow_gold_quality. However, there is nothing in the ontology that records the fact that this is a color concept, nor is there any information about colors; for example, that gold is complementary to blue.

Such classifications of concepts and relations between them are what OWL is designed to record, but putting this information into the Concept Map requires more work to be done. We expect this to be an incremental process, since the most useful structure of classifications and concept relationships can emerge only by a process of induction applied to a number of Concept Maps. An interim product resulting from our single Cmap is shown in Figure 4. The top part is an OWL-ized version of the Cmap itself, while the other parts of the Cmap record a variety of background information and links to other, existing, ontologies written using COE notational conventions. Concepts labeled with ‘naïve’ phrase transcriptions (like `_this`) are intended as placeholders which will require more attention later.

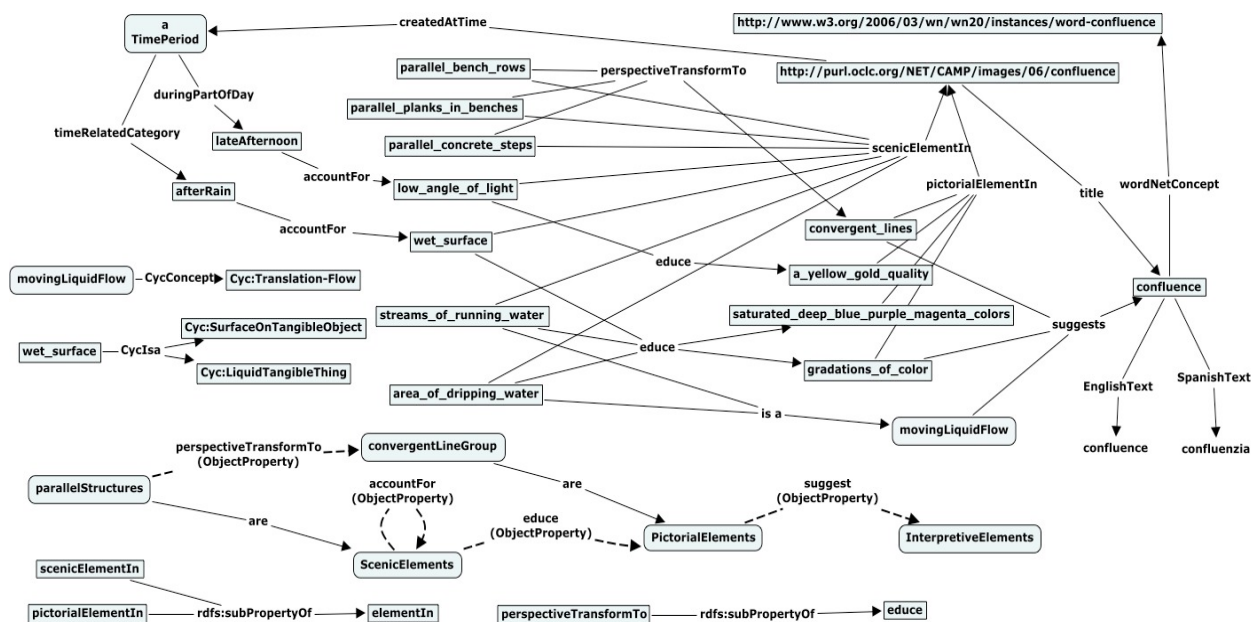


Figure 4. The formalized "Confluence" Concept Map.

Finding what one might call ‘OWL structure’ in the Concept Map requires paying close attention to categories (or classes) and properties. OWL is based on a global picture which regards *individuals* as categorized or classified in various ways into *classes*, and which have *properties* whose values are other individuals. In the above example, the fact that *yellow_gold* is a color would be a classification, and color-complement would be a property (of yellow, with value blue; of red, with value green, etc.) Much of the craft of writing adequate ontologies lies in making a deft choice of classifications and properties; and this is still very much an art rather than a science. For example, our conceptual analysis seems to have identified categories of ‘pictorial element’ and ‘scenic element,’ suggesting that ‘elements’ are the primary individuals in this domain. But what exactly *are* these “elements”? They seem to include things which one would intuitively think of as properties or facts (that some lines are parallel), and they seem to be very eclectic, including color, geometry and properties of surfaces. We anticipate that finer and more detailed categories will emerge as we proceed - they are in any case of use only insofar as they serve to state some general regularities which apply to more than one image - but the appropriate ontological strategy at this early stage is a kind

of principled agnosticism, to try to only make very ‘large’, weak generalizations, such as in our case that between *picture* and *scene*, and then expect more detailed structure to emerge from more examples. Already, we have noted that there seemed to be a single relationship called *educe*. (OWL allows us to use a property as a kind of classification of other properties, called its subproperties, and we expect *educe* to have many subproperties which have more precise meanings, such as *perspectiveTransformTo*.)

The same intuitive content can be encoded into a notation like OWL in a large number of different ways, and the choice among them is not always obvious (van Ossenbruggen, Troncy et al. 2006). For example, in an early version we had classes of ‘picture elements in the image,’ until we realized that (of course) such classes are of use only if they can be related to the particular image. Therefore, we decided to treat these as properties of the elements, with the image as value. This allows a single element to figure in more than one image. The art of compiling useful ontological markup lies in making such decisions well; but these go beyond the usual process of building a conceptual model, and require attention to much finer-grained conceptual distinctions.

The ‘background’ information in the final Cmap included general facts about categories and properties and also links to similar (perhaps identical) concepts in other published ontologies. The general facts, which have many variations, allow a Semantic Web reasoner to draw simple conclusions from the facts at its immediate disposal. For example, our Cmap does not assert directly that *parallel_bench_rows* is classified as a *ParallelStructure*, but this follows from the background fact that the latter is the ‘domain’ of the property *perspectiveTransformTo*, which is applied to the former in the Cmap (indicated by a bold dashed line.) Other facts establish subcategory relations and relationships between classes and property values. The relatively few background facts we have described here are sufficient to generate enough conclusions to increase the size of the Concept Map by a factor of approximately five.

Making links to existing published ontologies is a critical for the success of the Semantic Web. Such re-use means that concepts of broad utility become articulated into more and more pieces of available knowledge, increasing their utility for search and inference. In our example, the links are to two widely cited ontologies: WordNet (W3C 2004), which contains synonym and hypernym relationships for all English nouns, and the OpenCyc knowledge base, which contains a large number of ‘natural’ concepts under a large number of headings. It is notable that many of our concepts, even from this one image, seem to have no exact correspondence to any published concept corpora. There are, for example, several ‘standard’ temporal ontologies, which provide for detailed time attributions and define concepts such as millisecond or international time zone; but none of them contain concepts such as ‘late afternoon’ or ‘early morning,’ which are qualitative rather than numerical, and which are relevant to appearance (as in ‘early morning light’) rather than time-charting. Cyc has an elaborate classification of physical shapes and surfaces, but it is oriented to physical thinking, rather than to optical qualities of surfaces: Cyc for example ‘knows’ that touching a wet surface makes the toucher wet, but has no way to infer that a wet surface is typically reflective. We expect that we will be obliged to develop ‘artist-quality’ ontologies for such matters as kinds of lighting, colors, relevant atmospheric and surface forms, etc., as the project proceeds.

The export facility of COE will output as much of the Concept Map as possible into structured OWL, and translate the rest into unstructured RDF when possible, or simply ignore it. The result of exporting our map is an RDF/XML file. A Concept Map also contains a great deal of other information, including the actual layout of the nodes on the page, and may contain annotations in the form of isolated nodes containing text, or ‘sticky notes’. We call Concept Maps that contain both such informal Concept Map information and ontology information *Combination Cmaps*. The informal information may be placeholders that are awaiting translation into OWL/RDF, or they may be segments of Concept Maps that are used to explain particular pieces of OWL/RDF representation. The explanations may also be used to contextualize the ontology for human readers, so that potential users can get a better understanding of the original intent of the ontology author, and where and why the ontology may be useful. We are just beginning to discover the utility of Combination Cmaps, but we are already anticipating that they will benefit ontology development through greater understandability and collaboration.

5 Conclusions

Building a bridge between knowledge models, which reflect human learning, and formal ontologies, which can support machine inference, is challenging but also necessary and important. We cannot claim to have fully bridged this gap. While we have found that people are able to construct COE maps after a few hours of training, exporting these maps to useful OWL/RDF generally requires iterative development and specialized background knowledge to

interpret properly. In addition, our longer-term goal requires us to formalize even more human intuition in the form of ontologies that define the meanings of vocabularies which can describe aspects of pictures and artwork that are relevant to artist's concerns. Such ontologies do not yet exist, but experience with other areas of ontology development leads us to be guardedly optimistic that it is possible in a reasonable timescale.

One aspect of the CAMP project is as a test-bed for educing improvements to the GUI design of tools such as COE. Already, for example, it is clear that it would be desirable to make assertions about areas or regions of a picture rather than relating all markup to the entire image. Experience with CAMP also suggests several improvements to the COE interface to allow the compact representation of commonly used descriptive patterns in RDF/OWL.

The CAMP project is also a 'forcing function' in a larger sense. There have been many attempts to create 'common sense' or 'high-level' ontologies intended to provide concepts for representing intuitive or informal knowledge. However, even the most ambitious of these (such as Cyc) contain many concepts with a distinctly 'technical' flavor, more suited to expressing ideas in engineering or mathematics than those that seem to form the everyday substance of human experience and intuitive thought. Our intended domain of application is stubbornly resistant to any pressures to 'technicalize' its conceptual vocabulary, which is both part of its interest and what makes it challenging.

Finally, we expect and hope that the combination of tools, vocabularies, ontologies and representational ideas which emerge from CAMP and COE will be a useful resource for artists, archivists and anyone concerned with cataloging, writing about or searching for images on the Web.

6 References

- Berners-Lee, T., J. Hendler, et al. (2001). The semantic web. *Scientific American*. 284: 34-43.
- 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.
- Coffey, J. W., A. J. Cañas, et al. (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.
- Coffey, J. W., T. C. Eskridge, et al. (2004). A Case Study in Knowledge Elicitation for Institutional Memory Preservation Using Concept Maps. In A. J. Cañas, J. D. Novak & F. M. González (Eds.), *Concept Maps: Theory, Methodology, Technology. Proceedings of the First International Conference on Concept Mapping*. Pamplona, Spain: Universidad Pública de Navarra.
- Coffey, J. W. and R. R. Hoffman (2003). Knowledge modeling for the preservation of institutional memory. *Journal of Knowledge Management* 7(3): 38-49.
- Committee, T. S. (2002). *Exchangeable image file format for digital still cameras: Exif Version 2.2*. from <http://www.exif.org/Exif2-2.PDF>.
- Hayes, P., T. C. Eskridge, et al. (2005). Collaborative knowledge capture in ontologies. *Proceedings of the 3rd International Conference on Knowledge Capture*. Banff, Alberta, Canada, ACM Press.
- Hoffman, R. R., J. W. Coffey, et al. (2006). A method for eliciting, preserving, and sharing the knowledge of forecasters. *Weather & Forecasting* 21: 416-428.
- Hoffman, R. R., N. Shadbolt, et al. (1995). Eliciting knowledge from experts: A methodological analysis. *Organizational Behavior and Human Decision Processes* 62: 129-158.
- McNeese, M., B. S. Zaff, et al. (1995). AKADAM: Eliciting user knowledge to support participatory ergonomics. *International Journal of Industrial Ergonomics* 15: 345-363.
- Niles, I. and A. Pease (2001). Towards a standard upper ontology. *Proceedings of the International Conference on Formal Ontology in Information Systems*, Volume 2001. Ogunquit, Maine, USA, ACM Press.
- Schreiber, A. T., B. Dubbeldam, et al. (2001). Ontology-based Photo Annotation. *IEEE Intelligent Systems* 16(3).
- van Ossenbruggen, J., R. Troncy, et al. (2006). *Image annotation on the Semantic Web*. from http://www.w3.org/2001/sw/BestPractices/MM/image_annotation.html.
- W3C (2004). *Wordnet in RDF and OWL*. <http://www.w3.org/2001/sw/BestPractices/WNET/wordnet-sw-20040713.html>.