Knowledge Capture, Representation and Visualization

Alex Ivanov
School of Interactive Arts and Technology
Simon Fraser University, Surrey, BC, Canada
For IAT881: Ontological Systems
and Semantic Web Technologies
alexi@sfu.ca

ABSTRACT

This paper reviews the developments in the past three years in the topics of Knowledge Capture, Knowledge Representation, and Knowledge Visualization, from a semantic Web ontology perspective. The paper tries to show that these three topics blend or even overlap one another. Concept Mapping is one particular unifying theme. The paper will try to shed light on this by reviewing several prototypes, leading to a discussion of research directions that aims to conclude that graphical representations will play a key role in KC, KR, and KV and the semantic Web. Moreover, the future of these fields will make use of both semiotics as well as the design of collaborative spaces—in addition to the technology that underlies them.

KEYWORDS

Knowledge acquisition, information capture, knowledge representation, collaborative work, visualization, semantic Web, ontologies, concept maps, information systems.

1. INTRODUCTION

Despite great advances in IT over the past decade and ever-growing pools of data on the Web, the quality of this information has remained essentially the same. The advent of the semantic Web promises qualitative improvements, and great strides have been made in document processing, data mining, ontology mapping, and interoperable data-sharing. This paper focuses more on the content aspect of knowledge on the Web. It will first provide a brief review of the latest developments in knowledge capture (KC), knowledge representation (KR), and knowledge visualization (KV) over the past three years. Second, it will try to make evident that the three concepts of KC, KR, and KV blend or even overlap into one another. This is the case in the most promising of prototypes, which will be reviewed. The topic of Concept Maps, in particular, serves this unifying function between KC, KR and KV. This is not surprising, as indeed the term "concept" is at the core of both knowledge ontologies and most definitions of knowledge itself. One such definition is by J. D. Novak, who states that the primary elements of knowledge are concepts and relationships between concepts called propositions. He defines concepts as "perceived regularities in events or objects designated by a label", while propositions consist of "two or more concept labels connected by a linking relationship that forms a semantic unit." [3, p. 206]. The paper will try to shed light on this definition by reviewing several prototypes, leading to a discussion of research directions, aiming to conclude that graphical representations will play a key role in KC, KR, and KV and the semantic Web.

2. KNOWLEDGE CAPTURE

The use of ontologies allows for a deep semantic description in each domain, where a group of people share a common view on the structure of knowledge. Yet this still does not solve the knowledge acquisition challenge, or how to instantiate ontologies with instances of the concepts, and produce machine-useable formalized ontology content, especially without conforming to highly unnatural syntactic conventions. Several developments in the last three years deal with this challenge.

The Ontology Forge allows human experts to create taxonomies and axioms, and by providing a small set of annotated examples, machine learning can take over the role of instance capturing though information extraction technology. [6] Of course knowledge does not only reside in experts who create taxonomies. Much can be extracted from the millions of online forums that have mushroomed in recent years. A post-, rather than page-based, web forum extraction technology by Limanto, et al [12] generates wrappers by studying the structure of the forum HTML files and extracts data inside the pages using the created wrappers. The extraction enables advanced searching and improves the effectiveness of ranking during retrieval. Since this methodology pertains to the field of *data*[1] mining, it will not be elaborated on here.

Another vein of research in data capture is worth pointing out, however, since it shows the transition from KC to KR. This new area dismisses internal representations, but uses the world as its own best representation. As S.D. Larson says:

> Our simpler evolutionary ancestors did not have the luxury of high-level concepts as humans do to help them craft clever taxonomies. Their only view of the world were the low-level patterns of activation they received from arrays of sensor cells. How can we model the processes that enabled them to distinguish food from poison and mate from attacker? For such simple creatures, we must boil the definition of representation down to its most basic and pure: representation is about grouping similar things together and placing dissimilar things farther apart. [10, p. 204]

Larson's system, sketched in Fig. 1, demonstrates the ability of the system to ground symbols in sensory data. The system acquires sensory data and organizes it into classes in an unsupervised fashion. Later, this data is associated with symbols in a supervised fashion.

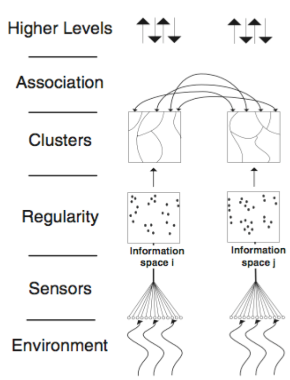

Fig. 1 The model of intrinsic representation [10]

3. KNOWLEDGE REPRESENTATION

According to Donald Norman, the ability to represent the representations of thoughts and concepts is the essence of reflection and of higher-order thought, and further, it is through meta-representations that we generate new knowledge. [14] Rouet et al [17] point out that comprehension involves the construction of a multilayered mental representation. The first level of representation is called a surface representation, consisting in encoding the explicit form of the message: visual features, letters, words and sentences. During the reading process, the surface representation is quickly subsumed by a propositional representation, in which content words are connected into micro and macro-propositions. Finally, propositions from the text are integrated with the reader's prior knowledge in order to form the situation model.

As we know, knowledge representation is a surrogate to something that exists in the real world. This by default makes KR imperfect, for a perfect representation can only be the actual thing being represented. This inherent compromise is exacerbated considering the well-known tradeoff between the representational expressivity of KR versus the efficiency of reasoning with that KR. Researchers have long debated various choices in formalisms, such as Rules, Frames, Semantic nets, and Formal logic. [7] One of the latest methodologies, the Postgres RDB system, is a set of transformations that can be used to move knowledge across two fundamentally different KR formalisms: Frame-based systems and Relational database systems (RDBs). As the authors point out:

> At a coarse level, the transformation from Protégé to RDBs is simple: Classes become tables, slots become attributes, and individuals become rows. However, this simple view omits a great deal from the source Protégé model. In fact, because the expressivity of frame-based systems is greater than that of relational databases, transformations in this direction must necessarily lose some information. Even if loss-less transformation across different KRs is impossible, enough information can be retained to make data transformations worthwhile. [7, p.197]

PostgresRDF is currently being tested in the bioinformatics domain. Another recent KR technique is semantic thumbnailing, by Sengupta et al [18]. While *image* thumbnails summarize structural layout, and are adequate for the purpose of human viewing and browsing, they are not appropriate for deriving any semantic content from the thumbnail. BioKnOT, on the other hand, creates document thumbnails designed for the consumption of software such as search engines, and other content processing systems.

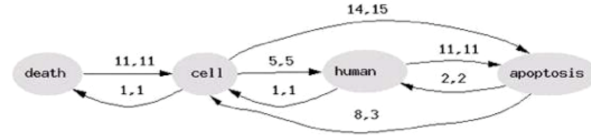

Fig. 2 Screenshot of BioKnOT showing the user's semantic thumbnail of a relevant document [18]

[1] Data has been considered the most basic term in the ladder of information: data->information->knowledge->wisdom.

Another issue in knowledge representation is well summarized by Stutt and Motta:

> What we lack are tags that can be used to indicate to a learner how the learning object may be contextualized, how an object should be interpreted, or how it fits into the central debates in the field, and how to navigate the space composed of the far more important structures of relations which knit topics, concepts, examples and so on, into the fabric of the disciplinary field. [19, p.136]

The idea of the Conceptual Web has recently been expressed, as a layer above the Semantic Web intended to make it more accessible to humans using graphical context maps, which include concepts and relations among concepts. Conzilla (www.conzilla.org) is one such downloadable application, but since this paper does not focus on information retrieval, semantic browsers per se will not be reviewed.

4. KNOWLEDGE VISUALIZATION

The knowledge engineering and KR communities have for years been primarily concerned with knowledge elicitation methodologies and the formal notation used to represent knowledge. [3] The recent interest in Knowledge Management on the part of the business community has brought attention to the effective *portrayal* as well as collaborative construction of knowledge. This has led to a forming of a field now called Knowledge Visualization (KV), defined as "the use of visual representations to transfer knowledge between at least two persons" [1]. Visualization techniques have typically dealt with managing the amount of large data sets that is visible at any given time. Fish-eye view visualization techniques are a well-known approach to solving this problem by maintaining a balance between local detail and global structure [9]. Related to the semantic Web, semantic fisheye views are based on this framework, with one or more "distances" defined semantically, rather than spatially or structurally. Once this semantic *degree of interest* distance is calculated, the visual representation is modified using one or more emphasis algorithms to allow the user to perceive the relative. [9]

Fig. 3. An example session of browsing the semantic relationships between images [9]

Another development that will empower KV has been the Scalable Vector Graphics (SVG) specification, an XML technology developed by the W3C with the purpose to create a standardized vector graphics format for the Web environment. SVG features a set of graphic techniques that are on par with the best graphics design software, integration with the Document Object Model and JavaScript standards, and allowing programs to access and update SVG content dynamically. The extensibility of SVG allows graphics elements to be combined with elements from other domains to form SVG documents, which implies the potential powerful capability of SVG to the collaborative sharing of semantics. Below are a screenshot from an SVG application in geography, and then parts of the related code. [20]

Fig. 4 An SVG rendering of a map [20]

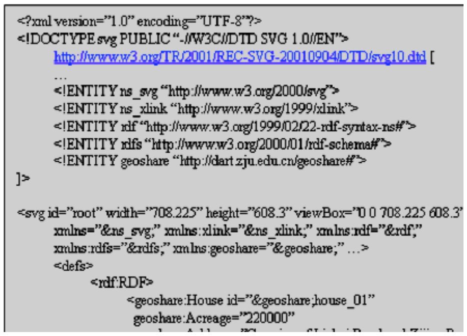

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE svg PUBLIC "-//W3C//DTD SVG 1.0//EN">
      http://www.w3.org/TR/2001/REC-SVG-20010904/DTD/svg10.dtd [
      ...
      <!ENTITY ns_svg "http://www.w3.org/2000/svg">
      <!ENTITY ns_xlink "http://www.w3.org/1999/xlink">
      <!ENTITY rdf "http://www.w3.org/1999/02/22-rdf-syntax-ns#">
      <!ENTITY rdfs "http://www.w3.org/2000/01/rdf-schema#">
      <!ENTITY geoshare "http://dart.zju.edu.cn/geoshare#">
]>

<svg id="root" width="708.225" height="608.3" viewBox="0 0 708.225 608.3"
      xmlns="&ns_svg;" xmlns:xlink="&ns_xlink;" xmlns:rdf="&rdf;"
      xmlns:rdfs="&rdfs;" xmlns:geoshare="&geoshare;" ...>
      <defs>
            <rdf:RDF>
                  <geoshare:House id="&geoshare;house_01"
                   geoshare:Acreage="220000"
```

Fig. 5 A fragment of the SVG document embedded with RDF semantics [20]

Two other visualization formats that deserve mention are Knowledge Charts and Causal Graphs. Knowledge Charts are ontologically permeated high-level representations of a community's knowledge or point of view. The charts represent structures—such as narratives and arguments—using ontologies and provide access to them using graphical representations. [19]

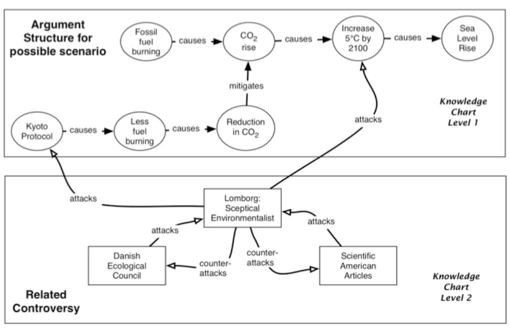

Fig. 6 A web of argumentation related to Global Warming [19]

Causal Graphs, on the other hand, together with their causal semantics for seeing and setting, have the potential to be as powerful a data visualization tool as line graphs or pie charts, argue Neufeld and Kristtorn. [13]

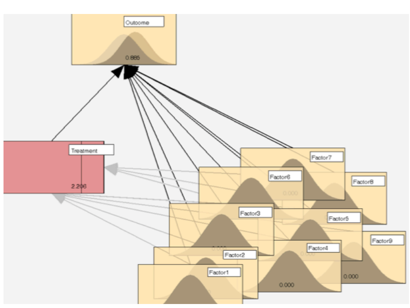

Fig. 7 Causal relationships in a scientific experiment [13]

5. CONCEPT MAPS

Novak and Gowin's research into human learning and knowledge construction has led to the development of Concept Maps: a graphical tool that enables anybody to express their knowledge in a form that is easily understood by others. [3] Concept Maps (CM) are a graphical two-dimensional display of knowledge that is comprised of concepts (usually represented within boxes or circles), connected by directed arcs encoding brief relationships between pairs of concepts. These relationships usually consist of verbs, forming propositions or phrases for each pair of concepts. The simplest CM would consist of two nodes connected by an arc representing a simple sentence such as 'grass is green'. As Neufeld and Kristtorn point out:

> From the perspective of information visualization, it is serendipitous that so simple a structure as a node-link diagram would be consistent with such a careful definition of causality. Node-link diagrams themselves bring semiotic power. Node-link representations form an important class of perceptual graphical codes: a circular line, though just a mark on paper or a monitor, is interpreted as an object at some level in the visual system. Similarly, linking lines imply relationships between objects. [13, p. 257]

As for arrows—250 meanings of them have been categorized! [13]

According to Hayes et al. [8] a Web-oriented knowledge capture tool must be more than simply a user interface to an ontology editor; it must in addition provide intuitive mechanisms for locating appropriate formalized concepts in previously published Web ontologies. Ontology maps, then, can be described as both CMs (allowing the gradual crystallization of thinking while constructing) and graphical representations of groups of OWL description logic axioms.

Hayes et al's "COE" interface provides the ability to output strictly correct OWL/RDF/XML syntax suitable for input to mechanical processors. Yet its graphical nature allows users to rely on spatial and geometric layout to "segregate concepts into meaningful 'chunks' which can be rapidly selected, copied and pasted, and which can be arranged to form meaningful and memorable visual patterns." The interface also suppresses some nodes and arcs that would be visible in a rendering of the RDF graph and represent 'obvious' information. Such information would normally never be mentioned by human users but are required by formal reasoners. This diminishes the perceptual overhead on human readers. In a usability study with a set of published OWL ontologies ranging in size from 11 to 12,260 RDF triples, the COE heuristics reduced the number of nodes by 48% and number of links by 67%. [8]

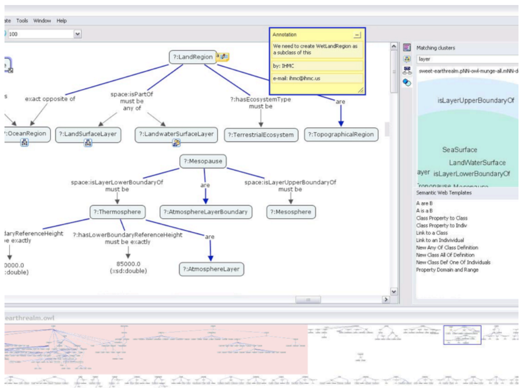

Fig. 8 COE's *definitional view* (left) provides the user with the ontological details necessary to define and modify OWL concepts. The *vicinity concepts view* (right) exposes a number of contextually related concepts (SeaSurface, WaterSurface, etc.), which users can click on to open the ontology map, in which the concept is defined, to investigate the concept further. [8]

Other similar visualizations have been plug-ins for Protégé such as ezOWL or Jambalaya, [20] but stand-alone 'offline' ontology editors are not the main focus of this paper. Another application similar to CEO will be mentioned, however, as it shows the flexibility of CMs. SHAKEN allows domain experts to author knowledge bases with minimal training in KR. [4] Traditional semantic network representations, and the graphical editors supporting them, have been dominated primarily by binary relationships, the authors claim. As Fig. 9 illustrates, ternary relations can be captured, and conditionals can also be represented. In a military experiment, experts quickly became adept at using the SHAKEN interface for capturing knowledge, collectively authored 56 additions to the critiquing knowledge, including 13 new classes. [4]

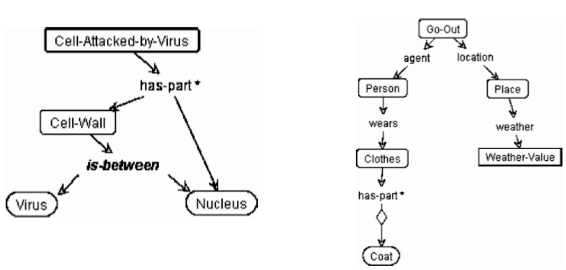

Fig. 9 (left) Graphical presentation of a ternary relation. The cell wall is in between the nucleus of the cell and the attacking virus.

The is-between hyper-edge represents an atomic statement with a ternary predicate: as usual, the incoming arc indicates the first argument; the two outgoing arcs, from left to right, are the second and third arguments. And in Fig. 9 (right) the diamond indicates that the value of an edge is conditional. the user can roll-over the diamond and get a description of the condition. "if it is snowing while going out, one should wear a coat." [4]

The role of ontologies on the semantic Web is undergoing a fundamental change: "rather than being input to specialized data handlers, it is seen as a form of public markup." [8] The term Collaborative Knowledge Visualization has appropriately been coined is Jasminko Novak and Michael Wurst [15], exemplified by their Knowledge Explorer.

This is how their application works. The user is first presented with an agent-generated knowledge map created by means of methods for autonomous machine clustering. This map serves as an initial context and navigation guide for the users exploration of the document space. As the user explores s/he identifies relevant documents and relationships between them which s/he can express by selecting individual items into personal collections and (re-)arranging them according to her personal understanding of their meaning. In this way the user creates a personal map as a natural result of her exploration of information. This template can then be learned by a personal information agent, which can semantically structure arbitrary information.

To achieve this, the Knowledge Explorer the two main elements of the knowledge map visualization: the Content Map and the Concept Map.

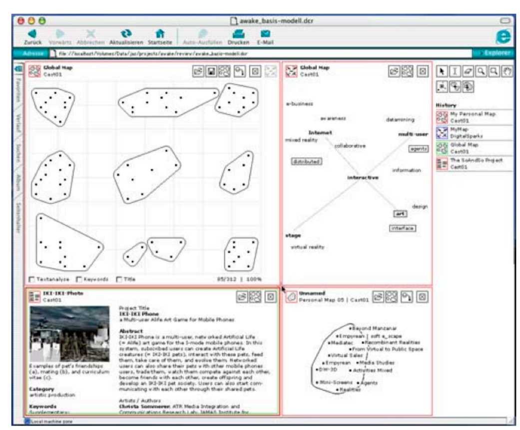

Fig. 10 The Content Map (left) provides an overview of the information space structured according to semantic relationships between items.

The clusters of related documents offers insight into implicit relationships between their content. The Concept Map, on the right, visualizes a concept network extracted from the document pool and redefined by the users. This provides a navigation structure and insight into the criteria that have determined the semantic structuring in the Content Map. [15]

As the authors say:

> Since the personalized map templates have been produced by a user as an effect of his interaction with information and can be dynamically applied to reflect his point of view, they are a form of representation of the user's knowledge that has previously not been expressed. Visualizing the personalized maps and the related concept structures, and making them available to other users is a way of making the users knowledge perceivable and available to others. [14]

Content analysis uses properties of items (word vectors, authors, etc.) to measure the similarity of these items, whereas in context analysis, if two objects appear together in many user-edited clusters, then we can assume that these objects are in some way similar. This is a very interesting feature of the Knowledge Explorer, as items are not only rated by users, like in "collaborative filtering" systems, but are put into the context of other items.

6. FUTURE DIRECTIONS

Even though using CM for KC, KR, and KV is clearly a good idea, are CMs a natural representation of knowledge as it occurs in the real world? The real world knowledge is found in form of text, diagrams, pictures, tables, formulas, etc., and so CM may need to be complemented where more compelling visualizations. [4] As early as 1999, Leake and Wilson developed a tool for capturing expert design knowledge and cases through concept mapping in the field of aerospace design that made use of graphics and pictures alongside CM. [11]

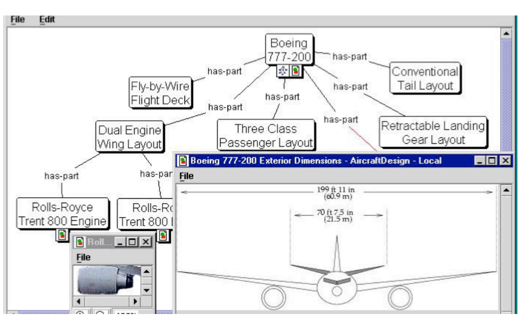

Fig. 11 Sample screen of the Cmap editor [11]

But graphical representations do not need to fit the node-link format. More recently, the "Collaboratory" is a good example of a complete rich media application. [5]

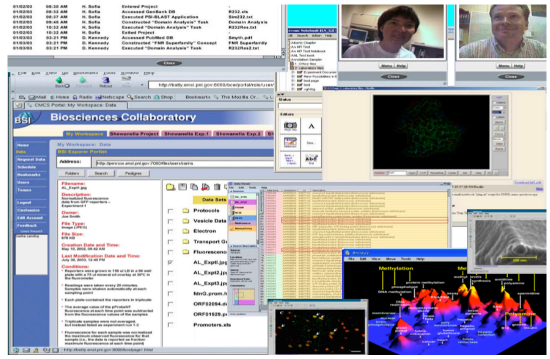

Figure 12. The *Collaboratory* provides an integrated set of tools – allowing researchers to share data, applications, and communications. [5]

It is important, however, to realize that virtual reality and 3D simulations will not necessarily be the ultimate in knowledge building. More than ten years go, Donald Norman distinguished between experiential and reflective cognition and artifacts. Experiential artifacts mediate between the mind and the world, while reflective artifacts allow us to ignore the real world and concentrate only upon artificial, representing worlds. In reflection, one wants to contemplate the experience and go beyond, finding new interpretations or testing alternative courses of action. Norman argues:

> Rich, dynamic, continually present environments can interfere with reflection: These environments lead one toward the experiential mode, driving the cognition by the perceptions of event-driven processing, thereby not leaving sufficient mental resources for the concentration required for reflection…. Until we learn how to provide stable external representations that can be examined, contrasted, and transformed into higher-order, more powerful representations, these new technologies will remain devices of exploration and entertainment and fail in their power to enhance cognition. [14, p.25]

One such true artifact for reflection, *mind mapping*, utilizes the way the brain works, and associative memory, visual memory, imagery, and Gestalt principles are exploited to reflect one's internal structure and processes. [2]

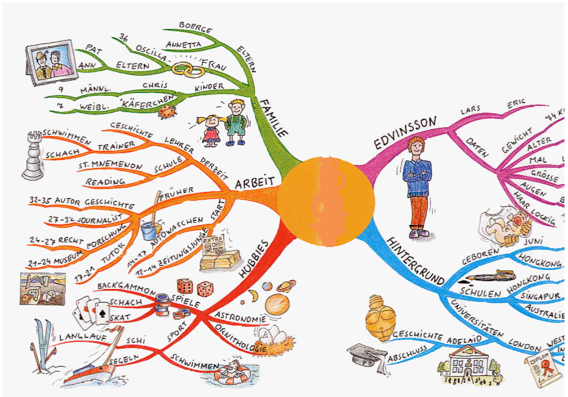

Fig. 13 An example of a mind-map [2]

In examples such as the mind-map, we see how a rich form interplays with content, and how these interplay with context and space. According to Stutt and Motta, it is likely that the Semantic Web for Learning, like Ancient Greece or Medieval Italy, will be composed of loosely related Knowledge Neighbourhoods. These will be locations on the Web where communities collaborate to create and use representations of their knowledge. [19] And as Price & Coulter-Smith say in their theory of pervasive information spaces:

> Our framework for representation of pervasive information is based upon the metaphor of space. We may take the approach of the town-planner, who must juggle issues of local economy with social and political forces. We may take the approach of the architect, who chooses to structure space according to his beloved metaphor (such as the Centre Pompidou or La Parc de La Villete). Or we may choose the mathematical certainty of a chessboard, a city-block of intersecting roads, or else the organic development of a European city, from medieval to contemporary. (16, p.99)

The future of knowledge capture, representation, and visualization, then, will clearly be an endeavor that makes use of both rich graphics as well as environmental design—in addition to the technology that underlies them.

REFERENCES

[1] R. Burkhard. Towards a framework and a model for Knowledge Visualization, in: Tergan and Keller (eds.) Knowledge and information visualization: searching for synergies, Heidelberg / New York, Springer Lecture Notes in Computer Science, 2004.

[2] T. Buzan, and B. Buzan. The mind map book: How to use radiant thinking to maximize your brain's untapped potential. New York: Plume Books, 1993.

[3] A.J. Cañas, R. Carff, G. Hill, M. Carvalho, M. Arguedas, T.C. Eskridge, J. Lott, and R. Carvajal. Concept Maps: Integrating Knowledge and Information Visualization, S.-O. Tergan and T. Keller (Eds.): Knowledge and Information Visualization, LNCS 3426, pp. 205– 219, 2005.

[4] V. Chaudhri, K. Murray, J. Pacheco, P. Clark, B. Porter, and P. Hayes. Graph-Based Acquisition of Expressive Knowledge, E. Motta et al. (Eds.): EKAW 2004, LNAI 3257, pp. 231–247, 2004.

[5] G. Chin Jr. and C.S. Lansing, Capturing and Supporting Contexts for Scientific Data: Sharing via the Biological Sciences Collaboratory. CSCW'04, November 6-10, 2004, Chicago, Illinois, USA.

[6] N. Collier, K. Takeuchi, A. Kawazoe, T. Mullen, and T. Wattarujeekrit. A Framework for Integrating Deep and Shallow Semantic Structures in Text Mining, V. Palade, R.J. Howlett, and L.C. Jain (Eds.): KES 2003, LNAI 2773, pp. 824–834, 2003.

[7] J. Gennari, P. Mork, and L. Hao. Knowledge Transformations between Frame Systems and RDB Systems, K-CAP'05, October 2–5, 2005, Banff, Alberta, Canada., 2005.

[8] P. Hayes, T. Eskridge, R. Saavedra, T. Reichherzer, M. Mehrotra, and D. Bobrovnikoff. Collaborative Knowledge Capture in Ontologies K-CAP'05, October 2–5, 2005, Banff, Alberta, Canada, 2005.

[9] P. Janecek, and P. Pu. Searching with Semantics: An Interactive Visualization Technique for Exploring an Annotated Image Collection, R. Meersman and Z. Tari (Eds.): OTM Workshops 2003, LNCS 2889, pp. 185–196, 2003.

[10] S.D. Larson. Intrinsic Representation: Bootstrapping Symbols from Experience, A.Y. Tawfik and S.D. Goodwin (Eds.): Canadian AI 2004, LNAI 3060, pp. 202–216, 2004.

[11] D. Leake, and D. Wilson. Combining CBR with Interactive Knowledge Acquisition, Manipulation, and Reuse K.-D. Althoff, R. Bergmann, L.K. Branting (Eds.), ICCBR-99, LNAI 1650, pp. 203-217, 1999

[12] H. Limanto, N.G. Nguyen, V.T. Trung, Q.H. Nguyen, J. Zhang, and Q. He. An Information

Extraction Engine for Web Discussion Forums, WWW 2005, May 10-14, Chiba, Japan, 2005.

[13] E. Neufeld, and S. Kristtorn. Picturing Causality – The Serendipitous Semiotics of Causal Graphs, A. Butz et al. (Eds.): SG 2005, LNCS 3638, pp. 252 – 262, 2005.

[14] D. Norman. Things That Make Us Smart: Defending Human Attributes in the Age of the Machine. Addison Wesley Publishing Company, 1994.

[15] J. Novak, M. Wurst, M. Fleischmann, and W. Strauss. Discovering, Visualizing, and Sharing Knowledge through Personalized Learning Knowledge Maps, L. van Elst, V. Dignum, and A. Abecker (Eds.): AMKM 2003, LNAI 2926, pp. 213–228, 2003.

[16] C. Price, and E. Coulter-Smith. Developing a Theory and Practice of Pervasive Information Capture, Processing, Visualization, and Documentation, SIGDOC'05, Sept 21-23, Coventry, England, 2005.

[17] J.-F. Rouet, H. Potelle, and A. Goumi. The Role of Content Representations in Hypermedia Learning: Effects of Task and Learner Variables S.-O. Tergan and T. Keller (Eds.): Knowledge and Information Visualization, LNCS 3426, pp. 343 – 354, 2005

[18] A. Sengupta, M. Dalkilic, and J. Costello. Semantic Thumbnails - A Novel Method for Summarizing Document Collections," SIGDOC'04, October 10-13, Memphis, TN, USA, 2004.

[19] A. Stutt, and E. Motta. Semantic Webs for Learning: A Vision and Its Realization, E. Motta et al. (Eds.): EKAW 2004, LNAI 3257, pp. 132–143, 2004.

[20] Z. Xu, H. Chen, H., and Z. Wu. Applying Semantic Web Technologies for Geodata Integration and Visualization, J. Akoka et al. (Eds.); ER Workshops 2005, LNCS 3770, pp. 320-329, 2005.

[21] G. Yolanda, and V. Ratnakar. IKRAFT: Interactive Knowledge Representation and Acquisition from Text, A. GÛmez-PÈrez and V.R. Benjamins (Eds.): EKAW 2002, LNAI 2473, pp. 27-36, 2002.