Modelling Naturalistic Argumentation in Research Literatures

Victoria Uren, Simon Buckingham Shum, Clara Mancini and **Gangmin Li** [1]

Abstract. This paper characterises weaknesses in the ability of current digital libraries to support scholarly inquiry, and as a way to address these, proposes services grounded in semiformal models of the naturalistic argumentation commonly found in research literatures. It is argued that a design priority is to balance formal expressiveness with usability. We summarise the requirements for an argument modelling scheme for use by untrained researchers, describe the resulting scholarly discourse taxonomy, contrast it with other domain modelling and semantic web approaches, before focusing on examples of computational services to support the filtering and analysis of the repository.

1 INTRODUCTION

This workshop issue brings together work which investigates the implications of modelling, with computational support, naturally occurring arguments as formulated in the course of everyday work. This paper contributes with respect to several issues raised by this challenge:

- A domain application of natural argumentation modelling, namely, to scholarly electronic publishing and discourse, is presented;
- This application is accomplished through Web-mediated computer supported collaborative argumentation, for modelling the specific types of argumentation found in research literatures;
- Tools are provided for interacting with structures of argument, including visualisation tools and interfaces supporting structured dialogue.

We start by characterising some weaknesses in current scholarly/scientific publishing infrastructures, and as a way to address these, propose computational services grounded in semiformal models of the naturalistic argumentation found in research literatures. Let us begin with a question to focus the imagination:

In 2010, will scientific knowledge still be published solely in prose, or can we imagine a complementary infrastructure that is 'native' to the emerging semantic, collaborative web, enabling more effective dissemination and analysis of ideas?

It is important to say that we are seeking neither to replace textual narrative as an expressive medium, nor its products such as books and peer reviewed publications. We seek instead to augment them by exploiting globally networked information in ways that – precisely because of its historical pedigree – the prose publication cannot support. Conventional scholarly publications are the result of a long co-evolution of notational form with print publishing technology, but are not designed to take advantage of today's information infrastructure. While information retrieval and human language engineering research seek to extract structure of different sorts from these texts, the strategy pursued here is to question why this structure is lost in the first place? Instead, we are investigating the interdependent representational and usability challenges in capturing and publishing the conceptual structure of a research article as a human and machine readable, semiformal structure.

In the following sections, we set out the rationale for this work (Section 2), and then focus on associated challenges, with an approach derived from the research into Hypertext, Semantic Web, Human-Computer Interaction, Computer-Supported Collaborative Work and Computational Linguistics. Section 3 specifies the particular requirements for an argument modelling scheme which will be usable by researchers untrained in conceptual modelling or argumentation theory, Section 4 describes the modelling scheme, before Sections 5 and 6 consider user interfaces for modelling arguments in publications, and computational services for analysing the repository as contributions are made to it.

2 LIMITS OF DIGITAL LIBRARIES

Researchers are benefiting from more rapid access to research documents as resources such as new digital libraries and e-print archives go online almost by the week, but researchers (like almost all other professions) are also drowning in this ocean, with less time to track growing numbers of conferences, journals and reports. But beyond tracking new results, there is the whole dimension of analysing a literature. Researchers are concerned with the significance of a contribution to the literature, but no digital library can answer common – but complex – questions which are fundamental to critical inquiry, and which we seek to instill in our students, such as:

- Which work supports or challenges this?
- What is the intellectual lineage of this idea?
- What data is there to support this specific claim or prediction?
- Who else is working on this problem?
- Has this approach been used in other fields?
- What logical or analogical connections have been made between these ideas?

Such questions self-evidently require complex interpretative work, and moreover, there may be disagreements of different sorts. The above questions require semantic annotation at a different level from that addressed by conventional metadata or ontologically-based markup in semantic web research, which seek to iron out inconsistency, ambiguity and incompleteness in the way resources are characterised (clearly these are undesirable if the domain is uncontentious). In contrast, principled disagreement about the significance of a contribution, conflicting perspectives, new evidence that changes the world to be modelled, and the resulting ambiguities and inconsistencies are precisely what define a field as research; they are the objects of explicit inquiry.

In sum, there remains a gap in the researcher's digital toolkit: tools to track (claimed) *contributions* in a field, and to express,

[1] Knowledge Media Institute, The Open University, Milton Keynes, MK7 6AA, UK

CMNA IV, 4th *Workshop on Computational Models of Natural Argument*
Edited by Floriana Grasso, Chris Reed and Giuseppe Carenini
ECAI 2004

analyse and contest their *significance*. It in this context that structured argumentation has a contribution to make to help individuals and research teams construct a picture of the key arguments in the literature from their particular standpoint. Let us now consider the detailed requirements for such a research tool.

3 MODELLING NATURALISTIC SCHOLARLY DISCOURSE

"Ontologies" are the term used in knowledge modelling and agent research, and increasingly within the semantic web community, to describe an abstract (implementation-independent) specification of concepts, attributes and relationships [3]. Typical semantic web applications develop an ontology to control interpretation or semantic annotation in a specific domain of inquiry (such as an ontology of problem-solving methods) or to model a particular aspect of the world (such as organisational functions), enabling machine-to-machine interoperability and interpretation. In contrast, we propose a semiformal ontology for scholarly discourse, primarily for *humans* to communicate through as a medium for publishing and discourse (although we envisage agents as protagonists and claim-makers at some point), with the express goal of supporting multiple (often contradictory) perspectives. In this sense it is an ontology for *principled disagreement*. It still requires consensus in the sense that participants subscribe to the ontology as a reasonable language for "making and taking perspectives" [1], but in contrast to most existing ontology applications, stakeholders need not agree at all on the structure of the field being modelled. All modelling is interpretation, but when there is meant to be consensus, the end-user community is not given the option of disputing the ontology or the way in which it has been applied. In contrast, our modelling scheme makes it explicit that every contribution can be contested. This emphasis is carried through into the language of the user interface and help information, which talks about "claims", and makes clear that the system's function is to serve as a medium for agreeing and disagreeing in various ways.

A representation scheme for modelling arguments in papers needs to achieve a fine balance between completeness and usability. It would be possible to produce an elegant formal ontology that could perform reasoning of the type supported in other computational argument modelling systems. However, if the database is to be populated by domain experts from fields outside knowledge engineering it seems implausible that a critical mass of readers of research papers would feel inclined to learn such a scheme or have the confidence to publish the argument maps they built using it. Conversely, too weak a scheme will not deliver sufficient services to make it worth the readers' while to use it.

4 MODELLING SCHOLARLY DISCOURSE

4.1 Data model

Our modelling scheme comprises nodes and links. As we now explain, nodes may be atomic or composite at the end user's discretion. *Atomic nodes*[2] are expressed as short pieces of free text succinctly summarising a 'contribution' (at whatever granularity the researcher wishes to express this). For instance, an (optionally untyped) atomic node might simply be the name of a new algorithm that the researcher wishes to add to the network as a contribution, e.g.: *PageRank.* A different, typed atomic node might summarise an empirical result: <*Data*> *Undergraduate chemistry exam performance is doubled after training on the ChemVR system.* These are now objects (loosely analogous to published websites with URLs) which others can link to in their own work, whether positively or negatively.

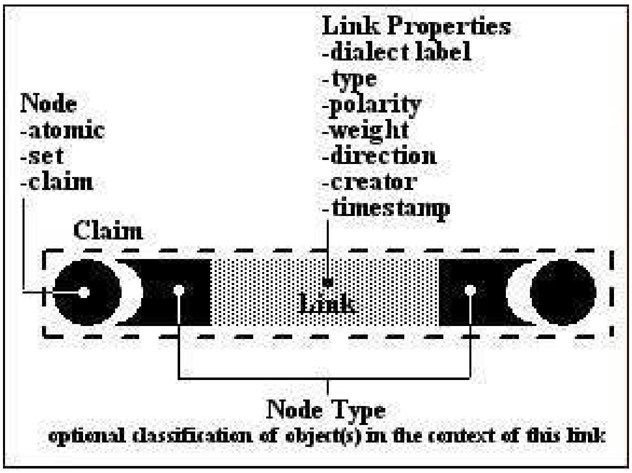

Figure 1. Structure of a *Claim* in the discourse ontology.

As shown in Figure 1, an object may optionally be assigned a type (e.g. *Data, Language, Theory*), stored as part of the link connecting it. By storing the node type in the link, rather than binding it intrinsically to the node, the typing of nodes is made context dependent: objects may play different *roles* in different contexts, since researchers may disagree on the node's type: e.g. is this *Language* also a *Theory*? Is this based on *Opinion* or *Data*? One person's underlying *Theory* may be someone else's *Problem.*

In addition to atomic nodes, two kinds of composite object can be used as the nodes in Claims. A *Set* is a group of objects (atomic nodes, Sets or Claims) declared by the user to share a common theme and enabling them to be referenced by a single named node (e.g. *Constructivist Theories of Learning*). *Claim triples* themselves can also be linked from or to other atomic nodes, Sets or Claims. This nesting allows users to build complex conceptual and argument structures.

To illustrate claim triples, consider the following:

[Decision Forest Classifier] *(uses/applies/is enabled by)*
[Decision tree learning]

This uses one of the General relations *uses/applies/is enabled by* to assert that the *Decision Forest classifier* studied in the paper uses a well known method, *Decision tree learning*. The latter node was introduced in a different document, so this link has a contextual role: it locates the paper near similar claims.

[Decision Forest classifier improves on C4.5 and kNN]
(is inconsistent with) [SVM and kNN outperform other classifiers]

This claim uses the negative, Supports/Challenges relation *is inconsistent with* to link one of the experimental results of this paper to a result in a third paper. In addition to its contextual role, locating the claim near other comparisons of classifiers, this claim has a rhetorical role: it contrasts pieces of evidence that make contradictory assertions.

The priority of the system in supporting multiple perspectives means that it does not add the kinds of constraints that would be

[2] In previous papers we have referred to atomic nodes as "Concepts" but have found that it is more helpful to refer to them in semantic hypertext language as nodes.

expected when one can safely assume a single worldview. One researcher may think that X is an example of Y, but a peer may argue the opposite. This is the substance of research discourse, but limits the scope for automated reasoning. However, we are focusing on the argumentation level primarily, with the domain model emerging as a secondary product; other modelling efforts could focus on fields where there is consensus (or where only consistent views are modelled), and build richer, more constrained representations that can support correspondingly more advanced reasoning.

4.2 Link semantics

A *link* between two nodes is typed with a natural language label from a discipline-specific *dialect*, which in turn is a member of a generic, discipline-independent *class* (e.g. *Problem-related; Taxonomic; Causal*).

Elsewhere we have described the iteration from the first to the current version of the ontology [2]. The current relational scheme is summarised in Table 1. Our goal is to provide a given research community with a dialect that will cover the most common claims that they make (there may well be exceptional kinds of contributions that fall outside the expressiveness of the vocabulary, but the generic *Other Link* is available for those situations).

4.3 Theoretical basis of link semantics

The link taxonomy evolved through a combination of theoretical and data-driven processes. The theory-driven approach derived from psycholinguistics and computational research on Cognitive Coherence Relations (CCR), combined with a semiotic perspective on representation, which emphasises the interpretive act of modelling [6].

According to CCR theory, discourse coherence is a cognitive phenomenon that goes beyond any linguistic expression. It depends on the interpreter's ability to create a coherent cognitive representation of the discourse content, by establishing coherent connections between its parts. The categories of discourse connectivity are expressed in natural language by specific indicators, but these are evidence of the deeper cognitive processes that natural language is optimised to express [8].

As proposed by comprehensive parametrical descriptions [5, 9], two discourse units can be related by additiveness, temporality (sequentiality) or causality; additiveness can be conjunctive or comparative (similarity); causality can be actual or hypothetical (conditionality); both causal and additive relations can be semantic (e.g. cause-effect) or pragmatic (e.g. argument-claim); they can have positive or negative polarity (e.g. similarity or contrast); the order of the related units can be forward (e.g. cause-effect), backward (e.g. effect-cause) or bi-directional (e.g. list).

Structuring and using our link taxonomy based on these cognitive primitives offers valuable advantages. First of all, this grounds the taxonomy in what - from experimental evidence - appears to be psychological reality, which in principle gives the taxonomy stability and applicability across different disciplines, media and discourse types. Secondly, this ensures that the main categories of this psychological reality are represented in the taxonomy, which ensures a more balanced expression of different kinds of connection. Thirdly, at any level of articulation and specialisation, this ensures that the taxonomy's links are accountable for by a small number of primitives, which allows for consistent discourse modelling, processing and searching at very different levels of granularity. Finally, defining our relation taxonomy based on the parametrical description proposed by CCR theory has the potential for underpinning the design of complex services in the future.

Table 1. The revised discourse taxonomy following a first iteration and use analysis.

Relation Class	Dialect label	Polarity/ Weight
General	is about	+/1
	uses/applies/is enabled by	+/1
	improves on	+/2
	impairs	– /2
	other link	+/1
Problem Related	addresses	+/1
	solves	+/2
Supports/ Challenges	proves	+/2
	refutes	–/2
	is evidence for	+/1
	is evidence against	–/1
	agrees with	+/1
	disagrees with	–/1
	is consistent with	+/1
	is inconsistent with	–/1
Causal	predicts	+/1
	envisages	+/1
	causes	+/2
	is capable of causing	+/1
	is prerequisite for	+/1
	prevents	–/2
	is unlikely to affect	–/1
Similarity	is identical to	+/2
	is similar to	+/1
	is different to	–/1
	is the opposite of	–/2
	shares issues with	+/1
	has nothing to do with	–/1
	is analogous to	+/1
	is not analogous to	–/1
Taxonomic	part of	+/1
	example of	+/1
	subclass of	+/1
	not part of	–/1
	not example of	–/1
	not subclass of	–/1

The relations of our taxonomy (Table 1) can easily be described in terms of CCR parameters [7]. For instance, the *General* relation 'is-about' derives from *elaboration*, a positive pragmatic additive relation. Elaboration is a relation between two discourse units (atomic or composite nodes in the data model), one of which has the rhetorical function of explaining, expanding, articulating the content of the other unit. Elaboration has a lot in common with another positive pragmatic additive relation of comparative nature: 'agrees-with', whose rhetorical function is reinforcing the content expressed in one discourse unit by adding up more content expressing the same perspective. Now, if the user was to search for all the discourse units that 'are-about' the discourse unit X, the system would know that all the discourse units that 'agree-with' unit X are also relevant, and it can propose them to the user as a secondary result of their search. A further example of this kind of analysis for an implemented service is given in section 6.2.

If the functionalities of our relations are CCR-based, however, their apparent organisation within the taxonomy is informed by practical considerations dictated by a more data-driven approach. The data-driven approach was, in fact, to model argumentation as we found it in a range of research domains, including computer supported collaborative work, text categorization, and literary criticism. Relations common to several domains were identified that are used in argumentation practice. We found we could classify these into groups with similar rhetorical implications: Supports/Challenges, Problem Related, Taxonomic, Causality, Similarity, and General. Each relation belongs to one group. Consistently with what observed by CCR theory, we also found that some relations occurred in pairs of opposites, as *proves* and *refutes*, where one has positive and the other negative implications. Refutes has negative polarity in our taxonomy since it implies *dis*proof and it is commonly used to relate an *"argument against"* to a claim, from an argumentative practice viewpoint, 'refutes' is considered a negative relation.

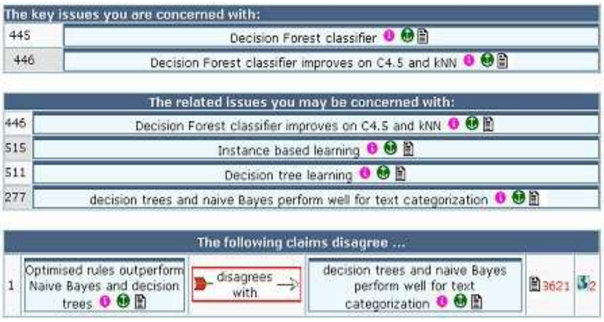

Figure 2. Arguments that contrast with the nodes in a research paper. Key: clicking displays node metadata; sets the node as the focus, to show incoming and outgoing relations; links to the document metadata/URL. links to information about the node's creator.

By defining relations in terms of type and polarity we can reason with them at a higher level of granularity than individual relations; it is not just the claims made using the *refutes* relation that represent *"arguments against"* something, but any claims made using links that have negative polarity. Furthermore the same taxonomy of relations can be employed by research communities, which speak different "dialects", or even different languages, by changing the labels of the relations, without changing the underlying functionality of ScholOnto. To summarise, thus far, our goal is to provide a given research community with a dialect that will cover the most common, significant kinds of 'claims' made in their literature (there may well be exceptional kinds of contributions that fall outside the expressiveness of the vocabulary, but a generic *Other Link* is available for those situations). The taxonomy could be much more expressive, rigorous and formal, but we are walking the tightrope between usability and formal rigour. We propose that these kinds of connections are expressed at a level, which most researchers would not only recognise, but indeed, would naturally use when summarising part of a literature.

5 INTERFACES FOR CONSTRUCTING ARGUMENT MODELS

As a research vehicle for developing these ideas, we have implemented a client-server system called ClaiMaker which enables distributed modelling of documents in a literature, and provides a variety of services for browsing and analysing the emergent conceptual graphs. Infrastructure details are given in Li, *et al.*[4], and are not of primary concern here.

Our interest in working with non-modellers has placed particular emphasis on user interfaces intended to support naturalistic argument modelling by non-experts, and ways to pursue the technology deployment strategies listed above. The ClaiMaker prototype has evolved as our understanding of the modelling process has grown.

For those interested in the user interface design issues, please see [2] and [10]. However, in this paper we focus on the extent to which computational services can still be deployed, even when the representation's expressiveness has been constrained by design criteria of simplicity and learnability.

6 COMPUTATIONAL SERVICES TO FILTER/ANALYSE CLAIM NETWORKS

6.1 Example 1: *Perspective* analysis

Consider a common question that many researchers bring to a literature: *"What arguments are there against this paper?"* Despite the centrality of such a notion, there is not even a language in which to articulate such a query to a library catalogue system, because there are no indexing schemes with a model of the world of scholarly discourse. There is no way to express the basic idea that *researchers disagree*. If we can improve on this, then we have a good example of the argumentation taxonomy adding value over existing retrieval methods.

How can we realise such a query? First, we are looking for *arguments against*, which map to the taxonomy as negative relations of any type (recall that all relations have positive polarity or negative polarity). At a trivial level, *this paper* corresponds to the currently selected document in ClaiMaker.[3] More substantively, *this paper* refers to the *claims* that researchers have made about the document, specifically, the nodes linked to it. Moreover, we can extend this to *related nodes*, using the following definition: *the extended set of nodes linked by a positive relation to/from the document's immediate nodes*.

For the given document, this discovery service does the following:

- finds the nodes associated with that paper;
- extends the set of nodes by adding positively linked nodes from other papers;
- returns claims against this extended node set.

Typical results are presented in Figure 2. ClaiMaker then supports further structured browsing; for instance, having discovered that one of the nodes related to the article is challenged by *Optimized rules outperform Naïve Bayes and decision trees*, clicking on the icon sets this as the focal node of interest, showing its immediate neighbourhood.

6.2 Example 2: *Lineage* analysis

A common activity in research is clarifying the lineage behind an idea. Lineage is essentially ancestry and (with its inverse, the descendant) focuses on the notion that ideas build on each other. Where the paths have faded over time or been confused, uncovering unexpected or surprising lineage is of course a major scholarly contribution. We have a more modest goal to start with in ClaiMaker: to provide a tool to pick out from the "spaghetti" of claims, candidate streams of ideas that conceptually appear to be building on each other. Our lineage tool tracks back (semantically, not in time) from a node to see how it evolved, whereas the descendants tool tracks forward from a node to see what new ideas

[3] If not already in the database (e.g. we are working with journal publishers), one can manually enter document metadata, or more conveniently, upload one's personal library of bibliographic metadata in a standard format such as Refer or Bib.

evolved from it. Since descendants are the inverse of lineage (and are implemented as its literal inverse) we will only discuss lineage.

So, let us consider a new query: *Where did this idea come from?* A claims network can be treated as a graph, with nodes as vertices, and the links between nodes as edges. A path in a graph is a sequence of connected edges. A lineage can be conceptualised as a path in which the links suggest development or improvement. The problem of finding lineage in ClaiMaker can then be formulated as a path matching problem, a well known problem in graph theory for which algorithms exist.[4]

To provide lineage analysis as a ClaiMaker service, path queries are constructed from link-types using a set of primitives. For example, we can search for paths that may be of any length, and which contain (in any order) any of the positive links that have type *similarity* in either direction, or the two general links *uses/applies/is enabled by* or *improves on*, going in the direction away from the target node of the query.

The *improves on* link type is included to reflect the notion of progress implicit in lineage, while *uses/applies/is enabled by* has a weaker implication of "building upon". In CCR terms these are both positive semantic causal relations: in the first case, one phenomenon causes its own improvement by the other in the same way that a problem calls for being given a solution; in the second case, one phenomenon is a direct cause or condition for the other to take place.

The *similarity* links - which constitute positive semantic comparative additive relations in CCR terms - are included because if a new node is like another that *improves on* a third, then the new node may well also be an improvement. *Similarity* links are acceptable in either direction because comparative relations are bi-directional (if A is like B, then B is like A).

Summarising, from the CCR viewpoint, the functionality of lineage needs to always follow positive relations, and they need to be either causal or comparative: either they denote a step forward along a development line, or a convergence across different lines. Figure 3 shows examples of acceptable paths that could be returned by this lineage analysis.

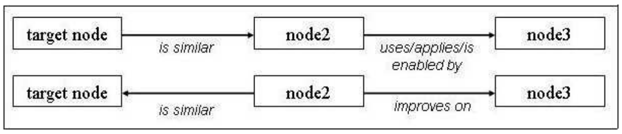

Figure 3. Examples of paths that could be returned by a lineage analysis on a target node (see text for the specification of the query).

The search can be tightened by filtering the paths returned to ensure they contain the *improves on* relation, after which only the second of the paths in Figure 3 would be retained. Conversely, one can relax the conditions to broaden the search, for instance, to permit the inclusion of any Problem-related links (see Table 1), since *addressing* or *solving* a known problem usually represents progress of some sort. One could also include Taxonomic links, since if a *part of* some innovation *improves on* another approach then it implies there may be improvement overall. Note that in these cases, the direction of the link is fundamental: it is only problems that the new node *solves* that are of interest, and even if a whole innovation is an improvement, there is no reason to assume that every *part of* it is also. One advantage of the path matching approach is that it facilitates the use of directional elements in queries.

The results of this kind of structural query can then be rendered in a variety of forms back to the user. Figure 4 shows a visualization of the structure extracted from the claims network in response to a lineage query about a node.

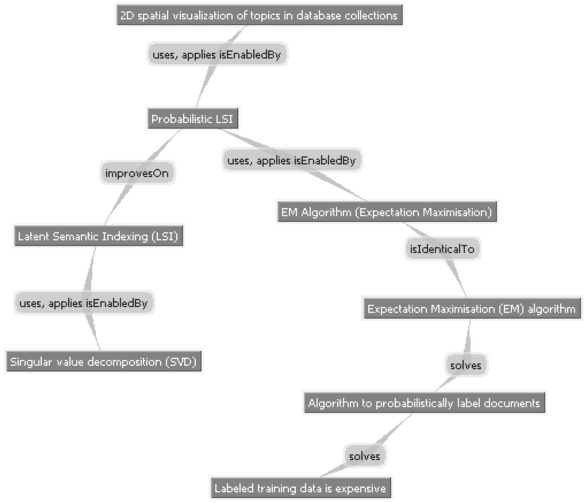

Figure 4. Visualization of the results of a *lineage analysis*, a representation of the claims in the network on which the top node explicitly and implicitly builds, or alternatively, a guide to the local context in which a node is embedded.[5]

The lineage function (and its inverse, descendants) can be thought of as providing an analytical tool to excavate the *foundation* under an idea (or conversely, an indicator of its *impact*). From a navigational perspective, they can be thought of as offering *focused browsing tools*. In response to a *"Where am I?"* question, they give answers in terms of developmental context, positioning ideas in the literature in terms of their evolution.

To summarise, term-based information retrieval handles documents as isolated entities defined by the words in them. Citations in a document give no indication of authors' intentions in referring to other work; we cannot even tell if a paper is referenced because the authors support or are diametrically opposed to it. The examples of *Perspective Analysis* and *Lineage Analysis* demonstrate how the discourse taxonomy can make the connections between ideas in different documents explicit, enabling novel and powerful kinds of query.

7 CONCLUSIONS

We have reported our approach to semiformal modelling of the argumentation commonly found in research literatures. The scholarly discourse taxonomy at the root of this approach was analysed in terms of cognitive coherence relations theory. Finally we focused on examples of computational services to support the filtering and analysis of the repository.

We argue that any scheme which aims to help researchers model naturalistic argumentation needs to balance usability against formal expressiveness. Cognitive coherence relations help in creating this balance since they allowed us to analyse the taxonomy, created

[4] A semantic web standard based on graphs is the *Resource Description Framework* <www.w3.org/RDF>. In the analysis presented here we use the *Ivanhoe* path matching tool available in the *Wilbur* RDF toolkit <wilbur-rdf.sourceforge.net>.

[5] Graph visualization courtesy the *Ceryle Project* by Murray Altheim, Knowledge Media Institute, Open University <www.kmi.open.ac.uk/projects/ceryle/>.

originally in a data driven manner, in a formal way. Analysis of existing services suggests that CCR parameters are used in them in a systematic manner. This leads us to propose that services based on CCR parameters should be a focus of our future research.

8 ACKNOWLEDGEMENTS

This work was funded by the UK EPSRC Distributed Information Management Programme (GR/N35885/01).

9 REFERENCES

[1]Boland, R. J. J., & Tenkasi, R. V. (1995). Perspective making and perspective taking in communities of knowing. Organization Science, 6(4), 350-372.

[2]Buckingham Shum, S., Uren, V., Li, G., Domingue, J., Motta, E., & Mancini, C. (2002). Designing representational coherence into an infrastructure for collective sensemaking.. Invited contribution to: National Science Foundation Workshop on Infrastructures for Distributed Communities of Practice, San Diego, CA.

[3]Gruber, T.R. (1993). A translation approach to portable ontologies. Knowledge Acquisition, 5(2):199-220

[4]Li, G., Uren, V., Motta, E., Buckingham Shum, S., Domingue, J. (2002) ClaiMaker: weaving a semantic web of research papers. In ISWC2002, 1st International Semantic Web Conference, Sardinia.

[5]Louwerse, M. (2001). An Analytic and Cognitive Parametrization of Coherence Relations. Cognitive Linguistics, 12 (3), pp. 291-315

[6]Mancini, C., & Buckingham Shum, S. (2001). Cognitive coherence relations and hypertext: From cinematic patterns to scholarly discourse. Proc. ACM Hypertext 2001, (Aug. 14-18, Århus, Denmark), 165-174. New York:. ACM Press

[7]Mancini, C. (2003). Towards Cinematic Hypertext: a Theoretical and Empirical Investigation. Ph.D. thesis. Knowledge media Institute, The Open University, Milton Keynes, UK

[8] Sanders, T.J.M, Noordman, L.G.M. (2000). The Role of Coherence Relations and Their Linguistic Markers in Text Processing. Discourse Processes, 29(1), pp.37-60

[9]Sanders, T.J.M., Spooren, W.P.M., Noordman, L.G.M. (1993). Coherence Relations in a Cognitive Theory of Discourse Representation. Cognitive Linguistics, 4(2), pp.93-133

[10]Uren, V., Buckingham Shum, S., Sereno, B. and Li, G.. (2003) Interfaces for capturing interpretations of research literature. Distributed and CollectiveKnowledge Capture Workshop, Knowledge Capture Conference, Florida, Oct. 2003