

Modeling Naturalistic Argumentation in Research Literatures: Representation and Interaction Design Issues

Simon J. Buckingham Shum,^{*} Victoria Uren,[†] Gangmin Li,[‡]
Bertrand Sereno,[§] Clara Mancini[¶]
*Knowledge Media Institute, The Open University,
Milton Keynes MK7 6AA, United Kingdom*

This article characterizes key weaknesses in the ability of current digital libraries to support scholarly inquiry, and as a way to address these, proposes computational 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, making it critical to coevolve the modeling scheme with appropriate user interfaces for argument construction and analysis. We specify the requirements for an argument modeling scheme for use by untrained researchers and describe the resulting ontology, contrasting it with other domain modeling and semantic web approaches, before discussing passive and intelligent user interfaces designed to support analysts in the construction, navigation, and analysis of scholarly argument structures in a Web-based environment. © 2007 Wiley Periodicals, Inc.

1. INTRODUCTION: ARGUMENT MODELING AS SCHOLARLY PUBLISHING

This journal special issue brings together work that investigates the implications of modeling, with computational support, naturally occurring arguments as formulated in the course of everyday work. This article contributes with respect to several issues raised by this challenge:

- A domain application of natural argumentation modeling, namely, to scholarly electronic publishing and discourse.
- This is accomplished through Web-mediated computer supported collaborative argumentation, for modeling the specific types of argumentation found in research literature.

^{*}Author to whom all correspondence should be addressed: e-mail: sbs@acm.org.

[†]e-mail: v.s.uren@open.ac.uk.

[‡]e-mail: g.li@ecs.soton.ac.uk.

[§]e-mail: bertrand.sereno@gmail.com.

[¶]e-mail: c.mancini@open.ac.uk.

INTERNATIONAL JOURNAL OF INTELLIGENT SYSTEMS, VOL. 22, 17–47 (2007)
© 2007 Wiley Periodicals, Inc. Published online in Wiley InterScience
(www.interscience.wiley.com). • DOI 10.1002/int.20188



- Tools are provided for interacting with structures of argument, including visualization tools and interfaces supporting structured dialogue.

We start by characterizing 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 to replace neither 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 coevolution of notational form with print publishing technology, but are not designed to take advantage of today’s information infrastructure. Whereas 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, Human–Computer Interaction, Computer-Supported Collaborative Work, and Computational Linguistics. Section 3 specifies the particular requirements for an argument modeling scheme that will be usable by researchers untrained in conceptual modeling or argumentation theory, and Section 4 describes the modeling scheme, before Sections 5 and 6 describe a series of user interfaces designed to support the variety of user tasks in the modeling environment. This article extends our previous work by contextualizing our approach to the specific field of naturalistic argument modeling, consolidating previously presented but unpublished material,^{1,2} updating the description of the user interfaces from earlier papers,^{3,4} expanding the theoretical rationale behind the representational scheme,⁵ and illustrating new literature modeling case studies and computational services from those already reported.^{2,6}

2. LIMITS OF DIGITAL LIBRARIES IN SUPPORTING SCHOLARLY INQUIRY

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 analyzing a literature. Researchers are concerned with the significance of a contribution to the literature, but no digital library can answer the obvious—but complex—questions that are fundamental to critical inquiry and that we seek to instill in our students:

- Which publications support and challenge this document?
- What is the intellectual lineage of this idea?
- What data are 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 characterized (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 modeled, 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, analyze, and contest their *significance*. It is in this context that structured argumentation has a contribution to make to support individuals and research teams in constructing a picture of the key arguments in the literature from their particular standpoints. Let us now consider the detailed requirements for such a research tool.

3. REQUIREMENTS FOR AN ARGUMENTATION SCHEME TO MODEL NATURALISTIC SCHOLARLY DISCOURSE

“Ontologies” are the term used in knowledge modeling and agent research, and increasingly within the semantic web community, to describe a specification of concepts, attributes, and relationships.⁷ Typical ontology-based 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 organizational 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 as much 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,”⁸ but in contrast to most existing ontology applications, stakeholders need not agree at all on the structure of the field being modeled. All modeling 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 modeling 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 supporting and contesting ideas in various ways.

A representation scheme for the arguments in papers needs to achieve a fine balance between expressiveness and usability. Although our ontology could be designed to support automated reasoning and verification of argumentation structures of the sort offered in other computational argument modeling research systems, 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. We do not yet claim to have delivered a system with a large user base, such is the difficulty of negotiating this trade-off, but it is a primary constraint in the design iterations we present here.

Our work derives from the tradition of hypertextual argumentation (see Ref. 9 for a historical account of the field’s emergence). This places an emphasis on interactive semiformal representations, often with graphical renderings, which are processable by both human and software agents. We describe later how specific hypertext functionality supports argument modeling in a concept mapping tool.

3.1. Data Model

Our modeling scheme comprises nodes and links. Nodes may be atomic or composite at the end user’s discretion. *Atomic nodes*^a are expressed as short pieces of free text succinctly summarizing 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, for example, *PageRank*. A different, typed atomic node might summarize an empirical result: *<Data> Undergraduate chemistry exam performance is doubled after training on the ChemVR system*. These are now objects (loosely analogous to published web sites with URLs) that others can link to in their own work (but unlike the web) using a semantically typed link.

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

^aWe refer later to nodes as “concepts,” but in explaining the data model, have found that it is more helpful to refer to them in semantic hypertext language as nodes.

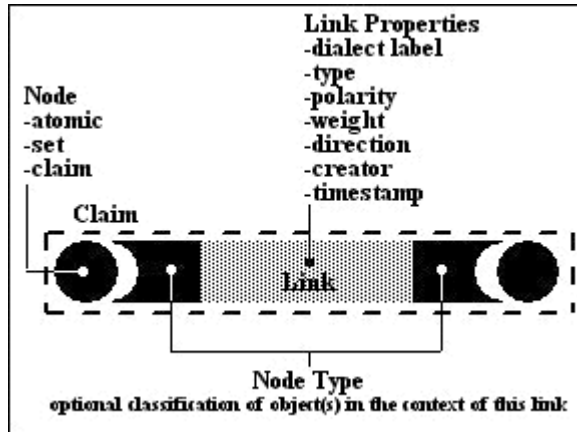


Figure 1. Structure of a Claim in the discourse ontology.

is made context dependent: Objects may play different *roles* in different contexts, because researchers may disagree on the node's type: for example, 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, applies, or is enabled by 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 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 modeling efforts could focus on fields where there is consensus (or where only consistent views are modeled), and build richer, more constrained representations that can support correspondingly more advanced reasoning.

3.2. 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*). The structure of the current discourse scheme is shown schematically in Figure 2. 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).

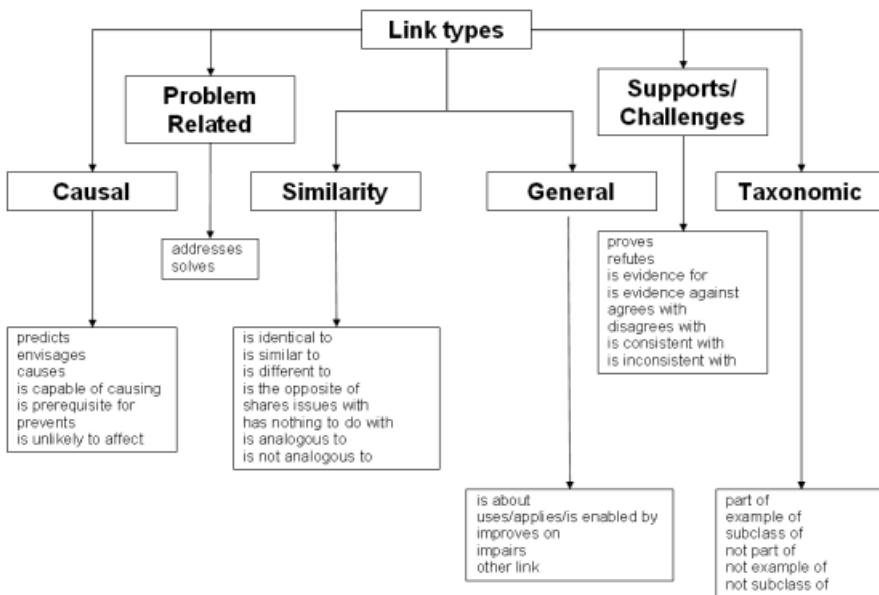


Figure 2. Class structure of the scholarly discourse ontology.

Defining relations in terms of class and dialect means the same classes can be employed by research communities who speak different “dialects” or even different languages: One can change the dialect labels of the relations without changing the underlying relational classes. Looking at the *Supports/Challenges* class, *refutes* is clearly a stronger term than *is inconsistent with*. Authors would be careful in their usage, particularly of stronger relations, but clearly they also both express the notion of a negative relationship between two nodes. We therefore add the explicit notions of polarity and weight (Table I), which are predefined and used by

Table I. The discourse ontology with polarity and weightings.

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
Causal	is consistent with	+/1
	is inconsistent with	-/1
	predicts	+/1
	envisages	+/1
	causes	+/2
	is capable of causing	+/1
	is prerequisite for	+/1
prevents	-/2	
Similarity	is unlikely to affect	-/1
	is identical to	+/2
	is similar to	+/1
	is different from	-/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 system, but end-users are not asked to provide these. Combining *classes*, *dialects*, *polarity*, and *weight* means we can reason at a higher level of granularity than individual relations, in delivering services by working with relations that share combinations of these properties (Section 5).

Elsewhere we have described the iteration from the first to the current version of the ontology.¹ The relational classes were originally derived from a data-driven approach of modeling naturalistic argumentation as we found it in a range of research domains, including computer-supported collaborative work, text categorization, literary criticism, genetics, philosophy of computing, applied ethics of technology, and film theory. Relations common to several domains were identified that we could classify in the classes shown above: Supports/Challenges, Problem Related, Taxonomic, Causality, Similarity, and General. Interwoven with this bottom-up approach was a theoretical strand of work, which we found enabled us to critique and validate the classes we had derived. Cognitive Coherence Relations theory (described next) provides a grounding for the relational classes and conceives relations in pairs of opposites, such as *proves* and *refutes*, where one has positive and the other negative implications.

3.3. Theoretical Basis of the Discourse Ontology Relations

The discourse ontology 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 that emphasizes the interpretive act of modeling.^{10,11}

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 optimized to express.¹²

Comprehensive sets of parameters have been proposed,^{13,14} defining a space of relational primitives by which two discourse units can be related. The basic relations are *additiveness*, *temporality* (sequentiality), and *causality*. Each of these is then parameterized: 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 bidirectional (e.g., list). Table II summarizes this scheme.

Grounding discourse relationships in a cognitive theory of coherence affords a number of interesting properties for building a system designed to support naturalistic argumentation. First, we have used the CCR typology as a tool to verify that the main relationships are represented in the taxonomy.¹⁰ Second, it grounds the discourse ontology in a set of relations which Sanders et al.'s¹³ experimental

Table II. Parametrical description of the main relations accounted for in CCR theory.^{13,14}

Source of coherence	Type	Polarity	Directionality	Hypotheticality	Comparativeness	Relations	Connectives
semantic	causal	positive	backward	actual		consequence-cause	B, since A
			forward	hypothetical		consequence-condition	B, if A
		negative	backward	actual		cause-consequence	A, hence B
			forward	hypothetical		condition-consequence	if A, then B
	temporal	positive	backward	actual		consequence-contrastive cause	C, despite A
			forward	hypothetical		consequence-contrastive condition	C, even if A
		negative	backward	actual		contrastive cause-consequence	despite A, C
			forward	hypothetical		contrastive condition-consequence	even if A, C
			backward	actual		backward sequence	A before B
			forward	hypothetical		sequence	B after A
	additive	positive	backward	actual		simultaneity	A while B
			forward	hypothetical		backward negative sequence	A until B
		negative	backward	actual		negative sequence	until A, B
			forward	hypothetical		conjunction, list	A and B
pragmatic	causal	positive	backward	actual		similarity	A like B
			forward	hypothetical		opposition	A but B
		negative	backward	actual		alternative	A or B
			forward	hypothetical		contrast, exception	A unlike B
	additive	positive	backward	actual		claim-argument	B, because A
		negative	backward	actual		argument-claim	because A, B
	additive	positive	backward	actual		contrastive claim-argument	B, although A
		negative	backward	actual		contrastive argument-claim	although A, B
		positive	forward	actual		enumeration, elaboration	A, moreover B
		negative	forward	actual		concession	A, however B

evidence substantiates as having psychological reality. In principle, this gives the taxonomy stability and applicability across different disciplines, media, and discourse types, and empirically, we have indeed modeled a wide variety of domains (see previous section).

Third, CCR makes it possible for the discourse relations we use to be resolvable back to a small number of relational primitives and their parameters and sheds light on the relationships between them. Representationally this is elegant (while also validating CCR's generalizability). For instance, the *General* relation *is-about* can be reexpressed as the CCR relation *elaboration* (whose parametrical values are positive/pragmatic/additive). *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, and 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. The current version of our server delivers a variety of services (see Section 5), but does not yet have a CCR-representational layer implemented; CCR has served more as a theoretical reference point and analytical validation tool in the system's development. Once such a layer was implemented, if the user were to search for all the discourse units that *agree* with node X, the system would know that all the discourse units that are *about* node X may also be of interest. At present, these relationships can be "hard-coded" in, but not inferred from CCR constraints.

To summarize, 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). We propose that these kinds of connections are expressed at a level that most researchers would not only recognize, but indeed, would naturally use when summarizing part of a literature.

4. 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 modeling of documents in a literature and provides a variety of services for browsing and analyzing the emergent conceptual graphs. Infrastructure details are given in Ref. 6 and are not of primary concern here. The focus is on the demands placed on user interfaces intended to support naturalistic argument modeling by nonexperts and ways to pursue the technology deployment strategies listed above.

As the ClaiMaker prototype has evolved and we have learned more about the problems users encounter with modeling, we have prototyped different interfaces for constructing models:

- a Web-based forms interface
- a plug-in for authors to produce concepts while writing in a word processor

- an argument map sketching interface to edit and manipulate claim structures
- a text annotation interface to view and edit candidate concepts identified in a research article.

4.1. Form Filling Interface for Claim Construction

The first version of ClaiMaker used forms with basic features such as keyboard input, text search, and dropdown lists. Its aim was to allow the project team to start inputting data as quickly as possible in order to populate a test collection that could be used for designing services. It took a stepwise approach to creating claim networks: first the user had to nominate the article they were modeling, then one form allowed her to create Concepts, another could be used to assemble Sets by searching for and selecting groups of Concepts, a series of other forms allowed claims to be made by selecting preexisting Concepts and joining them (see Figure 3). Capture was broken down into subprocesses, which meant that the user needed to understand the process as a whole in order to decide which step to take next and also had to know where in the menu system the appropriate form was located.

Although the “power users” on the project team did become reasonably fluent with the interface, even they had difficulty holding a gestalt view of the model in their heads as they went through the dissociated steps of building Concepts and

Figure 3. The ClaiMaker forms interface for creating a claim. The bottom bar gives details of the paper the reader is modeling. The user has already selected the concept to be linked from and given it the optional type “Evidence.” She is currently selecting a link from the drop down list of options. The next step will be to select the search button to look for the third component of the Claim triple.

then assembling them into Claims. It was clear that some radical changes were needed to make capture interfaces better support the cognitive processes involved in modeling.

4.2. Microsoft Word Plug-in for Claim Construction

One approach to tool deployment is to integrate any new tasks (in this case, argument construction and submission) with existing tools. When one needs to model the arguments in one's own, new papers, we hypothesize that claim construction might be best done as one is thinking about the conceptual structure of one's paper, that is, during writing, to minimize the delay between the expression of the idea in conventional prose and its formalization. As a first step, we have implemented a Microsoft Word plug-in (see Figure 4), which authors can launch directly from the Word toolbar.

The toolbar button opens a "semantic annotation" form for authors to enter the major types of Concepts in a paper as they write it. These can be classified in response to some prompts: Problem? Contributions? Uses/Applies? Improves on? Contrasts/Critiques? These prompts foreground the most important relational links in the ontology for summarizing an article's contribution, in other words, "promoting" them from the longer menu of relational types available in the more complex ClaiMaker forms interface (Figure 3) and turning them into questions. Once the concepts have been saved (as an XML file), the idea is that the Concepts will then be imported into ClaiMaker and used as a basis for further Claim building.

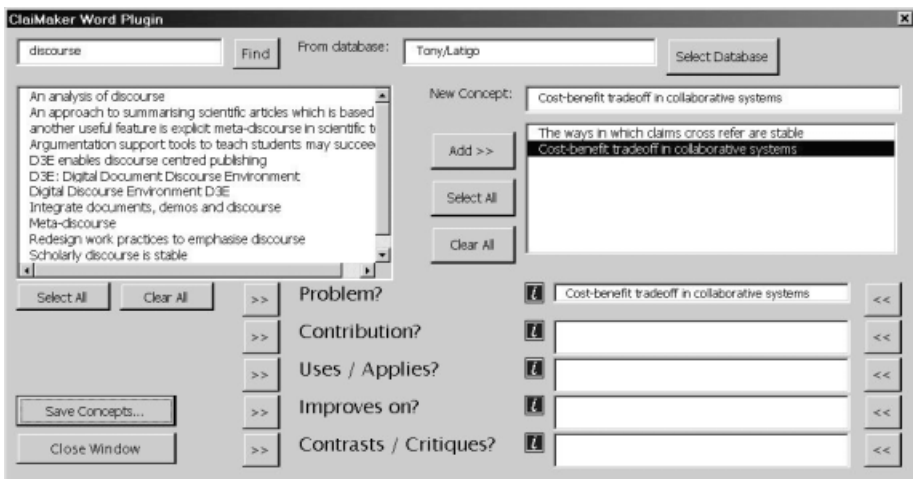


Figure 4. ClaiMaker Word plug-in. Existing Concepts on the web server can be searched and displayed in the panel top left. New concepts are displayed on the right and can be assigned types using the five prompts in the lower part of the screen.

4.3. ClaiMapper: Sketching Claim Structures

To overcome the problems of holding complex models in memory, the team found themselves resorting to pen and paper for sketching drafts of argument maps. Figure 5 shows the typical kind of sketch produced as one works out the structure of the literature, prior to entry in ClaiMaker.

The use of pen and paper with a software tool is a telling indicator that it is providing inadequate cognitive support for users, and it is well established that sketching is a fundamental activity in many forms of creative and conceptual representation.¹⁵ The sketching was mainly driven by a desire to consolidate one's own interpretation before committing it to the knowledge base. In the terms of Green's¹⁶ Cognitive Dimensions framework, the form-filling interfaces led to "premature commitment" by requiring users to commit to a structure before they have been able to validate it more broadly in the context of the overall structure. Consequently, a concept mapping tool has been developed, called ClaiMapper (Figure 6).

ClaiMapper is a stand-alone tool, based on the Compendium¹⁷ visual hypertext system.¹⁸ Instead of filling in a new form for each bipartite connection, the user can simply draw links between nodes, specifying the link type when prompted. Of particular use is the hypertext facility in which copying and pasting a node across the maps for multiple documents (whether a Concept, Set, or Map) does not

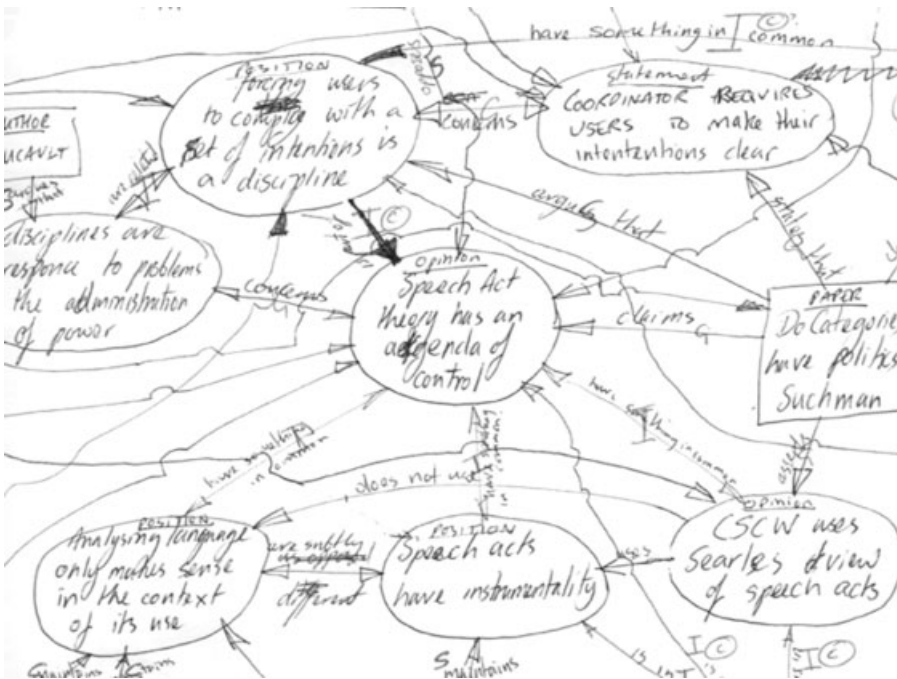


Figure 5. The typical kind of sketch produced as one works out the structure of the literature.

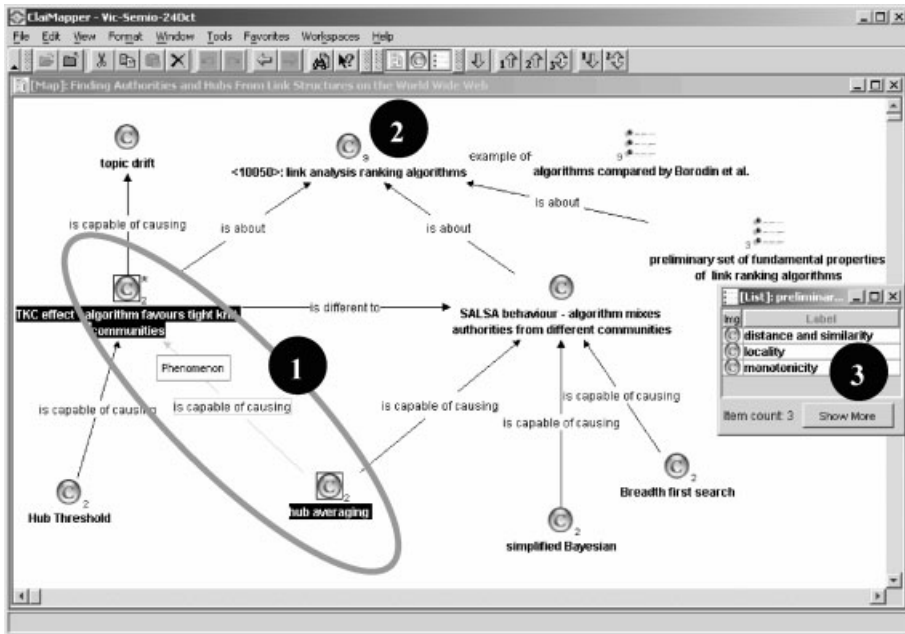


Figure 6. Sketching ClaiMaker-compatible models using the ClaiMapper tool. (1) In the circled Claim, the node *TKC effect* has the type (i.e., plays the role of) *Phenomenon*. (2) The Concept *link analysis ranking algorithms* is shown as being used in nine different contexts. (3) On the right is a Set named *preliminary set of fundamental properties of link ranking algorithms*, which, when opened, lists three concepts that the analyst has found.

literally clone it in the ClaiMapper’s local database, but simply creates a new pointer to the node: the interface updates the node’s display to indicate how many argument models the node is used in, whose names the user can display and jump to (e.g., the Concept *link analysis ranking algorithms* in Figure 6 is shown as having *nine* occurrences). Users can search the ClaiMaker server for existing concepts matching a selected node in a map, and can import or simply drag and drop search “hits” directly into ClaiMapper, creating nodes with full database metadata, ready to be reused through connection to new structures.

Based on our experiences to date, ClaiMapper has proven to be a significant advance in supporting the cognitive demands of modeling, seeing the “bigger picture,” and more rapidly creating claim structures, and the tool can, of course, be used for analysis and note taking without ever uploading the model to the server. However, there are still usability problems. The ClaiMaker server accepts XML exported from ClaiMapper and checks for duplicate node labels and illegal structures. Particularly for new users, the checking routines threw up significant numbers of errors and warnings. Warnings included duplicate Concepts in the database, which it wished the user to confirm were identical or to change. Some errors were simple, for example, misspelling of a link label. Others were structural, for instance, putting the whole of a model for a paper as an element of a claim.

It seems that, although we do not want the sketching interface to enforce premature structure, we do want it to give positive assistance to the user to build models that are valid and can be painlessly imported into the ClaiMaker database. The forms interface enforced legal structures because the range of operations on each form was limited to legal actions, and invalid inputs such as incomplete Claims were discarded. As we continue with our development of ClaiMapper, we need to tackle the question of how it can more actively communicate to a user what a syntactically “good” model ought to look like. One possibility, described elsewhere,³ is to provide readers with claim-making templates for stereotypical “genres” of papers in a field.

We have begun to investigate active support of a *semantic* nature in the context of modeling articles in the literature, described next.

4.4. ClaimSpotter: Document Analysis and Annotation for Claim Formalization and Reuse

The ClaimSpotter interface tackles the “chunking” problem identified by Buckingham Shum¹⁹ in a cognitive analysis of the use of graphical argumentation schemes. In essence, the user is faced with deciding what should be made into a Concept/Set/Claim for linking: what granularity should there be, how succinct or verbose should the label and detail be, and how should it be categorized (if at all)? In the context of modeling claims in a paper, this question clearly depends on the reader’s interpretation of the paper, and, therefore, the use of the original text as the basis for semiautomatic assistance in formulating claims is not straightforward.

ClaimSpotter is our first step toward an active user interface with concept suggestion and identification of potentially relevant areas in the source text. There are three elements:

- identification of the areas where the author presents and defends her argument, combined with approaches to break up the text into potential concepts
- provision of additional services to promote collaboration and reuse within a group of readers/annotators
- provision of an interface to support the capture/editing/construction of claims based on the candidate concepts that the tool has extracted.

4.4.1. Enhancing a Document

The first step of our approach is to identify areas where authors present and defend their argument. Because authors have to defend their position and their contributions and relate them (through support or criticism) to the positions of their peers (an account of this strategy can be found in the Create A Research Space Model²⁰), we believe that the ability to guess the role played by a sentence in this defense, using text analysis methods, provides a valuable resource in the task of interpretation, which can be seen as the task of positioning oneself with respect to the author’s assertions.

We have started to tackle this problem by using text patterns that can be consistently associated with certain kinds of assertion to identify and categorize

statements that signal stages of the argument. For example, our discourse ontology has natural language labels, which can be changed to fit the dialect of the domain, so the simplest approach is to identify locations where the labels appear or synonyms as defined in a user-editable thesaurus. This gives us an indication of where (and how) the author defends her argument. Another category of interest is statements about contributions made by the authors. These are identified using references to the document itself (e.g., “Section 2 describes . . .”) and references to the authors (e.g., “We have proposed . . .”). Once patterns such as these are combined with approaches to identify potential components of Concepts, such as noun-group identification, the system can propose a number of elements ready to use as a part of a Claim, while still leaving the reader free to edit them.

Relying on such a limited number of text patterns, although useful, does not account for the richness of expression one can use in defending one’s position. In a CARS-derived approach,²¹ the role played by a sentence (e.g., introducing the authors’ work, providing background information, or supporting a cited work) is guessed from a number of annotated examples described in terms of a much more exhaustive range of features including (among many others) sentence content and position in the document. We have reimplemented a simpler version of that approach; details of the different document filters can be found in Ref. 22.

To complement this approach, one could look at further means to enrich a document, for instance, the inclusion of hyperlinks between topically coherent passages²³ or between a term and its definition.²⁴ Figure 7 shows how candidate relations and some specific areas of a research paper are highlighted in the ClaimSpotter interface. Sereno et al.²⁵ report an empirical study into how researchers annotate a

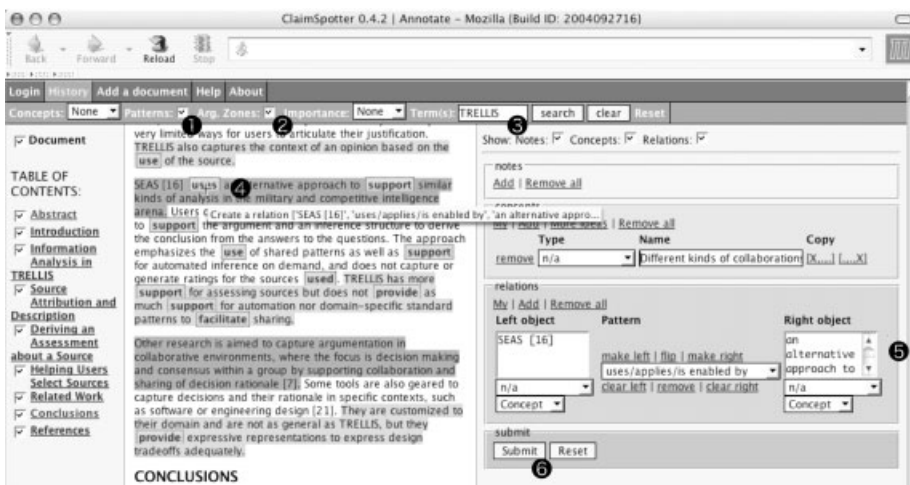


Figure 7. The ClaimSpotter interface attempts to reduce the “chunking” problem by helping the user focus on subsets of the original text. In this example, the user has combined the candidate relations (1), the rhetorically coherent areas (2), and a user-defined filter to help focus on subsets of the original text that are deemed interesting. Candidate relations found (4) can be clicked on and split into claim triples (5) and submitted immediately to the database (6) if desired.

research paper informed the design of ClaimSpotter, and a formative usability evaluation study of the interface.

4.4.2. Promoting Collaboration and Reuse

The second element of our approach aims at incorporating and making use of the Claims encoded by fellow readers and the Concepts they connect. Displaying the position defended by fellow annotators as a set of Claims indicates what has been said already about the document, including readings that are different in emphasis or focus from the author’s primary narrative and argument. Figure 8 shows the usage of a Concept over the corpus of documents. The Claims in which it has been used and the documents which it has annotated can be accessed from there. In this way, documents become connected through common Concepts, even if they do not directly reference each other. This provides a form of extended “semantic cocitation” that exploits the web of structured annotations and extends the citations of a document.

Finally, our investigations into user interfaces for authoring conceptual representations have led us to investigate the emerging phenomenon of web logging (or “blogging”). We are considering this as a network-centric paradigm for

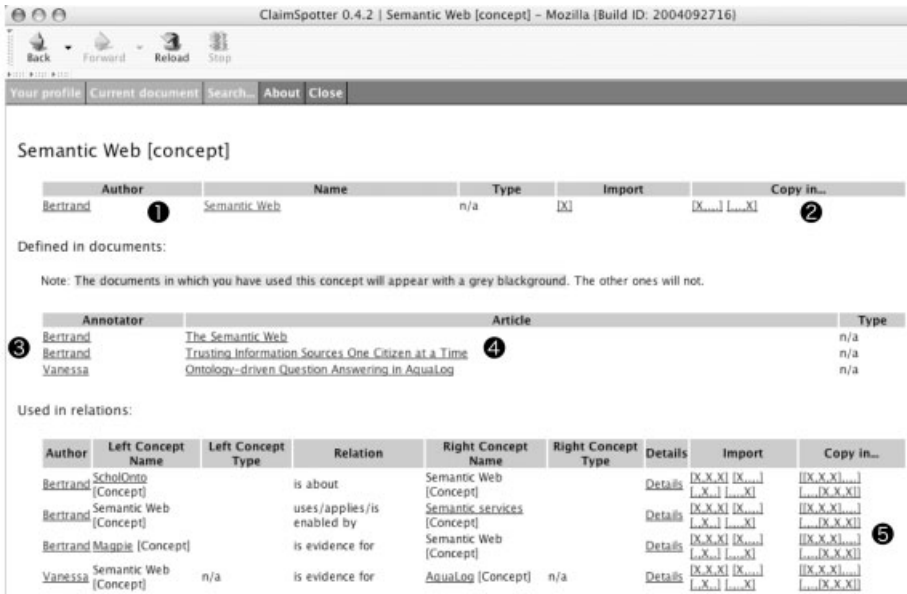


Figure 8. The user can access a History window for a Concept which displays, for instance, the author (1) and the different uses of that Concept over the corpus of documents (3) (4). It can be copied in the current document with a single click, if a user decides so (2). In a similar way, the different relations in which it is used can be imported in the document being currently annotated or copied in order to discuss them (5). Multiple links within the History window allow a quick navigation within the annotation repository.

publishing interlinked commentaries that has potential applications for research discourse. “Semantic Blogs”²⁶ add machine-processable semantics to undifferentiated webs of connections, which, in the context of our work, takes the form of establishing discourse relations between documents and blog entries (Figure 9).

4.5. Discussion: The Interplay of User Interface and Representation Design

The different interfaces that we have presented here were designed with the intention of facilitating claim construction, for example, by integrating it into other work activities such as reading (ClaimSpotter) and writing (the Word plug-in). In developing these interfaces, we have mainly tackled the usability side of the usability–formality balance, but we are seeing indications that improving the usability of input interfaces can affect the kinds of models that are built.

One important formal notion is “normalization,” that is, ensuring that there is only one entity in a model representing a particular concept. In the ScholOnto approach, we have never tried to enforce normalization. It is intended as a collaborative system with no “master view,” leaving open the possibility that if one user considers that his notion of, for instance, “ontology” is different from an existing one, there should be no restriction on him creating an identically named node:

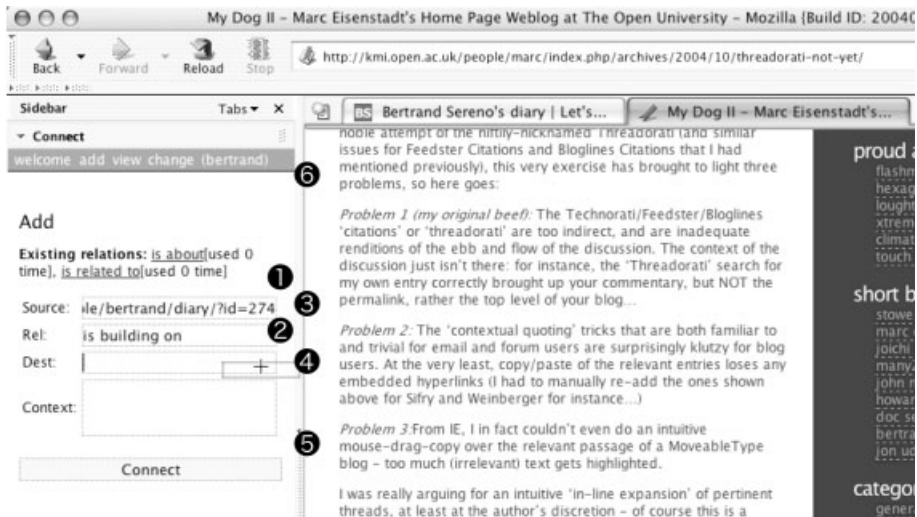


Figure 9. The “semantic blogging” interface enables the user to drag and drop links to create relational triples between blogs or documents.²⁷ Directed connections can be drawn between Web pages by selecting an existing relation (1) or creating a new one by typing it (2). The source and destination URLs can then be typed (3) or dragged and dropped (4) from any web page. An additional Context field allows one to add a note on the connection to give it more nuanced meaning (5). Additional functions are available to view and modify existing connections.

competition over the definition and ownership of terms is a natural part of research and not a practice that we could or should suppress. However, we do have some mechanisms for avoiding unintended duplication of nodes that would impair the usability of models. ClaimSpotter detects existing node labels and highlights them where they occur in the text of the document being analyzed. Users of ClaimSpotter appreciate this feature because it saves them the work of creating new nodes that they wish to reuse, presenting the information proactively in contrast to a time-consuming database search on potentially unknown keywords. The XML upload facility from ClaiMapper also checks for duplicates and allows users to substitute an existing node into their model. However, this comes after the point of creation and does not provide the low-cost insight into existing models that ClaimSpotter gives. An important advance on these features is to identify close-matches, work on which is under way.

Some of the interfaces guided users to make particular kinds of relations. This was explicit in the case of the Word plug-in, which offered a very limited palette of link types to the user, forcing them to concentrate on claims about problems, contributions, etc. (see Figure 4). A similar effect was observed with the ClaimSpotter interface but caused implicitly by the underlying approach to text chunking rather than explicitly by the interface itself. A key part of the development of ClaimSpotter focused on highlighting chunks of the text where an author asserts or defends her position. Pointing users at these places appeared to encourage them to make more claims that use “addresses” links.

One user who tried both ClaiMapper and ClaimSpotter observed that the latter led her to focus on concepts whereas the former encouraged the building of webs of relations. This seems a reasonable remark on the basic affordances of the two systems. ClaimSpotter takes a text and highlights interesting chunks. Chunks of text look like concepts, leading users to naturally think of that facet of the process. In contrast, ClaiMapper provides a canvas for users to lay out concepts, offering tools to organize and link them. The primacy given to a physical representation of the network may encourage users to craft an interconnected model. Other users of ClaimSpotter commented that they would have liked to have had such a visualization of the claims they were building (which has led to subsequent work to generate graph structures from the individual claims they construct using the traditional web form interface).

These early observations lead us to believe that the design of interfaces for creating claim networks, and possibly argument models in general, may influence the kinds and quality of models produced. Comparative studies are needed to analyze whether different interfaces bias users to produce different styles of model. However, we have found that users benefit from automatic support of the modeling process, particularly through text analysis. When support encourages good practice, for example, by highlighting existing concepts and making it easy to reuse them, users welcome this and take advantage of the functionality. This indicates that if other kinds of support were incorporated into an interface, such as suggesting appropriate link types for concepts of a particular type, users would find this valuable “scaffolding” as they sought to build rigorous, elegant models. Although the formality/usability balance may be hard to define, good interfaces that provide

users with support for key aspects of a representation can allow untrained users to push the balance toward formality.

We move now from user interfaces for *constructing* conceptual networks of ideas to tools that enable useful navigation around and interrogation of those structures.

5. NAVIGATING AND ANALYZING LARGE ARGUMENT STRUCTURES

In previous papers we have detailed a variety of mechanisms for delivering computational services over the conceptual graph of claims that is built as researchers submit their annotation models to the ClaiMaker server.^{1–3,5,6} We summarize these here to convey the end user’s interactional experience, as enabled by the underlying discourse ontology and refer the reader to the above papers for implementation details.

“Discovery Services” that users can access fall into two broad classes:

- graph theoretic analysis of claims networks by exploring the topography of networks
—example: *Cluster Analysis* identifying dense networks of concepts suggesting a coherent topic
- semantic analysis of claims networks that exploits the relational types
—examples: *Perspective Analysis*, which generates a report of supporting or challenging papers; *Lineage*, which traces the work on which the current paper directly builds, and its converse, *Descendants* (i.e., measures of semantic impact, including but going beyond citations).

Although users can still access services via the original ClaiMapper interface (recall that this was the first generation interface, primarily for the research team), we are moving toward a more accessible search tool called ClaimFinder, which delivers the services as tabs on a web page, rather than as items embedded in a drop-down menu in ClaiMaker. The default page provides a simple, single-field form for users to do keyword searching, with “advanced” search tabs delivering encapsulated services such as *Perspective Analysis* and *Lineage* (Figure 10).

On invoking one of the above ClaimFinder services, instead of returning a list of results, the tool generates interactive visualizations (currently in two possible formats) of the argumentative claim structures in which the relevant Concepts/Sets/Claims are embedded (Figures 11–13). These can be browsed by selecting a node to see the underlying detail, the source document it originates from, or to reveal/hide structure by zooming, rotating, or filtering the number of links from the selected node.

The visualization tool illustrated in Figure 12 is delivered via a Java applet when generated in response to a query, but it is also available as a self-contained Java application. The advantage of this is that if the user wants to save an argument map layout for future use, the application version can be used to open and display it, off-line if necessary. We anticipate that this will be particularly useful



(a)



(b)



(c)

Figure 10. The ClaimFinder search interface. ClaimFinder delivers the original ClaiMaker's Discovery Services via the tabbed search interface style with which Web users are most familiar. The default (a) is a simple, single-field form, whereas the others (b,c) lead the user into more advanced services, with the fourth, most complex, being the full ClaiMaker tool for adding new claims. Default search terms on each page invite the user to test the tool to see examples of visualized claim structures.

when crafting map layouts as instructional aids or “portal maps” for students and research peers.^b Let us now consider two examples to illustrate how combinations of relational class, dialect, weight, and polarity can be used in answering queries that cannot be articulated in our current digital libraries.

^bOf the interfaces presented, the ClaiMaker forms-based system is available for interested parties to test, with a variety of analysis services available to interrogate claim structures. The ClaiMapper sketching tool is available as a stand-alone application on request and is currently being integrated more tightly with ClaiMaker. ClaimSpotter and the Word plug-in are running prototypes, although in a preliminary state. Screen recordings with commentary illustrate the tools' interactivity more effectively than static screens and text: <http://claimmaker.open.ac.uk>

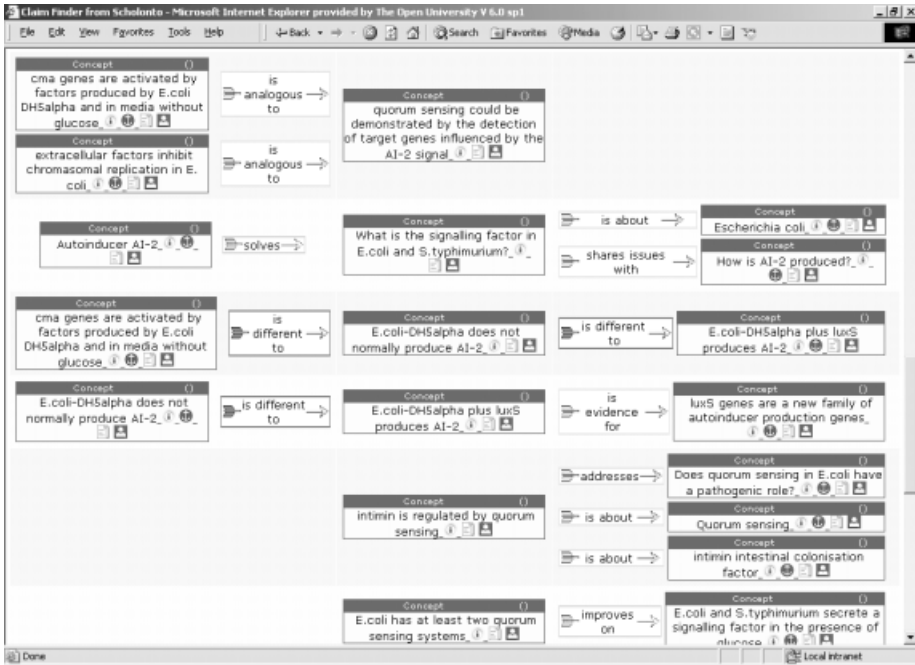


Figure 11. ClaimFinder generates interactive visualizations of argument structures in response to queries. In this rendering, a three-column tabular layout shows each Concept/Set in the search results, with incoming and outgoing links to Concepts/Sets in the left and right columns. This example is taken from modeling part of the test data set released from the *Proceedings of the National Academy of Sciences*, as part of a domain visualization symposium.²⁸

5.1. 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 catalog 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 realize 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.^c More substantively, *this paper* refers to the *claims* that researchers have made about the document,

^cIf 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.

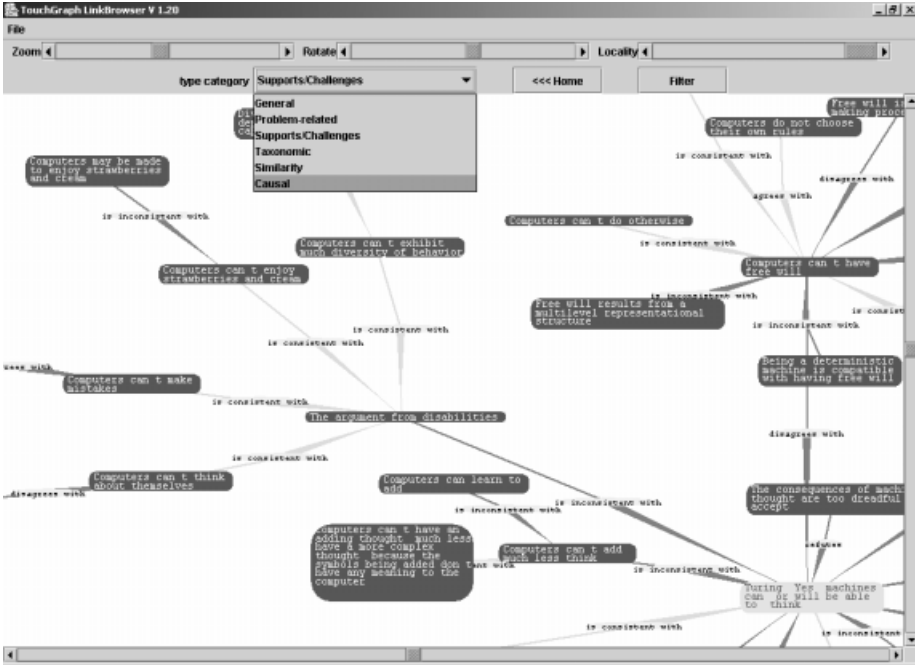



Figure 12. This shows a different format, making use of the TouchGraph Linkbrowser Java classes.²⁹ This uses a self-organizing graph algorithm to generate an interactive two-dimensional map, with controls to zoom, rotate, and limit how much is displayed at once from the currently selected node. We have added additional controls to semantically filter the graph around a target node, using the argumentation ontology’s relational classes. This example is taken from modeling the Philosophy of AI literature in the Turing debate, converting the large paper argument maps published by Horn³⁰ into interactive Web versions.

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 13.

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 neighborhood.

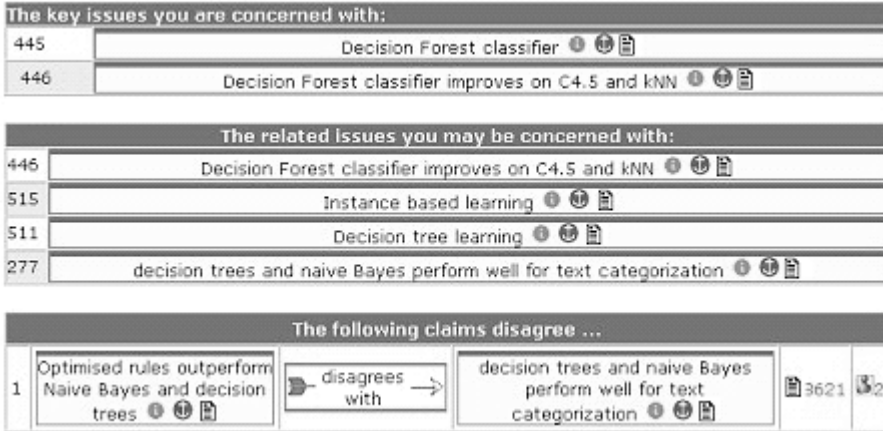


Figure 13. Arguments that contrast with the nodes in a research paper. Key: clicking ⓘ displays node metadata; ⌕ sets the node as the focus, to show incoming and outgoing relations; 📄 links to the document metadata/URL. 👤 links to information about the node’s creator.

5.2. 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 evolved from it. Because 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 conceptualized 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.^d

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

^dA semantic web standard based on graphs is the *Resource Description Framework*, available at: www.w3.org/RDF. In the analysis presented here we use the *Ivanhoe* path-matching tool available in the *Wilbur* RDF toolkit, available at: wilbur-rdf.sourceforge.net.

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, whereas *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 bidirectional (if A is like B, then B is like A).

Summarizing, 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 14 shows examples of acceptable paths that could be returned by this lineage analysis.

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 14 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 I), because *addressing* or *solving* a known problem usually represents progress of some sort. One could also include Taxonomic links, because 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 15 shows a visualization of the structure extracted from the claims network in response to a lineage query about a node.




Figure 14. Examples of paths that could be returned by a lineage analysis on a target node (see text for the specification of the query).

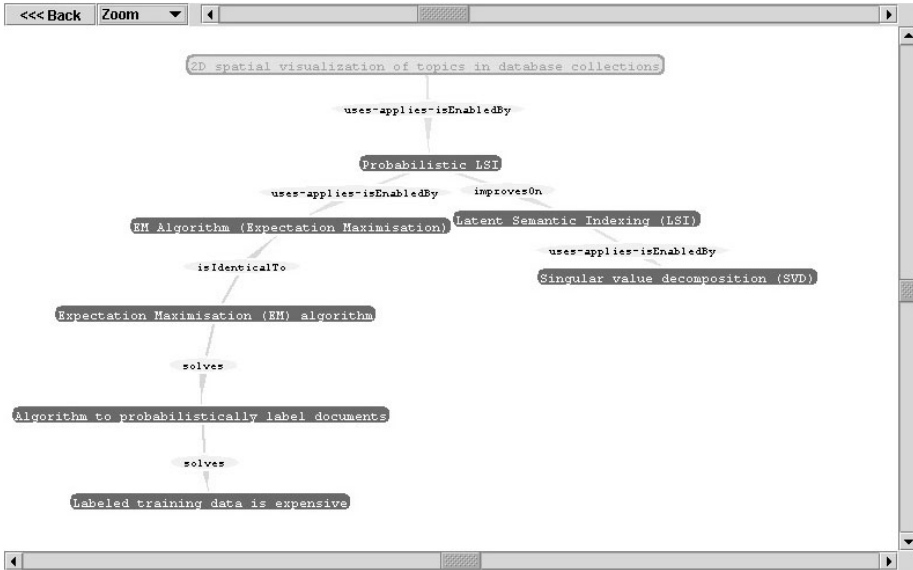


Figure 15. The *Lineage* service was conceived as a way to trace the “intellectual roots” of a concept, displayed at the top. The network is analyzed and filtered to show potentially significant relational types such as *uses/applies/is enabled by*, *improves on*, and *solves*. The *Descendants* service traverses the graph in the opposite direction to show impact of a concept.

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 summarize, 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 queries.

6. RELATED WORK

Research related to this work can be broadly grouped into the following categories:

- research into modeling natural argumentation
- research into Web-based annotation
- research into concept mapping
- research into modeling scientific discovery.

First, the research community represented by the series of workshops on Computational Models of Natural Argument (CMNA), and this special issue, is an obvious source of comparative approaches. Here we find theoretical analyses of naturally occurring argumentation and systems that support argument modeling and reasoning in applied fields with well-defined rules such as law (e.g., Ref. 31). The emphasis in this field to date has been on the scope for computational reasoning even in the face of the informality found in natural argumentation, and we are now considering how the lessons learned from this artificial intelligence research strand can be integrated with our own infrastructure, to add computational services when patterns can be detected in the claims networks. However, our philosophy of imposing minimal constraints on the degree to which analysts structure their work places our system at the informal end of the spectrum compared to other CMNA research. As a counterbalance, however, we note with interest that strong critics of formalization in interactive systems maintain that our approach is still too formal.³² Our efforts to negotiate the “formalization tightrope” will continue, with potential benefits to be accrued both through judiciously adding computational services while remaining acutely aware of the dangers of overstructuring interaction.

The approach presented here shares some of the aims of annotation technologies. Ovsianikov et al.³³ analyze common practices of traditional hand-written annotation and identify its primary uses as “to remember, to think, to clarify and to share.” They observe that the first three are predominant for traditional annotation, which, with the exception of reviewing, is a largely private affair, but that sharing becomes more important for software annotation systems that facilitate collaborative annotation. However the decisive benefit of annotation technology over traditional annotation is searchability. This reinforces our view that developing the search interface and services of the ClaiMaker system is central to encouraging and supporting knowledge capture.

The aspects of sharing and searching are prominent in collaborative Semantic Web annotation technologies, such as Annotea, being developed by the W3C.³⁴ The Semantic Web approach to annotation regards it as searchable metadata stored on web servers with Xpointers to original documents, but provides no semantic for relating annotations to documents or to each other. Our work can be framed as providing a relational semantics that make it possible for large numbers of annotations to remain manageable.

The TRELLIS system is a rare example of a system that adds a semantic element to annotation by linking statements drawn from web documents using a set of discourse, logical, and temporal connectives.³⁵ TRELLIS is designed to assist analysis of multiple documents, but does not consider multiple users collaborating and does not use the semantic relations to enable computational services to support the analysis of the data.

Concept mapping tools for teaching sense making and argument construction are well established. Our ClaiMapper tool (Figure 8) and the conceptual visualizations (Figures 10–12) draw inspiration in part from this research tradition, reviewed in the context of argument mapping by Buckingham Shum.⁹

Finally, Thagard’s work³⁶ on modeling scientific revolutions complements our work. Using a knowledge representation scheme focused on the conceptual

structures behind competing theories, he adds parameters to provide a quantitative indication of the “explanatory coherence” of a given theory, given the available evidence and competing theories. Thagard’s work contrasts with ours in its dependence on an expert modeler codifying theories at a finer granularity and with greater care than we can assume with our envisaged end users. The target of his modeling is complementary in the sense that our discourse ontology is designed to support the collaborative construction of claims—a form of computer-supported collaborative work—in contrast to the modeling of a well-understood debate, in which it is clear whether, for instance, a hypothesis has been refuted. ClaiMaker enables peers to contest this claim, rather than take it for granted. As with the more formal CMNA work, there is potential for integrating the two approaches.

7. RECENT AND FUTURE WORK

In this article, we have motivated the design of a Computer-Supported Collaborative Argumentation system for researchers to model, publish, and analyze “claims,” as a possible paradigm for scholarly publishing that exploits the properties of conceptual networks and the internet. We have drawn particular attention to the requirements implied by a system in which we cannot assume that end users have any training in the underlying semiformal argument modeling ontology, detailing our consequent approach to ontology and interaction design. This has implications for the computational reasoning that the model can support compared to other systems, but we argue that formalization comes at a high price for many users. The cost–benefit trade-off must deliver rapid enough benefits for the effort of modeling arguments. We are at a relatively early stage in the development of this infrastructure, and we cannot yet claim widespread adoption. However, we contend that the environment in its current state shows potential as a cognitive tool exemplifying how we may “read and write” ideas in a network-centric paradigm.

The most recent work has been to complete formal user testing. First, Sereno et al.²⁵ report on an evaluation study of ClaimSpotter. Second, an evaluation study has been conducted in which the same literature was reviewed using ClaiMapper and ClaiMaker. The resulting claims network was then studied by other researchers, using either ClaiMaker and ClaimFinder or reading a traditional literature review article.³⁷ One future strand of work concerns user interfaces, as we develop our semantic weblog environment to explore the properties of this as a user-friendly medium for constructing networks of commentary. Another future strand concerns more powerful reasoning to enhance usability either by imposing constraints on users that they find productive or by making helpful suggestions about argument structures. One approach is to embed CCR more deeply in the system in order to investigate the kinds of reasoning that it enables and another is to explore the possibility of integrating finer-grained approaches to argumentation modeling as being developed by other CMNA researchers.

Finally, scholarship and research is clearly not the only domain in which it is important to capture contrasting interpretations, and we are interested to investigate the potential of this work to support analysts in other domains of collective knowledge management and sensemaking.

8. CONCLUSION

If in the late 1980s, a visionary had painted a scenario of the explosive adoption of a global standard for information publishing and rendering that would overtake all known internet standards, s/he would have been treated with some skepticism, to say the least. The idea of “normal people” doing structured mark-up of their work in their own personal time would have sounded dubious. However, the World Wide Web successfully negotiated the cost/benefit trade-off for many people, who discovered the power of simple hypertext. We do not of course claim to be inventing the next Web, but have sought in our work to learn from its lessons and build on the shift in “network literacy” that is taking place. The internet and the Web were always envisaged as powerful tools for researchers, and although communications and distributed computation are revolutionizing some aspects, the way in which new knowledge is published and contested has remained almost untouched. We have painted a scenario of scholarly publishing and debate in which the Web paradigm of publishing resources to which others can link is taken the next step, with the specific needs of researchers in mind. Having developed a prototype environment to explore this space, we are now beginning to generate evaluation data as the tools become robust and usable. There is however much more to do.

Acknowledgments

We gratefully acknowledge the support of the EPSRC (Distributed Information Management Programme GR/N35885/01) in funding the Scholarly Ontologies Project and the programming by Michelle Bachler on ClaiMapper, Haibo Jia on TouchGraph, and Tony Brush on the ClaiMaker Word plug-in. We are also indebted to Robert Horn for providing digital copies of his Turing debate argumentation posters. Feedback from reviewers and participants at the CMNA 2004 ECAI workshop also helped to improve this article.

References

1. Buckingham Shum S, Uren V, Li G, Domingue J, Motta E, Mancini C. 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; 2002.
2. Uren V, Buckingham Shum S, Li G, Domingue J, Motta E. Scholarly publishing and argument in hyperspace. In: WWW 2003: 12th Int World Wide Web Conf, Budapest, Hungary. ACM Press; 2003. pp 244–250.
3. Buckingham Shum S, Uren V, Li G, Domingue J, Motta E. Visualising internetworked argumentation. In: Kirschner PA, Buckingham Shum S, Carr C, editors. Visualizing argumentation: Software tools for collaborative and educational sense-making. London: Springer-Verlag; 2003. pp 185–204.
4. Uren V, Buckingham Shum S, Sereno B, Li G. Interfaces for capturing interpretations of research literature. In: Distributed and Collaborative Knowledge Capture Workshop, ACM Knowledge Capture Conf, Florida, October 2003.
5. Uren V, Buckingham Shum S, Mancini C, Li G. Modelling naturalistic argumentation in research literatures. In: Fourth Workshop on Computational Models of Natural Argument, European Conf on Artificial Intelligence, Valencia, August 22–27, 2004.

6. Li G, Uren V, Motta E, Buckingham Shum S, Domingue J. ClaiMaker: Weaving a semantic web of research papers. In: First Int Semantic Web Conf (ISWC2002), Sardinia; 2002. pp 436–441.
7. Gruber TR. A translation approach to portable ontologies. *Knowl Acquis* 1993;5:199–220.
8. Boland RJJ, Tenkasi RV. Perspective making and perspective taking in communities of knowing. *Organ Sci* 1995;6:350–372.
9. Buckingham Shum S. The roots of computer supported argument visualization. In: Kirschner PA, Buckingham Shum S, Carr C, editors. *Visualizing argumentation: Software tools for collaborative and educational sense-making*. London: Springer-Verlag; 2003. pp 3–24.
10. Mancini C, Buckingham Shum S. Cognitive coherence relations and hypertext: From cinematic patterns to scholarly discourse. In: *Proc ACM Hypertext 2001*, August 14–18, Århus, Denmark. New York: ACM Press; 2001. pp 165–174.
11. Mancini C. Towards cinematic hypertext: A theoretical and empirical investigation. Technical Report KMI-04-6. Knowledge Media Institute, The Open University, Milton Keynes, UK; 2003.
12. Sanders TJM, Noordman LGM. The role of coherence relations and their linguistic markers in text processing. *Discourse Process* 2000;29:37–60.
13. Sanders TJM, Spooren WPM, Noordman LGM. Coherence relations in a cognitive theory of discourse representation. *Cogn Ling* 1993;4:93–133.
14. Louwse M. An analytic and cognitive parametrization of coherence relations. *Cogn Ling* 2001;12:291–315.
15. Goel V. *Sketches of thought*. Cambridge, MA: MIT Press; 1995.
16. Green TRG. Cognitive dimensions of notations. In: Sutcliffe A, Macaulay L, editors. *People and computers V*. Cambridge, UK: Cambridge University Press; 1989. pp 443–460.
17. Compendium: Visual hypertext concept mapping tool. Available at: www.CompendiumInstitute.org.
18. Selvin AM, Buckingham Shum S. Rapid knowledge construction: A case study in corporate contingency planning using collaborative hypermedia. *Knowl Process Manag* 2002;9:119–128.
19. Buckingham Shum S. Analyzing the usability of a design rationale notation. In: Moran TP, Carroll JM, editors. *Design rationale: Concepts, techniques, and use*. Hillsdale, NJ: Lawrence Erlbaum Associates; 1996. pp 185–215.
20. Swales JM. *Genre analysis: English in academic and research settings*. Cambridge: Cambridge University Press; 1990.
21. Teufel S, Moens M. Summarizing scientific articles: Experiments with relevance and rhetorical status. *Comput Ling* 2002;28:409–445.
22. Sereno B, Buckingham Shum S, Motta E. From scholarly documents to interpretative claims: An approach to bridge the formalisation gulf. In: *European Semantic Web Symp*, May 10–12, 2004, Heraklion, Greece; 2004. pp 357–371.
23. Hearst MA, Plaunt C. Subtopic structuring for full-length document access. In: *Proc ACM SIGIR'93 Conf*, Pittsburgh, PA; 1993. pp 59–68.
24. Blustein J. Automatically generated hypertext versions of scholarly articles and their evaluation. In: *Proc 11th ACM Hypertext Conf*, San Antonio, TX; 2000. pp 201–210.
25. Sereno B, Buckingham Shum S, Motta E. An environment to support sense-making with knowledge triples. In: *Proc IUI'05: ACM Conf on Intelligent User Interfaces*, San Diego, California, January 9–12, 2005. New York: ACM Press; 2005. pp 199–206.
26. *Semantic Blogging for Bibliography Management*. Hewlett-Packard Research Labs. Available at: <http://www.hpl.hp.com/semweb/biblio>.
27. *Semantic Blogging Project*, Knowledge Media Institute, Open University, UK. Available at: <http://kmi.open.ac.uk/projects/semanticblog>.
28. Shiffrin R, Börner K. Mapping knowledge domains. *Proc Nat Acad Sciences* 2004;101:5183–5185.
29. *TouchGraph home page*. Available at: www.touchgraph.com.

30. Horn RE. Infrastructure for navigating interdisciplinary debates: Critical decisions for representing argumentation. In: Kirschner PA, Buckingham Shum S, Carr C, editors. *Visualizing argumentation: Software tools for collaborative and educational sense-making*. London: Springer-Verlag; 2003.
31. Prakken H, Vreeswijk G. Encoding schemes for a discourse support system for legal argument. In: *CMNA2: 2nd Workshop on Computational Models of Natural Argument, ECAI 2002*, July 22, Lyon, France; 2002.
32. Marshall C, Shipman FM. Which semantic web? In: *Proc Hypertext 2003*, Nottingham, August 2003. New York: ACM Press; 2003. pp 57–66.
33. Ovsianikov IA, Arbib MA, Mcneill TH. Annotation technology. *Int J Hum Comput Stud* 1999;50:287–362.
34. Kahan J, Koivunen M, Prud'Hommeaux E, Swick R. Annotea: An open RDF infrastructure for shared Web annotations. In: *10th Int World Wide Web Conf (WWW2001)*, Hong Kong; 2001. pp 623–632.
35. Gil Y, Ratnakar V. TRELLIS: An interactive tool for capturing information analysis and decision making. In: *Proc EKAW 2002, Lecture Notes in Artificial Intelligence 2473*. Berlin: Springer; 2002. pp 37–42.
36. Thagard P. *Conceptual revolutions*. Princeton, NJ: Princeton University Press; 1992.
37. Uren V, Buckingham Shum S, Bachler M, Li G. Sensemaking tools for understanding research literatures: Design, implementation and user evaluation. *Int J Hum Comput Stud* 2006;64:420–445.