Knowledge Federation 2008 (Oct 20-22, 2008, Dubrovnik):
http://knowledgefederation.org/Knowledge_Federation_2008

Knowledge Federation as Hypermedia Discourse

Simon Buckingham Shum
Knowledge Media Institute, The Open University, Milton Keynes, MK7 6AA, UK.
http://kmi.open.ac.uk/people/sbs

Abstract: The motivation behind the concept of *Knowledge Federation* resonates closely with the orientation of the *Hypermedia Discourse* project[1] and the *Knowledge Cartography* focus of a forthcoming book.[2] We are concerned with the co-evolution of new tools and practices for "bringing together" ideas, sharing the KF orientation that meaning is always contextualised, that truth is often contested, and that social processes are central to the task of meaning-making in non-trivial domains. We try to add a sharper focus to what is meant by "social processes" by focusing specifically on the form of discourse by which problems are framed, and meaning is constructed in teams of analysts, e.g. deliberation over alternatives, dialogue seeking common ground, or rational debate and argumentation. Our work thus draws on the conceptual foundations offered by fields such as argumentation, cognitive coherence relations and organisational sensemaking. Hypermedia points to the engineering and aesthetics of managing webs of meaningful connections as visualizable networks of claims, issues, potential solutions, evidence, and so forth. One of our Hypermedia Discourse tools is Compendium[3] which has an established user community whom we support online and with an annual workshop, and around which has developed a relatively well developed set of practices for effective use in sensemaking.[4] A recently launched tool is Cohere[5] which is a Web 2.0 platform for making connections between ideas, including argumentation. We are designing Cohere to be as open as possible, interoperable and extensible via REST services and web feeds. These tools are finding application in a wide ranging contexts, with sensemaking around climate change being a particular focus of our current efforts.[6] The remainder of this paper introduces KF participants to our orientation to our shared concerns. It opens with a recent keynote address[7] which motivated the concept of Hypermedia Discourse, followed by a more detailed description of the rationale behind the design of the Cohere web application.[8]

---

1 *Hypermedia Discourse* project: http://kmi.open.ac.uk/projects/hyperdiscourse

2 *Knowledge Cartography: Software Tools and Mapping Techniques.* (Eds.) Okada, A., Buckingham Shum, S. and Sherborne, T. Springer: Advanced Information and Knowledge Processing Series (2008, in press). http://kmi.open.ac.uk/books/knowledge-cartography

3 *Compendium Institute*: http://CompendiumInstitute.org

4 Conklin, J. (2006) *Dialogue Mapping*. Wiley: http://cognexus.org/dm_book.htm

5 *Cohere:* http://cohereweb.net

6 *ESSENCE:* http://projects.kmi.open.ac.uk/essence

7 Buckingham Shum, S. (2007). Hypermedia Discourse: Contesting Networks of Ideas and Arguments. Keynote Address, *15th International Conference on Conceptual Structures*, Sheffield, July 2007. Lecture Notes in Computer Science, Volume 4604/2007, pp.29-44. Springer: Berlin. http://kmi.open.ac.uk/projects/hyperdiscourse/docs/SBS.ICCS2007.pdf

8 Buckingham Shum, S. (2008). Cohere: Towards Web 2.0 Argumentation. *2nd International Conference on Computational Models of Argument*, 28-30 May 2008, Toulouse. IOS Press: Amsterdam. http://kmi.open.ac.uk/projects/osc/docs/Cohere.COMMA2008.pdf

Hypermedia Discourse: Contesting Networks of Ideas and Arguments

Simon Buckingham Shum

Knowledge Media Institute, The Open University, Milton Keynes, UK
www.kmi.open.ac.uk/sbs

Abstract. This invited contribution motivates the *Hypermedia Discourse* research programme, investigating the reading, writing and contesting of ideas as hypermedia networks grounded in discourse schemes. We are striving for *cognitively and computationally tractable conceptual structures*: fluid enough to serve as augmentations to group working memory, yet structured enough to support long term memory. I will describe how such networks can be (i) mapped by multiple analysts to visualize and interrogate the claims and arguments in a literature, and (ii) mapped in real time to manage a team's information sources, competing interpretations, arguments and decisions, particularly in time- pressured scenarios where harnessing collective intelligence is a priority. Given the current geo-political and environmental context, the growth in distributed teamwork, and the need for multidisciplinary approaches to wicked problems, there has never been a greater need for sensemaking tools to help diverse stakeholders build common ground.

1 Introduction

I want to talk about the challenge of our generation. […] Our challenge, our generation's unique challenge, is learning to live peacefully and sustainably in an extraordinarily crowded world. [...] The way of solving problems requires one fundamental change, a big one, and that is learning that the challenges of our generation are not us versus them, they are not us versus Islam, us versus the terrorists, us versus Iran, they are us, all of us together on this planet against a set of shared and increasingly urgent problems. [...] But we are living in a cloud of confusion, where we have been told that the greatest challenge on the planet is us versus them, a throwback to a tribalism that we must escape for our own survival.

Jeffrey Sachs: 2007 Reith Lectures: http://www.bbc.co.uk/radio4/reith2007

With these "minds", a person will be well equipped to deal with what is expected, as well as with what cannot be anticipated; without these minds, a person will be at the mercy of forces that he or she can't understand, let alone control. [...] The disciplined mind… the synthesizing mind… the creating mind… the respectful mind… the ethical mind.

Howard Gardner: *Five Minds for the Future.* Harvard Univ. Press, 2006: p.2

U. Priss, S. Polovina, and R. Hill (Eds.): ICCS 2007, LNAI 4604, pp. 29–44, 2007.

The context in which we find ourselves presents problems on a global scale which will require negotiation and collaboration across national, cultural and intellectual boundaries. At the same time we are in a climate which questions claims to knowledge, and in which the quality of discourse is often poor. This, I suggest, presents both major challenges and unique opportunities for us as a community dedicated to understanding how to provide computational support for negotiating the construction of coherent, conceptual structures. We have choices about the kinds of problems we work on, the way in which we do our modelling, and the functionalities of the systems we offer. What do we have to offer?

My thesis is that part of the solution could be discourse-oriented tools to help capture, comprehend, and manage competing interpretations and arguments for action. There is a particular need to provide languages for communities to *agree and disagree* in principled ways. This paper considers the challenge of evolving interactive tools that are flexible enough to mediate and capture discourse between stakeholders with different perspectives, yet introduce sufficient structure to provide computational services. The *Hypermedia Discourse* research programme[1] is focused on co-evolving the semantics, user interfaces, technical infrastructure, and human work practices to embed such tools in highly pressured, real time sensemaking scenarios, face-to-face and over the internet, as well as to support extended, asynchronous discourse lasting from a few days to many years.

Discourse means different things in different fields. It is used here in a broad sense to cover the diversity of verbal and written workplace communication that we want to support, which would include the framing of problems, review of solutions, and argumentation. *Discourse communities* refers to communities of practice [15] and other networks of people who "make and take perspectives" [2].

The paper is organised as follows. I start by motivating the need for tools to assist with sensemaking in socially complex scenarios, in particular, to manage discourse when tackling wicked problems [22]. The attributes required of tools to support the expression, exploration and contesting of perspectives in shifting, contentious domains defines a new class of tool for Hypermedia Discourse. The Compendium methodology and tool is then introduced as a relatively mature exemplar, before concluding with directions for future research.

2 Sensemaking

The world, indeed our lives, make sense to the extent that we can sustain a coherent narrative about who we are and why we matter. If the story fragments, our identity crumbles if we cannot re-integrate it into our narrative [3]. When we are confronted by breaches in normality, Karl Weick draws our attention to *sensemaking* as literally "the making of sense": sharing interpretations using different representations of the situation. He proposes that: *Sensemaking is about such things as placement of items into frameworks, comprehending, redressing surprise, constructing meaning, interacting in pursuit of mutual understanding, and patterning. [30], p.6*

[1] Hypermedia Discourse project: http://kmi.open.ac.uk/projects/hyperdiscourse

Weick's concern is to characterise what people do in socially complex situations, when confronted by incomplete evidence and competing interpretations : *The point we want to make here is that sensemaking is about plausibility, coherence, and reasonableness. Sensemaking is about accounts that are socially acceptable and credible. […] It would be nice if these accounts were also accurate. But in an equivocal, postmodern world, infused with the politics of interpretation and conflicting interests and inhabited by people with multiple shifting identities, an obsession with accuracy seems fruitless, and not of much practical help, either.* [30], p.61

In other words, when there is uncertainty, *what else is there to do* but through discourse, construct a narrative to fill in the gaps?

3 Argumentative Discourse

Sensemaking wrestles with conflicting interpretations, tracks technical facts with emerging issues and ideas as the problem is reframed, and tries to reconcile socio-political arguments. This is a formidable functional requirements specification for a software tool to satisfy. Elsewhere [4, 5] we trace the work of design and policy planning theorist Horst Rittel, whose characterisation in the 1970's of "wicked problems" has continued to resonate since: *Wicked and incorrigible [problems]...defy efforts to delineate their boundaries and to identify their causes, and thus to expose their problematic nature.* [22]

Rittel concluded that many problems confronting policy planners and designers were qualitatively different to those that could be solved by formal models or methodologies, classed as the 'first-generation' design methodologies. Instead, an *argumentative* approach to such problems was required: *First generation methods seem to start once all the truly difficult questions have been dealt with. …[Argumentative design] means that the statements are systematically challenged in order to expose them to the viewpoints of the different sides, and the structure of the process becomes one of alternating steps on the micro-level; that means the generation of solution specifications towards end statements, and subjecting them to discussion of their pros and cons.* [22]

This intersects with Doug Engelbart's 40+ year mission to develop software tools to augment human intellect, our "collective capability for coping with complex, urgent problems" [14]. Our work in a variety of domains has led to the definition of a class of 'augmentation system' to assist argumentative design in Rittel's terms, and other modes of workplace discourse more broadly.

4 Hypermedia Discourse

Discourse modelling is at once both useful and limited. It is limited in the sense that, like any model, it captures only key features of the world's richness, in our case, the

richness of textual prose and verbal discourse.[2] However – if done appropriately – stripping out detail to focus on underlying structure can yield cognitive, computational and theoretical benefits:

- **Cognitive:** a well designed external representation exploits the human perceptual and cognitive system to direct attention to relevant information;
- **Computational:** a formal model also provides machines with structure to reason with;
- **Theoretical:** the removal of detail may assist in identifying generalisable patterns across diverse contexts (see discussion of Cognitive Coherence Relations later).

The function of a *medium* is to make it possible for people to express, and work with, structure. Sensemaking calls for a particular kind of discourse, expressed through one or more media. *Hypermedia* can be thought of as the craft, art, science and engineering of managing structure, specifically, relationships, making it the primary discourse modelling medium for several reasons:

- **Modelling discourse relations:** an utterance only has meaning in a context, that is, when juxtaposed with others before and after it, and in relation to other possible utterances that make its selection significant.
- **Expressing different perspectives on a conceptual space:** diverse stakeholders are usually needed to define and resolve wicked problems, so support tools need to provide support for modelling flexibly, to show agreements and differences between viewpoints.
- **Supporting the incremental formalization of ideas:** as understanding develops, so that patterns can be captured using representations that are intuitive, fast in real time usage scenarios, and expressive enough to enable computational support.
- **Rendering structural visualizations:** to assist users in grasping complex interconnections between ideas and information.
- **Connecting heterogeneous content:** the *content* that stakeholders refer to during sensemaking can range from media fragments which offer little or no obvious structure, to material sufficiently structured to support forms of machine reasoning; similarly, *relationships* may range from associations expressed spatially or as untyped links, to being formally grounded in a known semantic schema.

4.1 Key Characteristics

Bringing these concepts together, we can define a class of tools designed to model discourse as hypermedia networks, with the objective of making the process and product of discourse tangible and manipulable through the combination of:

- **A discourse ontology:** A set of explicit constructs that express a subset of the richness of human verbal or written communication. An example (discussed

[2] As described later, there are ways to compensate for the terseness of modelling by integrating source texts, audio and video as richer resources for humans (and possibly machines) to supplement the discourse model.

below) is IBIS; another that we have been developing is the ScholOnto discourse schema [7].

- **One or more notations:** Symbol system(s) for rendering the ontology. For instance, IBIS can be rendered as a textual outline, and as a directed graph flowing from left to right, or from top to bottom. Each has different affordances which can complement each other as coupled visualizations.
- **An intuitive user interface:** These tools are intended for knowledge workers in diverse sectors of society, not only for discourse modellers, knowledge engineers or information scientists. The notations are therefore just part of designing the overall cognitive and aesthetic experience of working with the tool.
- **Computational services:** The above come together as augmentation of human capability through software implementation. For instance, "services" would include more efficient *capture, interpretation, sharing, retrieval, discovery and integration* of discourse modelled in the 'knowledge repository'. Interoperability not only with other relevant tools, but also compatibility with existing work practices will contribute to the overall service augmentation.
- **Literacy and fluency:** The tool's functionality is only part of the story, however. We must also examine the capabilities assumed on the part of the user, which we will do under the heading of *literacy*, the ability to read and write ideas in the new medium in a manner appropriate to the context, ideally moving towards *fluency*.

5 Compendium

Having defined the key characteristics of a Hypermedia Discourse system, we focus now on the most mature approach we have developed, in terms of its dissemination and breadth of use. This has provided a longitudinal case study to reflect on issues of knowledge technology adoption and practice [9].

Compendium is a dialogical medium for modelling the discourse around problems. We are aiming for a tool which in the hands of skilled users, can facilitate the capture and structuring ideas, not only to model discourse, but also to model problem domains in a manner that invites and structures contributions, whether this is in a synchronous or asynchronous discussion. It is optimised for use in what is arguably the most demanding context of deployment for a knowledge representation tool: real time collaborative modelling. The software is a free Java application for all platforms, including the source code. Downloads and other community resources are coordinated via the not-for-profit Compendium Institute: www.CompendiumInstitute.org

5.1 Ontology

Compendium is a direct descendent of Conklin's gIBIS prototype [13] and the 1990's QuestMap product. Its ontology expresses Rittel's IBIS and similar Design Rationale schemes such as MacLean *et al's* Questions-Options-Criteria (QOC) [16]. The focus is on capturing key issues, possible responses to these, and relevant arguments. Users can define their own ontology if they wish, or map concepts in a completely

unconstrained manner. Entities are described in free text, while labels may be free text or grounded in a predefined scheme. Additional semantics can be expressed textually by defining one or more *Tag* groups, which operate as flat keyword spaces, analogous to web-based tagging, whereby tag combinations can be used to define different searchable views of the database. Semantics can, additionally, be expressed visually, either by predefining a palette of icons, or by selecting images to reflect ideas as they emerge in discussion (eg. from a library, or by searching the Web).

5.2 Notation

Some people use Compendium to support their preferred style of concept mapping [20]. However, following the gIBIS system, Compendium is designed specifically to render IBIS as a directed graph, normally with a root issue on the left, with the structure of the developing conversation about this issue growing to the right of the screen. User customizable icons distinguish different entities, and link colours with optional labels indicate relational semantics. Links typically point from right to left, to reflect the conversational dynamic that new contributions (added to the right) *respond-to* existing ones.

The discourse-orientation of the approach, and the demands of real time participatory modelling to capture the progress of meetings, have led to a number of notational strategies. A root *Issue* (signalled with a question mark icon) provides the orientation to a map, establishing the problematic context for the discussion: *Why are we here? To tackle this issue.* Two discourse modelling methodologies have developed around the capabilities of Compendium. *Dialogue Mapping* is a set of skills developed by Conklin [12] for mapping IBIS structures in real time during a meeting in order to support the analysis of wicked problems, as defined by Rittel. In Dialogue Mapping, Issues are usually unconstrained freetext expressions summarising an agenda item or a participant's contribution, with Ideas responding to them, and any associated arguments (Fig. 1).

Conversational Modelling [23] incorporates and extends Dialogue Mapping by deriving Issues from a modelling methodology (or for instance, an organizational procedure/best practice). Issue nodes can be saved as reusable *issue-template* structures to seed different kinds of discussions. Fig. 2 shows a fragment of one template, with *Idea* icons serving as placeholders for responses. These lead to consequent *Issues* to be considered (on the right).

In addition, the modelling methodology specifies that the placeholder Ideas appear in three different views, indicated by the numeral 3 on each Idea icon. Rolling the mouse over this numeral displays a menu of hyperlinks to these other views. When views are labelled informatively, this facility provides rich context at a glance to the different 'conversations' in which a node is being discussed. Node label auto-completion assists the reuse of these granular chunks, offering users a menu of existing nodes which they can select from as they type.

With the addition of *catalogues* of reusable nodes, metadata *tagging* and multiple linked *issue-templates*, Compendium provides generic building blocks to construct a discourse-oriented modelling environment for team deliberation (Tate et al [28] document the customisation of Compendium in an hour from receipt of a planning

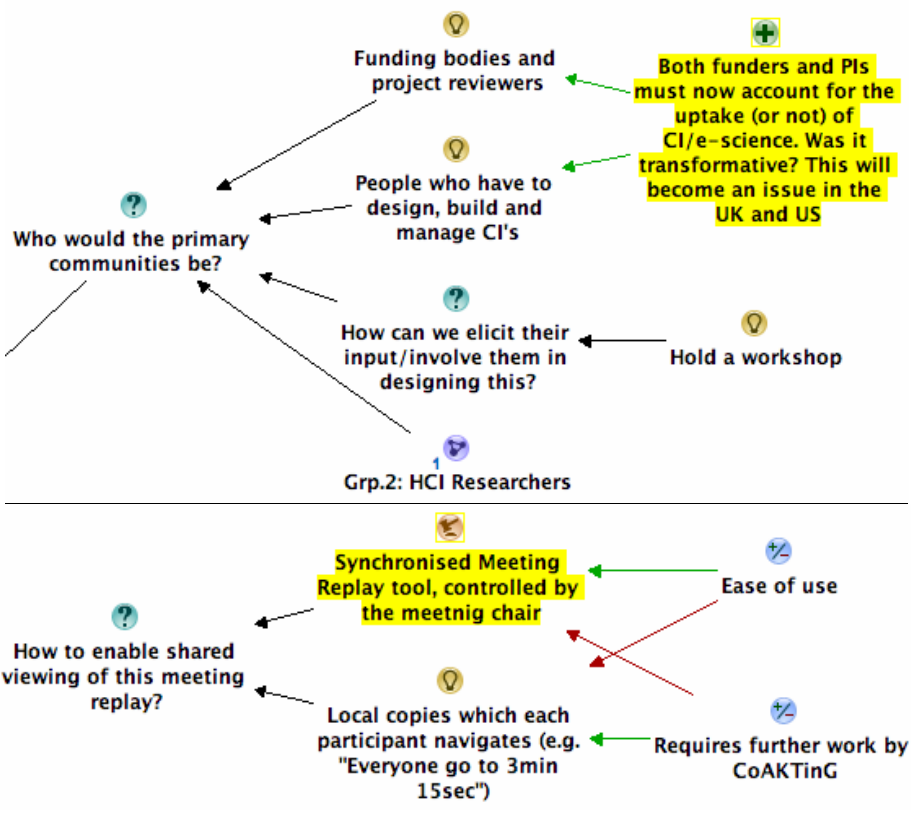

Fig. 1. Fragments from two Dialogue Maps using IBIS. In the top example exploring requirements for a website, a Pro argument of a political nature is highlighted, backing two Idea nodes. In the lower example, a QOC-style design discussion examines Option tradeoffs against more formally expressed design Criteria.

methodology). Conversational Modelling enables the real time capture of both expected, well-structured information through the use of issue templates, with the flexibility to capture unexpected, ad hoc information and discussions as they arise.

From a more formal knowledge representation perspective, we represent semantics using a variety of conventions. In a NASA field trial (Fig. 3), science metadata was represented using templates which look like visual forms, with each Issue inviting the team to answer (or if necessary debate) the values of the 'slots'.

An issue-template such as this provides a user-friendly way to engage in participatory modelling which permits argumentation if necessary, and results in a set of semantic assertions amenable to automated analysis (data entry into a simulation engine in this case). Each *Issue* in fact embodies the relational semantic connecting its answer to the entity represented by the containing map. However, rather than ask the team to complete sets of semantic triples, they are offered a set of question mark icons to which they need to link lightbulb icons. Thus, Fig. 2 provides an interface to elicit

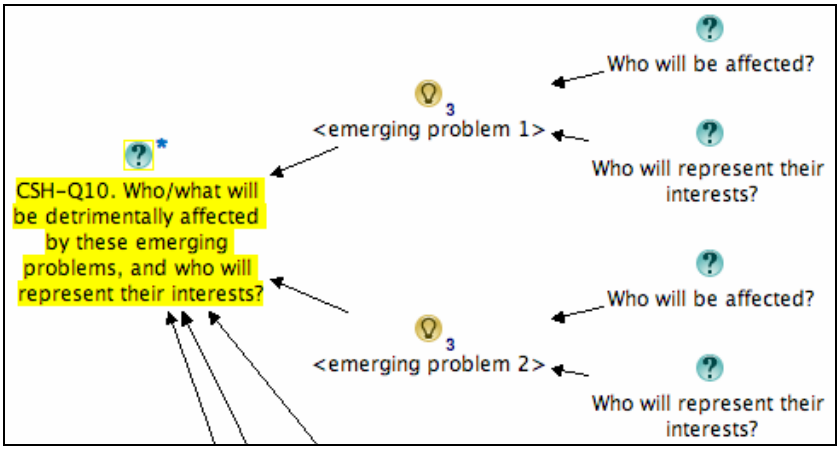

Fig. 2. An Issue-template used in Conversational Modelling. For each answer, there are two subsequent Issues.

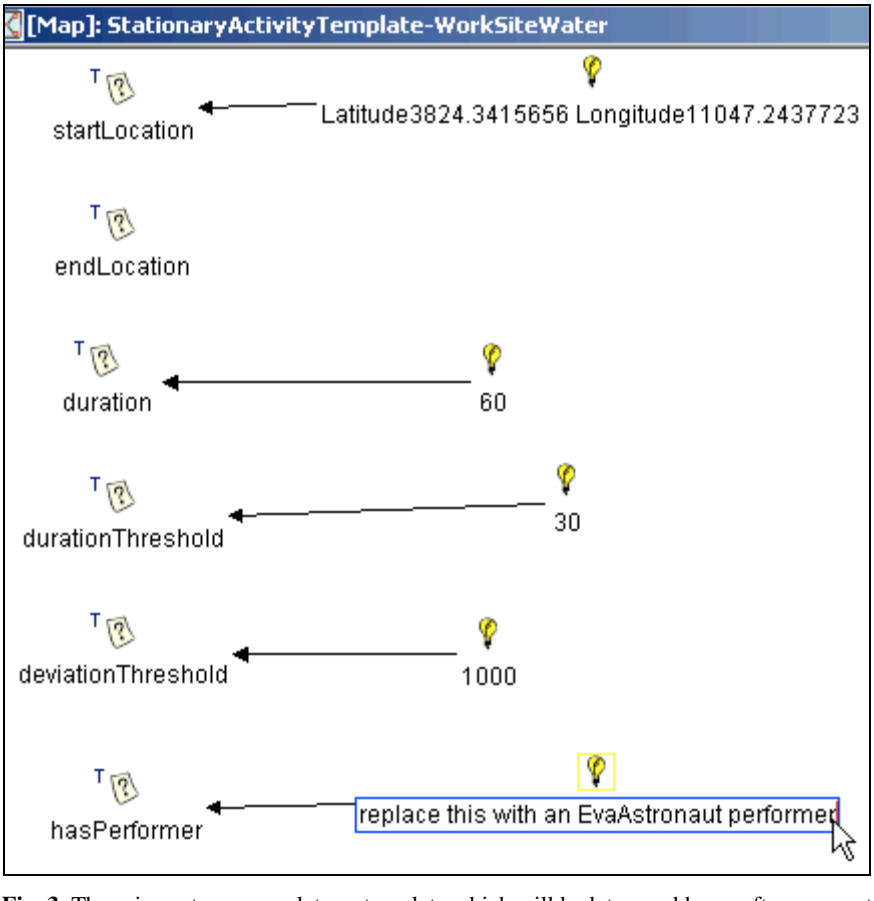

Fig. 3. The science team completes a template which will be later read by a software agent

the structured assertion *<user's answer> will_be_affected_by <emerging problem 1>,* while FIG. 3 will elicit *<WorkSiteWater> hasPerformer <user's answer>.*

Relational semantics are also expressed in the link types, but for speed – a key requirement in real time mapping under pressure – link types are set to be unlabelled by default, with the semantics loaded on the nodes' iconic language. Every link can be classified and labelled if desired using the default IBIS linkset, or a user defined linkset.

5.3 Intuitive User Interface

There are many improvements that could be made to Compendium, but as the preceding figures show, it looks familar to users of concept mapping or graph-editing applications. It comes with IBIS preloaded, and hypermedia functionality which makes it simple to (i) create navigational links to a given database view, and (ii) reuse a hypertext node simultaneously in different views by copying and pasting. A keyword tagging scheme combined with search assists with filtering nodes across many maps.

Complete beginners can learn to map simple but well-formed IBIS structures after working through a tutorial on the Compendium Institute website. End users can express quite sophistcated data and relationships without needing to perform complicated technical actions or remember arcane commands. The user feedback on the website reflects the personal sense of satisfaction that users have with the tool.

5.4 Computational Services

We earlier defined "services" as the set of affordances at the intersection of ontology, notation, user interface, and the human and machine reasoning these enable. Compendium's display has a number of visual affordances which enable one to read off information about the state of an analysis that is not immediately obvious, either in conventional text documents or other concept mapping approaches. This includes unresolved issues, competing ideas, the extent to which explicit evidence is used to back ideas, and the 'depth' of node reuse and tagging (an indicator of the degree of modelling utilised).

When Compendium is interfaced to other tools, its database can be automatically populated or reasoned about. Examples include the use of software agents to autonomously read data and pass this to a simulation and planning engine, and also to populate the database with multimedia data for subsequent analysis by scientists [10]; the exchange of issues with a planning tool which could analyse the option space exhaustively or raise new issues [28]; the export of populated issue templates to different notational formats for other stakeholders to work on [26].

Most recently, we have automated the exchange of Compendium data with an RDF triplestore, in order to deliver a video conferencing capture and semantic replay tool [8]. Fig. 4 illustrates the complementary use of video from meetings to 'fill in the gaps' that a terse conceptual graph cannot possibly express; conversely, Compendium provides semantic indexing within and across meetings, enabling users to jump to the point in a meeting when, for instance, an argument was made.

5.5 Literacy and Fluency

Advanced tools are more effective when used expertly. The concept of services must, therefore, be qualified by the degree of literacy and fluency that the user brings. Our research agenda is directed towards understanding the whole learning curve associated with reading and writing in this new medium. We have analysed the cognitive tasks that a beginner must learn [6] and there are training programmes to help with initial adoption of the tool, but equally, we need to characterise expert, 'fluent' use of the tool in the most demanding contexts we work in, namely, supporting real time sensemaking in time pressured teams (e.g. [10, 28]). Constructing a language for fluency should help to expand the boundaries of expertise, improve the apprenticing of new practitioners, foreground new functionalities that the tool should provide, and illuminate an emerging literacy in this new medium.

Selvin [24, 25] has begun to explore the nature of fluency in what he terms *Participatory Hypermedia Construction.* Detailed analysis of screen recordings from teleconferences and face-to-face meetings is providing an account of the representational moves that Compendium mappers make, and the different roles they can play in meetings.

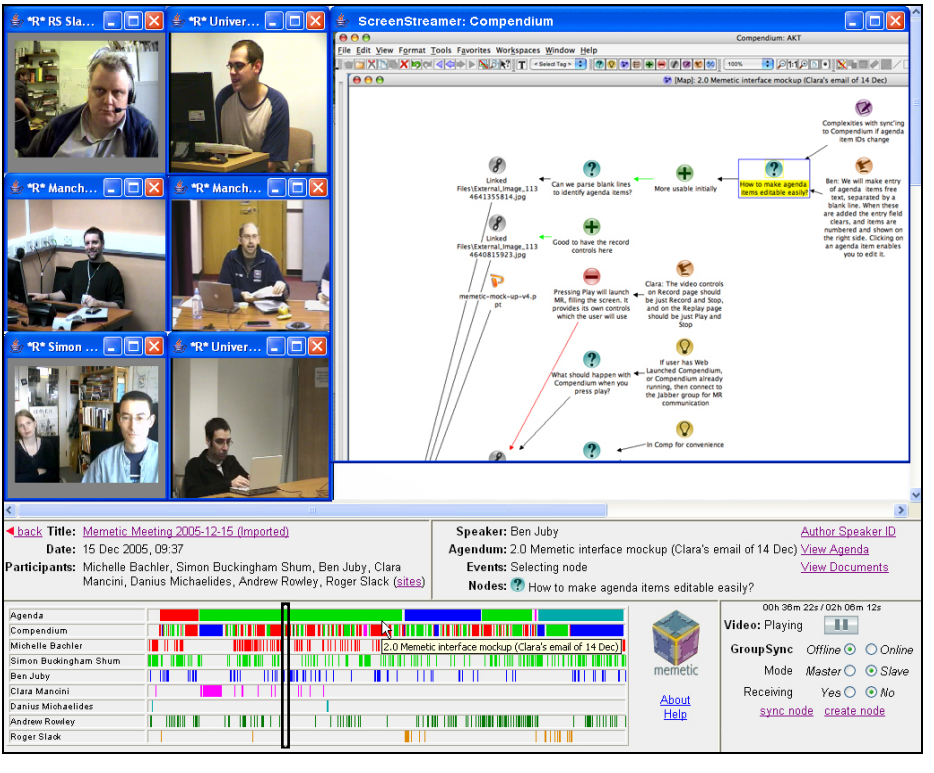

Fig. 4. The Memetic Meeting Replay tool, using Compendium nodes as a means of indexing and navigating meeting videos

6 Semantic Scholarly Publishing and Annotation

A second instantiation of the Hypermedia Discourse concept is the suite of tools developed in the Scholarly Ontologies project.[3] Unlike Compendium, which simply offers Web exports and supports the embedding of websites in IBIS conversational models, these tools were conceived from the start as distributed Web applications. The design rationale is the need for representational infrastructure to evolve the current prose document and associated practices for publishing and contesting research results and – equally significant – authors' *interpretations of their significance*. Within current research into 'e-Science' (UK) and 'Grid/cyberinfrastructure' (USA), this is a neglected part of the scholarly lifecycle, which is ironic: we engage in research in order to substantiate *knowledge level claims*. Perhaps, however, the absence of activity in this latter stage of research should not surprise us, because we are of course dealing with the difficult issue of computational support for an intrinsically *pragmatic* process, by which a discourse community (in this case, research peers) negotiates what some reported facts should be taken to *mean.* The emerging Pragmatic Web community has as a primary focus the interplay between formal representation and context, conversations and commitments to action, and it will be interesting to see how this takes shape.

We detail elsewhere [27, 29] the design and evaluation of ClaiMaker and the associated suite of tools for authoring (ClaiMapper), querying (ClaimFinder) and the collaborative, semantic annotation (ClaimSpotter) of research claims and argumentation. These are less mature than Compendium, proof of concept research tools which are not publicly available. Space precludes as detailed a treatment as Compendium, but ClaiMaker's 'hypermedia discourse profile' below conveys the essence of the approach:

- **Discourse ontology:** A two-layer relational taxonomy which provides base relational classes in which 'dialects' from different discourse communities are grounded (Fig. 5).
- **Notation:** A conceptual graph of claims that can be visualized using different schemes to show discourse connections between concepts annotated onto the literature.
- **User interface:** We have investigated a variety of interaction paradigms for annotation tools, in order to help untrained users create semantic annotations.
- **Computational services:** The use of a richer discourse scheme than IBIS enables us to offer more powerful services. For instance, the semantic citation maps can be filtered in response to queries such as, *What documents report data that challenges this author's hypothesis?What is the lineage of this concept: the key ideas on which this work builds?* (Fig. 6)
- **Literacy and fluency:** Being less mature than Compendium, we do not yet have a large enough user community to provide a good description of what it means to read and write such argumentative networks, particularly beyond initial learning. Our empirical studies provide insight into how untrained and more expert users construct and query claim networks [27, 29].

[3] Scholarly Ontologies project: http://kmi.open.ac.uk/projects/scholonto

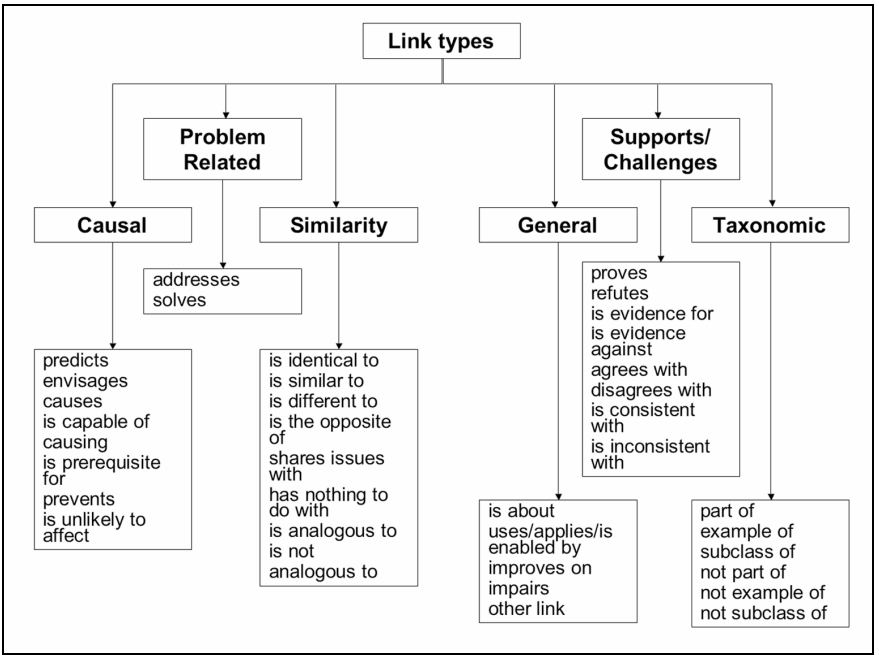

Fig. 5. ClaiMaker's discourse scheme, which groups the 'dialect' of a discourse community under more primitive relational classes

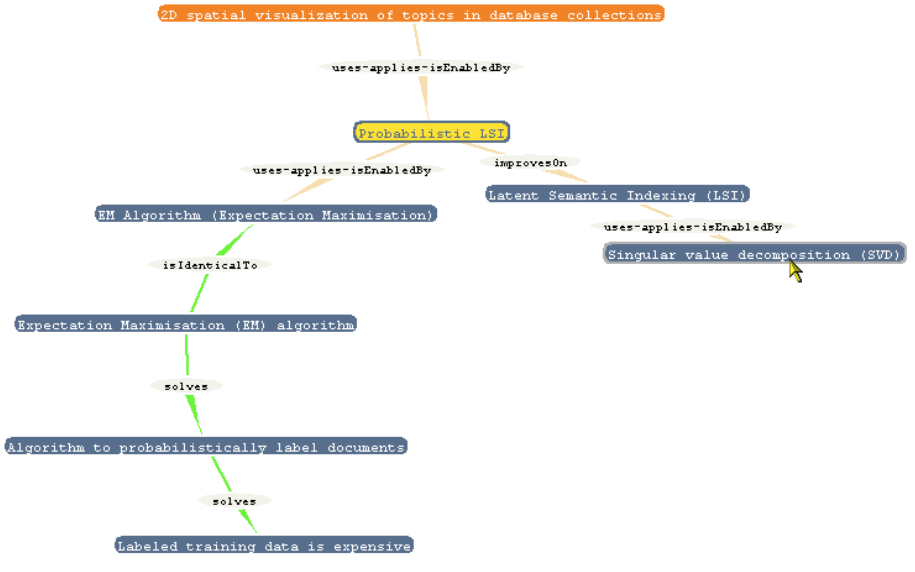

Fig. 6. ClaimFinder's *Lineage* query traces the 'intellectual roots' of a concept. displayed at the top. The conceptual graph is analysed and filtered to show potentially significant relational types such as *uses/applies/is enabled by, improves on*, and *solves.*

7 Conclusions and Future Work

The complexity of the dilemmas we face at an organizational, societal and global scale forces us into sensemaking activity. The requirements on tools to support such work have motivated basic and applied action research into a new class of Hypermedia Discourse tool to mediate, structure and augment the expressing and contesting of perspectives that may agree and disagree in principled ways. Such tools are hybrids borrowing from concept mapping, information visualization, discourse relations and decision-support. We need tools flexible enough for real time use in meetings, structured enough to help manage longer term memory, and powerful enough to filter the complexity of extended deliberation and debate on an organizational or global scale.

I suggest that this focus on the intersection of discourse and hypermedia provides insights into a number of pressing problems:

- **We have to talk.** The only way that anything is accomplished in this world is by people talking, building trust and sufficient common ground that they can frame problems in mutually meaningful ways, and commit to action in mutually acceptable ways. The challenge for a community such as ours is understand how to weave software support into the social fabric without ripping it, but possibly in the process, enriching that fabric to exploit the new threads we have to offer. The work summarised here points to possible ways to evolve network-native infrastructures for synchronous and asynchronous discourse, that step out of the shadow of the printing press and conventional meetings (building on their strengths, but transcending their limitations).
- **Modelling in the absence of consensus.** Knowledge-based systems (including for our purposes the data models and ontologies underpinning the Semantic Web) encapsulate *consensus models* of the problem domain, and how to reason about it. How can we provide computational services *in the absence of consensus*, when one group's assumption is another group's problem? This is the domain of discourse, especially argumentation, in which we provide a language for stakeholders to agree and disagree in principled ways. Compendium uses a semiformal network representation optimised for real time use. ClaiMaker uses finer grained semantics for modelling asynchronously in a more detailed manner.
- **Negotiating the knowledge capture bottleneck.** In knowledge engineering, but also in less formal approaches to Knowledge Management (KM), Organizational Memory and Design Rationale (DR), the cost/benefit tradeoff must be negotiated to acquire useful abstractions of naturally occurring activity, and experts' descriptions thereof. The Compendium approach emphasises the collaborative modelling of information, ideas and argument in order to add immediate value to the users (useful *working memory*), as well as seeding the *long term memory* required for KM. This has, for instance, provided a way of tackling the DR capture bottleneck [9].

Future work will continue to co-evolve tools and practices, study the skills associated with high performance discourse modelling, and develop conceptual

frameworks that recognise the complexity of modelling, mediating and mapping real discourse about wicked problems. Specific challenges we are working on include:

- **Distributed, online apprenticeship in hypermedia discourse.** The Compendium community now has members who are recognised 'expert mappers', but they are a scarce resource. A very applied concern is how to use the internet to spread this literacy through the creation of e-learning resources and 'e-apprenticeship'.
- **Social networks and folksonomic tagging.** Behind a conceptual structure are people. We are integrating our social networking tools with our conceptual networking tools to support *Open Sensemaking Communities*, learners and educators who must self-organise around open source learning resources, but by extension, any epistemic community on the internet. Based on the ScholOnto project, we have prototyped and formatively evaluated a next generation social bookmarking tool for linking tags via discourse connectives, moving from the annotation of isolated keywords on web reosources, to a mode knowledge construction and negotiation: *from tag clouds to tag webs* [27].
- **Hypermedia discourse engines as computational theory.** We are investigating the potential of modelling and reasoning over an upper level relational ontology, derived from linguistics coherence relations research [18]. If it is the case that we perceive 'coherence' in a medium because it structures elements according to a small, bounded set of relational primitives, then it should be possible to model and reason over such structures in a manner which is 'coherent' across different domains of discourse, languages and even cultures. Such an engine would be a formal expression, and test, of the hypotheses generated by this theory.

To return to our opening quote from Gardner's *Five Minds for the Future*, perhaps Hypermedia Discourse tools provide a way to move fluidly between the different minds: a way to provide representational scaffolding for disciplined modelling, but permitting the creative breaking of patterns when needed and the forging of new syntheses; a way to show respect for diverse stakeholders' concerns by explicitly integrating them into the conversation; a way to bring into an analysis 'messy' requirements such as ethical principles, as well as hard data and constraints. We have some evidence from our case studies that we're on the right track, but there remains much to do.

Acknowledgements. I am grateful to Al Selvin, Clara Mancini, Jack Park and David Kolb whose comments improved earlier versions of this paper. The evolution of Compendium has been a long term action research programme with Al Selvin, Maarten Sierhuis and Jeff Conklin, with programming by Michelle Bachler. Its development has been funded by the UK's research councils EPSRC, ESRC and JISC. The Meeting Replay tool is joint work with the Universities of Manchester and Southampton as part of the JISC Memetic project. The Scholarly Ontologies project was funded by the EPSRC, and is the product of work with Victoria Uren, Gangmin Li, Clara Mancini, Bertrand Sereno, Enrico Motta and John Domingue. The William and Flora Hewlett Foundation is now supporting the work through the Open University's *OpenLearn* initiative.

References

1. Boden, D.: The Business of Talk. Polity, Cambridge (1994)
2. Boland, R.J.J., Tenkasi, R.V.: Perspective Making and Perspective Taking in Communities of Knowing. Organization Science 6(4), 350–372 (1995)
3. Bruner, J.S.: Acts of Meaning. Harvard University Press, Cambridge, MA (1990)
4. Buckingham Shum, S.: The Roots of Computer Supported Argument Visualization. In: Kirschner, P.A., Buckingham, S., Carr, C. (eds.) Visualizing Argumentation, pp. 3–24. Springer, London (2003)
5. Buckingham Shum, S., Hammond, N.: Argumentation-Based Design Rationale: What Use at What Cost? Int. J. Human-Computer Studies 40(4), 603–652 (1994)
6. Buckingham Shum, S., et al.: Graphical Argumentation and Design Cognition. Human-Computer Interaction 12(3), 267–300 (1997)
7. Buckingham Shum, S., et al.: Modelling Naturalistic Argumentation in Research Literatures: Representation and Interaction Design Issues. Int. J. Intelligent Systems (Special Issue on Computational Modelling of Natural Argument) (in Press)
8. Buckingham Shum, S., et al.: Memetic: An Infrastructure for Meeting Memory. In: Proc. 7th Int. Conf. on the Design of Cooperative Systems. 2006 of Conf. Carry-le-Rouet (2006)
9. Buckingham Shum, S.J., et al.: Hypermedia Support for Argumentation-Based Rationale: 15 Years on from gIBIS and QOC. In: Dutoit, A., et al. (eds.) Rationale Management in Software Engineering, pp. 111–132. Springer, Heidelberg (2006)
10. Clancey, W.J., et al.: Automating CapCom Using Mobile Agents and Robotic Assistants. In: 1st Space Exploration Conf. 2005 of Conf. Orlando, FL (2005), http:// eprints.aktors.org/375
11. Conklin, J.: Dialogue Mapping: Reflections on an Industrial Strength Case Study. In: Kirschner, P.A., Buckingham, S., Carr, C. (eds.) Visualizing Argumentation, Springer, London (2003)
12. Conklin, J.: Dialogue Mapping: Building Shared Understanding of Wicked Problems. Wiley, Chichester (2005)
13. Conklin, J., Begeman, M.L.: gIBIS: A Hypertext Tool for Exploratory Policy Discussion. ACM Transactions on Office Information Systems 4(6), 303–331 (1988)
14. Engelbart, D.C.: A Conceptual Framework for the Augmentation of Man's Intellect. In: Howerton, P., Weeks. (eds.) Vistas in Information Handling, pp. 1–29. Spartan Books, Washington, DC, London (1963)
15. Lave, J., Wenger, E.: Situated Learning: Legitimate Peripheral Participation. Cambridge University Press, Cambridge (1991)
16. MacLean, A., et al.: Questions, Options, and Criteria: Elements of Design Space Analysis. Human-Computer Interaction 6(3 & 4), 201–250 (1991)
17. Mancini, C.: Cinematic Hypertext. Investigating a New Paradigm. In: Frontiers in Artificial Intelligence and Applications. IOS Press, Amsterdam (2005)
18. Mancini, C., Buckingham Shum, S.: Modelling Discourse in Contested Domains: A Semiotic and Cognitive Framework. Int. J. Human-Computer Studies 64(11), 1154–1171
19. McCall, R.: Fundamentals - Rationale Representation, Capture and Use. In: Dutoit, A., et al. (eds.) Rationale Management in Software Engineering, pp. 49–52. Springer, Heidelberg (2006)
20. Novak, J.D.: Learning, Creating, and Using Knowledge: Concept Maps as Facilitative Tools in Schools and Corporations. LEA, Mawah, NJ (1998)
21. Park, J.: Topic Mapping: A View of the Road Ahead. In: Maicher, L., Park, J. (eds.) TMRA 2005. LNCS (LNAI), vol. 3873, Springer, Heidelberg (2006)

22. Rittel, H.W.J.: Second Generation Design Methods. Interview in: Design Methods Group 5th Anniversary Report: DMG Occasional Paper, vol. 1, pp. 5–10 (1984). In: Cross, N. (ed.) Developments in Design Methodology (reprinted), pp. 317–327, J. Wiley & Sons, Chichester (1972)
23. Selvin, A.: Supporting Collaborative Analysis and Design with Hypertext Functionality. Journal of Digital Information 1(4) (1999)
24. Selvin, A.: Fostering Collective Intelligence: Helping Groups Use Visualized Argumentation. In: Kirschner, P.A., Buckingham, S., Carr, C. (eds.) Visualizing Argumentation, Springer, London (2003)
25. Selvin, A.: Aesthetic and Ethical Implications of Participatory Hypermedia Practice, Technical Report KMI-05-17, Knowledge Media Institute, Open University (2006)
26. Selvin, A.M., Buckingham Shum, S.J.: Rapid Knowledge Construction: A Case Study in Corporate Contingency Planning Using Collaborative Hypermedia. Knowledge and Process Management 9(2), 119–128 (2002)
27. Sereno, B., Buckingham Shum, S., Motta, E.: Formalization, User Strategy and Interaction Design: Users' Behaviour with Discourse Tagging Semantics. In: Workshop on Social and Collaborative Construction of Structured Knowledge. 16th Int. World Wide Web Conference, Banff, Canada, May 8-12, 2007 (2007), [http://www2007.org/workshops/paper_30.pdf]
28. Tate, A., et al.: Co-OPR: Design and Evaluation of Collaborative Sensemaking and Planning Tools for Personnel Recovery, Technical Report KMI-06-07, Knowledge Media Institute, Open University (2006)
29. Uren, V., et al.: Sensemaking Tools for Understanding Research Literatures: Design, Implementation and User Evaluation. Int. J. Human Computer Studies 64(5), 420–445 (2006)
30. Weick, K.E.: Sensemaking in Organizations. Sage Publications, Thousand Oaks, CA (1995)

Cohere: Towards Web 2.0 Argumentation

Simon BUCKINGHAM SHUM*
Knowledge Media Institute, The Open University, UK

Abstract: Students, researchers and professional analysts lack effective tools to make personal and collective sense of problems while working in distributed teams. Central to this work is the process of sharing—and contesting—interpretations via different forms of argument. How does the "Web 2.0" paradigm challenge us to deliver useful, usable tools for online argumentation? This paper reviews the current state of the art in Web Argumentation, describes key features of the Web 2.0 orientation, and identifies some of the tensions that must be negotiated in bringing these worlds together. It then describes how these design principles are interpreted in *Cohere,* a web tool for social bookmarking, idea-linking, and argument visualization.

Keywords: argumentation tools; argument visualization; usability; Web 2.0

1. Introduction: The Need for Distributed, Collective Sensemaking Tools

The societal, organizational, scientific and political contexts in which we find ourselves present problems on a global scale which will require negotiation and collaboration across national, cultural and intellectual boundaries. This, I suggest, presents both major challenges and unique opportunities for us, as the community dedicated to understanding computational support for argumentation: our challenge is to work with relevant stakeholders to co-evolve new practices with flexible, usable tools for communities to express how they *agree and disagree* in principled ways, as part of building common ground and mutual understanding.

While our previous work has focused on the real time mapping of issues, dialogue and argument in contexts such as e-science teams [1] and personnel rescue [2], this paper focuses specifically on the challenge of designing engaging, powerful tools for distributed, primarily asynchronous work which, in particular, exploits the strengths of the "Web 2.0 paradigm". The paper begins by reflecting on the kinds of expectations that Web users and developers now bring, before surveying the current state of the art in Web Argumentation tools. We then describe how Web 2.0 principles, as introduced, have informed the design of a prototype tool called Cohere, concluding with a vision of how COMMA researchers might extend or interoperate with it as it moves towards a Web services platform.

* Correspondence to: Knowledge Media Institute, The Open University, Milton Keynes, MK7 6AA, UK. Tel: +44 (0) 1908 655723; Fax: +44 (0) 1908 653196; http://kmi.open.ac.uk/people/sbs

2. The Web 2.0 Paradigm

A lot is being written about the Web 2.0 paradigm, a term first dubbed in 2004 [3]. While some dismiss it as marketing hype, it does serve as a useful umbrella term to cover significant new patterns of behaviour on the Web. There are many lists of the key characteristics of Web 2.0, not all of which are relevant to our concerns (e.g. e-business models). In this section we select several characteristics for their impact on the user experience of collective information structuring. Together these present a challenge to the design of practical Web Argumentation tools, given the expectations that users now have from their everyday experience of the Web. If we cannot create tools within the new landscape, argumentation tools will remain a much smaller niche than they should be — and as this paper seeks to demonstrate, need be.

2.1. Simple but Engaging Multimedia User Interfaces

The World Wide Web has established itself as the default platform for delivering interactive information systems to professionals and the public. Although early Web applications lacked the elegance and interactivity of desktop applications due to the need for the client to communicate every state change to the server, the gap is closing rapidly with the emergence of good graphic design principles, controlled layout and stylesheet management, and critically, so-called Rich Internet Applications: interactive multimedia capabilities such as Adobe Flash embedded as standard browser plugins, and approaches such as AJAX (Asynchronous JavaScript And XML) for caching local data to increase the responsiveness of the user interface [4]. Users increasingly expect Web applications to have a clean, uncluttered look, and to be as responsive as offline tools. Given a choice of Web offerings, the user experience can determine whether or not a tool is adopted.

2.2. Emergent, Not Predefined, Structure and Semantics

Argumentation focuses on a particular kind of semantic structure for organising elements. Of central interest, therefore, is the Web 2.0 emphasis away from predefined information organizing schemes, towards self-organised, community indexing ('tagging') of elements, resulting in so-called "folksonomies" that can be rendered as tag clouds and other visualizations. Persuading 'normal people' (in contrast to skilled information scientists or ontology engineers) to create structured, sometimes high quality, metadata was previously thought impossible, and the success and limits of this approach is now the subject of a new research field that studies collaborative tagging patterns, e.g. [5].

Another way in which this emphasis expresses itself is in the new generation of tools that make it easy to publish one's opinion of the world. Free, remotely hosted blogging tools such as Wordpress and Blogger make it very easy for non-technical users to create a personally tailored journal or diary and syndicate their ideas. Blogs demonstrate one way to negotiate the formality gulf successfully, providing expressive freedom (essentially, traditional prose and graphics), with just enough structure to reap some benefits of hypertext (entries are addressable as URLs, timestamped, tagged, and syndicated as web feeds – see below). The limitation of blogging at present is that like the Web at large, there are no semantics on the links between postings, thus failing to

provide any support for an analyst who wants to gain an overview of the moves in a debate, or indeed, any kind of inter-post relationship.

2.3. Social Networks

Web 2.0 applications are dominated, although not exclusively restricted to, sites that either seek explicitly to connect people with people, often via the artifacts that they share. They are designed such that the greater the numbers participating, the higher the return on effort invested. Social tools provide a range of ways in which users are made aware of peer activity, for instance, alerting when another user 'touches' your material (e.g. by reusing it, making it a favourite, tagging it), or by mining social network structure to suggest contacts in a professional network. Social tools also provide mechanisms for building reputation, from the trivial (how many "friends" one has), to potentially more meaningful indices, such as authority based on the quality of material or feedback that a user posts, or professional endorsements.

2.4. Data Interoperability, Mashups and Embedded Content

A core idea behind the Web 2.0 paradigm is access to data over the web from multiple applications. Web feeds using RSS and Atom have become the lingua franca for publishing and subscribing to XML data in a simple manner that many non-technical users now handle daily. Public APIs and web services enable the more sophisticated access that enterprise architectures require, while semantic web services promise to overlay ontologies on these layers so that they can be configured according to function. "Mashups" of data sources fuse disparate datasets around common elements (e.g. geo-location, person, date, product), often accessed via customisable user interfaces such as Google Maps [6]. While many mashups typically need to be crafted by a programmer, others can be generated by end-users, given a sufficiently flexible environment. The results of a search may bring together data in new ways.

The phenomenal growth of web applications such as Google Maps, YouTube, Flickr and Slideshare is in part due to the ease with which users can embed remotely hosted material in their own websites. By providing users with the 'snippet' code (which may be HTML or JavaScript), such applications empower users to in turn provide their readers with attractively presented access to the material, which can in turn be embedded by those readers in their sites. The material thus spreads 'virally', as the links to a resource increase: it is no longer necessary to visit a web page to access its content.

3. Web Argumentation Tools

A significant strand in COMMA research focuses on the design, implementation and evaluation of practical software tools for creating and analysing arguments. Following the entity-relationship modelling paradigm that lends itself so well to software, as well as the work of pioneering argument and evidence mapping theorists such as Wigmore and Toulmin, these tools provide a way to construct arguments as structures comprising semantically linked elements taken from one or more argumentation schemes. The argument structures may be left implicit behind text-centric user interfaces, or rendered explicitly as trees or networks to help the author and reader visualize and edit the

argument [7]. The intended users of such tools include members of the public engaged in a public consultations and societal debate [8], students or educators in a learning context [9], lawyers [10], and analysts in many other walks of professional life such as public policy [11] and scholarly publishing [12]. Research in this field examines issues including the translation of argumentation theory into computable representations [13], the nature of expert fluency with such tools [14, 15], and empirical studies of the tools' usage in all of the above domains.

In light of the high design standards and new possibilities that the Web 2.0 paradigm sets, it is clear that existing tools have limitations. First, there are desktop applications like Compendium [30] and Rationale [16] with high quality user interfaces refined through the feedback from their extensive user communities: however, these are limited to publishing read-only maps to the Web, either as JPEG images, or as interactive image maps. Single user applications like CmapTools which have been migrated to 'groupware' versions provide remote editing of maps, but do not exploit the Web 2.0 functions described above.

Finally and most relevant, there are a number of Web-native applications, designed from the start to support large scale, multi-user construction. Some websites now provide a very simple structure for structuring the two sides of a debate, while others provide a more articulated argumentation language. Beginning with the least structured, we see the emergence of sites such as Debatepedia, which is modelled on Wikipedia, providing a debating resource showing unstructured prose arguments for and against a particular proposal, demarcated in two columns [17]. CoPe_it! [18] is designed for community deliberation, and provides a way to synchronise views between IBIS graphs (it also integrates with Compendium in this respect), an IBIS outline tree, and a conventional threaded discussion forum. CoPe_it! also provides a mechanism to evaluate the strength of a position, and so represents another interesting development. Its interaction design is at present rather rudimentary compared to Web 2.0 interfaces. It does not have an end-user customisable semantics, interoperability with existing Web data sources, or mechanisms to syndicate content outside the application.

Parmenides is designed to support web-based policy consultation with the public, and incorporates a formal model of argumentation [19]. It provides a forms-based, questionnaire interface to elicit views from the user, populating an argumentation structure, which it then reasons over to elicit further views. Parmenides enforces a particular argument ontology (it was not designed as a social web application) and does not appear to support any other Web 2.0 characteristics.

ClaiMaker [20] was a Web 1.0 era application, developed in our own prior work modelling the claims and arguments in research literatures. ClaiMaker, and its sister tool ClaimSpotter [21], provided vehicles for us to validate empirically the usability of the data model and a number of user interface paradigms. This has led us to carry the core data model through into Cohere, while relaxing the constraint that restricted users to the predefined classifications of nodes and links. Cohere's visualizations are also versions of those first prototyped in ClaiMaker.

TruthMapping goes much further than this, aiming specifically at tackling some of the limitations of threaded discussion forums, with a clear distinction between unsupported premises, which when supported become claims, and a way to post rebuttals and responses to each of these [22]. DebateMapper uses a combined graphical and outline structure to map debates using the IBIS scheme, with contributions tagged as issues, positions and arguments [23]. DebateMapper perhaps illustrates most clearly some the Web 2.0 interaction design principles, but provides no open semantics, or an

open architecture to enable services on the data. The ArgDF system [24] is the first argumentation tool to adopt a Semantic Web architecture based around the W3C standard Resource Description Framework (RDF) for distributed data modelling and interchange. Moreover, ArgDF is probably the first interactive tool to ground its argument representation in the recently proposed Argument Interchange Format (AIF) [25]. This combination of AIF+RDF is a notable advance. However, while proving the conceptual and technical feasibility of a semantic web orientation for argumentation, it does not yet have a user community, and it cannot be regarded as a Web 2.0 application as defined above.

4. The Cohere system

We now describe how we are trying to incorporate the Web 2.0 principles introduced above to create an environment called *Cohere* [cohereweb.net] which aims to be semantically and technically open, provide an engaging user experience and social network, but provide enough structure to support argument analysis and visualization.

4.1. Emergent Semantics: Negotiating the Formalization Gulf

In any user-driven content website, the challenge is to keep entry barriers as low as possible to promote the growth of the community, yet maintain coherence of navigation and search, through the careful design of the data model and user interface. The combination of data model and user interface must seek the right balance between constraint and freedom. This Web 2.0 orientation might seem to be in tension with an environment designed to promote rigorous thinking and argumentation. Our approach is to start with relaxed constraints in order to foster engagement with the idea of structuring ideas in general, but provide tools to incrementally add structure as the user recognises the value that it adds in a given context.

Cohere is, therefore, styled to invite playful testing by people who may not first and foremost be interested in argumentation. Instead, the website invites them to *make connections between ideas.* This broader framing aims to meet the need of many sensemaking communities to express how ideas or resources are related (whether or not this is argumentative) in a way that goes beyond plain text blog postings, wikis or discussion forums. A typical pair of connected Ideas in Cohere is illustrated in Figure 1.

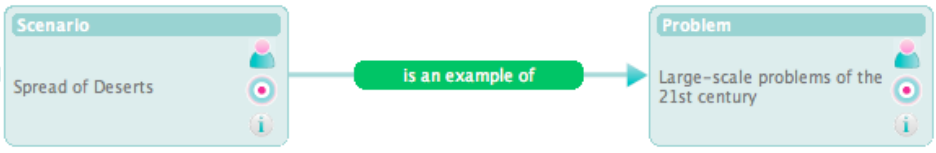

Figure 1: Example of a user-defined connection between two Ideas in the Cohere system

In Cohere, users are free to enter any text as an Idea and its detailed description. The examples seeding the database convey implicitly that Idea labels are generally short and succinct. Users are encouraged by the user interface to reuse existing Ideas, with an autocomplete menu dropping down as they type to show matching Ideas already published: as far as possible, we want them to describe the same Idea using the same label.

Users must, however, constrain their contributions by:

- creating labelled connections between Ideas (e.g. *is an example of*)
- reusing, or creating, a connection from a list of either positive, neutral or negative connections

Users can optionally:

- assign roles to Ideas (e.g. *Scenario; Problem*)
- add descriptive details (displayed when the *Info* icon is clicked)
- assign websites to Ideas (listed when the Idea is selected)

The Cohere data model is inherited from the ClaiMaker prototype [11]. The essence is summarised informally in Figure 2.

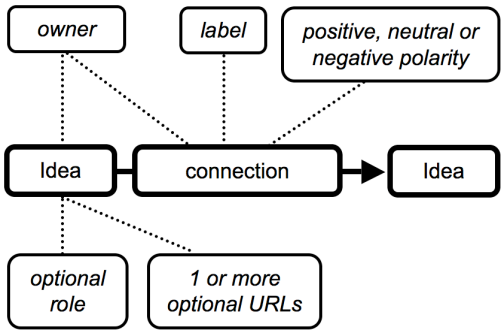

Figure 2: Cohere's data model

The provision of mechanisms to enable flexible linking of web resources around what we are calling Ideas is a goal shared by the Topic Maps community [26], whose data model is very close to that of Cohere. Intruigingly, the two were developed entirely independently, yet arrived at the same core data model, which we take to be a form of empirical validation. In the more formally defined, and more wide-ranging Topic Map Data Model, *topics* (=Cohere Ideas) point to one or more *resources* (=websites); topics can be linked by *associations* (=connections), and topics may play one or more *roles* within a given association (=roles). A Web 2.0 application called Fuzzzy [27] is based on the Topic Map standard and shares some similarities with Cohere, as does the HyperTopic system [28]; neither, however, provide support for argumentation.

While not mandating that the user engage in argumentation, the language of deliberation and argument is nonetheless at the heart of Cohere: (i) the *roles* that Ideas can play in a connection include the IBIS scheme's *Question, Answer, Pro, Con* and the user can define new ones (e.g. *Datum, Claim, Warrant* for Toulmin); (ii) the connection types offered to users are clustered by positive, neutral or negative polarity, with defaults including discourse moves such as *proves, is consistent with, challenges, refutes*. These default connection types are also leveraged in the predefined visualization filters offered by the Connection Net tab, described later (Figure 6). While the interface makes it clear that users may choose to ignore the defaults and create their own connection language, and the roles Ideas can play, the fact that all connections are classed as broadly positive, neutral or negative provides a way to express not only disagreement in the world of discourse, but could signify inhibitory influence (e.g. in biological or ecological systems modelling), or antagonistic

relationships (e.g. in social networks). It is entirely up to the individual or team to define their modelling scheme.

4.2. Visualizing IBIS-Based Dialogue Mapping in Cohere

The default roles that an Idea can play in a connection are Questions, Answers, Pros and Cons, derived from the Issue-Based Information System (IBIS) developed by Rittel [29], and implemented in the Compendium tool referred to earlier. This is used to model what Walton and Krabbe [30] classified as deliberation dialogues over the pros and cons of possible courses of action to address a dilemma.

Our previous work has demonstrated the value of real time IBIS dialogue mapping in meetings, and the use of IBIS as an organising scheme around which an analyst can map, asynchronously, the structure of public policy debates which can then be published as read-only maps on the Web [31]. Cohere now provides a platform for collaborative deliberation and debate mapping over the internet, with primarily asynchronous use in mind to start with. (Real time mapping requires a tool like Compendium which has a user interface optimised for rapid mapping. However, it is our intention to optimise for real time mapping in the longer term, perhaps by adapting Compendium as an applet for Cohere).

4.3. Visualizing Argumentation Schemes and Critical Questions in Cohere

In related work [32], we have demonstrated how Walton's argumentation schemes and associated Critical Questions, rendered as XML files in the Argument Markup Language [33], can be transformed into Compendium XML and expressed as IBIS structures in Compendium. The resulting argumentation scheme templates can now be modelled in Cohere as illustrated in Figure 3.

4.4. Social Networking and Reputation

All Ideas other than one's own have their owner clearly indicated iconically. Clicking this displays the user profile, making it possible to learn more about the person behind the ideas. We are beginning to add metrics to make users aware when they arrive at the site how many active users there are, and what the most recently posted, reused and linked Ideas are. Web feeds in the future will enable users to request notification whenever one of their Ideas is embedded in someone else's connection, or in someone else's website (see below).

4.5. Interoperability: Web Data as Platform

Central to O'Reilly's notion of Web 2.0 is the notion of web data as the platform on which many applications can compute. Cohere exposes and consumes data in a variety of ways:

- Publishing and importing XML Web feeds
- Importing XML data from the Compendium offline dialogue and argument mapping tool
- Embedding pointers to its data in other applications as URLs and HTML 'snippets'
- Exposing data in a variety of standards to engage different communities

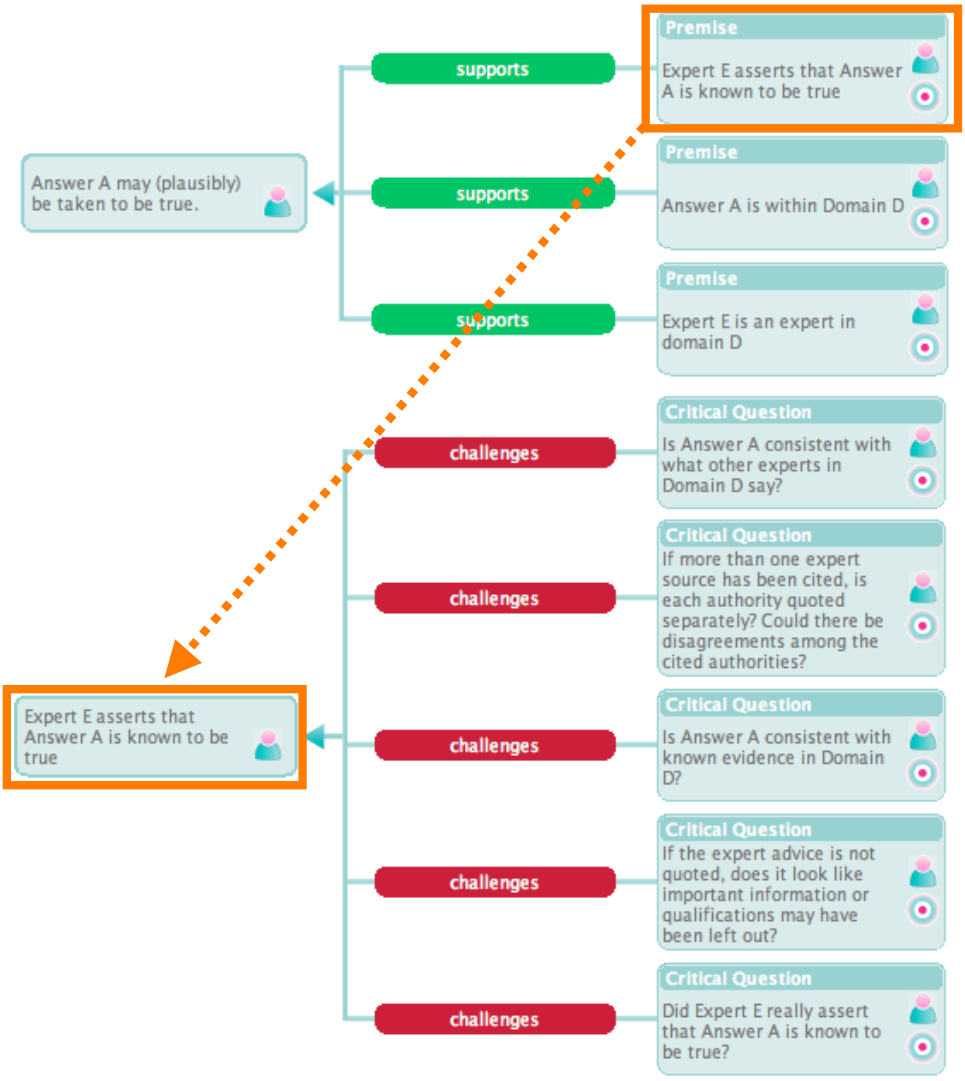

Figure 3: Rendering Walton's Critical Questions on the *Argument from Expert Opinion* scheme, as an IBIS

Web feeds: Cohere seeks to build on the significant effort that many users already invest in social bookmarking with folksonomic tagging tools such as del.icio.us, or in blogging with tools such as Blogger or Wordpress. These are currently two of the most dominant ways in which users share their views, and Cohere aims to leverage this by enabling users to import/refresh the Web feed (RSS or Atom) for any bookmarking or blogging site. Entries are converted into *Ideas* and annotated with the relevant URL, ready for linking. We are in the process of implementing an RSS feed so that users can track new Ideas as they are published. We plan to develop this capability, so that individual Ideas can be subscribed to, with alerts everytime someone connects to or from them.

Ideas and views as URLs: It is increasingly hard to find an artifact or building these days without a URL on it. The web depends on the URL as a way for non-

technical users to connect web pages, save searches, and disseminate sites of interest via standard tools such as email, slides and wordprocessors. The design of URLs goes beyond cool top level domain names, to the art of URLs that communicate their content to people, in contrast to machine-generated addresses that have no obvious pattern.

It was considered essential, therefore, to make Cohere's content addressable and accessible as URLs. This required the creation of a guest login status for non-registered users to successfully reach an address, and the design of a URL syntax that specified the visualization type and search criteria. The URL for an Idea, a triple, or a Connection List/Net is accessed by the user in what has become the conventional manner, by clicking on a URL icon to copy and paste the address that pops up.

Embedding ideas and views in other websites: Once a URL addressing scheme is in place, it becomes possible to provide such embeddable snippets for users, as introduced above. Pasting this <iframe> code into a web page creates an embedded, interactive view onto the Cohere database, which reproduces the buttons to get the URL and snippet code, to encourage further dissemination (Figure 4).

Figure 4: A Cohere connection embedded as a snippet in another web page. The three buttons below take the user to the connection within the Cohere website, provide the URL to this link, and provide the HTML embed code. Users can embed single Ideas, or whole interactive maps.

Multiple Import/Export data formats: By the time of the conference, we will have implemented further data formats for importing and exporting Cohere structures. A priority is to provide Argument Interchange Format compatibility, with other candidates being Topic Maps, Conceptual Graphs, and OWL.

4.6. Mashup Visualizations

Our objective is to help forge links not only between Ideas, but between the people publishing them. As Cohere starts to be used, it is inevitable that popular Ideas will be duplicated: if the site is successful, we can expect many people to be working on the Idea *Global Warming*, or making reference to everyday concepts such as *Capitalism* or *World Wide Web*. We have started to design views that help render the structures that will result from many users working on common Ideas. This is a long term challenge, but Figure 5 shows the first tool called Connection Net, which uses a self-organising graph layout algorithm that can render all of one's personal data, or filtered views of the world's data. In particular, Ideas with a border are used by more than one person, and as shown, right-clicking on it enables the user to view all the owners of that Idea. In this way, just as folksonomies enable disparate users to discover related resources and people, Cohere aims to reveal new connections and users working on the same Idea, or perhaps more interestingly, making similar or contrasting connections.

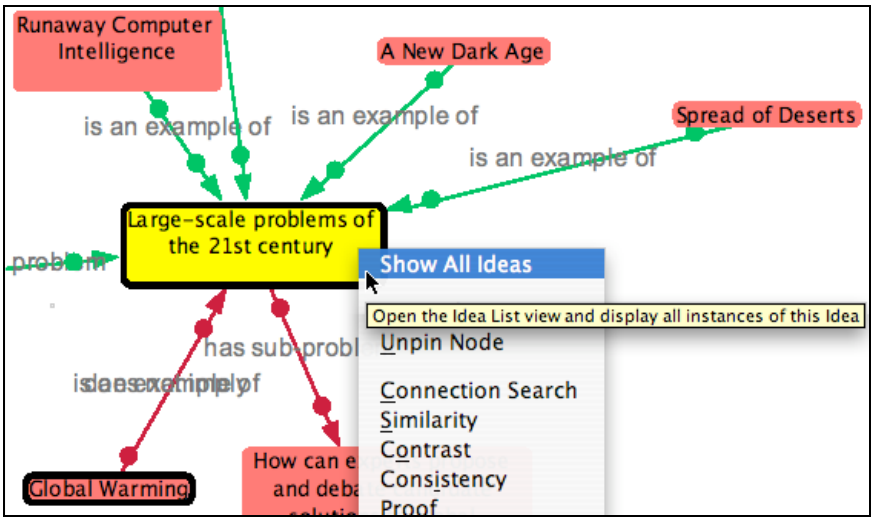

Figure 5: The Connection Net view merges all matching Ideas in a single node, and lays out the graph automatically

Filter buttons in the Connection Net view make use of the connection types, as shown in Figure 6. A number of saved filters are shown, for example, *Contrast* searches the database from a focal Idea on a specific subset of connection types of a contrasting nature, e.g. *challenges, has counterexample, is inconsistent with, refutes.* Users can define their own custom searches, and in the future will be able to save them as shown in the example buttons.

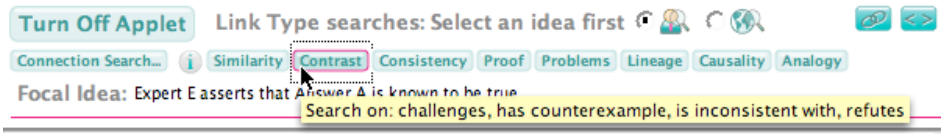

Figure 6: Semantic filter buttons that show only a subset of connection types from a focal Idea. The example shown is a Contrast search: rolling over it displays the connection semantics it will search on. User defined searches are issued from the Connection Search button on the left.

4.7. Implementation

Cohere is implemented on Linux, Apache HTTP server, MySQL database, and PHP. The user interface exploits the AJAX approach to caching data in the browser to create a highly responsive interface, with few delays between action and feedback. Cascading Style Sheets are used extensively to control presentation. In addition, a Java applet from the Prefuse information visualization classes [34] has been customised to provide self-organising, interactive graph visualizations under the Connection Net tab. Compendium (op cit) serves as an offline mapping tool (a cross-platform Java desktop application with Apache Derby or MySQL database). Data is uploaded to Cohere currently using the Compendium XML scheme. Cohere is currently a freely hosted application, and an open source release is planned by end of 2008.

5. Present Limitations, and Future Work

This project is tackling a complex but important challenge: to create tools providing a compelling user experience by harnessing two forces that seem on first inspection to pull in opposite directions: on the one hand, informal social media with low entry thresholds and few interaction constraints, and on the other, mechanisms for structuring ideas and discourse. We have presented the design rationale behind Cohere, a web application for structuring and visualizing information and arguments, publishing Ideas, and discovering new intellectual peers. In order to balance the informal+formal design criteria, we bias to the informal, with interface and architectural mechanisms to add structure as desired. Understandably, Cohere is being trialled initially by individuals testing it for personal reflection and information management, but ultimately, we hope to augment distributed communities engaged in intentional, collective sensemaking.

We are now in disussion with other COMMA research groups to explore how their argument modelling approaches could be integrated with Cohere. We are moving to a web services architecture [cf. 35] and plan to enable data exchange via the W3C's RDF and OWL, and the proposed Argument Interchange Format [25]. These developments are designed to evolve Cohere from a closed web application, towards a collaboration platform for structured, social argumentation and deliberation for wider experimentation by both end-users, developers and argumentation researchers. Future technical work will support different argument layouts, more flexible visualizations, group permissions, and the management of template 'pattern libraries' (currently managed via the offline Compendium tool). Our pilot usability evaluations are leading to interface changes that are being added at the time of writing, and will be followed by more in depth studies in the lab (cf. [12]), and with end-user communities.

Finally, our goal with Cohere is to provide an environment for the emergence of social and argument structures. While not currently exposed in the user interface, Cohere has built into its data model the Cognitive Coherence Relations modelling scheme described in the COMMA'08 debate modelling paper by Benn *et al.* [36], which seeks an integrated approach to modelling social networks and argument networks. A future objective is to investigate this modelling paradigm within Cohere.

Acknowledgements: Cohere is implemented by Michelle Bachler, with contributions from Michele Pasin and Alex Little, to whom I am indebted. The *Open Sensemaking Communities* Project [www.kmi.open.ac.uk/projects/osc] is funded by the William and Flora Hewlett Foundation, as part of the Open University's *OpenLearn* Initiative [www.open.ac.uk/openlearn].

6. References

[1] Sierhuis, M. and Buckingham Shum, S. (2008). Human-Agent Knowledge Cartography for e-Science. In: *Knowledge Cartography,* (Eds.) Okada, A., Buckingham Shum, S. and Sherborne, T. Springer

[2] Tate, A., Buckingham Shum, S., Dalton, J., Mancini, C. and Selvin, A. (2006). Co-OPR: Design and Evaluation of Collaborative Sensemaking and Planning Tools for Personnel Recovery. *Technical Report KMI-06-07* , Open University, UK. www.kmi.open.ac.uk/publications/pdf/KMI-TR-06-07.pdf

[3] O'Reilly, T. and Battelle, J. (2004). The State of the Internet Industry. Opening Presentation, *Web 2.0 Conference*, Oct. 5-7 2004, San Francisco, CA.

[4] Garrett, J.J. (2005). Ajax: A New Approach to Web Applications. *Adaptive Path* (Feb. 18, 2005): www.adaptivepath.com/ideas/essays/archives/000385.php

[5] Golder, S. and Huberman, B.A. (2006). Usage Patterns of Collaborative Tagging Systems. *Journal of Information Science*, 32(2). 198-208.

[6] See for instance the *Crisis in Darfur* project by the The United States Holocaust Memorial Museum and Google Earth: www.ushmm.org/googleearth

[7] On argument visualization see for instance the volume in [15] and the recent conference *Graphic and Visual Representations of Evidence and Inference in Legal Settings*: www.tillers.net/conference.html

[8] *Argumentation Support Systems for eParticipation.* EU DEMO-net Project: www.demo-net.org

[9] Andriessen, J., Baker, M. and Suthers, D. (2003). *Arguing to Learn*. Kluwer

[10] Verheij, B. (2005). *Virtual Arguments. On the Design of Argument Assistants for Lawyers and Other Arguers.* T.M.C. Asser Press, The Hague

[11] Lowrance, J. (2007) Graphical Manipulation of Evidence in Structured Arguments. *Oxford Journal of Law, Probability and Risk*

[12] Uren, V., Buckingham Shum, S.J., Li, G. and Bachler, M. (2006) Sensemaking Tools for Understanding Research Literatures: Design, Implementation and User Evaluation. *Int. Jnl. Human Computer Studies*, **64**, (5), 420-445.

[13] Reed, C., Walton, D. & Macagno, F. (2007). Argument diagramming in logic, law and artificial intelligence. *Knowledge Engineering Review*, **22** (1), 87-109.

[14] Conklin, J. (2006). *Dialogue Mapping.* Chichester: Wiley

[15] Selvin, A. (2003). Fostering Collective Intelligence: Helping Groups Use Visualized Argumentation. In *Visualizing Argumentation*, (Eds.) P. Kirschner, S. Buckingham Shum, and C. Carr. Springer: London

[16] van Gelder, T. (2007). The Rationale for RationaleTM. *Oxford Journal of Law, Probability and Risk.*

[17] *Debatepedia*: International Debate Education Association: wiki.idebate.org

[18] Karacapilidis, N. and Tzagarakis, M. (2007). Web-based collaboration and decision making support: A multi-disciplinary approach. *Int. J. Web-Based Learning and Teaching Technologies*, 2, (4), 12-23.

[19] Atkinson, K., Bench-Capon, T. and McBurney, P. (2006): PARMENIDES: Facilitating deliberation in democracies. In: *Artificial Intelligence and Law*, **14** (4), pp. 261-275.

[20] Buckingham Shum, S.J., Uren, V., Li, G., Sereno, B. and Mancini, C. (2007). Modelling Naturalistic Argumentation in Research Literatures: Representation and Interaction Design Issues. *International Journal of Intelligent Systems*, **22**, (1), pp.17-47.

[21] Sereno, B., Buckingham Shum, S. and Motta, E. (2007). Formalization, User Strategy and Interaction Design: Users' Behaviour with Discourse Tagging Semantics. Workshop on Social and Collaborative Construction of Structured Knowledge, *16th Int. World Wide Web Conference*, Banff, AB, Canada

[22] *TruthMapping*: www.truthmapping.com

[23] *DebateMapper*: www.debatemapper.com

[24] Rahwan, I., Zablith, F. and Reed, C. (2007). Laying the Foundations for a World Wide Argument Web. Artificial Intelligence, **171**, (10-15), pp 897-921

[25] Chesnevar, C., McGinnis, J., Modgil, S., Rahwan, I., Reed, C., Simari, G., South, M., Vreeswijk, G., Willmott, S. (2006) Towards an Argument Interchange Format, *Knowledge Engineering Review*, **21** (4), 293-316.

[26] ISO 13250-2: Topic Maps — Data Model: www.isotopicmaps.org/sam

[27] Lachica, R. and Karabeg, D. (2007). Towards holistic knowledge creation and interchange Part I: Socio-semantic collaborative tagging. *Proc. Third International Conference on Topic Maps Research and Application*, Leipzig. Lecture Notes in Artificial Intelligence, Springer: Berlin

[28] Zaher, L., Cahier, J-P., and Guittard, C. (2007) Cooperative Building of Multi-Points of view Topic Maps with Hypertopic. *Proc. Third Int. Conf. Topic Maps Research and Application*, Leipzig. Springer

[29] Rittel, H.W.J., Second Generation Design Methods. In: *Developments in Design Methodology*, N. Cross (Ed.), 1984, pp. 317-327, J. Wiley & Sons: Chichester

[30] Walton, D.N. and Krabbe, E.C.W. (1995). *Commitment in Dialogue*, SUNY Press.

[31] Ohl, R. (2008). Argument Visualisation Modelling in Consultative Democracy around Wicked Problems In: *Knowledge Cartography,* (Eds.) Okada, A., Buckingham Shum, S. and Sherborne, T. Springer: London

[32] Buckingham Shum, S. and Okada, A. (2008, in press). Knowledge Cartography for Controversies: The Iraq Debate. In: *Knowledge Cartography,* (Eds.) Okada, A., Buckingham Shum, S. and Sherborne, T. Springer: London

[33] Walton, D.N. & Reed, C.A. (2002) Diagramming, Argumentation Schemes and Critical Questions. *Proc. 5th International Conference on Argumentation (ISSA'2002)*, SicSat, Amsterdam, pp. 881-885.

[34] *Prefuse*: interactive information visualization toolkit: www.prefuse.org

[35] EU *ASPIC* Project: *Argument Service Platform with Integrated Components*: www.argumentation.org

[36] Benn, N., Buckingham Shum, S., Domingue, J. and Mancini, C. (2008). Ontological Foundations for Scholarly Debate Mapping Technology. *Proc. 2nd Int. Conf. on Computational Models of Argument*. Toulouse, May 2008. IOS Press: Amsterdam.