

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

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 programme¹ 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

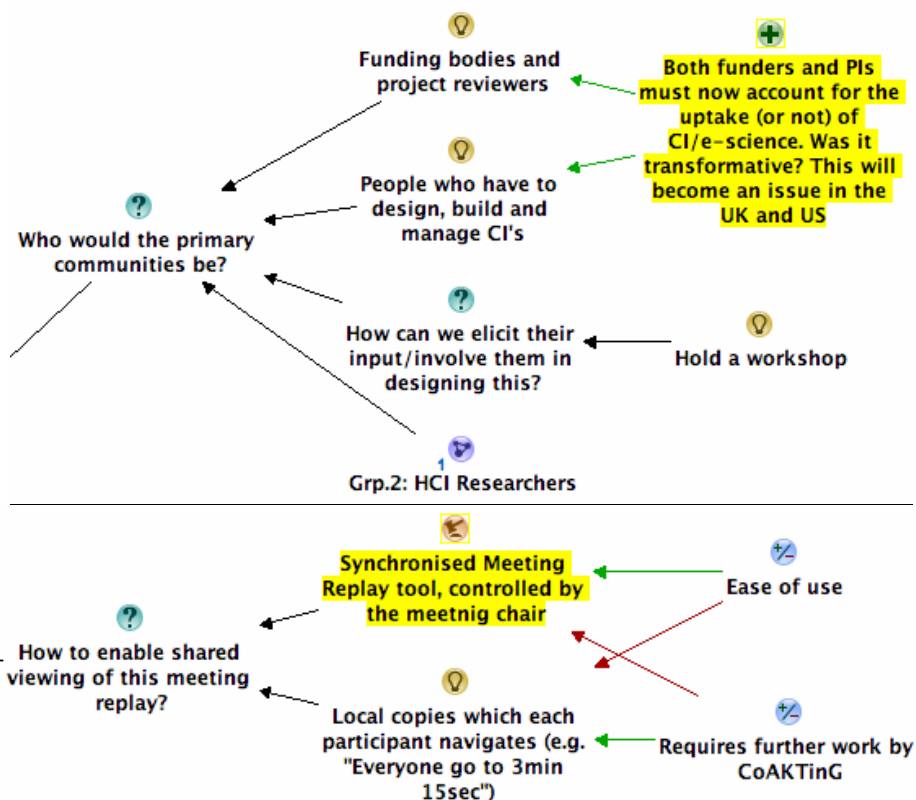


Fig. 1. Fragments from two Dialogue Maps using IBIS. In the top example exploring requirements for a website, a Pro argument of a political nature is highlighted, backing two Idea nodes. In the lower example, a QOC-style design discussion examines Option tradeoffs against more formally expressed design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

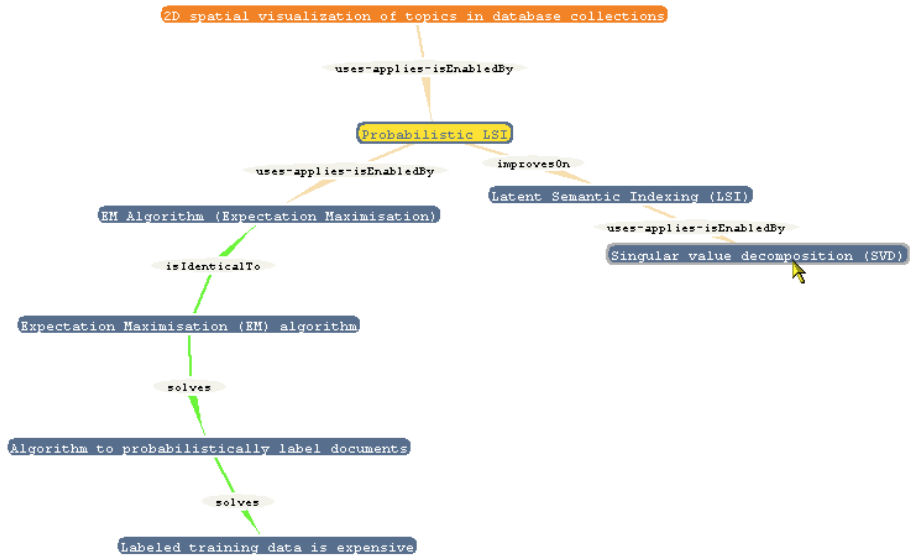


Fig. 6. ClaimFinder’s Lineage query traces the ‘intellectual roots’ of a concept, displayed at the top. The conceptual graph is analysed and filtered to show potentially significant relational types such as *uses/applies/is enabled by*, *improves on*, and *solves*.

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 resources, 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, Mahwah, 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)