

Sensemaking on the Pragmatic Web: A Hypermedia Discourse Perspective

Simon J. Buckingham Shum
Knowledge Media Institute & Centre for Research in Computing
The Open University, Milton Keynes, MK7 6AA, UK
<http://kmi.open.ac.uk/people/sbs>

Abstract: The complexity of the dilemmas we face on an organizational, societal and global scale forces us into sensemaking activity. We need tools for expressing and contesting perspectives 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. This has been the motivation for a programme of basic and applied action research into *Hypermedia Discourse*, which draws on research in hypertext, information visualization, argumentation, modelling, and meeting facilitation. This paper proposes that this strand of work shares a key principle behind the Pragmatic Web concept, namely, the need to take seriously diverse perspectives and the processes of meaning negotiation. Moreover, it is argued that the hypermedia discourse tools described instantiate this principle in practical tools which permit end-user control over modelling approaches in the absence of consensus.

1 Introduction

The complexity of the dilemmas we face at an organizational, societal and global scale forces us into *sensemaking* activity. To manage such complexity, this paper argues that there is an important role for discourse-oriented tools to help capture, comprehend, integrate and manage competing interpretations and arguments for action. There is a particular need to provide languages for communities to express and contest perspectives in principled ways, using tools that are flexible enough to mediate, capture and add value to discourse between stakeholders, yet introduce sufficient structure to provide computational services without straitjacketing the discourse. This has motivated a research programme 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 in scenarios of extended, asynchronous discourse lasting from a few days to months or years (projects; research discourse).

This paper seeks to clarify the intellectual and technical threads with which we work, in the hope that they add some distinctive colour and texture. In a nutshell, we seek to provide practical tools for discourse communities to add their own interpretive layer over whatever media, syntax and predefined semantics they may be using. This agenda intersects with the concept of the Pragmatic Web, as recently

articulated [SDD06]. There is, we suggest, common ground in our understanding of how the world is, which motivates requirements for a new kind of representational palette: placing centre-stage (rather than finessing) the pragmatics of meaning negotiation, for instance, in order to adapt (rather than unproblematically “re-use”) ontologies in different contexts; the reality of diverse perspectives (hard or impossible to formally model) when it comes to tackling complex socio-technical problems; the importance of tools affording end-user manipulation to reflect perspectives; the collective evolution of semantics in situ rather than being straitjacketed by an abstract worldview.

The paper is organised as follows. I start by reflecting briefly on the global and organizational context in which new approaches to augmenting discourse must operate. This motivates the need for tools to assist with sensemaking in socially complex scenarios (§3), in particular, to manage discourse when tackling wicked problems. 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 (§4). The Compendium methodology and tool is then introduced as a relatively mature exemplar (§5), followed by the second example of ClaiMaker (§6). §7 reflects briefly on the value of the Hypermedia Discourse concept, before concluding with discussion of how this relates to other work articulating a Pragmatic Web orientation (§8).

2 The context in which we find ourselves...

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 [LW91] and other networks of people who “make and take perspectives” [BT95]. Later, two discourse representation schemes are introduced.

2.1 Discourse: global context

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 those of us working on the next generation of interactive, computational aids for managing information, ideas and arguments. 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.

2.2 Discourse: organizational context

“Meetings are where organizations come together. (They) remain the essential mechanism through which organizations create and maintain

the practical activity of organizing. They are, in other words, the interaction order of management, the occasioned expression of management-in-action, that very social action through which institutions produce and reproduce themselves.” [Bo94]

Given the centrality of meetings in knowledge work, they provide a context in which tools for mediating discourse could make a difference. Some meetings can, of course, revolve around issues of identity and status that may lead participants to undermine tools which they do not think will play to their strengths. That being said, not all meetings are so politically charged. In the knowledge-intensive organizations with which we work, the focus is typically on the practical challenges of improving discourse across disciplines, organisational boundaries, geography and timezones. Coupled with the common complaints about ineffectual meetings, sterile debates and ubiquitous information overload, there is much interest in how things might change for the better.

3 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. 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.” [We95] p.6

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.” [We95], 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?

3.1 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 [Bu94] 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." [Ri72]

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." [Ri72]

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" [En63]. 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.¹ However – if done appro-

¹ 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.

priately – 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 is 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 tool 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.

- **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 modellers with research interests. 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*.

We turn now to two instantiations of the hypermedia discourse concept: Compendium and ClaiMaker.

5 Compendium

Compendium² is a tool 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 can be used as a personal tool, or asynchronously in a group, or in what is arguably the most demanding context of deployment for a knowledge representation tool, real time collaborative modelling.

5.1 Ontology

Compendium is a direct descendent of Conklin’s gIBIS prototype [CB88] 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) [MBY91]. Compendium has therefore provided a longitudinal case study to reflect on issues of knowledge technology adoption and practice [BSS06]. The representational 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

² The software is a free Java application for all platforms, including the source code, available from the Compendium Institute: www.CompendiumInstitute.org

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 ‘folksonomies’. 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

Compendium renders IBIS structures as a directed graph, with a root issue on the left or at the top, 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. Two discourse modelling methodologies have developed around the capabilities of Compendium. *Dialogue Mapping* is a set of skills developed by Conklin [Co05] 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 (Figure 1).

Conversational Modelling [Se99] 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. Figure 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). 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. Such templates are essentially *patterns* that can seed conversations of different sorts with issues, ideas, arguments, constraints and information. 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.

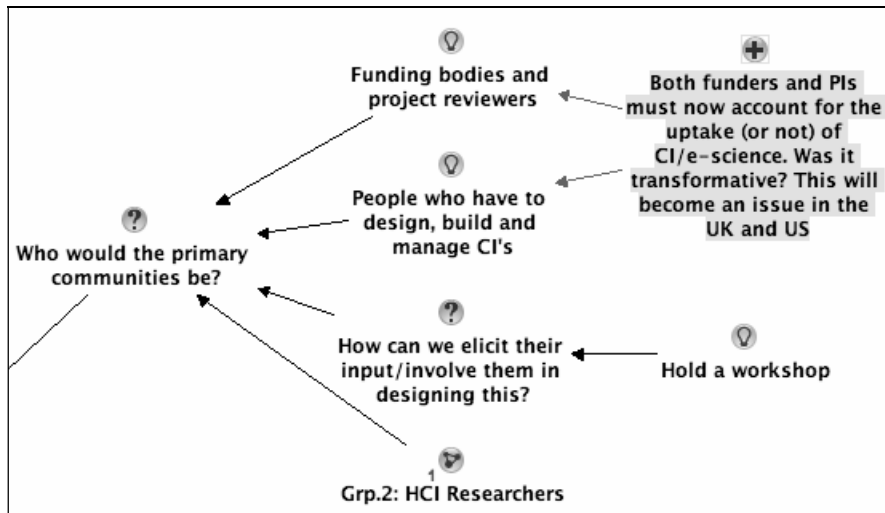


Figure 1: Fragment from a Dialogue Map exploring requirements for a website. Discussion is organised around the issue *Who would the primary communities be?* A Pro argument of a political nature is highlighted, backing two nodes representing possible Ideas/Positions/Options responding to the issue.

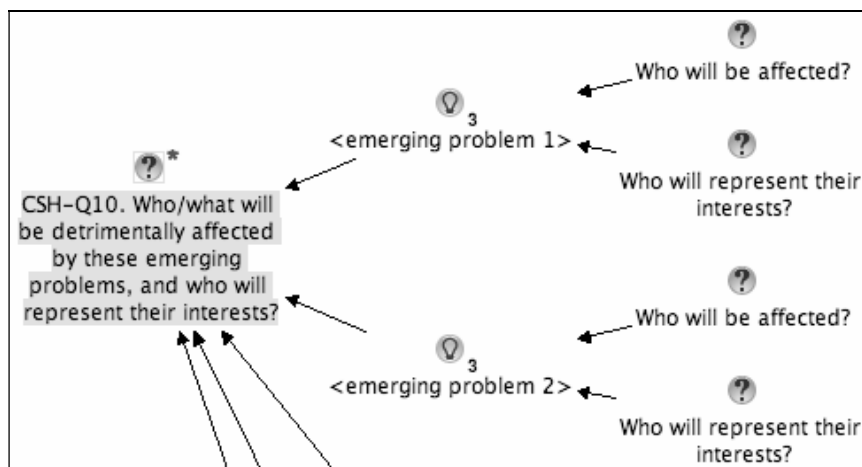


Figure 2: An Issue-template used in Conversational Modelling. For each answer, there are two subsequent Issues.

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.

From a more formal knowledge representation perspective, we represent semantics using a variety of conventions. In a NASA field trial (Figure 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.

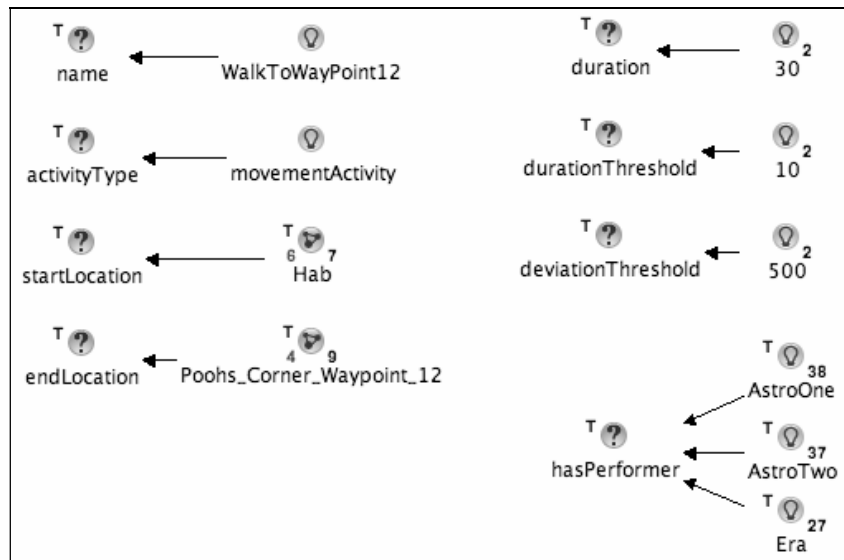


Figure 3: Science metadata in Compendium, negotiated by humans but readable by agents. Systematic tagging and structure are used to specialise the nodes. There could be a Dialogue Map behind the answer to each “issue”, which captured important context or rationale for each element in the final model.

5.3 Intuitive user interface

We have a long list of feature requests from our user community (>10,000 downloads, with active users across all organisational sectors). However, as the preceding figures show, Compendium looks familiar to users of concept mapping or graph-editing applications, and is simple to learn. It comes with IBIS preloaded, and hypermedia functionality which makes it simple to create navigational links to a given database view, and reuse a hypertext node simultaneously in different

views by copying and pasting. The 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 reported.

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 a 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 [CSA05]; the exchange of issues with a planning tool which could analyse the option space exhaustively or raise new issues [TBD05]; the export of populated issue templates to different notational formats for other stakeholders to work on [SB02]. Most recently, we have deepened the integration between Compendium and the Semantic Web, automating the exchange of Compendium data with an RDF triplestore, in order to deliver a videoconferencing capture and semantic replay tool [BSD06]. Video from meetings ‘fills in the gaps’ that terse IBIS graphs cannot possibly express; conversely, the extended-IBIS provides semantic indexing within and across meetings, enabling users to jump to the point in a meeting, for instance, when an argument was made.

5.5 Literacy and fluency

We take it as given that 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 is directed towards understanding the whole learning curve associated with reading and writing in this new medium. There are training resources to help with initial adoption of the tool, but equally, we wish 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. [CSA05][TBD05]). 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 liter-

acy in this new medium. Selvin [Se05] has begun to explore the nature of fluency in what he terms *Participatory Hypermedia Construction*. Detailed analysis of screen recordings is providing an account of the representational moves that Compendium mappers make, and the different roles they can play in meetings.

6 ClaiMaker: semantic scholarly publishing and annotation

A second instantiation of the Hypermedia Discourse concept is ClaiMaker. Unlike Compendium, which simply offers Web exports and supports the embedding of websites in IBIS conversational models, ClaiMaker was conceived from the start as a distributed Web application. Its 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 life-cycle, 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*.

We detail elsewhere [UB06] the design and evaluation of ClaiMaker and the associated suite of tools for authoring (ClaiMapper) and querying (ClaimFinder) research claims and argumentation. These are less mature than Compendium, proof of concept research tools which are not yet 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 (Figure 4).
- **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 [UB06].
- **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?* (Figure 5)

- **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. Empirical studies have demonstrated that untrained users can construct and query claim networks [SBM05][UB06].

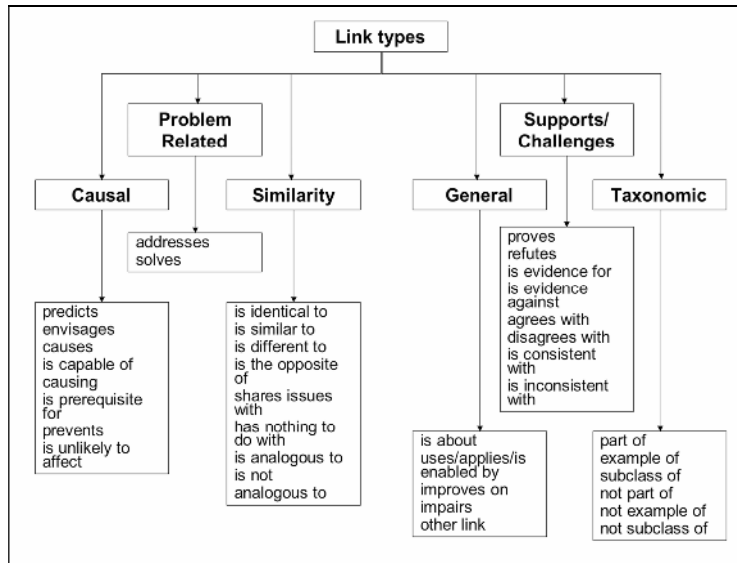


Figure 4: ClaiMaker’s discourse scheme, which groups the ‘dialect’ of a discourse community under more primitive relational classes.

7 What does Hypermedia Discourse buy us?

How does this dual focus on discourse and hypermedia make an advance? We suggest they provide insights into a number of problems of interest to the Pragmatic Research community.

Modelling in the absence of consensus. Knowledge-based systems (including for our purposes 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.

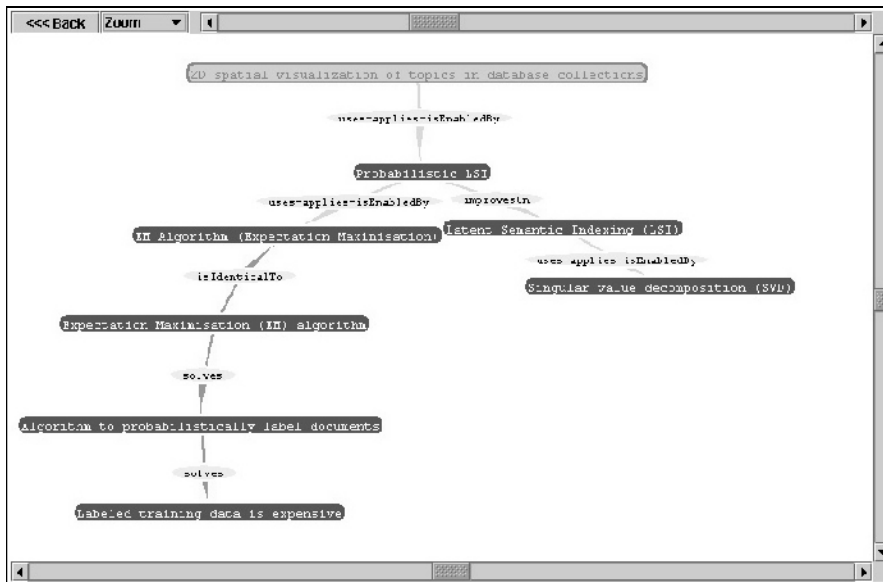


Figure 5: ClaimFinder’s *Lineage* query traces the ‘intellectual roots’ of a concept, displayed at the top. The conceptual graph is analysed and filtered to show potentially significant relational types such as *uses/applies/is enabled by*, *improves on*, and *solves*.

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 [BSS06].

Knowledge modelling from a linguistic, semiotic perspective. While there is substantial activity investigating upper level ontologies of conceptual primitives (entities, time, etc), our focus on discourse has led to work on the *relational primitives* from which domain-specific relational semantics can be derived. We adopt a relation-centric approach to reflect the process of semiosis by which meaning derives from, and shifts with, the context of use. This draws on linguistic research into Cognitive Coherence Relations to map between relational grammars across fields and modalities as diverse as scholarly discourse, cinema and text [Ma05][MB06]

8 Hypermedia Discourse and the Pragmatic Web

We have shown that the rationale for Hypermedia Discourse as a concept, and its technical implementation in one widely used tool (Compendium) and a more experimental prototype (ClaiMaker) appear to exemplify several strands in the emerging Pragmatic Web perspective: the need to co-evolve tools and work practices that recognise the realities of helping different stakeholders work together to model the world, whether in real time or at the more sedate (though increasing) pace of scholarly discourse. We suggest that the work we have been doing on what it might mean to become *literate* with tools for negotiating meaning such as Compendium [Co05][Se05], is an important part of a wholistic perspective that recognises that ‘cool tools’ alone do not change *practice* (a transition highlighted by de Moor and Aakhus [DA06]). They need to be embedded in the skillset of at least one person, and embedded in the matrix of tools that constitute the work setting.

Although we claim that Hypermedia Discourse (perspective and tools) exemplifies aspects of the Pragmatic Web, somewhat reflexively (but not surprisingly), we would also claim that they can be applied to help achieve specific practical goals that one would anticipate as relevant to the Pragmatic Web agenda. For example, Compendium can be used to collectively design semantic web applications, helping to merge existing ontologies into one that all stakeholders agree on. [BMS02] It can be used to map discussions in which both ‘hard’ technical data as well as ‘soft’ factors such as cultural or political factors can be integrated [TBD05].

Looking at other published work articulating a Pragmatic Web perspective, one can distinguish between the *stakeholders* (humans and agents) and the desired *objective* (which varies). Thus, we can see an interest in *agent-agent negotiation* in order to achieve *shared semantics for interoperability* [De05][Si02], while others [RS03] are interested in *human-agent negotiation* in order to create *user interfaces that adapt* to device and end-user characteristics.

Our work on Hypermedia Discourse has focused largely on *human-human negotiation of common semantics for human processing*, although in principle there are many scenarios for agents to operate on the conceptual graphs underpinning Compendium and ClaiMaker. However, we have demonstrated the use of Compendium as a medium for *human-agent collaboration* [CSA05], although the emphasis we place on the quality of the user interface to enrich human-processing of the representations (especially in real time pressured contexts) has moved us to the semi-formal rather than the formal end of knowledge representation.

To conclude, and return to the opening metaphor, this paper has described the threads with which we weave in Hypermedia Discourse. It is clear that these resonate with hues and textures in the Pragmatic Web fabric as envisaged thus far, and we hope that this work contributes to the emerging picture.

Acknowledgements. The Hypermedia Discourse concept is emerging from joint work with many valued colleagues. The evolution of Compendium has been a long term action research programme with Al Selvin, Maarten Sierhuis and Jeff Conklin, with Michelle Bachler leading the programming. Its development has been funded most recently by the UK's research councils EPSRC, ESRC and JISC. 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. My thanks to Al Selvin, Clara Mancini, Jack Park and David Kolb whose comments improved this paper.

References

- [Bo94] D. Boden. *The Business of Talk*. 1994, Polity: Cambridge.
- [BT95] R.J.J Boland, R.V. Tenkasi. Perspective Making and Perspective Taking in Communities of Knowing. *Organization Science*, 1995. 6 (4): p. 350-372.
- [Bu94] S. Buckingham Shum, N. Hammond. Argumentation-Based Design Rationale: What Use at What Cost? *Int. J. Human-Computer Studies*, 1994. 40 (4): p. 603-652.
- [BSD06] S. Buckingham Shum 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.
- [BSS06] S. Buckingham Shum et al. Hypermedia Support for Argumentation-Based Rationale: 15 Years on from gIBIS and QOC, in *Rationale Management in Software Engineering*, A. Dutoit, et al., Editors. 2006, Springer-Verlag: Berlin. p. 111-132.
- [BMS02] S. Buckingham Shum, E. Motta, J. Domingue (2002). Augmenting Design Deliberation with Compendium: The Case of Collaborative Ontology Design. in *Workshop on Facilitating Hypertext-Augmented Collaborative Modelling, ACM Hypertext Conference*. [PrePrint: <http://cognexus.org/ht02>].
- [CSA05] W.J. Clancey et al. Automating CapCom Using Mobile Agents and Robotic Assistants. *1st Space Exploration Conf.*. 2005 of Conf.. Orlando, FL: <http://eprints.aktors.org/375>
- [Co05] J. Conklin. *Dialogue Mapping: Building Shared Understanding of Wicked Problems*. 2005, Chichester: Wiley.
- [CB88] J. Conklin, M.L. Begeman, gIBIS: A Hypertext Tool for Exploratory Policy Discussion. *ACM Transactions on Office Information Systems*, 1988. 4 (6): p. 303-331.
- [DA06] A. de Moor, M. Aakhus. Argumentation Support: From Technologies to Tools. *Communications of the ACM*, 2006. 49 (3): p. 93-98.
- [De05] A. de Moor. Patterns for the Pragmatic Web, in *International Conference on Conceptual Structures*, F. Dau, M.-L. Mugnier, and G. Stumme, Editors. 2005, Springer, Lecture Notes in Computer Science (Vol. 3596). p. 1--18.
- [En63] D.C. Engelbart. A Conceptual Framework for the Augmentation of Man's Intellect, in *Vistas in Information Handling*, P. Howerton and Weeks, Editors. 1963, Spartan Books: Washington, DC: London. p. 1-29.

- [LW91] J. Lave, E. Wenger. *Situated Learning: Legitimate Peripheral Participation*. 1991, Cambridge: Cambridge University Press.
- [MBY91] A. MacLean et al. Questions, Options, and Criteria: Elements of Design Space Analysis. *Human-Computer Interaction*, 1991. 6 (3 & 4): p. 201-250.
- [Ma05] C. Mancini. *Cinematic Hypertext. Investigating a New Paradigm*. Frontiers in Artificial Intelligence and Applications. 2005, Amsterdam: IOS Press.
- [MB06] C. Mancini, S. Buckingham Shum, Modelling Discourse in Contested Domains: A Semiotic and Cognitive Framework. In Press. *Int. J. Human Computer Studies*.
- [RS03] A. Repenning, J. Sullivan. The Pragmatic Web: Agent-Based Multimodal Web Interaction with no Browser in Sight. in *Proc. INTERACT'03*. 2003 of Conference: IOS Press, IFIP, pp. 212-219.
- [Ri72] H.W.J. Rittel. Second Generation Design Methods. *Interview in: Design Methods Group 5th Anniversary Report: DMG Occasional Paper*, 1972. 1: p. 5-10.
- [SDD06] M. Schoop, A. De Moor, J. Dietz. The Pragmatic Web: A Manifesto. *Communications of the ACM*, 2006. 49 (5): p. 75-76.
- [Se99] A. Selvin. Supporting Collaborative Analysis and Design with Hypertext Functionality. *Journal of Digital Information*, 1999. 1 (4).
- [Se05] A. Selvin. Aesthetic and Ethical Implications of Participatory Hypermedia Practice. 2006, *Technical Report KMI-05-17*, Knowledge Media Institute, Open University.
- [SB02] A. Selvin, S.J. Buckingham Shum. Rapid Knowledge Construction: A Case Study in Corporate Contingency Planning Using Collaborative Hypermedia. *Knowledge and Process Management*, 2002. 9 (2): p. 119-128.
- [SBM05] B. Sereno, S. Buckingham Shum, E. Motta. ClaimSpotter: An Environment to Support Sensemaking with Knowledge Triples. *10th ACM Conf. on Intelligent User Interfaces*. 2005 of Conf.. San Diego: ACM Press: NY.
- [Si02] M.P. Singh. The Pragmatic Web. *IEEE Internet Computing*, 2002 (May/June): p. 4-5.
- [TBD05] A. Tate et al. Co-OPR: Design and Evaluation of Collaborative Sensemaking and Planning Tools for Personnel Recovery. 2006, *Technical Report KMI-06-07*, Knowledge Media Institute, Open University.
- [UB06] V. Uren et al. Sensemaking Tools for Understanding Research Literatures: Design, Implementation and User Evaluation. *Int. J. Human Computer Studies*, 2006. 64 (5): p. 420-445.
- [We95] K.E. Weick. *Sensemaking in Organizations*. 1995, Thousand Oaks, CA: Sage Publications.