



Knowledge Media Institute

Structuring Discourse for Collective Interpretation

Simon J. Buckingham Shum and Albert M. Selvin

KMI-TR-98

January, 2001

<http://kmi.open.ac.uk/tr/papers/kmi-tr-98.pdf>

Distributed Collective Practices 2000: Conference on Collective Cognition and Memory Practices.
Paris, 19-20 Sept., 2000. www.limsi.fr/WkG/PCD2000



Structuring Discourse for Collective Interpretation

Simon Buckingham Shum

Knowledge Media Institute
The Open University
Milton Keynes
MK7 6AA
UK

sbs@acm.org
www.kmi.open.ac.uk/sbs

Albert M. Selvin

Verizon
White Plains
New York
NY 10604
USA

albert.m.selvin@verizon.com

Abstract

This paper reflects on three examples of a discourse-oriented approach to supporting *collective interpretation*. By this, we mean activities involving two or more people who are trying to make sense of an issue. The common theme linking the examples is that each mediates interpretive activity via a software environment which structures discourse: participants construct their interpretation within a representational framework which in return provides computational services. As a by-product, this persistent trace of the sensemaking process can serve as a collective memory resource for subsequent *reinterpretation*. Based on the three examples, we draw attention to specific challenges that discourse-structuring technologies raise, and strategies for tackling them. A generic issue emerging from this work is the design of ontologies (representational schemes) by and for communities of practice.

Keywords: collective interpretation, collective memory, structured discourse, hypertext

1. Introduction

Digital library research addresses new possibilities for indexing and recovering physical and digital artifacts. A collective or organizational *memory* system may be seen as a digital library specifically concerned with the capture and recovery of *conceptual content* such as ideas, rationale, activity history, or lessons learned. These of course are still reified in a form that must be stored and recovered.

Current indexing approaches focus on the encoding of stable *primary content attributes* to improve retrieval and interoperability. These attributes may be manually codified according to a metadata scheme or controlled vocabulary, or extracted automatically. Inconsistencies and personal interpretations in encoding at this level are considered undesirable in order to simplify resource discovery and interoperability between networked repositories. In contrast, inconsistencies and personal interpretations are precisely the attributes that a system needs to reify and support collective *sensemaking* — the work of articulating and possibly contesting the *meaning and significance* of an artifact or idea. Our focus is on this latter class of system: computational support for constructing explicit, interpretive levels of indexing which embed artifacts and ideas in a context for subsequent interpretation. Thus, on encountering an artifact or idea, one has an additional view of how it has been understood by others.

This paper reflects on three different examples of systems for structuring interpretation-oriented discourse: journal peer review, sensemaking in meetings, and scholarly debate. The three examples *mediate and structure* discourse in particular ways, thus offering a collective memory resource for subsequent reinterpretation. In the following section, we present the rationale for focusing on *structured interpretive discourse*, before introducing the three examples, the challenges that they raise, and strategies for addressing them.

2. Structuring interpretive discourse

dis·course (dskôrs, -krs) *n.*

Verbal expression in speech or writing.

Verbal exchange; conversation.

A formal, lengthy discussion of a subject, either written or spoken.

Archaic. The process or power of reasoning.

Discourse-structuring technologies take as their point of departure the fact that the collective interpretation of an artifact or idea is invariably accomplished through discourse, taking one or more of the forms in the above definition. Given the joint challenge of supporting interpretive activity and building a collective memory resource, it seems natural to build on the particular forms of interpretive discourse that already exist. An hypothesis, therefore, is that making discourse the basis of computational support will be intuitive to end users in two respects: as the means of ‘data capture’ (since they have to communicate anyway), and in subsequent reuse (since one way to recover material will be by using conversations/writings as indices into the repository).

Another role that a discourse structuring scheme can play is to reduce complexity in order to help participants tackle an ill-structured problem systematically. An effective scheme serves as a filter on the universe of possible issues by focusing attention on a subset and providing a vocabulary in which to conduct this interpretive discourse. (Naturally we must ask who has designed the scheme and whether participants are willing to use it.) Standard email systems, whilst clearly capable of mediating and structuring interpretive discourse, would not count as discourse *structuring* systems in these terms. They provide no explicit semantics, that is, a way to classify interpretive contributions to a discourse with respect to either their *content* (what is this contribution about?) or their *rhetorical relationship* to the discourse (what level of contribution is this? how does it build on existing work?).

To summarise this distinction another way, there are two main strategies to constructing discourse-oriented memory resources, of which we are focusing on the second:

- Minimise the change in discourse practices as far as possible, by using technologies that record and extract structure from naturalistic discourse in order to make the record more reusable. Examples: analyse emails [22,42], video [18] or meeting capture data [23].
- Structure discourse to make it more effective by providing a specific vocabulary and computational services, and as a by-product, more reusable. One can constrain factors such as *what* is discussed *when*, *by whom*, *for how long*, and *how* it is reified.

Making discourse structured and persistent changes the context and nature of that discourse (e.g. see recent workshops [17]), and such interventions must be introduced with care. Structuring makes explicit certain elements of the discourse that were previously implicit. The sense of “implicit” varies depending on the technology and context of use. It may be *verbal* and ephemeral, written but *private*, or *uninterpretable* to a machine. Given these changes, a key challenge is to design the structuring technologies such that the changes enhance the discourse. Participants must *themselves* reap benefits in return for adopting new practices and engaging in more formalized discourse. As a *by-product* of the primary interaction to accomplish the immediate task, this must also seed a structured, reusable memory resource for future use and users. Ignoring this cost-benefit tradeoff can result in groupware failures of the sort identified by Grudin [15]; users cannot be expected to use a system completely altruistically for the benefit of others, or hypothetical future scenarios in which they might benefit. The three examples negotiate trade-offs on a range of human dimensions that are explored shortly.

3. Overview of the three examples

JIME: E-journal peer review

The Journal of Interactive Media in Education (JIME) is a peer reviewed, electronic journal (ejournal), published since 1996, to promote interdisciplinary dialogue through the use of a Web-based peer review process. JIME articles are published in a purpose-designed Web document-discussion user interface, which tightly links the article to an area for review comments and discussion. Reviewers can post comments under threads based on the

journal's review criteria (e.g. Originality of Ideas), or they can make section-specific comments. The review process is designed to enable authors, reviewers and the wider community engage in constructive discussion as opposed to the conventional anonymous 'issuing of a verdict'. Authors have the right of reply, and reviewers (non-anonymous) are accountable for what they say. This intellectual history is preserved with the final publication in the form of an edited version with the most significant comments and replies, which remains an open forum for authors (e.g. to post updates) and readers to comment.

There are hundreds of links in a given article, which are generated by a Web publishing toolkit called D3E (Digital Document Discourse Environment). D3E has found many applications beyond JIME, is under development on an open source basis, and is in use by a variety of groups. Elsewhere, we present detailed descriptions of the rationale underlying D3E and JIME [9,38-40].

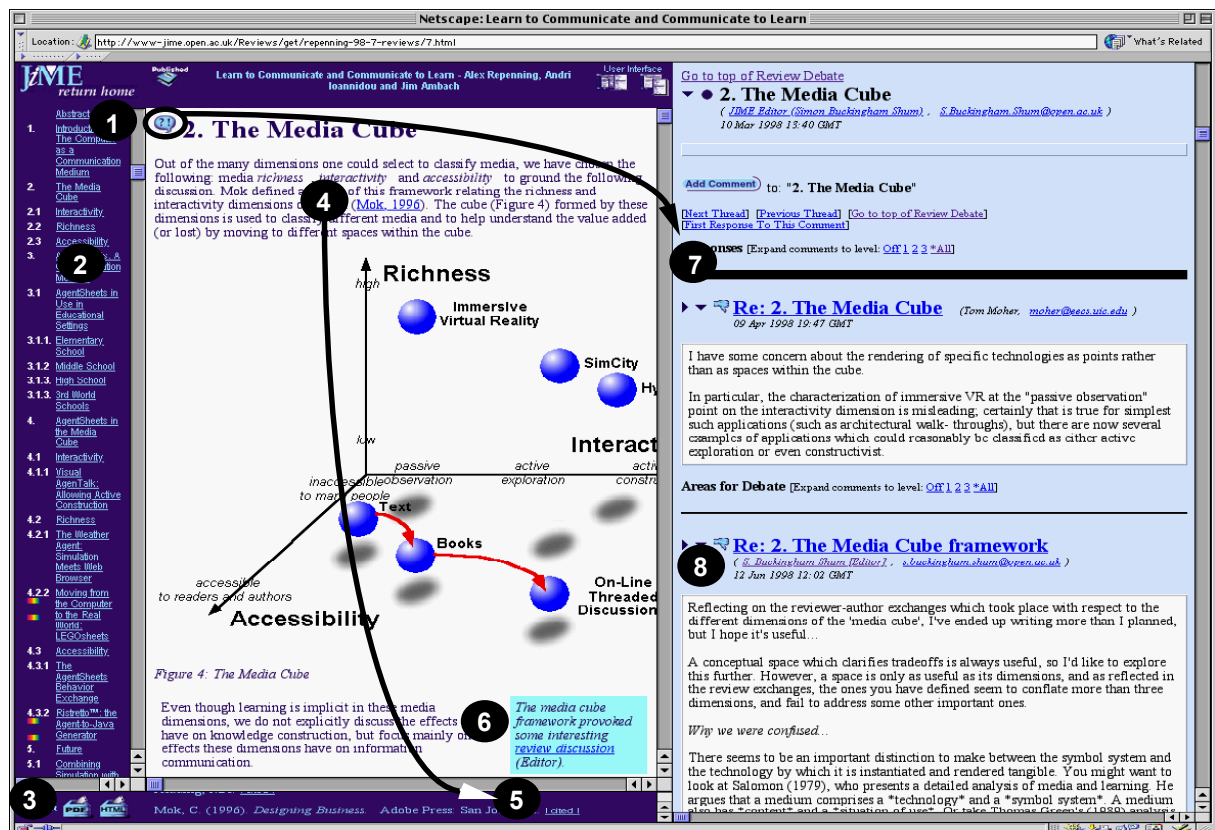


Figure 1: JIME's user interface, generated by the D3E toolkit from a submitted HTML file. On the left is the Article Window, on the right the Commentaries Window showing the top level outline view of discussion about the document. Key: [1] Comment icon embedded in each section heading; clicking displays section-specific comments; [2] active contents list extracted from the section headings; [3] print versions as HTML and PDF; [4] numeric or author/date citation automatically linked to corresponding reference in footnote window; [5] a reverse hyperlink is inserted for each citation of a reference; [6] an editorial note to draw attention to a controversial issue in the author-reviewer debate that 'made it' into the published version; [7] section-specific review comment; [8] an editorial comment summarising the review discussion and specifying change requirements. (Note that there are two versions of the user interface: one as shown, and for smaller displays, the document and discussion are placed in separate browser windows.)

Compendium: Collaborative modelling and discourse mapping

Compendium, first developed in 1993 as an approach to aid cross-functional business process redesign (BPR) teams, has been applied in several dozen projects in both industry and academic settings [28,29,31]. Its origins lie in the problems attending teams working over weeks or months to design business processes: keeping track of the plethora of ideas, issues, and conceptual interrelationships without needing to sift through piles of easel sheets, surfacing and tracking design rationale, and staying on track and "bought-in" to the project's overall structure and goals.

The set of techniques which represent the *Compendium* approach revolve around a graphical hypermedia system¹ for the development and application of (i) question-oriented *templates* which serve as semiformal ontologies to structure the subject matter of a particular project (Figure 2), and (ii) a set of *metadata tags* that can be assigned to any concept in the database (Figure 3). A hallmark of the approach is the ability to move between formal and prescribed representations and informal, ad hoc communication, incorporating both in the same view if that is helpful to the participants. Hypertext nodes and links can be added either in accordance with templates or in an opportunistic fashion.

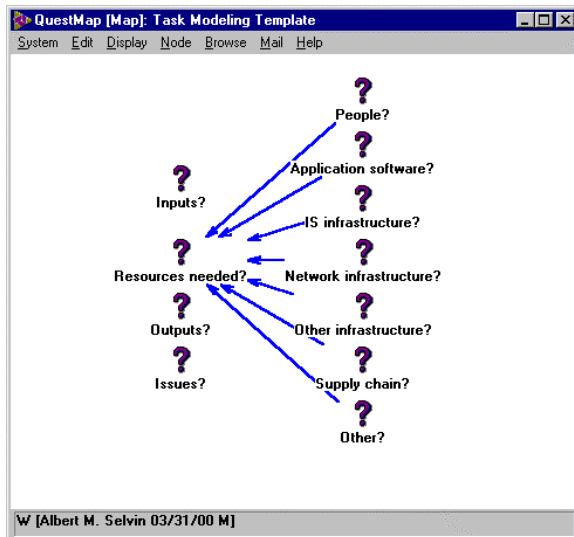


Figure 2: A Compendium question template representing the concerns of a particular stakeholder group (application in Figure 4).

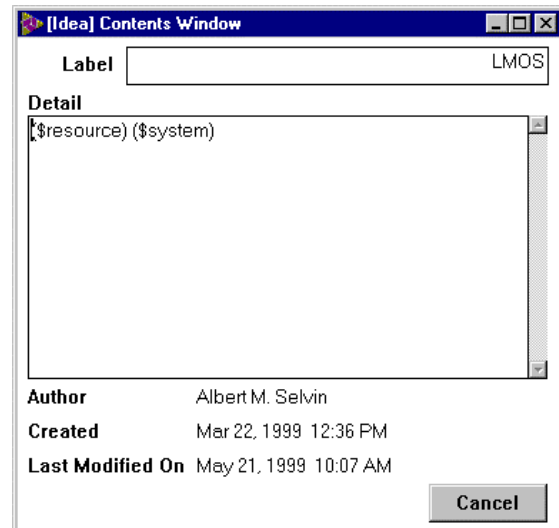


Figure 3: Optional metadata tags added to the content of a node, assisting subsequent harvesting and analysis of elements

The key feature of the early approach was the combination of an Issue-Based Information System (IBIS) concept-mapping tool [12], which supported informal and exploratory conversation and facilitation, with a structured modeling approach [30]. This allowed teams to move along the spectrum of formal to informal communication as well as prescribed to spontaneous approaches, as their needs dictated. It also let them incrementally formalize data [35] over the life of the project. As the approach was tested and refined over the course of several years, additional modelling methods were added, plus tools to transform Compendium's hypertext models into established organizational document forms, and vice-versa [32].

¹ In the examples presented in this paper, the tool is GDSS Inc.'s QuestMap™ product <www.gdss.com>. Verizon is currently developing a Java-based product with more comprehensive support for the Compendium methodology.

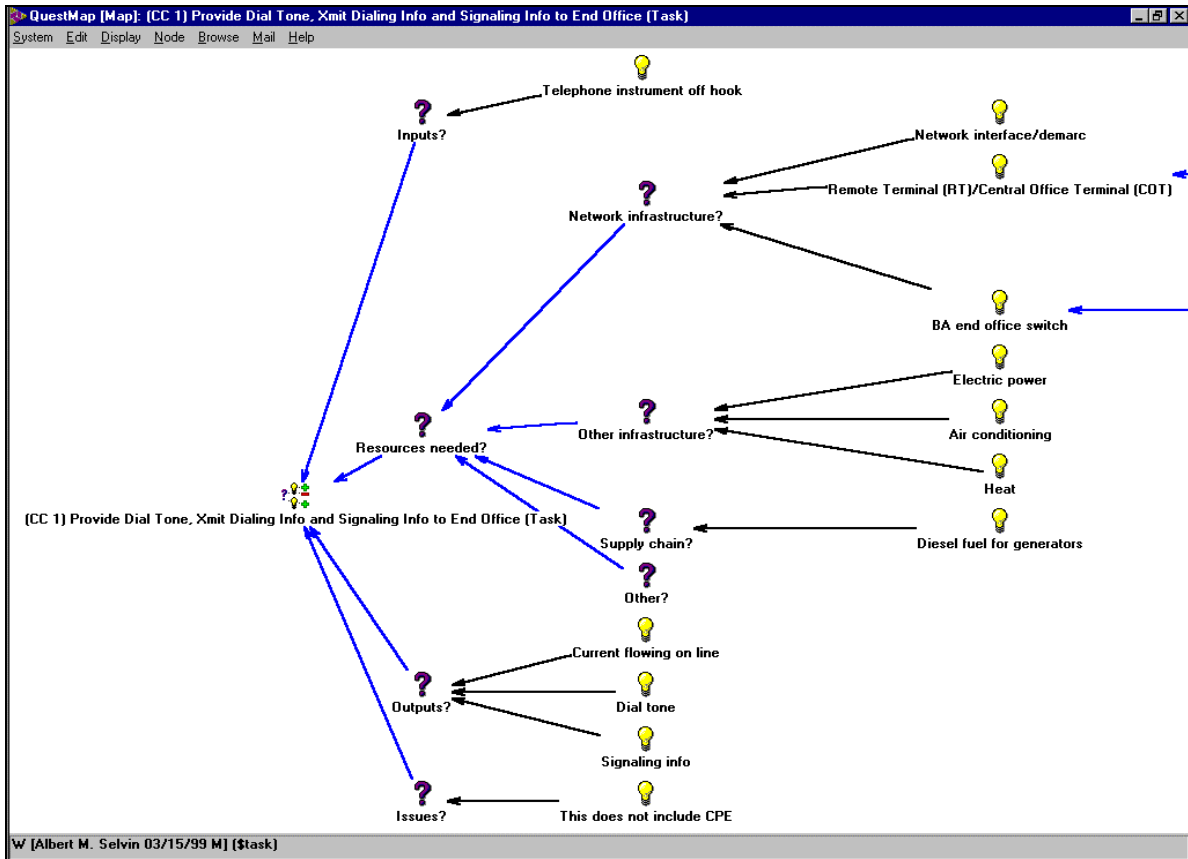


Figure 4: The Compendium template from Figure 2 being instantiated [33].

ScholOnto: Modelling research discourse in literatures

The Scholarly Ontologies (ScholOnto) Project is the youngest of the three examples, aiming to provide researchers with digital library services based on structured interpretive discourse. Researchers are typically interested in phenomena such as:

- the *intellectual lineage* of ideas: e.g. where has this come from, and has it already been done? (“Are there any arguments against the framework on which this paper builds?”)
- the *impact* of ideas: e.g. what reaction was there to this, and has anyone built on it? (“Has anyone generalised method M to another domain?” “Has anyone extended Language L?”)
- *perspectives*: are there distinctive schools of thought on this issue? (“Has anyone proposed a similar solution to Problem P but from a different theoretical perspective?”)
- *inconsistencies*: e.g. is an approach consistent with its espoused theoretical foundations?; is there contradictory evidence to a claim? (“Are there groups building on Theory T, but who contradict each other?”)
- *convergences*: are different streams of research mutually reinforcing in interesting ways? (“Who else uses Data X in their arguments?”)

Currently, there is no way to articulate such questions in a library, analogue or digital. The ScholOnto project seeks to address the fundamental requirement for an ontology capable of supporting scholarly research communities in interpreting and discussing evolving ideas: overlaying *interpretations* of content, and supporting the emergence of (possibly conflicting) *perspectives*.

Research disciplines are in constant flux and by definition lack consensus on many issues. Whilst this renders futile the idea of a ‘master ontology/taxonomy’ for a discipline, there does appear to be one stable dimension, namely scholarly discourse, by which we mean the *way* in which new work is expressed and contested. Thus, it is hard to envisage when researchers will no longer make claims about a document’s *contributions* (e.g. “this is a new theory, model, notation, software, evidence...”), or contest its *relationships* to other documents and ideas (e.g. “it applies, extends, predicts, refutes...”). Our approach provides an environment for scholars to make such claims in explicit, computer readable form (whereas normally they remain implicit in a document’s text, or undeclared). Decoupling concepts from *claims* about them is critical to supporting the emergence and co-existence of multiple *perspectives*. The kind of ontology we are working with is shown in Figure 5, suggesting concepts and relationships suitable for a wide range of disciplines. This generic scheme already enables interoperability between different domains, but we also envisage that different disciplines might re-express or specialize concept or relational types. Details and examples of modelling are presented in [8].

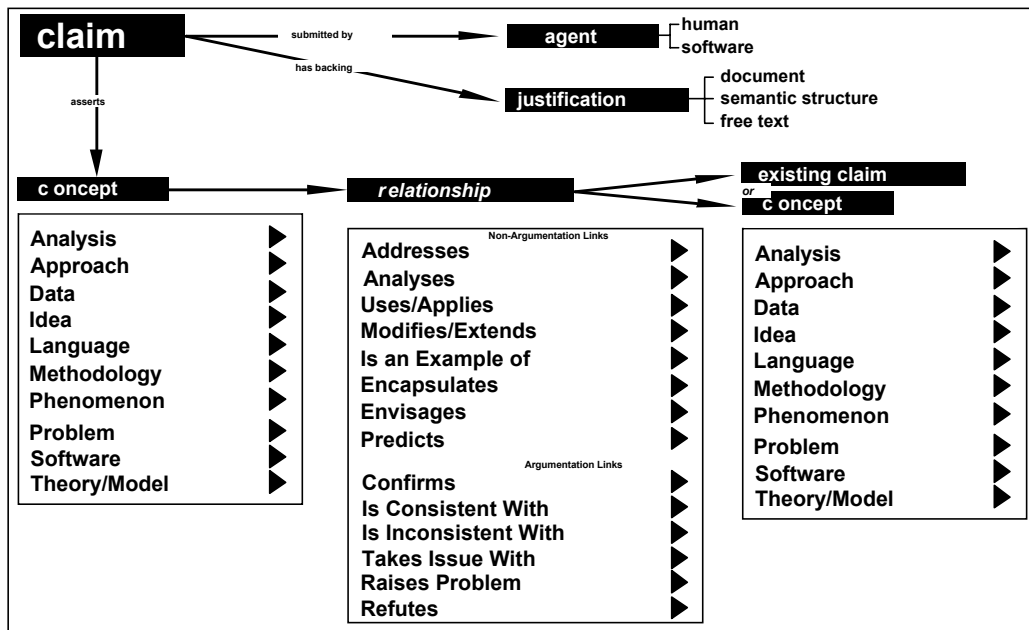


Figure 5: The structure of a scholarly *Claim* in the ScholOnto ontology. All claims are owned by an agent, and have some form of justification. Claims assert new relationships with other claims, or between concepts. It is hypothesised that such a scheme could be customised and adopted by research communities, whilst maintaining technical interoperability, and hence opening new possibilities for interdisciplinary discourse.

The ScholOnto Project is a skeletal prototype at present, but once developed further, its digital library server should deliver novel services to a research community based on its semantic model of the discourse in a literature. These services include agents and search based on semantic interest profiles, knowledge based filtering and visualizations of the literature, and the ability to analyse literatures for significant structural patterns [5].

4. Comparison of the three examples

Table 1 compares these three examples as follows:

- the collective activities for which they provide support
- their current status in terms of depth and breadth of usage
- the way in which they mediate and represent discourse
- the manner in which they function as a collective memory resources
- geographical and temporal spread of participants (same time/same place? etc.)
- social/institutional obstacles to acceptance/uptake
- cognitive obstacles to acceptance/uptake

Table 1: Summary of three technologies for structuring interpretive discourse.

Discourse-Oriented Technologies for Collective Memory			
	JIME/D3E	Compendium	ScholOnto
supports the following collective activities...	<ul style="list-style-type: none"> ● structured discussion of structured documents ● peer review and debate 	<ul style="list-style-type: none"> ● visual sensemaking ● domain modelling ● meeting facilitation ● project management 	<ul style="list-style-type: none"> ● research publishing/ dissemination ● research debate ● literature analysis
current status	<ul style="list-style-type: none"> ● JIME published since 1996 ● generic document-discussion interface used in many contexts by many groups 	<ul style="list-style-type: none"> ● evolving since 1993 ● used in >40 business projects ● usage spread within one company, and across several others 	<ul style="list-style-type: none"> ● project started 1999 ● modelling environment in place ● literature modelling underway
represents discourse as...	<ul style="list-style-type: none"> ● threaded discussions 	<ul style="list-style-type: none"> ● hypertexts rendered as concept maps 	<ul style="list-style-type: none"> ● knowledge based semantic nets constructed and rendered in different ways
functions as a collective memory resource by...	<ul style="list-style-type: none"> ● providing a document-specific interpretive space ● preserving intellectual history of document 	<ul style="list-style-type: none"> ● expressing stakeholder perspectives ● preserving contexts in which concepts are discussed ● capturing decision rationale ● tracing documents to underlying motivations/assumptions ● hypertextual linking, visual landmarks, bookmarks and guided tours through the above 	<ul style="list-style-type: none"> ● reifying a community's perceptions of a piece of work ● structuring ongoing debates ● services for searching and analysing literatures: e.g. structure; history

<p>geographical and temporal spread: discourse is conducted...</p>	<ul style="list-style-type: none"> ● asynchronously over the Net 	<ul style="list-style-type: none"> ● synchronously in meetings, or asynchronously over the Net 	<ul style="list-style-type: none"> ● asynchronously over the Net
<p>social/ institutional obstacles to acceptance/ uptake</p>	<ul style="list-style-type: none"> ● making the review process transparent ● changed roles and tone of discourse ● naming reviewers 	<ul style="list-style-type: none"> ● implications of making rationale explicit ● risk of opening up discourse that stakeholders want to keep closed ● resistance to changing habitual ways of speaking/holding meetings ● high degree of practitioner skill required to use in ill-structured contexts 	<ul style="list-style-type: none"> ● implications of making explicit claims and connections that might otherwise remain hidden or implicit ● agreement on core discourse ontology, and evolution of community-specific extensions
<p>cognitive obstacles to acceptance/ uptake</p>	<ul style="list-style-type: none"> ● learning to engage in threaded discussions with authors/ reviewers ● learning to formulate reflective contributions to debates in a timely manner 	<ul style="list-style-type: none"> ● representational literacy: discipline of chunking, naming, and linking ideas at appropriate granularity ● new group processes: e.g. learning to slow down and listen to others 	<ul style="list-style-type: none"> ● representational literacy: discipline of chunking, naming, and linking ideas at appropriate granularity

5. Socio-technical strategies for deployment

As indicated in Table 1, structuring and storing interpretive discourse creates various ‘ripples’. Several strategies are proving to be effective in the introduction of these new media. None of these strategies is especially novel in and of themselves, but are simply finding fruitful application to discourse structuring technologies. They are introduced below, before demonstrating their application to the examples (Table 2).

Practice-oriented strategies

Participatory design and evolution of discourse structuring schemes. Literacy with spoken and written language provides unlimited scope for nuanced expression. Introducing a discourse scheme in order to structure a memory resource and/or provide computational services is fraught with dangers such as poverty of expression, enforced commitment to hard categories, and inappropriate granularity of formalization (see below). Co-designing and evaluating discourse representations with the eventual users, with real deployment is the best way to satisfy evolving human and computational requirements.

Facilitators and mediators. Members of a user community rarely know how to exploit a new medium to its fullest when it first appears. Beyond learning the ‘mechanics’ of a new interface, users must learn, somehow, a new set of mappings from their intentions to expression in the new medium. For instance: How can a particular tone be set using this tool? What representational strategies are there for expressing established classes or genres of contribution? How might a codified, persistent addition to a hypertext be misinterpreted or change in meaning over time, and are there strategies for minimising such semantic drift (cf. [10])? The role of *facilitator/mediator* refers to the modelling of such practices by more users more expert in the medium, and/or the shaping of discussions through interventions. Orlikowski’s meta-structuring theory of ‘technology-use-mediation’ provides a systematic analysis of this process [25], which we have contextualised to the editorial mediation of new peer review practices in JIME [40].

Integration with existing practices. This is by now an established rubric, but one often ignored. As anthropological and other CSCW-related work reiterates, the material and social fabric of the workplace/community of practice is a finely balanced ecology. If a new tool is to carve a niche, it has not only to justify itself technically, but allow stakeholders to adjust their practices so that the often tacit forces that structure the workplace can be sustained.

Technical strategies

Significant computational ‘added value’. At the risk of stating the obvious, a new tool must deliver significant added value, notwithstanding the other factors listed. New technologies of any sort are only used over sustained periods because individuals gain personally in some way, and the computer’s role in the equation is to add some form of computational value. Without demoting the importance of interaction design, use of early personal computers, the Net, and digital libraries demonstrate that users will willingly suffer awkward interfaces and other inconveniences if a new tool provides (what is perceived as) a worthwhile increase in quality of life. On its own, raw computational power is not usually enough to guarantee adoption, but without it, other strategies lack motivation.

Integration with existing technical infrastructure. This follows from the point on integration with existing practices. As the design space of office applications converges on a small set of market leaders, and standards-based interoperability emerges to assist networked-computation and third party developers, there is less and less excuse on technical grounds for failing to integrate new technologies with the existing applications on users desktops, or providing Web/Java user interfaces that integrate into the intra/internet. It has never been easier to extend products with a new toolbar/button or menu item, or detect a user’s ID on a network rather than require yet another password. This level of integration can make all the difference in lowering the barriers to adoption.

Optional, incremental formalization. Historically, structured CMC systems (e.g. Coordinator [41]; Lens [21], and hypertext argumentation tools [13,16,27]) have encountered resistance for a variety of reasons. These include rejection due to the cognitive overhead of learning a new codification scheme [4,7,11,34], objections to limited expressiveness [20,24], and concerns over the implicit politics of categories and formalisms [2,3,37]. Research on the use of hypertext for knowledge representation [35] and design rationale/memory capture [19] has dubbed the term *incremental formalization* to refer to tools that enable users to choose when and how to add finer-grained, computer-readable codification to informal content (where “informal” means “semantically uninterpretable by machine”, which in most systems usually includes natural language, written text, and images). This strategy is adopted in all three discourse technologies in various ways.

Table 2: Practice-oriented and technical strategies applied to three discourse-structuring systems.

	JIME/D3E	Compendium	ScholOnto
	practice-oriented strategies		
Participatory design and evolution of the discourse structuring scheme	<ul style="list-style-type: none"> ● JIME peer review discourse scheme has minimal formality (optional agree/disagree icons for comments), and has not changed since journal launch ● other D3E applications introduce new categories specific to the domain 	<ul style="list-style-type: none"> ● stakeholders can understand the visual mapping notation without explanation ● template semantics are negotiated by stakeholder groups ● issue maps are created in real time for immediate verification and ownership of the record 	<ul style="list-style-type: none"> ● core ontology will be refined following public releases and use ● research groups and communities will evolve their own extensions to the core ontology
facilitation/mediation	<ul style="list-style-type: none"> ● JIME editors facilitate the review discussions and assist new editors in overseeing the review process (see [40]) 	<ul style="list-style-type: none"> ● in a well understood domain, the facilitator's main role is to capture discourse in the tool ● in more complex situations, more advanced facilitator and mediation skills (such as developing/adapting formalisms on the fly) help parties to make progress 	<ul style="list-style-type: none"> ● seed example literatures to demonstrate the range of codification and discourse possibilities ● as communities begin to use the system, more experienced members will model possible uses
integration with existing practices	<ul style="list-style-type: none"> ● the conventional textual scholarly document is the central focus, plus familiar threaded discussion ● JIME breaks from the traditional review model, but most reviewers appreciate the advantages of the new conversational model 	<ul style="list-style-type: none"> ● used to augment teams addressing real problems ● imports from, and exports to established organizational notations and document genres 	<p>possibilities:</p> <ul style="list-style-type: none"> ● integration with peer review processes (cf. JIME) ● integration with journal submission and publication (cf. integration of arXiv eprint server with physics journals [14]) ● integration into teaching contexts

technical strategies			
significant computational 'added value'	<ul style="list-style-type: none"> ● uses the Net to support scholarly debate in a reflective, but more timely, manner than possible in paper journals ● review debate on specific themes and sections is clustered together ● document and discourse integrated visually and navigationally ● authors, reviewers and wider community are alerted by email to new contributions to the debate ● time-based and interactive research data and artifacts can be assessed by reviewers 	<ul style="list-style-type: none"> ● granular reuse of hypertextual objects plus metadata supports flexible tracing and analyses of concepts ● import/export from/to other documents reduces mundane conversion tasks 	<p>possibilities:</p> <ul style="list-style-type: none"> ● knowledge-based services will assist in managing the discourse network ● semantic search agents ● semantic filtering ● literature visualizations
integration with existing technical infrastructure	<ul style="list-style-type: none"> ● accessed over the Net via a web browser, plus plugins for some interactive materials ● email alerting and submission of comments opens participation to lowest common denominator technology ● automatic HTML-conversions from most word processors can be processed by the D3E toolkit 	<ul style="list-style-type: none"> ● standard applications extended to translate to/from Compendium ● accessible over intra/internet ● exports to website 	<ul style="list-style-type: none"> ● accessed over the Net via a Web browser + Java applets <p>possibilities:</p> <ul style="list-style-type: none"> ● email submission of documents for codification ● email alerting of new contributions to the literature or discourse network
optional, incremental formalization	<ul style="list-style-type: none"> ● reviewers can choose to tag a comment with an agree/disagree icon, or leave it unclassified (most common) 	<ul style="list-style-type: none"> ● creation of new templates as needed ● (modification on the fly), choosing to tag nodes or not, etc.) ● creation of metadata scheme as needed 	<p>possibilities:</p> <ul style="list-style-type: none"> ● text analysis to extract concepts already registered in the library ● integration with publication and bibliographical tools

6. Discussion

With respect to the conference themes, this work raises some interesting issues about the interplay of discourse structuring and formalisms in communities of practice.

Discourse structuring can work

As noted under ‘incremental formalization’, attempts to structure discourse have a chequered history. However, our experience is that when practices are appropriately co-evolved with new technologies and representations, discourse structuring can work. In JIME and Compendium we see promising signs of discourse formalisms being adopted and successfully used within the targetted user communities (ScholOnto has yet to be used beyond the immediate project team). JIME builds on very familiar mechanisms to structure discourse (document structure and a threaded interface). Compendium, in contrast, introduces a radically new representation, facilitated by an expert, who passes on the skill to others by modelling it in practice, as well as through formal training. Our experiences with JIME and Compendium encourage us that ScholOnto can also succeed (elsewhere we have detailed why it may succeed where related systems have failed [8]). Indeed, we are experimenting with a Compendium-like graphical interface onto ScholOnto to build directly on the success of its user interface.

Ontologies, boundary objects and communities of practice

In the cases of Compendium and ScholOnto, we are in essence dealing with the introduction of a formal ontology into a community of practice (also discussed in [6]). Ontologies, by definition, require *consensus* on what should be included and how it should be structured. Once in place, ontologies serve to control interpretation with the goal of avoiding misunderstandings and confusions. Aligning *computational* agents with ontologies is one matter, but enforcing their use for communication between *people* is quite another. Arguably, sociological and associated CSCW perspectives on collective representation (e.g. [1,26]) build on radically different assumptions to those of the AI/knowledge modelling community (at least, as reflected in the level of attention given to such issues in typical knowledge modelling publications). “Formal ontologies meet communities of practice” as a scenario focuses attention on the intriguing challenge of forging perspectives and implicit worldviews from ethnography, artificial intelligence and human-computer interaction.

One response to this challenge is to focus on *boundary objects*, as first introduced by Star and Griesemer [36] from social analyses of scientific practice. “Boundary objects are objects which are both plastic enough to adapt to local needs and the constraints of several parties employing them, yet robust enough to maintain a common identity across sites” (p. 393). Each of the three examples presented expands the notion of a document beyond an artifact, linking it to an explicit discourse: in JIME, to a peer review; in Compendium, to the underlying network of issues and decisions; in ScholOnto, to a semantic web reflecting the claims and debates in which the document’s ideas are embedded. We are now dealing with *discourse representations as boundary objects* that are persistent, but being discourse, quite possibly growing.

Ontologies are in one sense an extreme case of boundary objects embodying a perspective; after all, an ontology’s *raison d’être* is the precise articulation of a perspective. One might assume that their brittleness and ontological assumptions make them poor candidates for supporting multi-perspective dialogue. There are two responses to this. Firstly, the fact that an ontology’s purpose is to make *transparent* the conceptual categories and relationships that are deemed important by a community opens them up for negotiation (if they are treated in this way), thus distinguishing them from more opaque representations that *conceal* the assumptions and perspectives they embody (e.g. many legal documents; a project review intended to secure further funding; a survey intended to motivate specific changes).

Secondly, we are focusing on ontologies not (primarily) for modelling the micro-worlds that are the normal focus of attention, but for reifying the normally invisible (or at least, unattended to) structure of *discourse*—*ways of talking about* these micro-worlds. The purpose of discourse structuring schemes is to support multi-perspectival discussion, not to close it down with the imposition of a master view. Of course a discourse ontology is still prey to the accusation that it is imposing a way of talking, but complex though this is, it is possible to do well. Compendium’s templates successfully operate as boundary objects between stakeholder groups. In Star and Griesemer’s terms, Compendium’s representations are *robust* enough to meaningfully capture discussions for subsequent reuse, yet *plastic* enough to allow negotiation over the meaning of ambiguous or controversial elements (using graphical argumentation structures if desired, or just visually mediated talk). The maps represent

multiple perspectives at a granularity and level of abstraction that is acceptable to all stakeholders (allowing them to preserve local control over detail), and can also transform one community's perspective into more appropriate formats (*boundary objects*) for others, thus assisting their *sense-making* (see [32] for details of 'representational morphing' into and out of Compendium). We hypothesise that ScholOnto will be able to replicate such 'collective ownership of ontologies' on a larger distributed scale.

7. Conclusion

To return to the distinction made at the start, the examples analysed here adopt a complementary approach to that normally found in the fields of digital libraries and corporate memory. In contrast to codifying 'primary artifacts' in any detail, the emphasis in all three examples is on codifying *discourse* about those artifacts. Firstly, this assists collective interpretation. Secondly, being digitally mediated, it seeds a new kind of collective memory that assists recovery of primary artifacts via new routes, and makes possible new kinds of discourse analysis. Primary metadata is still required, but 'meta-metadata' has been added: the interpretive layer.

As we are poignantly reminded by Bowker and Star [3], "classification systems provide both a warrant and a tool for forgetting" (p. 277)... "the classification system tells you what to forget and how to forget it" (p.278). This has ramifications for collective memory systems. We suggest that the kinds of systems reviewed in this paper offer possible ways to tackle this problem. Resource classification schemes typically impose a perspective, and leave little or no room for ambiguity, dissent or uncertainty. Deployed in conjunction with a conventional digital library, interpretive discourse systems such as D3E, Compendium and ScholOnto fill these 'borderline' niches, providing a place for multiple perspectives to co-exist, and for new ideas to emerge in discussions.

We regard collective memory resources as offering skeletons for an appropriately skilled 'paleontologist' to interpret and build a reasonable reconstruction of the living creature that once inhabited a particular environment. To pursue the metaphor a little further, the larger the population expected to make sense of the skeleton, the longer the time elapsed since it was buried, and the more the surrounding environment has changed, the more contextual information must be codified with the skeleton to assist analysis. The art and science of designing a collective memory resource might therefore be framed in the following terms: the co-design of *technologies* with individual and collective *practices* and *literacies*, to enable the *reading* and *writing* of material *sufficient* to support meaningful *reconstruction* as *living memories*, by an envisaged *population*, within an envisaged *timescale*, for a particular *purpose*.

Given this pervasive problem of reinterpretation, we propose that it is worthwhile to invest effort in understanding how interpretive layers can be effectively constructed by and for communities of practice. We have highlighted some of the problems that the introduction of interpretive discourse ontologies may encounter, but this should not deter us.

Acknowledgements

The examples outlined in this paper are the products of extremely stimulating collaborations with Tamara Sumner and Mike Wright (JIME/D3E), Maarten Sierhuis (Compendium). and John Domingue and Enrico Motta (ScholOnto).

8. References

1. Bannon, L.J. and Kuutti, K. Shifting Perspectives on Organizational Memory: From Storage to Active Remembering. In *Proc. HICSS'96: 29th Hawaii International Conference on System Sciences, Vol. III - Collaboration Systems and Technology*, Hawaii, Jan., 1996, 1996, IEEE Computer Society Press: Washington, pp. 156-167
2. Bowers, J. The Politics of Formalism. In *Contexts of Computer-Mediated Communication*, Lea, M., (Ed.), Harvester Wheatsheaf, 1991, pp. 232-261
3. Bowker, G.C. and Star, S.L. *Sorting Things Out: Classification and its Consequences*. MIT Press: Cambridge, Mass., 1999
4. Buckingham Shum, S. Analyzing the Usability of a Design Rationale Notation. In *Design Rationale: Concepts, Techniques, and Use*, Moran, T.P. and Carroll, J.M., (Ed.), Lawrence Erlbaum Associates: Hillsdale, NJ, 1996, pp. 185-215

5. Buckingham Shum, S., Domingue, J. and Motta, E. Scholarly Discourse as Computable Structure. In *Proc. SC2: Second International Workshop on Structural Computing*, San Antonio, Texas. June 3, 2000, Lecture Notes in Computer Science, Springer-Verlag: Berlin [<http://kmi.open.ac.uk/projects/scholonto/>].
6. Buckingham Shum, S., Domingue, J. and Motta, E. Ontology-Driven Document Enrichment: Philosophical, Modelling and Organizational Issues. In *Industrial Knowledge Management. A Micro Level Approach*, Roy, R., (Ed.), Springer-Verlag: London, in press
7. Buckingham Shum, S., MacLean, A., Bellotti, V. and Hammond, N. Graphical Argumentation and Design Cognition. *Human-Computer Interaction*, 12, 3, 1997, pp. 267-300 [<http://kmi.open.ac.uk/tr/papers/kmi-tr-25.pdf>].
8. Buckingham Shum, S., Motta, E. and Domingue, J. ScholOnto: An Ontology-Based Digital Library Server for Research Documents and Discourse. *International Journal on Digital Libraries*, Sept., 2000, in press
9. Buckingham Shum, S. and Sumner, T. New Scenarios in Scholarly Publishing and Debate. In *The Knowledge Web: Learning and Collaborating on the Net*, Eisenstadt, M. and Vincent, T., (Ed.), Kogan Page: London, 1998, pp. 135-151 [<http://kmi.open.ac.uk/knowledgeweb>].
10. Carter, L.M. Arguments in Hypertext: A Rhetorical Approach. In *Proc. Hypertext 2000*, San Antonio, TX, 2000, ACM: New York
11. Conklin, J. Hypertext: An Introduction and Survey. *IEEE Computer*, 20, 9, 1987, pp. 17-41
12. Conklin, J. and Burgess Yakemovic, K.C. A Process-Oriented Approach to Design Rationale. *Human-Computer Interaction*, 6, 3&4, 1991, pp. 357-391
13. Fischer, G., Morch, A.I. and McCall, R. Design Environments for Constructive and Argumentative Design. In *Proceedings of CHI'89: Human Factors in Computing Systems*, 1989, ACM: New York, pp. 269-275
14. Ginsparg, P. Winners and Losers in the Global Research Village (invited address). In *Conference on Electronic Publishing in Science, UNESCO: Session: Scientists' View of Electronic Publishing and Issues Raised*, Paris, 19-23 February, 1996, 1996 [<http://xxx.lanl.gov/blurb/pg96unesco.html>].
15. Grudin, J. Groupware and Social Dynamics: Eight Challenges for Developers. *Communications of the ACM*, 37, 1, 1994, pp. 92-105
16. Hashim, S.H. WHAT: An Argumentative Groupware Approach for Organizing and Documenting Research Activities. *Journal of Organizational Computing*, 1, 3, 1991, pp. 275-302
17. HICSS Workshops on Persistent Conversation (1999-2001). In *Hawai'i International Conferences on Systems Science* <<http://www.hicss.hawaii.edu/>>, 1999-01 [1999: <http://www.ascusc.org/jcmc/vol4/issue4/> or <http://computer.org/proceedings/hicss/0001/00012/0001toc.htm>].
18. Informedia: *Digital Video Library Research Project*. Carnegie Mellon University, USA [<http://www.informedia.cs.cmu.edu/>].
19. Lee, J. Incremental Acquisition and Formalization of Design Rationale. In *Working Notes, AAAI'92 Workshop on Design Rationale Capture and Use*, July 15, 1992, San Jose: CA, 1992, pp. 179-188
20. Lee, J. and Lai, K. What's in Design Rationale? *Human-Computer Interaction*, 6, 3&4, 1991, pp. 251-280
21. Malone, T.W., Grant, K.R., Lai, K.-Y., Rao, R. and Rosenblitt, D. Semistructured Messages are Surprisingly Useful for Computer-Supported Coordination. *ACM Transactions on Office Information Systems*, 5, 2, 1987, pp. 115-131
22. Masterton, S.J. Computer Support for Learners Using Intelligent Educational Agents: The Way Forward. In *Proceedings Sixth International Conference on Computers in Education*, Beijing, Oct., 1998, 1998
23. Minneman, S., Harrison, S., Janssen, B., Kurtenbach, G., Moran, T.P., Smith, L. and van Melle, W. A Confederation of Tools for Capturing and Accessing Collaborative Activity. In *Proceedings of the ACM Multimedia '95 Conference*, 1995, ACM: New York

24. Newman, S. and Marshall, C. Pushing Toulmin Too Far: Learning from an Argument Representation Scheme. Xerox Palo Alto Research Center, *Technical Report SSL 92-45*, 1991
25. Orlikowski, W., Yates, J., Okamura, K. and Fujimoto, M. Shaping Electronic Communication: The Metastructuring of Technology in the Context of Use. *Organization Science*, 6, 4, 1995, pp. 423-444
26. Robinson, M. and Bannon, L. Questioning Representations. In *Proc. of ECSCW'91: 2nd European Conference on Computer-Supported Collaborative Work*, Bannon, L., Robinson, M. and Schmidt, K., (Ed.), Kluwer: Amsterdam Sept 25-27, 1991, pp. 219-233
27. Schuler, W. and Smith, J. Author's Argumentation Assistant (AAA): A Hypertext-Based Authoring Tool for Argumentative Texts. In *Proc. ECHT'90: European Conference on Hypertext: Argumentation, Design & Knowledge Acquisition*, Cambridge, 1990, Cambridge University Press, pp. 137-151
28. Selvin, A. Leveraging Existing Hypertext Functionality to Create a Customized Environment for Team Analysis. In *Proceedings of the 2nd International Workshop on Incorporating Hypertext Functionality Into Software Systems*, 1996
29. Selvin, A. Facilitating Electronically: Using Technology to Help Maria. *The Facilitator (Special Issue on Automated Meeting Support)*, September 1998, 1998
30. Selvin, A. Supporting Collaborative Analysis and Design with Hypertext Functionality. *Journal of Digital Information*, 1, 4, 1999 [<http://jodi.ecs.soton.ac.uk/Articles/v01/i04/Selvin/>].
31. Selvin, A. and Sierhuis, M. Case Studies of Project Compendium in Different Organizations. In *Workshop on Computer-Supported Collaborative Argumentation, Conference on Computer-Supported Collaborative Learning*, Stanford, CA (12-15 Dec., 1999), 1999 [<http://kmi.open.ac.uk/sbs/csca/cscl99/>].
32. Selvin, A.M. and Buckingham Shum, S.J. Repurposing Requirements: Improving Collaborative Sensemaking over the Lifecycle. In *Profess'99: International Conference on Product Focused Software Process Improvement*, Oulu, Finland, June 22-24, 1999, 1999, pp. 539-559 [<http://www.inf.vtt.fi/pdf/symposiums/1999/S195.pdf>].
33. Selvin, A.M. and Buckingham Shum, S.J. Rapid Knowledge Construction: A Case Study in Corporate Contingency Planning Using Collaborative Hypermedia. In *Proc. KMaC 2000: Knowledge Management Beyond the Hype*, Birmingham, UK (16-19 July 2000), 2000, Operations Research Society [<http://kmi.open.ac.uk/kmi-abstracts/kmi-tr-92-abstract.html>].
34. Shipman, F.M. and Marshall, C.C. Formality Considered Harmful: Experiences, Emerging Themes, and Directions on the Use of Formal Representations in Interactive Systems. *Computer Supported Cooperative Work*, 8, 4, 1999, pp. 333-352 [<http://bush.cs.tamu.edu:80/~shipman/cscw.pdf>].
35. Shipman, F.M. and McCall, R. Supporting Knowledge-Base Evolution with Incremental Formalization. In *Proc. ACM CHI'94: Human Factors in Computing Systems*, Boston, Mass., 1994, ACM Press: New York, pp. 285-291
36. Star, S.L. and Griesemer, J. Institutional Ecology, "Translations," and Coherence: Amateurs and Professional in Berkeley's Museum of Vertebrate Zoology, 1907-1939. *Social Studies of Science*, 19, 1989, pp. 387-420
37. Suchman, L.A. Do Categories Have Politics? *Computer Supported Cooperative Work*, 2, 1994, pp. 177-190
38. Sumner, T. and Buckingham Shum, S. From Documents to Discourse: Shifting Conceptions of Scholarly Publishing. In *Proc. CHI 98: Human Factors in Computing Systems*, Los Angeles, CA, 1998, ACM Press: NY, pp. 95-102 [<http://d3e.open.ac.uk/general/d3e-chi98/>].
39. Sumner, T. and Buckingham Shum, S. A Toolkit for Publishing Web Discussion Documents: Design Principles and Case Studies. In *Proc. APCHI 98: Asia Pacific Computer Human Interaction Conference*, Kanagawa, Japan (15-17 July), 1998

40. Sumner, T., Buckingham Shum, S., Wright, M., Bonnardel, N. and Chevalier, A. Redesigning the Peer Review Process: A Developmental Theory-in-Action. In *Proc. COOP'2000: Fourth International Conference on the Design of Cooperative Systems*, Sophia Antipolis, France: 23-26 May, 2000, 2000 [<http://kmi.open.ac.uk/tr/papers/kmi-tr-96.pdf>].
 41. Winograd, T. A Language/Action Perspective on the Design of Cooperative Work. *Human-Computer Interaction*, 3, 1, 1987/88, pp. 3-30
 42. Zimmerman, B. and Selvin, A. A Framework for Assessing Group Memory Approaches for Software Design Projects. In *Proc. ACM DIS'97: Conference on Designing Interactive Systems*, Amsterdam, NE (18-20 August), 1997, ACM Press: New York, pp. 417-426
-