Preprint. Final version appears as:
ATTFIELD, S. & BLANDFORD, A. (2010) Discovery-led refinement in e-discovery investigations: sensemaking, cognitive ergonomics and system design. *Artificial Intelligence and Law*. 18.4. 387-412

Discovery-led refinement in e-discovery investigations: sensemaking, cognitive ergonomics and system design

Simon Attfield[1], Ann Blandford[2]

[1&2] UCL Interaction Centre, University College London, Gower Street, London, WC1 6BT, UK

[1] Tel: +44 (0)20 7678 0694

[1&2] Fax: +44(0)20 7679 0699

[1] s.attfield@cs.ucl.ac.uk

[1] http://www.uclic.ucl.ac.uk/people/s.attfield/url here

The work reported in this paper was funded under the EPSRC project Making Sense of Information (EP/D056268)

Keywords: electronic data disclosure, e-discovery, e-disclosure, investigations, sensemaking, information interaction, collaboration

Abstract: Given the very large numbers of documents involved in e-discovery investigations, lawyers face a considerable challenge of collaborative sensemaking. We report findings from three exploratory workplace studies which looked at different aspects of how this challenge was met. From a sociotechnical perspective, the studies aimed to understand how investigators collectively and individually worked with information to support sensemaking and decision making. Here, we focus on discovery-led refinement; specifically, how engaging with the materials of the investigations led to discoveries that supported refinement of the problems and new strategies for addressing them. These refinements were essential for tractability. We begin with observations which showing how new lines of enquiry were recursively embedded. We then analyse the conceptual structure of a line of enquiry and consider how reflecting this in e-discovery support systems might support scalability and group collaboration. We then focus on the individual activity of manual document review where refinement corresponded with the inductive identification of classes of irrelevant and relevant documents within a collection. Our observations point to the effects of priming on dealing with these efficiently and to issues of cognitive ergonomics at the human-computer interface. We use these observations to consider how visualisations might enable reviewers to deal with such refinements more efficiently.

Introduction

Electronic Data Disclosure (EDD, e-disclosure or e-discovery) is a process in which electronic data is sought, located, secured, and searched with the intent of

using it as evidence in civil or criminal proceedings, or as part of an inspection ordered by a court or sanctioned by a government (Conrad 2007). Lawyers involved in corporate litigations and regulatory investigations routinely face an immense challenge. Their aim is to identify and present documents relating to the activities of people within an organisation as these pertain to the aims of the investigation with the ultimate goal of telling a compelling tale (Socha-Gelbmann 2009). A key resource for this activity is a vast evidence-base of documents obtained through a large-scale recovery exercise. This will include a range of user-generated content, such as emails and office documents which record the everyday activities of the organisation under scrutiny. Once secured, this mass of documents must be subjected to extended and meticulous filtering and review in order to identify the relatively few documents that have a bearing on the case.

The reason why this task has become so onerous and expensive in recent years is because of the tectonic shift within organisations from paper to electronic documents. Paul and Baron (2007) describe this as a 'pulse' in the history of information resulting as it has in an information landscape in which information artefacts are created and communicated in quantities never seen before and which are increasing exponentially. Electronic discovery requests for email alone can result in thousands to millions and even tens of millions of documents (Baron et al. 2007). This presents a serious challenge to the legal system to effectively identify a complete evidentiary record (Paul and Baron 2007) within reasonable constraints of time and cost.

Testament to this problem and the speed with which it has come about was illustrated graphically by Jeane Thomas, a partner within Crowell & Moring's Antitrust Group, during her keynote address at the DESI II Workshop in 2008. Between 1996 and 2005 Crowell & Moring handled a series of Mergers and Acquisition on behalf of one of their clients. In each, potential competition issues led to document requests from the US Department of Justice. For the first, twelve to fifteen lawyers were required for the manual review; this resulted in a production of around three hundred boxes of paper. By 2004 the business had moved from being mostly paper-based to being mostly electronic. To fulfil a similar transaction, the firm employed 125 contract lawyers for three months. They reviewed 30 million pages and produced 12 million relevant pages. The following year a further transaction was conducted, only this time the firm needed a team of around 600 lawyers for the review. They read around 112 million pages and produced 17 million relevant pages.

The phenomenal increase in the number of documents created and held by institutions is referred to by Paul and Baron (2007) as 'information inflation'. They argue that this, combined with myriad and continually evolving forms of corporate writing (e.g. office documents, email, instant messaging, blogs, wikis, and potentially now Google Waves) held on multiple, and distributed forms of institutional digital memory (e.g. servers, personal computer hard drives, removable memory), has stressed the legal system to the point where change is essential. In addition the numbers of investigations are increasing. Within the EU, for example, regulatory investigations are expected to increase due to significant enhancements in the powers and resources available to regulatory authorities and their willingness to use them (Wildisen 2009). One effect of the 'credit crunch' has been to bring about a change in the organisational culture of the UK's Serious

Fraud Office (SFO) to more closely resemble the proactive stance of financial regulators in the US (Wildisen 2009). This combines with additional investigative and punative powers, such as the right afforded to the Office of Fair Trading (OFT) to mount "dawn raids" (Wildisen 2009).

Given the scale of such effort involved in conduction investigations there has been a natural and growing interest in the development of technologies and techniques that might help address them. Technologies attracting particular interest in this arena include media restoration tools, dedicated document management systems, information visualization, case analysis tools, and advanced information retrieval systems (such as concept search and information extraction). In particular, there has been an interest in the role of search and how it can be conducted to best effect. As a precursory step to review, search represents an essential step in mitigating the challenge of high review loads. Consequently, attention within the DESI community has been drawn to the need for search technologies and related techniques which can be shown to offer good performance in an e-discovery scenario (see for example, Brassil et al, 2009). A central initiative in this regard is the TREC Legal Track.

In addressing the question of how to design technology for e-discovery, however, we argue that it is important to recognise that e-discovery is an exercise in collaborative sensemaking. Sensemaking has been described as "the reciprocal interaction of information seeking, meaning ascription and action" (Thomas et al. 1993 p.240), and as "the deliberate effort to understand events" (Klein et al. 2007 p.114). It occurs when people face new problems in unfamiliar situations and their current knowledge is insufficient (Zhang et al. 2008). Characteristically, sensemaking involves a bi-directional interaction between engagement with data (i.e. bottom-up processing) and continually evolving representations and understanding that account that data (i.e. top-down processing) (Klein et al. 2006; Klein et al. 2007; Pirolli and Card 2005; Russell et al. 1993). We believe that understanding the details and dynamics of how legal staff individually and collaboratively perform e-discovery 'in the wild' is likely to provide important insights concerning the kinds of technological support they would find most useful.

The perspective we take is to view e-discovery as a collaborative, sociotechnical challenge. Given its scale and the need for resolution within a reasonable timeframe, e-discovery is typically conducted by teams of people working in close collaboration. Lawyers with different levels of experience and seniority work together and with paralegals, litigation support managers, records specialists and technologists (Kaplan 2008) using technology to manage recovery and review and ultimately make sense of the gathered evidence in a way that furthers the investigation. Frequently the e-discovery 'team' will also extend beyond the boundaries of a single organisation to include outside litigation service providers and e-discovery consultants. In this context the need for effective collaboration, including both the distribution of evidence and tasks and the integration of resulting knowledge, is particularly pressing.

Others have stressed the socio-technical perspective when considering design for e-discovery and argued for work-practice studies to explore this. Benedetti et al. (2008), for example, point out that when one examines how work is actually

organised and carried out, an emergent richness and variety becomes apparent, and that work-practice studies are an essential part of acquiring an understanding for designing useful and intelligent tools. In addition, we argue that investigating how work happens in context can make visible significant patterns in thinking, action and collaboration which can provide valuable insights for how to support that work more effectively.

In this paper we report results from three work-place case studies of large e-discovery investigations. The investigations were performed by lawyers and other staff within the London offices of an international law firm. The case studies were ethnographic and exploratory in nature. Our aim was to understand the ways in which the investigators individually and collectively worked with information to support sensemaking and decision making. Interviews with investigating legal staff and key artefacts they used provided a source of data for eliciting detailed reconstructions of the challenges that they faced and how their activities and thinking were structured in response. The data were analysed using inductive methods common to ethnographic studies as a source of reflection for technological requirements and future research.

In this paper we focus on two areas of our findings. The first takes a macroscopic perspective on collective problem structuring. It explores the decomposition of research problems during the investigations. Since this decomposition provided a basis for the distribution of labour, it is a significant issue in respect of collaboration. We describe the structuring we observed in terms of an ontological framework which we refer to as a *line-of-enquiry framework*. We then consider the implications of the framework for the design of collaborative, e-discovery support systems.

We then contrast this by considering an aspect of individual working in the context of this larger collaborative activity. We focus in particular on findings related to the task of manual document review. The need to manually review documents in e-discovery is widely recognised as presenting a considerable overhead in terms of cost and time. If anything, this present the most significant challenge to the legal system in terms of performing e-discovery matters effectively. Document review is a cognitively intense activity. At the centre of it are people, usually junior lawyers, who sit at computers and scan or read one document after another making judgements about relevance, typically inspecting thousands of documents over a period of weeks. We consider some aspects of this activity with particular reference to issues of cognitive ergonomics in relation to the design of document review system interfaces.

What binds these two issues together the prominence of *discovery-led refinement*. By this we mean the ways in which discoveries made during the investigations through engagement with the materials of the investigations result in new insights which allow investigators to re-frame the problems at hand and developing new goals and strategies in order to address them. Understanding how discovery led-refinement occurs has implications for understanding how to develop technologies which support the natural evolution of thinking during an investigation. In particular, we observe two kinds of discovery. The first are discoveries that investigators made about the domain under investigation as inferred from exposure to new evidence. The second kind of discovery concerns insights about

the evidence itself as a collection of documents which are worked with in the process of making the first kind of discovery.

The remainder of this paper is structured as follow: In the next section we outline the method used for gathering data in the case-studies. We then describe the line of enquiry framework, its motivation, and its implications. Following this we describe discovery-led refinement during the document review task and discuss its implications in relation to the design of interactive data visualisations.

Method

The case-study research method was interpretive and inductive (as described by Klien and Myers (1999)). Rather than being guided by hypotheses and predefined independent and dependent variables, we used the broader and more exploratory research question of understanding how corporate investigators structure and coordinate action. Our aim was to examine the situated performance of e-discovery in order to uncover the "complexity of humans sensemaking as the situation emerges" (Klein and Myers, 1999, p69). Klein and Myer's describe a number principles for conducting research of this kind, of which the most fundamental is that of the hermeneutic circle. According to this idea "all human understanding is achieved by iterating between considering the interdependent meaning of parts and the whole that they form" (Klien and Myers, 1999, p.72). In other words, we come to understand a complex whole by interpreting detail in terms of abstract interpretations and forming abstract interpretations based on interpretations of detail. This is itself a sensemaking process and characterises our data-gathering and analysis approach which aimed to generate abstract conceptualisations based on the data which would account for the data.

Our approach can also be described as *idiographic* (Luthans & Davis 1982) insofar as we were interested in considering individual experiences in a limited number of cases in depth. This is in contrast to a *nomothetic* approach which is concerned with deriving generalisable laws. This is not to say that generalisable laws are not useful, but rather that considering a few cases in detail is a good place to begin the process of abstraction.

Participants were recruited for 1:1 interviews from the London offices of a large, corporate law firm using a combination of theoretical (Strauss & Corbin 1998) and snowball sampling (Johnson 1990). Each participant had worked on one of three e-discovery investigations. Theoretical sampling was used to focus in on emerging issues and explore similarities and contrasts between investigations. Following Strauss & Corbin (1998) data gathering and analysis were interleaved.

Fourteen in-depth interviews were conducted. Interviews lasted from 45 minutes to 1hr 40 minutes. Although we would have liked to, confidentiality constraints made it impossible to conduct observations of investigation work. However, during and/or after interviews, key artefacts were made available for inspection including review software loaded with investigation data. The availability of such artefacts during interviews made it possible to conduct informal reconstructions of work activities which were used in order to explore the ways in which aspects of the work (such as document review) unfolded in some detail in relation to the tools and resources used and created in the investigations.

Interviewee roles included a technical coordinator (responsible for e-discovery support), two trainees, six associate lawyers, one senior associate lawyer and three partners. A senior associate who managed one investigation was interviewed twice. Ten interviews (including the two with the senior associate) pertained to a single investigation whose goal was the identification of a suspected fraud; one interview pertained to an earlier suspected fraud (chosen to test the generality of findings within one kind of legal matter); and three pertained to a matter concerning the origin of anomalies within a set of legal contracts (to test the generality of findings across contrasting types of matter).

The interviews were semi-structured with participants asked initially to provide a broad account of how the investigation had unfolded from the beginning of their involvement. During or after this account they were prompted to provide detail in relation to their interactions with evidential documents and external representations that the investigators created (either as hard-copy or mediated through software tools), and also how they coordinated their work with other team members. Participants were encouraged to contextualise these detailed descriptions in terms of their rationale, including the ongoing problems and questions of the respective investigation. In order to invite the participant to correct the researcher's understanding and provide additional detail, aspects of their accounts were summarised by the researcher at intervals during each interview.

Interviews were transcribed and analysed using open coding (Strauss and Corbin 1998) in order to generate a set of abstract themes or 'categories' that describe the data. These were refined on an ongoing basis through constant comparison against the data (Strauss and Corbin 1998).

One of the major themes emerging from the analysis related to the way in which discoveries from evidence prompted the decomposition of initially broad investigation issues into embedded sub-issues. In the following two sections we report describe how this happened, first in general terms across investigations as a whole, and following that in relation to the activity of document review.

Discovery led, recursive lines of enquiry

For each of the investigations a major source of evidence was a collection of documents (the 'document universe') resulting from ongoing document-recovery field. Some hard copy documents had been recovered from filing cabinets but by far the majority were electronic documents recovered from email servers and workstation hard-drives. Other sources of evidence included telephone records and interviews with witnesses and suspects. The investigations were both large and collaborative in nature with different tasks distributed across members of the respective teams.

One of our interests was to explore how the teams decomposed the problems that they tackled and how their results were integrated. This contrasts with other approaches to the study of sensemaking which have tended to focus on describing process (e.g. Pirolli and Card, 2005; Klein et al. 2006; Klein et al 2007). However, we begin by outlining the process to provide some context.

Figure 1 shows a very simple schematic to illustrate the process of the investigations. Recovered documents were added to a server and, in most cases, were searchable. Queries were then devised to retrieve documents relevant to evolving questions (document selection). The resulting documents were reviewed and electronically coded for relevance to any of the 'issues' currently active within the investigation (document review and classification). This had the effect of forming collections of relevant documents on which further work could be performed.

Information was then extracted from the relevant documents and re-represented within integrated analyses (schematisation). The most important of these were chronologies representing sequences of events, including details of meetings and email communications. A number of separate chronologies were created, and as these evolved important content was selected and consolidated into single master chronologies which provided an overview of known 'facts'.

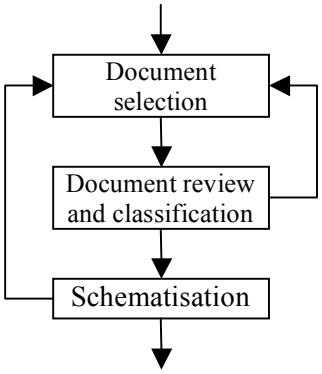

Figure 1. An overview of the investigation process

During the investigations the engagement with evidence that occurred through document review and schematisation had the effect of enhancing the investigators' understanding in a way that gave rise to new, more focussed issues and questions, resulting in the construction of information seeking strategies (e.g. search queries). These new issues and questions, however, tended not to be departures from the initial investigation issues and questions, but rather provided a re-specification of the original issues and questions in more detailed terms that were made available through new knowledge. This had the effect that new issues and question formed lower-level and recursively embedded lines of enquiry. We illustrate this with three examples each of which operates at a different level of investigation granularity.

From contract class to specific contracts

In one of the investigations, the high-level objective was to explore the possibility of fraud within a particular class of contract. A team of investigators were assigned to this task. Searches were constructed and run and documents resulting from these were passed to the team for review. However, the team had the initial

problem that they did not know what contracts there were within this class for the company they were investigating. This made constructing the initial queries difficult. Identifying characteristics such as contract names or associated employee names were unavailable. Details such as these would have provided multiple investigation foci with associated terms that might have been used to for searching. But without knowing the 'identities' of the contracts in question, it was only possible to specify a single broad issue, characterised at the level of the contract class. As a senior associate said:

> P4: Well actually what *[class]* contracts does the company have? And no one in the company knows or can tell you so you're then trying to piece that together. You know you're seeing references to *[contract a]*, you're seeing references to *[contract b]*, to *[contract c]*, to *[contract d]* and you've got no idea and you're trying to build up absolutely everything. I mean the scope of what you're trying to do is immense and you're having to define it as you go along…

'Defining' the investigation as you go along characterises the task well. The process of 'identifying' the different contracts had to be done by 'bootstrapping' inferences from information uncovered in documents and using this to re-specify issues and questions. Once the identity of a contract was known, then it could be defined as an investigation problem in its own right with associated information seeking strategies leading to evidence specific to it. Further, a subset of the team could be defined who would focus exclusively on this area, so consolidating effort and developing knowledge.

From contract focus to a time-period focus

Despite the foci provided by the identification of specific contracts of interest, the number of documents that were responsive to searches based on each of these was nevertheless large. Consequently, the investigators needed a way of focussing in on areas that would address the broader investigation questions most effectively. Given the nature of the allegations there were particular kinds of activity which were of interest and these would necessarily have occurred at specific periods within a contract lifecycle. However, when these periods had occurred for any one contract was initially unknown. As the investigators responsible for each contract reviewed documents and built their chronological representations of activities on the contracts, so these periods came to light:

> P5: …we'd be thinking, well if we're right on this, this is a really important build up […]. Or, we think money must have been sucked out of this business around this time. […] [Junior Partner] selected certain periods and posed certain questions in relation to those periods. And we would go back and interrogate the information further.

Importantly, the identification of particular and limited time periods of interest within a contract lifecycle allowed the investigators to develop new strategies for document retrieval involving date-delimitation.

> P5: If for example, three days were going to be really important, then we wouldn't worry about search terms. […] We would just say, give me every document that bears this date, created, edited, sent – anything. […]

Other information seeking strategies that took advantage of the identification of particular periods included the examination of telephone and expense records within certain time-windows. Telephone and expense records could provide useful and suggestive evidence about the kinds of activities of key protagonists. But

examining them was a slow process which was consequently expensive in terms of investigator time. Added to this telephone and expense records only made sense in relation to a backdrop of existing interpretations and theories.

From issue focus to event focus

Working on any of the issues involved the investigators in reviewing retrieved evidence and from this drawing inferences about events that had taken place, such as meetings and significant communications between protagonists. It was these inferences which were used as a basis for the representations that they created, and in particular, chronologies. For example, evidence for an event might take the form of an email between two people proposing a meeting. However, an email proposing a meeting would not present conclusive evidence that the meeting took place. And so an email like this could initiate a very specific set of theories and questions surrounding a single event on a given day,

> P4: […] So you put an entry down for November 20th and then you'd start looking for documents which relates, which might give evidence that that happened, that it actually happened […] and if it did happen who else was involved, who were they meeting, what were they doing, what were they saying to each other?

> P16: Let's take an example, like you're looking into a question as to whether someone was missold some securities and the relationship takes place over several months, various statements are made at various different times, what you may well find in that type of scenario is that when you look into it, 90% can be agreed you know there's no real dispute. But there will be a key meeting or a key conversation which took place for which there is no accurate records […] and what you're then trying to do is to work out exactly what happened at that meeting or during that call.

Faced with this situation, an investigator might focus in on this event and search for further evidence, or they might record the event as a conjecture in a chronology, and continue to review documents in the hope that they would come across further evidence relating to it.

These examples illustrate the way that new discoveries emerging from the investigation work prompted the decomposition and refinement of investigation issues into lower-level lines of enquiry. They are some common features which we will briefly explore:

1. Researching issues brought information to light that acted as a cue for more focused lines of enquiry. Without this knowledge these focused lines of enquiry would have been impossible;
2. New lines of enquiry were not complete departures but acted as sub-problems. Once the investigation of a sub-issue and all its embedded sub-issues had been exhausted, its outcomes could propagate back up to inform the outcomes of superordinate issues;
3. Despite 2, each new line of enquiry was independent insofar as it posed new questions and gave rise to new research strategies;

This discussion of the decomposition and focusing of research issues however, is incomplete without considering how in practice work on coordinate issues could inform each other. In addition to vertical information flow, it was also seen as essential for investigators working on different sub-issues to discuss their findings and theories and exchange information. One reason for this was the imprecision of information seeking strategies (such as search). A lawyer working on one local

area of enquiry could, and frequently did, turn up documents which could be of interest to a lawyer working on an unrelated area of enquiry somewhere else in the investigation.

To support the sharing of information, multiple communication mechanisms were put in place. These included daily review meetings in which investigators would be asked to summarise what they were finding, what they had inferred from this, and what sorts of information they were looking for. This had the effect of updating fellow investigators about their theories of relevance. During these reviews their inferences would be tested, alternative interpretations suggested, as well as documents and findings offered which might have a baring on their line of investigation. The communication mechanisms also included informal 'huddles' in which groups of investigators working in a similar area would discuss and exchange evidence. Also, knowing the interests of other investigators in the team, as new documents were uncovered these would be passed around on an ad hoc basis. As one lawyer explained "The amount of communication that has to go on in order to make that work is phenomenal".

The organisation of the investigation around embedded lines of enquiry in a collaborative context led us to consider how the emerging investigation structure might be reflected within systems for supporting large-scale collaborative sensemaking. A key motivation was to develop a way of representing an investigation in a way that mitigates information overload. One partner responsible for an investigation described the problem of information overload in this way,

> P6 Because, erm, my … the thing which was concerning me, coming at this, because all that I or the partner at [company] could take from this was a certain amount… there's only a certain amount of information that you can handle, from a personal perspective.

Sharing information and taking a 'horizontal' view was seen as important for identifying links between lines of enquiry, but this shear complexity of this could swamp the investigators,

> P12 … but often I think the problem was there was just too much going on, so you couldn't really draw any sort of themes from what was going on because there was just too much.

Hence controlling the quantity of information to which any one person would need to attend is an important requirement,

> P6 You could have a blog or some kind of an intranet or whatever, but there's a real risk of information overload. So targeting the right things to the right people… But the real balance… you see it in a lot of what we do… between giving people the information that they need to link all the pieces together, and not overloading, because then you just get paralysis.

Related to this, the investigators expressed the need for representations that supported filtering in order to eliminate extraneous information,

> P7 … where you have so many… 10, 20 issues whatever that you are looking at, or that the team as a whole are looking at, if you want to construct a theory about a subset of responsibility, it's a bit confusing if you see everything. So it would be quite helpful if you could somehow have maybe both… have the overview of everything… and then… only see events and document relating to a particular subset of issues that you are looking for. That might be helpful.

> P4 … we want to look at and analyse a certain event, you just want to be able to home in on five entries on a certain date, or on a event involving two or three people, so its really just the filtering of it just goes straight to what you want and because you just want the bare minimum that you need to get the answer.

Of particular interest was the idea of developing a generic and extensible representational framework which could be used as a basis for the conceptual design of an investigation system which would represent all the elements of an investigation whilst using the way in which an investigation structurally evolves as a means for defining and filtering information thematically into separate yet interconnected 'contexts' (approximately equivalent to 'issues'). An investigator responsible for a specific 'context' could then focus on elements relevant to *them* (e.g. questions, queries, evidence etc.) to the exclusion of extraneous information, and investigators responsible for larger, integrated parts of an investigation could take a similar and yet more broadly defined view. However, it would be important for coordinate or 'horizontal' contexts to available such that information could be passed from one to another.

Such a representation would allow:

- The gradual decomposition of areas of investigation as these occur naturally through exposure to evidence;
- The representation of 'contexts' corresponding to lines of enquiry at different level of granularity;
- The elimination of extraneous information (noise) for any context;
- Relating superordinate and subordinate contexts such that outcomes propagate up (and meaning propagates down);
- Relating coordinate contexts such that evidence can be passed from one context (and responsible investigator) to another;

In order to elaborate the requirements for such a representation further we re-examined our data in way that would help to reveal the conceptual elements that were common to a given line of enquiry. This would allow us to see what range elements of an investigation should be represented in any given 'context'. To do this we performed a Grounded Theory (Strauss and Corbin 1998) analysis using the concept of a *line-of-enquiry* as a core category in order to develop a framework of elements within a given line of enquiry independent of level of granularity. We describe the resulting framework in the next section.

A Line-of-enquiry framework

The framework takes a *line-of-enquiry* as a primary object. A *line-of-enquiry* has seven element types which represent those things that an investigator working within a line of enquiry generates and works with. They are: *theories*, *questions*, *information seeking strategies*, *evidence and evidence collections*, *knowledge representations, assigned investigators* and *lower-level lines of enquiry*. Significantly, given this last element type, lines of enquiry recursively embed such that one line of enquiry each can give meaning to one or more lower-level lines of enquiry, each with similar structure. In the following we describe each type of element:

Theories

Our data showed that theories or conjectures were central to a line of enquiry; they were theory-led. One senior associate expressed the centrality of theories in defining an issue (line of enquiry) in this way,

> P4: Well it's the theories that then define the issues you are coding for and looking for. […] we had lots of sub-issues and theories, well sub-theories that were helping to define the issues […]

> P7 I mean, your task would be to look at, say, contract so-and-so, so you would mostly be constructing a theory as to what went on there.

Theories were triggered by some kind of cue. This cue could be an allegation that had been made, or information that was revealed through the investigation process. For example, above we show how identifying a business activity of a certain type, or a key time period, or an event could provoke a more focused line of enquiry. Each of these was associated with a theory, however broad, about what could potentially have been the case (e.g. a contract involved fraudulent activity, fraudulent activity occurred within a particular time frame, a meeting took place).

Through the process of the investigation, theories were systematically investigated and eliminated when the evidence found was contradictory or unsupportive. When all the lower-level theories associated with a line of enquiry had been eliminated then a higher-level issue would become inactive.

Questions

The investigators made a natural move from theories to research questions, and in many cases these were explicitly recorded and shared across an investigation team. Research questions specified requirements for information that would test theories or simply elaborate their focus. This elaboration could then provide cues for further decomposition or could yield other unexpected findings.

> P6 You begin to ask yourself questions about, well, "What was really happening in this period of a week? This is slightly odd, because, of course, we can see *that* going on there, *that* going on there and *that* going on there. And this guy's flying from here to here to here, this guy's no where near the picture, but then he emerges there. OK, what I want to do is drill in and find out exactly what is happening, and these are the questions that I've got.

Information seeking strategies

Questions naturally gave rise to information seeking strategies. Most commonly these were keyword searches over the document universe designed to provide evidence responsive to the questions.

> P4 We ran keyword searches on all of that data and we ran I don't know how many, probably about 150/200 keyword searches.

> P4 Let's say if you know that Joe Bloggs was meant to be in *[location]* around *[date]*, it means that then you can on the server run a search for documents involving certain people around that week to actually see whether it did happen and if it did happen who else was involved, who were they meeting, what were they doing, what were they saying to each other?

Any given line of enquiry could have multiple queries developed iteratively and these could also be repeated periodically as new documents were added to the

main collection over time. The range of information seeking strategies, however, depended on the questions and the evidential resources available. In addition to search, and as already discussed, information seeking strategies might include the examination of telephone records, reviewing expense records, or asking questions of specific witnesses in interview.

Evidence and evidence collections

The information seeking strategies yielded information. In the case of searches this took the form of document collections (results sets).

> P12 So basically … [x] will come back with the search, it will get uploaded, we have a hundred search results set up,...

> P4. We were running these keyword searches […] they would throw let's say 10,000 hits […] and then we ended up with what we now have—130-odd thousand documents on our database […] and these are documents which each of them has been reviewed, each of them has been subjectively coded and that is the main source of information with the witness evidence.

Search results were manually reviewed for relevance by issue teams and relevant documents tagged. This then created a further collection of documents which was used for generating knowledge representations.

Knowledge representations

Within each line of enquiry, the investigators continually reviewed and collated evidence and recorded the inferences they drew from the evidence within different forms of analysis product. These included event chronologies, written narratives, social network diagrams, and organograms showing formal organisational structures. Knowledge representations were organised around two types of concept: The first were *people*. It was important to discover and maintain records of the central cast of characters for each line of enquiry and to record relationships between them. To do this the investigators created profiles of key protagonists and in some cases drew link charts to represent relationships.

> P4. ... and other things you would do is, create files on individual people, that would be a repository for key information. […] physical files. When I say physical, most of them were documents. But we would print them out and put them together with relevant documents and things like that. And they would often function as an index.

The second kind of concept corresponded to the events that told the story relevant to a given line of enquiry.

> P8. So yeah, the main thing we were doing was updating chronologies, keeping the big picture of what had happened, keeping that up to date and accurate, that was a general thing that we were always doing.

Each chronology event record included a date and time, a summary description, a list of people involved in the event, and references to the supporting evidential documents. These representations provided the basis for the evaluation of theories.

Assigned investigators

Given the team setting, a given line of enquiry could be allocated to one or more investigators. Knowing who was assigned to what area of the investigation provided a basis for lateral information sharing. Hence, these assignments formed part of the concepts associated with a line of enquiry.

> P16 We did have a team of probably about half a dozen associates working on it, looking at various different areas and we […] looked at different areas of the organisation so we had one team looking at how [x] had been working, another team looking at particular aspects of [y], another team looking at what the Chief Executive had been doing. […] we identified five I think it was areas, fairly disparate areas that we thought we needed to investigate as a starting point and then what we did is we then set up mini teams that focused on those areas and you then became masters of information in your specific area of investigation.

Lower-level lines of enquiry

Finally, and as already discussed, knowledge arising from investigation work on a line of enquiry could give rise to any number of more focussed problems which addressed more specifically defined aspects of a broader investigation issue.

The framework we have described provides an ontology of concepts associated with a line of enquiry. We have found these elements to occur irrespective of granularity. In some cases, a line of enquiry might concern a single relationship or a single event, whilst the investigation as a whole can be considered a line of enquiry.

When instantiated, the framework gives rise to a hierarchy of enquiry nodes, with elements represented at each node. By implementing this framework within a sensemaking support system we anticipate a number of advantages centring around the simultaneous decomposition and integration of multiple strands of an enquiry. By allowing investigators to selectively access information associated with a particular line of enquiry or ‘context’, the framework can support the elimination of extraneous information for focussed analysis. Conversely, with outcomes propagating up within the hierarchy, they it would be possible to integrate the elements of an enquiry at any higher level. This has implications for the filtering of knowledge representations such as chronologies and link charts used in schematisation. By associating the component elements of such representations with framework nodes, users could use node selection to view these different strands of an investigation in different combinations, thus enabling them to easily explore links between apparently separate issues.

Finally, integrating data and user-generated knowledge representations from multiple aspects of a collaborative investigation provides an opportunity for a system to automate the process of identify potential links between disparate parts of a large investigation which might otherwise have gone unnoticed. This could be based, for example, on matching common characters or travel locations across apparently unrelated lines of enquiry. Investigators alerted to these could then explore the extent to which they offer explanatory leverage. The details of this matching would depend upon specific user-needs and the details of data and knowledge representations within the system. However, the opportunity for automated matching may itself dictates requirements on how information is represented within the system.

Discovery-led refinement during document review

In this section we consider the effect of discoveries made during manual document review. The way that this affects the flow of work leads us to consider issues of cognitive ergonomics in relation to the design of document review system interfaces.

One of the major overheads in e-discovery arises from the need to employ appropriately knowledgeable and experienced staff people to individually review documents and record their relevance to one or more issues under investigation. These documents typically arise as a result of a broader information seeking strategy, such as search. Hence, document review is a particularly expensive information triage activity (Buchanan and Owen 2008).

Document review system interfaces tends to exploit a common information architecture design referred to by Tidwell (2006) as a Two Panel Selector. Using this design, which is common to many familiar programs such as email clients, folders are presented in a side-bar for users to select. Selecting a folder displays its contents as a list in the centre of the screen. As a further application of the pattern, a document text can be displayed by selecting list items. In addition, tick-boxes associated with each document allow the user to assign metadata codes to designate properties including privilege and relevance to defined issues in a matter.

Based on our interview data, we focus here on two issues of discovery-led refinement in relation to the use of these systems during document reviewer. These are: the identification of classes of irrelevant documents, and the identification of related relevant documents.

Identifying classes of irrelevant documents

We use the notion of *discovery-led refinement* to refer to the ways in which discoveries that arise through engagement with the materials of an investigation can result in investigators re-framing the problems that they are dealing with in such a way that they develop new goals and strategies to address them. We have also distinguished between two kind of discover: discoveries about the domain under investigation and discoveries about properties of the document collection which represent at least one part of the investigators' evidence. The identification of classes of irrelevant documents is knowledge of this second kinds.

Our interviewees who had been involved in document review reported that by far the majority of documents that they viewed were irrelevant to their investigation and that review could be a fairly tedious activity. Consequently, review typically involves reading irrelevant document after irrelevant document. However, they also said that as they progressed, they began to notice types of irrelevant document and familiarity with these could help them work more efficiently.

For example, one trainee assigned to a large document review described working through a "massive" folder of documents. She noticed that a number of documents significantly predated the events that were under investigation. This information provided all that she needed to know in order to judge them irrelevant. And so she described adapting her strategy; for each new document the

first things she looked at was the date (the documents did not have metadata denoting date and so she was unable to use an automated filter search). If she saw that the date was out of range the she could tag it as irrelevant without any further inspection. Using this strategy, and given the number of documents which fell outside the range of interest, she felt she was able to reduce overall time and effort necessary to review the folder.

Another trainee performing manual document review noticed that among the documents he was reviewing there were a significant number of invoices. He recognised that, in virtue of the matter that he was working on, invoices would simply not be relevant. Consequently if he could make the 'invoice' determination early on for each new document, he could work more quickly and with less effort. Given that invoices, and in particular invoices from a given company, have predictable surface features concerning layout, he visually scanned for these on any new document. He became accustomed to identifying invoices on low-level visual cues rather than detailed reading. Another recognition cue that this participant used was a pattern that he observed in the way these documents appeared in the sequence of documents,

> P12: … you would get the invoice followed by the cover letter, every time, and there was a whole series.

In both of the above cases, a reviewer becomes aware of the existence of a subset of documents within a wider set through exposure to subset instances. The process is one of induction. This induction, combined with recognising characteristic cues allowed these irrelevant documents to be identified quickly.

However, participants also reported that this identification was subject to a priming effect. Where multiple members of such a set were found in quick succession the strategy would be 'to hand', whereas temporal separation between exposures to subset members could slow the recognition process. This corresponds with a reduction in the priming effect over time. In other words, the greater the separation between two documents which were irrelevant on similar grounds, the greater time that would be taken to make that determination. We consider the implications of this after we consider as similar phenomenon: the identification of related relevant documents.

Identifying related relevant documents

Our participants reported a similar effect in relation to the identification of emergent subsets of relevant documents. In working through their allocated documents, reviewers became to familiar with the narrative or 'story' underlying the documents. They reported that understanding this narrative helped them to interpret subsequent documents relating to the same issue. However, the narratives could be complex and technical, and long lapses between exposures to documents related in this way could slow down the interpretation process. As one trainee said,

> P9 [...] it's easier if you've just, say if you've done this over the course of three weeks, it's much easier if you've just read the document that related to it, to read the next one and it makes it quicker to read it because you don't have to go, what was that about again? Why did I think that was relevant? […] so it's helpful if then the next

> document that's relevant to that tricky point is next to it because then you can just use the same knowledge as opposed to having to reconstruct it two weeks later.

The learning effect here is similar to that of recognising subsets of irrelevant documents. Familiarity with a subset supports more efficient decisions about members of that subset—only in this case the subsets are relevant documents. However, increasing the interval between exposures to members of the same subset increases the cognitive effort involved in recognition. To confound the task further, temporal separation between exposures to related documents also meant that multiple threads of narrative needed to be tracked simultaneously. Each may impose interference effects on the other, add additional cognitive load to the review, and impede the efficiency and effectiveness of relevance recognition.

Supporting the development of interests during review

These case-study examples lead us to formulate two hypotheses concerning the document review task—that the efficiency and effectiveness of reviewers' relevance judgements are adversely effected by:

a. temporal separation between exposures to similar, irrelevant documents.

and,

b. temporal separation between exposures to related, relevant documents.

These hypotheses are based on the reflections of document reviewers extracted from a series of unstructured interviews. They remain to be tested over a larger sample using objective performance metrics. However, they to draw attention to the issues of cognitive ergonomics, which, if we understand them better, might provide additional leverage for addressing the excessive overheads imposed by manual review in e-discovery.

Both hypotheses relate to the order in which documents are encountered and a proposed effect on performance. If they are correct, they might be explained through reference to cognitive momentum surrounding particular issues such that congruent stimuli are easer to process more quickly. Intervening time and tasks can fracture that momentum. Priming effects of this type are well known in psychology and have been studied extensively (for example, see McNamara, 2005).

As with the structural decomposition of lines of enquiry, an important part of the value of understanding these issues depends on the leverage they offer for the design of supporting tools. Since the question concerns the timing and order in which documents are experienced it is also a question about how a document review system leads the user from one document to another during the review process. Document review systems typically display documents in list form. Each new document that is viewed is simply the next one in the list. What we might consider, then, is whether alternative designs could help the user make strategically informed decisions about which document to inspect next, each time that decision is made.

A significant challenge to this, however, is that the classes of document emerge inductively; we cannot predict a priori what the interesting relationships between documents will be. Nevetheless, we can think in general terms about tools and representations which might respond to the dynamic development of interest. In considering these questions we suggest two possible approaches. The first, which we consider briefly, is to offer relevance feedback mechanisms at the review interface which allow the user to track down documents related to a given document exemplar, whether of the relevant or irrelevant type. Traditional relevance feedback mechanisms, however, may offer a rather limited option, based as they are on concept searching techniques, such as latent semantic indexing. These approaches characterise documents in terms of their lexical content and this may underestimate the richness of cues necessary for making the associations users want to make. For example, the fact that a user recognises yet another invoice by its structural cues leaves open the question of whether such documents could be discriminated in lexical terms, or whether it would be necessary to extend relevance feedback to address structural features.

We will develop the second approach in a little more detail. Hypothesis b. concerns separation between exposures to related, relevant documents. An approach, which has the potential to address this specific problem, is to represent documents at the interface using an interactive information visualisation. Information visualisations display document sets in ways that reveal properties and relationships between documents graphically. They can impose structure on a dataset. This can help the user shape and control the flow of information they receive (McNee and Arnette 2008). However, there are many properties and relationship that can be presented, and an open question is what would assist users in deciding where to go next during e-discovery review.

One solution is to use tools that automatically cluster documents on the basis of lexical similarity prior to the main review. Solomon and Baron (2009), for example, propose this strategy for exactly the reasons considered here; as a means of helping reviewers maintain ‘context’ and so improve review efficiency. An example commercial product of this type is the Attenex Patterns visualisation. Attenex Patterns displays documents as a series of embedded clusters according to relationships determined through the analysis of term distributions within the document collection. Documents with related content are shown in proximity and the user can exploit these associations to consider related documents together rather than in isolation. McNee and Arnette (2008) argue for the use of review productivity as a metric for assessing the value of visualisations such as this and using decisions per unit time as the measure of productivity. Accordingly, they claim improvements in excess of one order of magnitude compared to traditional systems. This, of course, does not include any assessment of decision quality.

Semantic proximity based on the words in a text is one possible way of relating documents, but there are others. An increasingly high proportion of documents recovered during e-discovery are emails. One way of associating emails in a potentially meaningful way is to represent them in terms of discussion threads as inferred from their subject lines. In Figure 2 we show a prototype visualisation designed to do this called ‘ThreadsVI’. Threads VI is shown populated with a set of emails which are derived from a keyword search over the Enron email collection (as collected and prepared by the CALO Project at SRI

[http://www.cs.cmu.edu/~enron/]). The search returns 88 emails sent between Jan 2000 and Oct 2001 relating to a research collaboration that took place between Enron and another organisation.

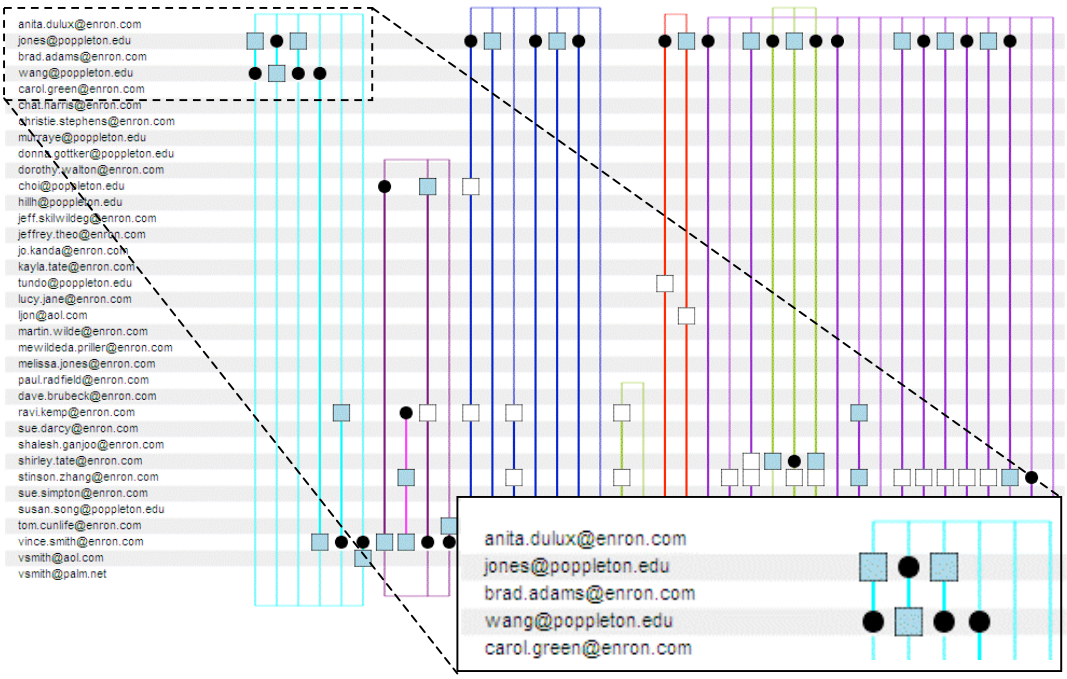

Figure 2. 'ThreadsVI', a prototype visualisation which shows emails linked by discussion thread. A dark 'blob' at the intersection between an email and an address shows the email sender; a coloured square shows who the email was sent to; and a white square represents a 'Cc:' recipient.

In the representation, each email is shown as a vertical line. Emails belonging to a common thread are represented in the same colour and are linked at the top and bottom. Email addresses are listed down the left hand side of the interface (these have been anonymised in the figure). A dark 'blob' at the intersection between an email and an address shows the email sender; a coloured square shows who the email was sent to; and a white square represents a 'Cc:' recipient. (The 'exploded' rectangular area, which is produced as part of figure 2 only, shows this more clearly). Clicking with the mouse on a blob or a square opens the email in another window.

The idea behind ThreadsVI is to present the user with a 'visual index' of an email collection that can inform choices about what to select next. If an email proves interesting then the user can identify other emails that are likely to reveal more of the underlying narrative. An additional benefit is that the interface makes global properties of a set of communications available at a glance, such as who is prominent in a discussion, who is more peripheral and who initiates communication.

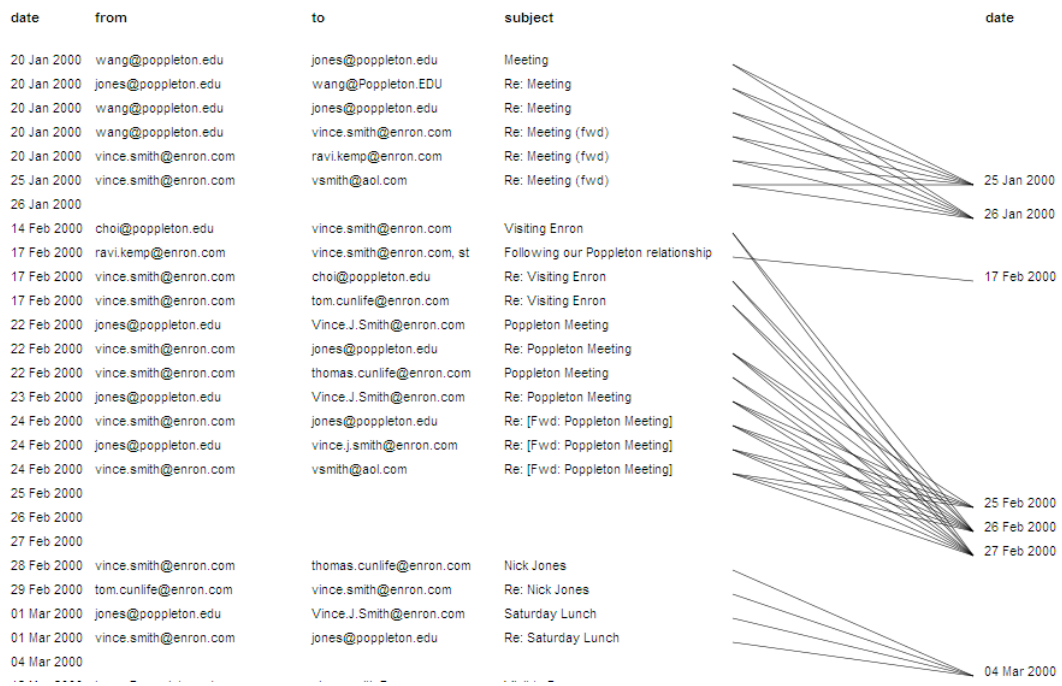

Figure 3. 'EventsVI', a prototype visualisation which shows emails linked to dates mentioned in the text. Emails are linked to date representations according to references within the email texts.

Another approach we are exploring is the 'EventsVI' visualisation, shown in figure 3. EventsVI is motivated by the observation that lawyers frequently construct event chronologies to help them make sense of documents in e-discovery (Attfield and Blandford 2008). However, information about a given event, such as a meeting between protagonists, can be distributed across multiple emails. Consequently, there should be value in drawing together emails that refer to common events. EventsVI does this by showing emails in a chronological list view (anonymised in figure 3) with individual emails linked to date representations according to references within the email texts. EventsVI was constructed using the same data set as ThreadsVI. Given that our aim is to consider its value as a representation, rather than to evaluate any particular date extraction technology, the representation was built around a hand-coded index. However, the visualisation is interactive, in that the user can inspect the full text of any email by clicking on its email. An advantage of EventsVI is that the user can see instantly which emails are linked by their discussion of a common event both before and after and, for that matter, which dates are subject to more discussion.

ThreadsVI, EventsVI and Attenex Patterns are discussed here to demonstrate the idea that interface design might play an important role in allowing users to maintain cognitive momentum. The basis for this is the idea of providing information that can help users to make informed decisions about what documents they look at and in what order. Further research is required to understand what designs work well and the scale of the impact that they can have on the e-discovery review challenge.

Discussion

In this paper we have reported findings from a series of socio-technical case-studies of e-discovery investigations conducted in a large law firm. These findings draw attention to the role of *discovery led refinement* concerning both the domain under investigation and an evidential document collection. Discoveries lead investigators to reframe their goals and restructure their tasks in the interests of efficient and effective working.

By their very nature, e-discovery investigations can be uncertain. However, by identifying the kinds of developments that can occur we are in a better position to design for them. Investigators need to establish effective ways of managing decomposition followed by coordinated integration, and the systems that they use can and should play an important role in this. Systems designed to support sensemaking, whether this be searching, filtering, extracting, constructing schematic representations, presenting a story, or integrated combinations of these, need to reflect the way that users naturally structure their problems. They need to support users in making sense of parts, and in making sense of the whole.

We began by showing how new discoveries can lead to new lower-level lines of enquiry. Essentially, these exploit new knowledge to form multiple re-specifications of the investigation problem which are more focused and more tractable. Our analysis of the structural composition of a line-of-enquiry reveals a recursive framework of conceptual entities which can be used to describe recurring elements associated with multiple, embedded lines of enquiry.

This recursive framework structures large-scale sensemaking challenges as ‘investigations within investigations’ based on supporting the definition of recursively embedded investigation contexts. By reflecting this framework in design we anticipate that it is possible to be responsive to the gradual focussing of an investigation through discovery and to support collaborative work by allowing investigators to focus on particular areas of investigation at different levels of granularity and to the exclusion extraneous information, whilst also allowing upward and lateral propagation of information from one context to another.

We have also considered *discovery-led refinement* as this occurs during document review and related this to design possibilities. Document review is a time-intensive and demanding task. A problem with traditional review system interfaces is that emergent goals of identifying classes of irrelevant documents and identifying classes of related, relevant documents are not well supported. Recognising document classes and their signature characteristics allows reviewers to employ strategies for increasing their efficiency. However, interfaces that enforce temporal separation between exposures to related documents and interfere with users’ cognitive momentum. In relation to this, we have reviewed potential solutions including visual representations of document collections which draw attention to features that are predictive of document relatedness.

Finally, the findings we have presented are drawn from exploratory studies of e-discovery undertaken in the field. We began with the broad aim of understanding how people doing e-discovery structure and coordinate action. Our study approach was to gather data about a complex, collaborative activity that would

help us to identify opportunities design approaches that might improve support for how that activity is conducted. In doing so we have outlined findings that we believe point in some valuable directions. More research is necessary to develop a detailed understanding of the problems and to test solutions. In the spirit of the special issue we have aimed to introduce some practical issues encountered during e-discovery investigation. We anticipate that what we present here provides some outline of these issues whilst also outlining the directions in which these issues point in terms of supporting system design.

References

Attfield S, Blandford A (2008) E-disclosure viewed as 'sensemaking' with computers: The challenge of 'frames'. Digital Evidence and Electronic Signature Law Review 5.

Baron J, Braman R, Withers K, Allman T, Daley M, Paul G, (2007) The Sedona Conference best practice commentary on the use of search and information retrieval methods in e-discovery. The Sedona Conference Journal 8:189-223

Benedetti V, Castellani S, Grasso A, Martin D, O'Neill J (2008) Towards an Expanded Model of Litigation, DESI Workshop on Supporting Search and Sensemaking for Electronically Stored Information in Discovery Proceedings. http://www.cs.ucl.ac.uk/staff/S.Attfield/desi/DESI_II_agenda.html. Accessed 14 December, 2009

Brassil D, Hogan C, Attfield S (2009) The Centrality of User Modelling to High Recall with High Precision Search. In Proc. of the IEEE International Conference on Systems, Man and Cybernetics 2009.

Buchanan G, Owen T (2008) Improving skim reading for document triage, Proc. 2nd Int. Symposium on Information Interaction in Context

Conrad, JG (2007) E-discovery revisited: A broader perspective for IR Researchers', DESI Workshop on Supporting Search and Sensemaking for Electronically Stored Information in Discovery Proceedings. http://www.umiacs.umd.edu/~oard/desi-ws/papers/conrad.pdf. Accessed 14 December, 2009

Johnson JC (1990) Selecting ethnographic informants. Sage, CA

Kaplan, A (2008) A Conversation with Corporate Counsel: e-Discovery Trends and Perspectives, Industry Research Report, http://pdfserver.amlaw.com/ltn/conversation_corp_counsel.pdf, Accessed 14 December, 2009

Klein G, Moon B, Hoffman R (2006) Making sense of sensemaking 2: A macrocognitive model. IEEE Intelligent Systems, 21(5):88-92.

Klein G, Phillips JK, Rall EL, Peluso DA (2007) A data-frame theory of sensemaking. In: Hoffman R(ed) Expertise out of context: Proc. of the Sixth International Conf. on Naturalistic Decision Making (Pensacola Beach, Florida, May 15-17, 2003). Lawrence Erlbaum, US pp 113-155

Klein HK, Myers MD (1999) A set of principles for conducting and evaluating interpretive field studies in information systems. MIS Quarterly 23(1):67-94

Luthans F, Davis TRV (1982) An idiographic approach to organizational behavior research: The use of single case experimental designs and direct measures. The Academy of Management Review 7(3): 380-391

McNamara RP (2005) Semantic priming (Essays in Cognitive Psychology), Taylor & Francis, NY

McNee SM, Arnette B (2008) Productivity as a Metric for Visual Analytics: Reflection on E-Discovery, Proc. of the 2008 Conference on Beyond Time and Errors: Novel Evaluation Methods for Information Visualization.

Paul GL, Baron JR (2007) Information inflation: Can the legal system adapt? Richmond Journal of Law and Technology, 13(3), http://law.richmond.edu/jolt/v13i3/article10.pdf. Accessed 14 December, 2009

Pirolli P, Card S, (2005) The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In: Proc. International Conference on Intelligence Analysis (McLean, VA, May 2-6, 2005). https://analysis.mitre.org/proceedings/index.html. Accessed 14 December, 2009

Russell, D. M., Stefik, M. J., Pirolli, P. and Card, S. K. (1993) The cost structure of sensemaking. In: Proc. of INTERACT '93 and CHI '93 Conf. on Hum. Factors in Comp. Sys. (Amsterdam, The Netherlands), ACM Press, New York, pp 269-276

Socha G, Gelbmann T (2009) Strange times: 2009 Socha Gelbmann report, Law Technology News, August 2009.

Solomon RD, Baron JR (2009) Bake offs, demos and kicking the tires: A Practical Litigator's Brief Guide to Evaluating Early Case Assessment Software and Search and Review Tools. http://www.kslaw.com/portal/server.pt. Accessed 14 December, 2009

Strauss A, Corbin J (1998) Basics of qualitative research: Techniques and procedures for developing grounded theory. (2nd ed). Sage Publications, London.

Thomas JB, Clark SM, Gioia, DA (1993) Strategic sensemaking and organisational performance: Linkages among scanning, interpretation, action and outcomes. Academy of Management Journal, 36:239-270

Tidwell J (2006) Designing interfaces: Patterns for effective interaction design, O'Reilly, CA.

Willisden G (2009) March of the regulators, New Law Journal, 159(7356) http://www.newlawjournal.co.uk/nlj/content/march-regulators. Accessed 14 December, 2009

Zhang X, Qu Y, Lee Giles C, Soong P (2008) CiteSense: Supporting Sensemaking of Research Literature. In: Proc. of the CHI '08 Conf. on Hum. Factors in Comp. Sys., (Florence, Italy), ACM Press, New York, pp 677-680.