

# ***Predictive Coding, Storytelling and God: Narrative Understanding in e-Discovery***

*Lawrence Chapin<sup>1</sup>, Simon Attfield<sup>2</sup>, Efeosasere Moibi Okoro<sup>2</sup>*

<sup>1</sup> Lawrence Chapin is an attorney in New York City engaged in managing complex e-Discovery projects for a global e-discovery provider.

<sup>2</sup> Department of Computer Science, School of Science and Technology, Middlesex University, The Burroughs, London, NW4 4BT, United Kingdom.

## ***Abstract***

The potential for gains that predictive coding can offer to the e-discovery review process lies in the power that it has to generalise human relevance decisions from a relatively small set of exemplars to a larger collection. These gains, however, depend as much on the theory of relevance developed by the human handler, and the review team in general, as on that developed by the machine. After differentiating human and machine theories of relevance and discussing the way in which these symbiotically co-evolve during predictive coding iterations, we examine the human theory in more detail. We draw on sources from philosophy, psychology and practice to make the case for the centrality of narrative, or story-telling, in constructing and anchoring the human theory of relevance. We then use this to inform a series of guidelines for structuring predictive coding based review around the notion of narrative. These include: narrative frame-working; differentiating sub-plots and episodes; designating a chief storyteller; using external representations; and narrative sharing that leaves no reviewer behind.

## ***1. Introduction***

Of all the challenges facing e-discovery practitioners, none is more daunting than that which Stuhldreher (2012) calls searching for that needle-in-the-haystack in masses of electronically stored information in all its new and evolving forms, and identifying that comparatively small set of documents that are relevant to the matter at hand, and from among those, finding the rarer documents that really matter, that truly mean something. Practitioners are asked to do all this, and do it well - effectively and efficiently.

The coming to pass of this workshop is as good a proof as any that the traditional solutions to these twin problems of volume and complexity (The Sedona Conference, 2009) are increasingly seen as inadequate, and that the time has come for advanced search techniques and practices that can actually deliver what an increasingly overburdened legal system requires. Keyword search, it is argued, exhibits poor performance and attorneys frequently resort to manual review (Stuhldreher, 2012). But manual review also often performs badly (Hogan, Bauer, & Brassil, 2010), with reported recall figures varying from 70% (Dysart, 2011) to as low as 20% (Grossman and Cormack, 2011). Manual review is also expensive with companies sometimes settling simply to avoid the cost of disclosure (Stuhldreher, 2012), leading to what has been called the “weapon of discovery” (Losey, 2013).

There is growing testimony that predictive coding, in its various forms, offers significant advantages over keyword search and exhaustive manual review. The much discussed watershed decision of the United States District Court in *da Silva Moore v. Publicis Group* (2012) is certainly more than just another voice

in that growing chorus. The enthusiasm is understandable. Predictive coding has been estimated to cut the cost of manual review by 50% to 70% (Stuhldreher, 2012), by leveraging the relevancy judgments of expert human reviewers concerning collection subsets. The human reviewers are asked first to render their judgments, whereupon the power of the predictive technologies to generalize them is loosed upon the wider collection.

This entire process hinges upon cognitive models. Documents chosen according to a human model are presented to a predictive coding system in order for the system to construct its own model, from which it identifies those documents in the larger set which it thinks are relevant. The results are then looped back to the reviewers who may widen the sample and perform further manual review, and re-present the results back to further refine the system model. So, an iterative process ensues, designed to narrow a very large collection of documents towards a subset which are responsive to the matter at hand - and at lower cost than other forms of review.

Important questions arise, however, about the deployment of predictive coding in practice. As an application of supervised machine learning, predictive coding explicitly places the human in critical phases of the loop. Effectively, the machine leans on its human handlers and their evolving grasp of the material at hand. It makes sense to us, therefore, that predictive coding may often depend on how humans are supported in developing and consolidating their own understanding. Further, we propose that the human understanding of what matters in any one case may very well rest upon how well the human reviewers apprehend the case as a *narrative* or a series of narratives. We argue that making sense in those settings entails elaborating and testing narratives constructed as plausible explanations for the facts uncovered — that, at its heart, such managed review is storytelling.

In this paper we draw upon sources which we believe make a strong case for structuring the processes of e-Discovery and predictive coding in particular around narrative representation of evidence. We then proceed to offer a number of recommendations for doing so. Our comments are most directly informed by the recent experience of the first and second authors in dealing with regulatory investigations - as a practitioner in the field in the first instance and an academic researcher in the second. We believe, however, that these principles bear strongly upon the field where regulators play because they are principles that hold wherever stories are being told to decision makers, including those in more traditional litigation.

In the next section we begin by exploring the distinction we draw between machine theories and human theories. In section three we draw on sources from a range of fields to elucidate the centrality of narrative in the human theory of relevance in e-discovery. We follow this section with some proposals for placing narrative at the heart of the deployment of predictive coding in e-discovery.

## ***2. Machine Theories and Human Theories***

### ***Finding God***

Not long ago in his widely read and influential blog, one of our leading practitioners of cutting edge electronic discovery, who is also one of the leading proponents of predictive coding, very publicly confessed that the computer he was using in one of his signature searches had been stumped. Several weeks before, Ralph Losey had begun a series of searches in the well-trodden Enron collection (Losey, 2012a). His purpose, he explained in the blog, was to encourage more attorneys and litigants to use

predictive coding. His chosen field was a nearly 700,000 document slice of the Enron database consisting of emails and attachments, organized in three separate data sets, and running from a time that the company was in high cotton, to the time of its dissolution. His chosen quest was the search for documents related to involuntary employee terminations, on the thinking that distinctions between voluntary and involuntary terminations involve "fine relevancy distinctions." To the best of his knowledge it was a first, Losey said: a "blow-by-blow" description "from the trenches" of a predictive coding project. Now, in week seven of that narrative undertaking, he said he ran into an unexpected limitation of his methodology. That is, his computer's focus was "too myopic to see God."

More specifically, one of the documents that the computer found to be a very close call was an email containing an inspirational message about God of the kind you find on the internet: God would ask not about the size of your home, but who you welcomed in; not about the kind of car you drove, but how many you helped by driving them. "It was kind of funny to see," Losey wrote, "that this email confused the computer, whereas any human could immediately recognize that this was a message about God, not employee terminations. It was obvious that the computer did not know God." (Losey, 2012b).

It is important, of course, not to make too strong a point here. As with any other subject matter, appeals to the divine could certainly be made the object of effective search methodologies. Nonetheless, it struck us as ironic that such a lawyerly observation on the apparent inaccessibility of the divine should come in the context of a novel embrace of the narrative form. After all, it is to narratives themselves that the people of many cultures and many times have resorted in their attempts to dig down into and construe the events of their lives. For reasons both psychological and philosophical that we will discuss here, answers to quests for meaning have long come in narrative form. In a sense, narrative has been a carrier of meaning. Indeed, western theology in the last part of the 20<sup>th</sup> century was marked by an epistemological approach still known as Narrative Theology. Yet, here paradoxically, at least one slice of an innovative effort to tell a modern story bumped into God and fell short.

Situational shortfalls, notwithstanding, of course, it is hard not to be impressed by the burgeoning technological capabilities of the offerings in the highly competitive e-discovery space. Still, even modern wisdoms about what is possible - and meaningful - can stand to benefit from reminders that come in more traditional garb. In that spirit, we offer a parable as a cautionary tale told by friends.

### ***The Parable of the Tank Detector***

We call it the Parable of the Tank Detector. It is a well-travelled story in the corridors of Artificial Intelligence departments everywhere, and we do not attest to its historicity. It does, however, say something about what can happen when we endeavor to teach machines to think.

Once upon a time, or so the story goes, the American military were developing a computer system that they could train to identify tanks on the battlefield. The approach involved connecting a 'neural network' to a camera. The training was to be done using photographs. So the design team went out into the field and took 100 photographs of scenes with tanks in various orientations - out in the open, hiding behind trees, and the like. They also took 100 photographs of scenes with no tanks present. The system would be taught using both positive and negative cases.

They split all the photographs into two sets, one for training and one for testing the system after training had taken place. Using the training set, they showed the system pictures of tanks and said, "Tank". They also showed the system pictures without tanks and said, "No tank". Each time the system would first

have a guess, and if shown to be wrong would adjust itself. A keen understanding would emerge, it was hoped, of the key features it needed to consider in making the right judgment. From entirely random beginnings the system's performance improved. It got so proficient that it could give a correct answer most of the time. The next step was to test the system on the remaining photos - the set that it had not yet seen. It behaved extremely well - perfectly in fact, categorizing every photo as either 'tank' or 'no tank' correctly. The designers decided to commission a further set of photos for more testing. The pictures came back and they were shown to the system. Only this time its performance was abysmal - no better than flipping a coin.

It took the designers a while to work out what was going on. It turned out that the original photographs with tanks and without tanks had been taken on different days. The 'tank' days happened to be sunny. The 'no tank' days had been cloudy. Each time the system was shown a photograph with a tank, it saw bright sunlight, blue skies and shadows. Each time it saw a photograph without a tank, it saw grey skies and an absence of shadows. This was the meaning of 'tank' it inferred. The designers had developed a sunny day detector, and a good one at that.

Thus we are reminded that when we teach by ostensive definition we cannot say what has actually been learned. By no means is this merely an academic caveat - it is, for example, what parents and teachers struggle with when it comes to knowing precisely what children have learned. Fine but important differences in understanding often find expression only in later behaviors. Such learning of ours is often soon enough, of course. Who cares, for example, if your understanding of 'fruit' differs from that of your child, so long as when you ask her to buy fruit, she comes back with something that you can happily put in your citrus salad. It's only when she comes back with an Idaho potato that you know that is time for another talk.

## ***A Convergence of Moving Targets***

During the predictive coding process a machine learning system uses exemplar documents to construct an internal, generalized model of document relevance. It then uses this model to discriminate relevant from non-relevant documents. Following Hogan, Bauer and Brassil (2010) we refer to the model as a *theory of relevance*. It is a theory because it extends beyond a finite set of judgments to a theoretically infinite set. Notably, the system constructs a theory of relevance based on exposure to documents and judgments about them, i.e. the theory evolves in the face of ostensive definitions. What is presented to the system as exemplars of relevance, however, depends in turn on prior human judgment by a document reviewer. Hence there is a second theory of relevance in play - that of the reviewer.

We understand the machine theory of relevance as very different in character from that of the human reviewer. The system characterizes relevance in terms that it understands - typically by mapping document characteristics, such as the presence and absence of text tokens, to locations within a multidimensional parameter space. In this parameter space relevance can be reduced to a calculation of spatial proximity. For the human reviewer, though, a theory of relevance is not so much about text tokens as a characterization of the subject matter of the case. Here judgments are concerned with whether documents are 'about' certain matters of human conduct and activity which form the focus for the case. They concern people and places, clubs and institutions, contracts and complaints - the full cornucopia of corporate life.

This mismatch is significant and arises because the two systems perform different kinds of abstraction over the raw data; effectively they have different internal languages. The question of addressing this

conceptual ‘mismatch’ or ‘misfit’ is arguably a fruitful direction and we briefly discuss two ways in which this may be done. One is for the human to understand, or translate, their interests into a language that the machine understands. This is actually what is done in keyword searching. Where the machine abstracts concepts from exemplars in terms of a multidimensional parameter space, the benefits of having users understand this would need to be offset against the costs training them on what goes on ‘under the bonnet’ and designing systems with sufficient transparency at the interface. A more user-friendly alternative might be to have the system abstract in terms that the user does. This is a perspective that has been adopted in respect of developing deeper user-centered evaluation methods (see for example, Blandford, Green, Furniss and Makri, 2008), and has been proposed as an approach for designing for more effective e-discovery (Attfield and Blandford, 2008). From a technical perspective, however, the challenges here are very big—perhaps as big as the problem of artificial intelligence itself, and certainly beyond the scope of this paper. For now, however, we note that linguistic approaches requiring limited *a priori* knowledge such as information extraction and ontology evolution may go some way to bridge the mismatch.

For now we are left with a mismatch. One way of understanding predictive coding is as a process that aims for convergence between the two, implicitly held, theories of relevance, and one that does this by systematically exposing and reconciling the relevance decisions of each. Through this process the machine’s theory of relevance evolves and is corrected in the light of human judgments. But just as the machine’s theory evolves, so the human theory is also dynamic. It too will evolve through exposure to documents and a developing understanding of the case (Hogan, Bauer and Brassil, 2010). One of the authors previously conducted a series of case studies of e-discovery regulatory investigations at a large London law firm (Attfield & Blandford, 2010). In the words of a Senior Associate interviewed in one of those case studies, “The scope of what you’re trying to do is immense and you’re having to define it as you go along”. A partner on the case similarly commented, “There’ll be constant refinement of who and what we thought was important.” Underlying the task is a complex sensemaking problem which involves the user in understanding what happened and what is important - an understanding that can be expected to be defined and refined as discovery progresses and the problem is better understood.

Sensemaking is often understood as a process of reciprocal interaction between bottom-up exposure to data on the one hand, and the top-down application of interpretations on the other (Klein, Phillips and Peluso, 2007; Pirolli and Card, 2005; Weick, 1995; Dervin, 1983). Interpretation gives meaning to data and also defines what counts as data (as opposed to ‘noise’). Exposure to data can elaborate, lend support to or challenge an interpretation, perhaps forcing it to be modified or abandoned. Sensemaking has been described as a process of placing stimuli into some kind of framework (e.g., a mental narrative, a model of others’ motives and intentions, a spatial “map,” etc., which then allows us to “comprehend, understand, explain, attribute, extrapolate and predict (Starbuck and Milliken, 1988). In e-discovery investigations, like other kinds of investigation, new discoveries cue new human theories which then become a theme for further investigation in a process of ‘discovery led refinement’ (Attfield and Blandford, 2011).

As the concept of what is important develops, so the human theory of relevance changes with it. These changes can be evolutionary or they can be revolutionary. Both, for example, are captured by Klein, Phillips and Peluso’s (2007) Data-Frame theory of sensemaking, in which sensemaking is articulated as an ongoing process of framing and re-framing in the light of data. A frame, or interpretation, which can be guided as much by our wealth of background knowledge and conditioning as the data, creates expectations. Violations challenge the frame and present a surprise, bringing the frame into question and provoking a re-assessment of current ‘understanding’.

Hence, we can understand the application of predictive coding to e-discovery as one that aims for the convergence of two moving targets, with each theory providing decisions and exemplars which help to refine the other. We now turn to the question of how to be understand the human theory in this process and from this, how to support its development and consolidation. We turn to the significance of narrative as a form for making sense of the human affairs that provide the ultimate subject-matter for e-discovery.

### ***3. The Significance of Narrative***

When faced with the question of what to do for the final argument at the close of trial, Jim McElhaney's fictional trial advocacy guru, and hero of his regular ABA Journal column, Angus responds, "People don't make their decisions with syllogisms and rational progressions of principle. Stories—not rules—are what really influence our thinking. Since the dawn of time, we have used stories to teach, explain, understand how the world works, memorialize events and instill moral values." (McElhaney, 2009). In Jim McElhaney's own words in an interview broadcast on YouTube, "Stories are what interest people" (McElhaney, 2012).

Research in cognitive psychology has yielded real evidence for what McElhaney's fictional hero knew so deep in his fictional gut. Pennington and Hastie conducted a series of studies into the comprehension of evidence in legal cases (Pennington and Hastie, 1991). In an initial study, they showed a group of experimental participants a filmed re-enactment of a murder trial. They then recorded debriefing interviews with them and analysed the recordings to see how the participants mentally structured the information they had heard. Of key interest was whether participants would internally organise the information by order of presentation, according to the legal arguments, in terms of character sketches of the various witnesses or as narratives. The analysis showed that participants organised the information into a story structure, even though it was not heard in that form. Indication for this came from observation such as events being described as causally linked (e.g. "Johnson was angry so he decided to kill him" etc). Interestingly, they also found that only 55% of the references to actions, mental states and goals were actually present in the testimony, whereas the remaining 45% had been inferred by the participants in order to form more coherent narrative episodes.

Pennington and Hastie went on to perform a number of studies designed to confirm and further explore what they had found. In one of these, they presented evidence from the trial in textual form, manipulating the order of presentation. Two orders were used: *story order*, in which the evidence was presented in a temporal, causal sequence; and *witness order*, in which the evidence was presented as it was in the original trial. This manipulation was applied independently to statements for the *defense* and statements for the *prosecution*.

A hundred and thirty college students listened to the statements. The results were startling. Of the participants who heard the prosecution in *story order* and the defense in *witness order*, 78% chose to convict. Of the participants who heard the defense in *story order* and the prosecution in *witness order*, only 31% chose to convict. The explanation offered was that the information was simply easier to understand when presented as a narrative. In the words of McElhaney, "Stories are how we understand the interrelationship of events."

This and other studies like it provided a basis for what Pennington and Hastie called the *Story Model*. According to the Story Model, people find it easiest to make sense of trial information by constructing a

narrative which accounts for and explains the evidence. It makes more sense that way. Importantly, the narrative is created not just from the evidence but also by reasoning from general beliefs and expectations about the social world. We are all deeply experienced in human practices, tendencies, motivations and reactions. Narrative allows us to leverage this social knowledge to make sense of what we see and hear.

But to borrow from the title of a book to which we are indebted in our exploration of these issues, why narrative? What is it about *story* that makes it crucial to understanding human life? (Hauerwas and Jones, 1997). How do we locate its significance in a wider range of diverse human affairs? Most important for us in this context, as a matter of practical reason, what is it about the narrative form that so penetrates to the core of what it means to be human?

If philosopher Stephen Crites is to be credited, the significance of narrative to the understanding of human affairs is that actual human experience—and by that he means the form of active consciousness itself as it operates on its horizon of events—is itself a story. He leans on no less an historical personage than Augustine for the idea that, “Consciousness grasps its objects in an inherently temporal way and that this temporality is retained in the unity of its experience as a whole” (Crites, 1971). Only narrative, Crites says, can contain the surprises, the disappointments and reversals and achievements of actual experience. It is narrative, we would say, that enables all these facets of each moment to “make sense.”

Alasdair MacIntyre (1981) makes a related claim for the importance of narrative to a deeper understanding of human life. “Modern” life and thought, he says, both tend toward a fragmentation of the self. Socially, we partition life—our life of work and our life of leisure, our public life and private. These separations, and the distinctiveness of each, are the terms in which we have come to think and feel. It is said that when we are able to understand our actions in their most basic and irreducible form, we touch the foundations of what it means to be human. But, MacIntyre says, more fundamental than even the most closely studied and cleanly defined actions to what it means to be human are what he calls “intelligible action, actions of a kind that enable us to see them as flowing from a human agent’s intentions, motives, passions, and purposes.” These are matters of context, the actions for which we may be held to account. It is when human actions escape such understanding that we are “baffled”, he says.

But such intentions can only be understood from a particular context, from the human setting in time and place; this was Klein and colleagues’ ‘frame’. And it is precisely this context that narrative provides and provokes. It occurs to the authors that the complex settings that tend to give rise to demands for electronic discovery are precisely those that can be baffling even to those who live and work within those settings. It can be hard to understand so many different actors, and their many individual actions. But according to MacIntyre, we are able to overcome our bafflement - the acts of others become intelligible - when they find a proper place in a narrative.

Take the setting of any large corporate enterprise, for example, operating a far flung network of integrated businesses where complexity and breadth of reach are marks of the prevailing culture. To share in any culture, MacIntyre (1977) says, is to share in a certain scheme of understanding and interpretations by which ordinary actions are made intelligible and the countless social transactions made possible. He reminds us how these vast enterprises can be like the realm of Shakespeare's Hamlet, challenged as he was to construct and construe the narrative according to which he would order and interpret events. Like the realm of that Denmark, our complex modern worlds can be unsettled by hitherto unsuspected truths – to use MacIntyre's phrase – precisely the kind of uncovering made

possible by the deep probing of our evolving technologies and their big data. Aha! moments can come even on the biggest stages, when things are seen not as before, but as they really are. This is the radical reframing of Klein et al – in MacIntyre the ‘epistemological crisis’ – that can only be resolved “by the construction of a new narrative...” in which the agent comes to understand how the criteria of truth and understanding must be reformulated.

Thus, we can even say, the players operating in those high stakes dramas of our world are not unlike the children of whom the psychologist Bruno Bettelheim wrote when they are challenged to revise what they think are settled and comforting understandings of the goings on around them. Bettelheim writes that stories are essential to a well-ordered childhood. (WN, 142) “Before and well into the oedipal period (roughly, the ages between three and six or seven), the child’s experience of the world is chaotic....During and because of the oedipal struggles, the outside world comes to hold more meaning for the child, and he begins to try to make sense of it....As a child listens to a fairy tale, he gets ideas about how he may create order out of the chaos that is his inner life”. MacIntyre adds here that it is Bettelheim's argument that it is from fairy tales “that the child learns how to engage himself with and perceive an order in social reality. It is the child who is deprived of the right kind of fairy tale at the right age who later on is apt to have to adopt strategies to evade a reality he has not learned how to interpret or handle.”

For Crites, MacIntyre and Bettelheim - and fabled lawyers like Angus apparently - it is through narrative that we integrate the complexity of human life into a coherent whole of intelligible action. All around us, in fact, our cultures are held in thrall by the power of stories - the stories upon which vast sums of money are won and lost every year. In his book *Story* (1997), which has been called the bible of screenwriting, Robert McKee asks why story as an art form “rivals all activities for our waking hours?” It is because, he says, quoting the critic Kenneth Burke, “stories are equipment for learning.” Our appetite for them “is a reflection of the profound human need to grasp the patterns of living.”

Given the significance of the narrative as a form for making sense of the complexity of human life, we might expect it to have a central role in investigations of human life. Here, we return briefly to the case studies of e-discovery investigations reported by Attfield and Blandford. In those studies the investigating lawyers organised information in a number of forms, but of those, chronologies played the most central role. On this topic the Senior Associate quoted earlier said, “I think it’s a very natural way for us to think here, we always use chronologies, our great organizing basis.” (Attfield and Blandford, 2011). Seeing events as a narrative allowed them to focus in on key time periods for more forensic investigation using more specific searches, develop more focused questions to ask witnesses during interviews, and to perceive events that seemed odd, inexplicable, unexplained or missing.

#### ***4. Applications for E-discovery***

Any decision to embrace an approach to managed document review will certainly mean a different operational landscape for those involved. First, and above all perhaps, it will mean a different mindset, one reflecting the grasp of relevancy as fundamentally derived from an understanding of a case as narrative. Ideally, this will be true of everyone involved, we believe, but it is especially important for those leading and managing the efforts of the group. Especially for those driving and directing the everyday affairs of the project, it is vital that there be an active and ongoing awareness of the role of storytelling in shaping both the processes and product of the team. The evolving narrative should be in everyone's clear line of sight, not lost amidst discussions of some detached notion of responsiveness.



Indeed, the notion of responsiveness should be driven by the context of narrative. The review plan should have ways of ensuring that this is the case.

So we turn now to the question that we initially set ourselves concerning implications for the deployment of predictive coding in practice. We do this in terms of a set of practice guidelines which are designed to locate narrative centrally within the thinking of a review team and within the deployment of predictive coding.

### ***Working the Frame***

As important as it is that project leadership embrace narrative as a tool in the carrying out of their mission, it is no less important that the narrative approach touch all team members – throughout the ranks and from top to bottom. This means narrative frame-working at all levels, from those who triage the general sets and shape the pools from which batches will be checked out, to those whose sometimes lonely task it is to sift through the occasional haystack in search of the elusive needle. We believe that for reasons we have discussed above, storytelling offers a heightened level of cognitive engagement, a certain traction for those involved in difficult search endeavors. The authors have seen how this works on the ground as reviewer eyes widen and light up upon securing a grasp of the connections between even obscure data points that narrative is so good at providing. In the wake of such revelation, pace quickens, and quality improves (Hamilton and Chapin, 2012).

The authors argue that there is always a place for narrative reflection, regardless of the shape of the review process, however lean it may be in turning to predictive coding to streamline data processing and reduce costs. One of the authors is acquainted with a proprietary predictive coding process now being deployed that is said to offer remarkable rates of efficiency in the expert hands of its developer. Of the six steps said to be critical to this predictive coding from inception all the way to quality assurance and production, three call for narrative engagement of some kind. The dialogues that help to launch the case, the initial reviews of samples that set further gears in motion, and the supervised learning iterations that refine the machine's product all have a role for storytelling. It is the belief of the authors that the place for narrative will always be a vital one.

### ***Threads of the Yarn: Differentiation in Sub-plots and Episodes.***

Investigations can be complex, involving multiple lines of enquiry. Typically these become differentiated as they emerge over time. These divisions can be leveraged in both the distribution of labor in the review team and in the deployment of predictive coding. On the first point, separate review teams can be assigned to and champion individual lines of enquiry. This can have the advantage of enhancing engagement of the review team and increasing their cognitive momentum by reducing task switching. On the second point, separate lines of enquiry will, by definition, have their own theories of relevance for both human and machine. For the human reviewer the theory will be characterized by a particular cast of characters, places and events. For the machine each will be characterized by a corresponding distribution of text tokens and specific areas of parameter space. By differentiating lines of enquiry during predictive coding iterations it may be possible to improve both the efficiency and effectiveness (in terms of precision and recall) of the process.

Each sub-plot may have multiple episodes of activity, or at least, can be thought of in that way. Similar to the differentiation of sub-plots, this can have implications for how review is handled in terms of both cognition and machine learning. Cognition is supported where complex phenomena are decomposed

into separate units, or 'Chunks' (Chase and Simon, 1973) which can be assigned to long term memory. Decomposing stories in this way is natural and has been understood as hierarchical (Pennington and Hastie, 1991). Naming these chunks can help memory through a process of verbal recoding (Miller, 1956). From the perspective of predictive coding, the arguments for the differentiating episodes are essentially the same as those for sub-plot differentiation. Different episodes may be characterized by different textual elements in documents and so advantage may be found in presenting them separately.

Sometimes these episodes and subplots emerge as follow up investigations into the particulars of specific events. We have found it remarkably helpful to offer the elements of storytelling as alternative ways of organizing these follow up efforts. Highly trained lawyers struggle sometimes to organize them along the formal lines of legal "issue." We have also seen them break free of such intellectual log jams when encouraged to think along the classic story lines of who, what, when, where, and why. As William Speros has very recently noted, they "are the stuff of which competent seed sets are made." (Speros, 2013). If reviewers are going to be asked to assist in the telling of a story, it makes sense, we believe, to offer structures – such as review protocols – that are themselves organized around the basic elements of narrative.

### ***The Storyteller-in-Chief***

Given the importance of narrative to the success of the overall project, one might consider designating someone as chief story teller, the person responsible for staying on top of the narrative as it evolves and for keeping that narrative in front of those working with the documents. The title Story Master has come to our minds, as a way of clarifying that aspect of a particular person's job. Although it may be advantageous to consider sub-plots separately, they may well crisscross in unexpected ways and so benefit from being related and perhaps integrated. Differentiation of sub-plots brings with it a need to enable channels for cross-fertilization. The chief story teller might take the role of what we have seen described as a 'bee' cross-pollinating between groups. Both roles were taken by a partner in an investigation we saw,

"I would go along to some of those meetings, probably once a week, and give them my overview from what I was hearing from interviews and the broader issues I was seeing, and then what I also would do whenever I had the chance was to spend some time in the room talking to people. So going to one of the groups and, "What have you seen?", "What's worrying you?", "What's this piece of paper you've got sat on the side there?", and saying "Oh well, you had better go and tell X about that because in an interview someone said Y and that feeds into that." So there's an element of informal cross-fertilization, a bit like a bee hopping from flower to flower."

This person can also take on the role of monitoring the flow of information. Are the important emails reaching their intended audience? Are they being read? What blockages are developing in the culture of that workplace? What parts of the ecosystem are threatened by collapse? In that regard, it is our experience, there is much to be said for the importance of managing by walking around, and keeping fingers on the pulse.

### ***Represent and Rehearse the Play***

The learning landscape should be marked by a keen awareness of the special role that external representations can play in enabling people to externalize and review their thinking. It is important to provide tools for representing and rehearsing parts of the story at every turn. Like laying out a jigsaw

puzzle, externalizing evidence as an ordered chronology allows people to engage in the creative practice of narrative interpretation, and in doing so, to ask ‘what if’ questions, to see gaps and to reframe their information needs.

One of the authors recently conducted a study looking into how representatives of the UK intelligence services would conduct a mock terrorist investigation using a novel interface design (Rooney, Wong, Attfield and Choudhury, Submitted). The interface allowed users to search for documents which appeared as cards on a large screen-based ‘canvas’. These could then be manipulated free-style or arranged automatically in chronological order. The participants were observed to arrange documents into chronologies and then repeatedly and carefully review them, constructing and reviewing possible explanatory narratives. This allowed them to consider the implications of what they had found and, in turn, to specify what further information they might require.

The story can be written formally as text, sketched out informally on whiteboards or paper or represented graphically as timelines. Consider the full range of options, all the way from cutting edge visual display to the old fashioned advantages of hard copy and whiteboards near at hand. The simple truth is that people learn best in different ways and so this should be characterized by an appeal to epistemological diversity. And where possible, link story elements to underlying evidence so as to attest to the provenance of the story.

### ***No Reviewer Left Behind***

It needs to be recognized that the twists and turns of any good story can be difficult to grasp. Accordingly, the project landscape should include an ongoing teaching function, an intentional effort to leave no reviewer behind, if we may use that paraphrase. This should mean formal group sessions when necessary, as well as the more frequent opportunities that arise when good, off the cuff questions are asked, and full, patient answers called for.

Bearing in mind the essentially creative nature of storytelling in all its forms, project life ought to be characterized by a thoughtful approach to the question of how reviewers are to be treated. Storytelling is at its greatest strength when used as a lever for the full range of talents, both rational and intuitive, that are likely to be found whenever groups of people are gathered. Such gifts cannot be yanked out like so many wisdom teeth. They can only be offered by those whose pleasure it is to do so (Chapin, 2011).

Added to this, different sensemakers in any enterprise come equipped with different experiences, knowledge and expectations, and these are directly implicated in any interpretive exercise, which narrative construction surely is. By offering up narratives and their supporting evidence for inspection these alternative interpretations can be brought to bear. They may not be accepted, but by exposing them there may be an increased chance of overcoming the epistemological crisis and reframing a narrative towards more fruitful directions (Hamilton and Chapin, 2012).

## ***5. Conclusions***

Perhaps it is because narrative is so much a part of what it means to be human, that the avenues for exploration of its use in e-discovery may well be endless. After all, both the producers and the consumers, both ultimate and penultimate, of narrative-driven intelligence stand to gain much from use of the narrative form. If our sources – philosophical, psychological and professional – are to be credited, we all share a natural bent towards the power of story. We all share a deeply engrained understanding

of stories as a vehicle for important values, such as "meaning." Our sources agree – stories are how we figure the world out, how we make sense of it. The footprints of millions of ordinary people as they make their way to movie theatres around the world are further witness. What we have tried to do here is to make the case that, along with predictive coding itself, the creative use of the narrative form can be a powerful lever of all the other resources brought to bear upon the challenges that electronic discovery is facing.

What we have also attempted to convey here about applying narrative skills is how natural it would seem to be for e-discovery practitioners to want to do so, indeed how almost strange it is that they have not done so yet, at least not as a general approach. After all, the end result of all the combined processes of electronic discovery is almost always storytelling of some kind – whether to another party sitting across a table, or a decision-maker sitting behind a bench, or, in some important cases, to a world hungry to know what has really gone on in important corners. Yet, explicit narrative thinking has tended to occur among only a select few.

By no means does the momentous industry turn to predictive coding signal that narrative thinking has lost its place. Rather, we argue that where such technologies are employed, it means that two parallel theories of relevance are at work, each helping to refine the other, both moving with each reiterative step towards a convergence theory. In convergence, each serves the other - the machine's grasp of the data on the one hand, and the narrative understanding of the real life context on the other.

If we are right that there is benefit in an orientation in e-discovery towards narrative then the benefits should be measurable. But what are they exactly and how might they be measured? At one level, document review is a classification task. Hence, given some benchmark 'ground truth' relevance decisions for a collection, standard IR measures of recall and precision can be used as assessments of effectiveness; time taken can be used to measure efficiency. In addition, objective measures of performance might be augmented by measures of more subjective, psychological responses. Our anecdotal evidence is that reviewers feel more engaged when there is strong sense of a narrative at the heart of an e-discovery review exercise, and with increased feelings of engagement we might anticipate greater understanding of the facts of the case. Hence psychological engagement can act as a proxy, at least, for the kinds of things we are interested in. The measurement of psychological engagement in technological environments has received increasing interest in recent years, and a number of approaches exist (Lalmas, Attfield, Kasai and Piwoworski, 2011), including a standardized questionnaire for measuring engagement developed by O'Brian and Toms (2010). Reviewer-knowledge might be additionally assessed by quizzing them about key facts after the task. In ongoing research by two of the authors, we are using measures of recall, precision, task-time, engagement and understanding to assess the impact of a narrative orientation in e-discovery type tasks. In this study, the narrative orientation extends to representations of evidence that participants are asked to create. We ask participants to perform a small investigation in which they review documents and organise selected information into graphical structures that are either chronological, argumentational or freeform (unconstrained). We intend to report the results of this study as future work.

It is certainly no little thing that the setting in which electronic discovery takes place is one in which the deconstruction and reconstruction of narratives can be dramatic, and involve high stakes, and occur under the tight constraints of court-ordered time and in the harsh glare of public light. If we are correct in our assessment of the power of narrative to address those challenges, then we are likely also correct in sensing that there is opportunity ahead for those willing and able to effectively deploy it. Further, we

even wonder, given the nature of the crises needing resolution in such matters, whether a decision to do so might not even be deemed a matter of professional responsibility.

## **References**

Attfield, S. & Blandford, A., 2011. Making Sense of Digital Footprints in Team-based Legal Investigations: The Acquisition of Focus. *Human–Computer Interaction - Special Issue on Sensemaking* 26(1-2).

Attfield, S. & Blandford, A. 2010. Discovery-led Refinement in e-Discovery Investigations: Sensemaking, Cognitive Ergonomics and System Design. *Artificial Intelligence and Law - Special Issue on e-Discovery*, 18(4) p.387.

Attfield, S. & Blandford, A. (2008) E-disclosure Viewed as ‘Sensemaking’ with Computers: The Challenge of ‘Frames’ Digital Evidence and Electronic Signature Law Review 5. Also in: G.Chandana (Ed) (2009) E-Discovery and Metadata: Evidentiary Issues, Amicus Books, India.

Bettelheim, Bruno, 1976, *The Uses of Enchantment*, New York: Alfred A. Knopf.

Blandford, A., Green, T. R. G., Furniss, D. & Makri, S. (2008) Evaluating System Utility and Conceptual Fit Using CASSM. *International Journal of Human–Computer Studies*. 66.

Chapin, L., 2011, Contract Coders: e-Discovery's "Wasting Asset"?, *E-discoveryTeam Blog* (Guest Entry) November 15 2011. Available on line at: <https://e-discoveryteam.com/2011/11/14/contract-coders-e-discoverys-wasting-asset/>

Chase, W.G. & Simon, H.A., 1973. Perception in chess. *Cognitive Psychology*, 4, p.55.

Crites, S., 1971. The Narrative Quality of Experience, *Journal of the American Academy of Religion*, 39(3), Reprinted in: Hauerwas, S., and Hones, L.G. (eds.), 1997, *Why Narrative?: Readings in Narrative Theology*, Oregon: Wipf and Stock.

Da Silva Moore v. Publicus Groupe, 2012, WL 607412 (S.D.N.Y), approved and adopted in DaSilva Moore v. Publicus Groups, 2012 WL 1446534.

Dervin, B., 1983. An Overview of Sense-making Research: Concepts, Methods, and Results to Date. *International Communications Association Annual Meeting*, Dallas, May, 1983. Available online at: <http://faculty.washington.edu/wpratt/MEBI598/Methods/An%20Overview%20of%20Sense-Making%20Research%201983a.htm>

Dysart, J. (2011). A New View of Review. *ABA Journal*, 97(10), p.26.

Grossman, M. R., and Cormack, G. V., 2011. Technology-Assisted Review in E-Discovery can be More Effective Than Exhaustive Manual Review. *Richmond Journal of Law and Technology*, 17(3).

Hauerwas, S., and Jones, L.G. (eds.), 1997, *Why Narrative?: Readings in Narrative Theology*, Oregon: Wipf and Stock.

Hogan, C., Bauer, R.S. & Brassil, D., 2010. Automation of legal sensemaking in e-discovery. *Artificial Intelligence and Law*, 18(4), p.431.

Hamilton, W., and Chapin, L., 2012, Storytelling: The Shared Quest for Excellence in Document Review, *E-discovery Team Blog* (Guest Entry), 8<sup>th</sup> January 2012. Available online at: <http://e-discoveryteam.com/2012/01/08/storytelling-the-shared-quest-for-excellence-in-document-review/>

Klein, G., Phillips, J.K. & Peluso, D.A., 2007. A Data-frame Theory of Sensemaking, in: *Expertise Out of Context: Proceedings of the Sixth International Conference on Naturalistic Decision Making*, Pensacola Beach, Florida, May 15-17, 2003, ed. Robert Hoffman, Lawrence Erlbaum Associates Inc, 2007 p.113.

Lalmas M., Attfield S., Kazai G. & Piwowarski D., 2011. Towards a Science of User Engagement, UMWA 2011 Workshop on User Modeling and Web Applications. Available online at: <http://www.dcs.gla.ac.uk/~mounia/Papers/engagement.pdf>

Losey, R., 2012a. Day One of a Predictive Coding Narrative: Searching for Relevance in the Ashes of Enron, *E-Discovery Team Blog*, 1<sup>st</sup> July 2012. Available online at: <http://e-discoveryteam.com/2012/07/01/day-one-of-a-predictive-coding-narrative-searching-for-relevance-in-the-ashes-of-enron/>.

Losey, R., 2012b. Days Seven and Eight of a Predictive Coding Narrative: Where I have another hybrid mind-meld and discover that the computer does not know God, *E-Discovery Team Blog*, 29<sup>th</sup> July 2012, Available online at: <http://e-discoveryteam.com/2012/07/29/days-seven-and-eight-of-a-predictive-coding-narrative-where-i-have-another-hybrid-mind-meld-and-discover-that-the-computer-does-not-know-god/>.

Losey, R., 2013. The Increasing Importance of Rule 26(g) to Control e-Discovery Abuses. *E-Discovery Team Blog*, 24<sup>th</sup> February 2013, Available online at: <http://e-discoveryteam.com/2013/02/24/the-increasing-importance-of-rule-26g-to-control-e-discovery-abuses/>

MacIntyre, A., 1977, Epistemological Crises, Dramatic Narrative, and the Philosophy of Science, *Monist* 60, 4, p.453. Reprinted in Hauerwas, S., and Hones, L.G. (eds.), 1997, *Why Narrative?: Readings in Narrative Theology*, Oregon: Wipf and Stock.

MacIntyre, A., 1981, The Virtues, The Unity of a Human Life, And The Concept Of A Tradition, in *After Virtue* . Notre Dame University Press.

McElhaney, J., 2009. The Arsenal of Persuasion: Make Your Closing Arguments a Really Good Story, *ABA Journal*, July 2008. Available online at: [http://www.abajournal.com/magazine/article/the\\_arsenal\\_of\\_persuasion/](http://www.abajournal.com/magazine/article/the_arsenal_of_persuasion/)

McElhaney, J., 2012. McElhaney Explains the Origin of Angus, *Youtube*, Available online at: <https://www.youtube.com/watch?v=T7Ujgwi3VRg>

McKee, R., 1997. *Story: Substance, Structure, Style and the Principles of Screenwriting*, New York, Harper-Collins.

Miller, G.A., 1956. The Magical Number Seven, Plus or Minus Two: Some Limits on our Capacity for Processing Information. *Psychological Review*, 63, p.81.

O'Brian, H, and Toms, E., 2010. The Development and Evaluation of a Survey to Measure User Engagement. *Journal of the American Society for Information Science and Technology*, 61(1), p.50.

Paul, G.L. & Baron, J.R. 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

Pennington, N. & Hastie, R., 1991. Cognitive Theory of Juror Decision Making: The Story Model, *A. Cardozo Law Review*, 13, p.519.

Speros, W., 2013. Predictive Coding's Erroneous Zones Are Emerging Junk Science. *E-discovery Team Blog* (Guest Entry), 28<sup>th</sup> April 2013. Available online at: <http://e-discoveryteam.com/2013/04/28/predictive-codings-erroneous-zones-are-emerging-junk-science/>

Starbuck, W.H. & Milliken, F.J., 1988. Executives' Perceptual Filters: What They Notice and How They Make Sense, in: *The Executive Effect: Concepts and Methods for Studying Top Managers*, Donald C. Hambrick, ed. Greenwich, CT, JAI Press, 1988, p35.

Stuhldreher, T., 2012. Predictive coding cuts discovery expenses. *Central Penn Business Journal*, 28(37), 19.

The Sedona Conference, 2009. Commentary on Achieving Quality in E-Discovery, Available online at <https://thesedonaconference.org//publication/sedona-conference%25C2%25AE-commentary-achieving-quality-e-discovery-process>

Pirolli, P. & Card, S., 2005. The Sensemaking Process and Leverage Points for Analyst Technology as Identified Through Cognitive Task Analysis, paper presented at the International Conference on Intelligence Analysis, McLean, VA, May 2-3, 2005. Available online at: [http://vadr.cc.gatech.edu/documents/2\\_\\_card-sensemaking.pdf](http://vadr.cc.gatech.edu/documents/2__card-sensemaking.pdf).

Rooney, C., Wong, W., Atfield, S & Choudhury, T., Submitted. INVISQUE as a Tool for Intelligence Analysis: The Construction of Explanatory Narratives, Paper Submitted to the 2013 IEEE Conference on Visual Analytics Science and Technology.

Weick, K., 1995. *Sensemaking in Organisations*, London; Sage.