

2.10 NATURALISING KNOWLEDGE MANAGEMENT

David Snowden, Cognitive Edge Pty

Introduction

This paper presents a science based approach to knowledge management, and related aspects of research in the fields of social and management science. The science of complex adaptive systems provides a structure within which we can understand high levels of uncertainty in which outcomes are inherently unknowable, but which can be understood and managed as evolutionary mechanisms. This requires a rethinking of risk management from probabilistic models based on possible outcomes, to an understanding of the degree of stability and volatility within the system. Cognitive science has transformed our understanding of how humans make decisions, challenging the model of human decision making as one of rational choices based on personal self interesting, to one based on matching patterns acquired through experience or narrative transfer.

The scientific use of narrative, and its fragmented and fractal nature offers considerable potential to transform knowledge management practice, and research methods. A new approach, originally developed in the field of risk assessment and horizon scanning in the field of anti-terrorism is described and illustrated. Links are made to the fragmented narrative form of social computing which has achieved considerable voluntary take up. A semi-structured indexing approach is taken to achieve a synthesis of narrative pattern analysis with human pattern recognition to reject empirical approaches to social research, within a modern scientific paradigm.

At the heart of this project is a view of meta-narrative as an emergent property or strange attractor arising from social interaction which is discoverable and actionable in the sense of quantum mechanics rather than the laws of motion. By taking narrative as a fragmented form of support for cognition, and using the ubiquity of the web and social computing, together with the representational and information processing capacity of computers we can considerably augment and enhance the natural pattern based intelligence that underpins human decision making, and more so to radically reduce interpretative conflict in the process. In consequence we can assume radically new and effective approaches to the management of risk.

What is it to manage knowledge?

The general focus here is to see knowledge management as centred on two principle objectives:

- To improve the way in which individuals and organisations make decisions by various means including but not limited to, sharing knowledge from previous experiences.
- Creating the conditions in which novel forms of knowledge can emerge, either in the context of immediate need or in research and development.

Debates on what is knowledge are as old as philosophy with multiple debates over the ages. Blackburn (1994) summarises the various debates as one between two

metaphors. The first is that of a building, which places a requirement on the philosopher to provide a secure foundation for claims as to what is known. This metaphor implies that there is some form of given together with a rational theory of confirmation and inference as a method of said construction. The second is that of a boat whose stability is due to interlocking nature of its parts which favours coherence and holism but is vulnerable to scepticism. Probably the best known is Plato's formulation in the *Theaetetus*, namely that knowledge is justified true belief. This paper takes a naturalising approach as its starting point, seeking to establish a basis for understanding knowledge based on understanding of how knowledge is formed and deployed by humans. Critically it seeks to ground that understanding in the natural sciences. It can be argued, although this is not the place, that such an approach can overcome the dichotomy of the two metaphors, seeing knowledge as paradoxically both a thing and a flow, as an electron is both a wave and a particle (Snowden 2002). As such, in discussing knowledge management we are focused on the two functions or objectives identified above. The argument here is that while definitions of knowledge remain controversial, it is possible to achieve a fair consensus on what means to manage knowledge within an organisation. A working definition, linked to the objectives above would be: *Knowledge management is process of creating sufficient shared context to allow for the effective exercise of judgement within an organisation.* That phrase is not intended or justified as a definitive definition, but more as an indicator of the approach being adopted.

This paper takes a theory based approach to the field. It uses insights from complexity and the cognitive sciences to create practice which is then reported. The paper does not attempt to derive theory from direct observation of good or bad practice in organisations; the dominant method in management science and practice. Indeed it is argued that such approaches have severe limitations. The confusion of *correlation* with *causation* represents a real issue for management science; indeed the whole issue of causation in social systems is problematic. Bennis & O'Tool (2005) attribute motivation by a not so subtle reference to *physics envy* which, while targeted at Business Schools, could equally well be applied to much social science. They accuse academics of writing increasingly obscure papers for increasing specialised audiences to achieve a pseudo-comparison with colleagues in the hard sciences. The extension of this pseudo-objectivity into the consultancy profession is endemic in the practice of knowledge management. The issue is well summarised in a delightful metaphor from Christensen & Raynor (2003) as follows:

Imagine going to your Doctor because you're not feeling well. Before you've had a chance to describe your symptoms, the doctor writes out a prescription and says "take two of these three times a day, and call me in a week."

"But – I haven't told you what's wrong," you say. "How do I know this will help me?"

"Why wouldn't it" says the doctor. "It worked for last two patients"
No competent doctors would ever practice medicine like this, nor would any sane patient accept it if they did. Yet professors and consultants routinely prescribe such generic advice, and managers routinely accept such therapy, in the naïve belief that if a particular course of action helped other companies to succeed, it ought to help theirs, too.

The metaphor represents a fundamental challenge to a case based, prescriptive approach using the benefits of hindsight (or retrospective coherence to use the appropriate CAS term). In particular it challenges one of the most common assumptions in knowledge management, namely that one of its purposes is to discover and disseminate *best practice*. One of the arguments in this paper is that avoidance of failure has had higher evolutionary value than imitation of success, and in consequence the human race is more inclined to learn from and distribute *worst practice*. Best practice has its place, but that place is restricted to highly ordered, stable situations. (Snowden 2002). A brief example will illustrate this; if your child burns his finger on a match then he learns more about fire and through many cases of avoiding such error; we learn from tolerated failure.

Equally in most organisations negative stories spread faster than those of success, and in greater number. There is at least an argument that evolutionary pressure to deal with high levels of uncertainty has meant that avoidance of mistakes is a more effective strategy than imitation of success. Narrative forms of knowledge in all human traditions have developed story forms, in particular the use of archetypes, to distribute failure without attribution of blame. One of the best illustrations of this is the Sufi's wise fool, the Mullah Nasrudin. (Shah 1985). Faced with personal failure, the approach is not to confess sin, but instead to create a Nasrudin story that will naturally spread the learning without traceability to source. One of the basic features of the naturalising approach to knowledge management is built on such natural human practices, rather than the more traditional content based approach which involves abstracting narrative into formal documents.

Avoiding the case based approach, this paper uses findings from other disciplines outside the normal range of management and social science to draw conclusions as to both the nature of knowledge management and its practice. The reference to *naturalising* places this paper within the naturalising tradition of epistemology (Kornblith 1985, Clark 1988) although it avoids a strict connectionist approach. The use of *naturalising* is also extended from its normal use in epistemology, to contrast with what will be termed an *idealistic* and normative approach. Idealistic approaches are typified by a strong dependence on case based reasoning and attempts to achieve predetermined objectives.

The current context of knowledge management

A previous paper by the author (Snowden 2002) identified three generations of thinking and practice in KM as follows:

- A pre 1995 period, characterised by a focus on growing use of technology to distribute information. The common language during this period was a Decision Support Systems (Lucas 1979), Data Warehousing and consolidation of application systems into what became known as *Enterprise Resource Planning* (ERP). The language of KM is not common in this period with some exceptions. Karl Wiig for example recalls using the term from 1984 and delivered a paper to the UN's ILO in 1986 entitled *Management of Knowledge: perspectives of a new opportunity*. This period coincides with, and is heavily influenced by increasing reliability and adoption of computing.
- The trigger point for KM as a mass movement or fad within the field of management can be linked to the publication of *The Knowledge Creating Company* (Nonaka

ka & Takeuchi 1995) which popularised the SECI model (Figure 1). This much quoted model with its distinction between tacit and explicit knowledge uses (but is not faithful to) language originally introduced by Polanyi (1962). This was not the only origin of KM. We also have the intellectual capital movement (Stewart 1997, Sveiby 1997), developments from information management (Davenport & Prusak 1998) and economics/strategy (Boisot 1998) coupled with major developments in collaborative software epitomised by Lotus Notes and the increasing commercial and private adoption of email and use of the Internet. However 1995 sees the creation of KM offerings in the large management consultants, the appointment of Chief Knowledge Officers (CKO) at or near board level and a sudden growth in conferences and trade magazines. The focus on this period, as for the first period remains on *content* and the creation and distribution of information; the SECI model is largely unchallenged. Knowledge is seen as an object existing in tacit or explicit forms and capable of transfer from one or other state.

- The third, emergent age which rejects the idea of knowledge as existing in tacit or explicit form: informed by the natural sciences, focused on *context* more than content and arguing that knowledge should be seen, paradoxically, as both a thing and flow. In 2002 this was a controversial theory but it has gained increasing acceptance, in part due to the failure of KM programs to achieve the various promises implied by the adoption of methods and software designed to make tacit knowledge explicit.

In 2007 technology has moved on with the rapid and pervasive growth of social computing, the principle advocates of which tend to reject the KM label wishing to avoid guilt by association with top down IT led initiatives. The number of KM conferences is radically reduced in number, and in attendee numbers. While some KM practices exist in consultants they are mainly focused on the tail end of a management fad life cycle (Government adoption of industrial best practice) and are strongly linked to technology implementations. There are few if any CKOs left, although there are many former ones in private practice as consultants. The need to provide knowledge to improve decision making, and to create new knowledge persist but they do so under different labels. This is part of the normal cyclical nature of management theory and the need to create novelty to gain attention but (in the contention of this author) there is something more significant in play and KM may in the future be seen as one of the last great fads, along with Business Process Engineering (BPR), Total Quality Management (TQM), the Balanced Score Card (BS) and The Learning Organisation (LO). KM in a sense may have been more resilient than most as it was not associated with single point of origin, and was also able to build on a substantial body of professionalism in the Library Sciences and elsewhere.

The reason for claiming a higher level of significance to the current changes is that all the quoted past movements have been based on a common set of assumptions about systems and human decision making which are under increasingly challenge from the natural sciences. Specifically these are:

- The assumption that humans make decisions by evaluating all available information, and are motivated by personal self-interest. This leads to a believe that if you want people to do something then you communicate well, and provide appropriate incentives. One can add some sophistication to this with game theory but the general

principle is remains. The adage of *the right information in the right place at the right time* underpins management science and practice. A computer model of information processing influences human psychology, sometimes to the point of farce; to quote “The mind is a limited capacity information-processing system that behaves in a law-like fashion” (Andrade & May 2004). With this goes the assumption that information can be interpreted consistently (in part the context-independent assumptions of first and second generation KM) through process definition, training and control.

- An assumption of repeatable and discoverable relationships between cause and effect within systems, which allow the definition of *best practice* and the creation of repeatable recipes for achieving organisational goals. The latter assumption is fundamental to the factory model of management consultancy which also grew and was maintained by BPR and linked software implementations. It also provides the basic assumption behind most research techniques in management science, in particular the validation of theory by retrospective selection of case based examples. This leads to a heavy dependency on correlation techniques and the subsequent danger of confusing correlation with causation. This was earlier referenced as a focus or concentration on the determination and then achievement of an ideal future state.

The challenge from the natural sciences has major implications for the future of management study and practice. KM was the first of the management movements to really focus on the role of humans within an organisational system other than as widgets requiring cultural change and motivation. LO also attempted this, but did not engage with the process aspects of the business, it stayed separate and did not in consequence face the problematic nature of interaction between human knowledge and the need for encapsulation with systems. This was the challenge taken up by KM. BPR in contrast attempted to confirm human practice to system need, LO simply ignored it. As a result some of the early application of new scientific thinking emerged from KM practice. In the next two sections we will look at the natural science challenge to both of the above assumptions before continuing to look at the implications for knowledge discovery and creation.

The pattern basis of human intelligence

Klein (1998) established in a seminal work on decision making that humans make decisions on a first fit pattern match either with past, or hypothesized future experience. Critically he established the choice of patterns is one of *satisfying* not *optimising*; it is not the *best fit*, but the *first fit* patterns which are used. This is radically different from the information processing, rational decision maker identified earlier. In a classic experiment (reference) created a video with six students, half dressed in white, half in black, who pass two basket balls between them. The viewer is instructed to count the number of times those dressed in white pass the basket ball. During the short sequence that follows a student dressed in a gorilla suit crosses the screen, beats her chest and exits stage left. On completion of the exercise observers offer ranges of answers as to the number of passes, few if any see the gorilla. The reason for this is of considerable relevance and corroborates and extends Klein’s original work. We do not scan all the information that is in front of us, typically 5–10 per cent. Based on this partial scan we match against patterns stored in our long term memory and perform a first fit pattern

match against those patterns. To do anything else would be to deny our evolutionary inheritance.

The position is even more complex as cognitive bias, or partial data scanning is linked into our cognitive development. All mammals have an extended period of post natal plasticity in the development of the brain which has profound influence on decision making and knowledge use in later life. During our early years imitation and other factors means that the basic patterns of our brain are profoundly influenced by those to whose behaviour we have to adapt namely family and peer groups; *One persons nature is another person's nurture* (Wexler 2006) While we maintain plasticity through out our lives it is at its most open to new patterns in the period before puberty. From that point onwards we seek to shape the world to our expectations, a practice which becomes stronger from our mid twenties. Our understanding of the phylogenetic emergence of human knowledge no longer permits the notion of the self-interested, atomistic and rational information-processor model of humankind which has formed the basis of most knowledge management practice.

Interestingly one group of people make rational decisions based on information processing, and in other than mild cases this is considered an educational handicap; they are autistic. In human terms computers are autistic, they are simply very fast at what they do. We do have limited capacity information processing capability (reference back) but it is not the basis of our intelligence. Our considerable capacity to utilise and blend patterns is the basis of our intelligence, and evolutionary adaptability is more associated with pattern utilisation than information processing. Our ability to link and blend patterns in unusual ways, known as conceptual blending (Fauconnier & Turner 2002) gives us ability to adapt rapidly to changing context and critically to innovate as well as to use that most powerful tool of explanation, knowledge transfer and teaching, metaphor.

By implication much of KM and Management Science have been operating off a false model of human decision making, worst still a model that that if instantiated in process would reduce cognitive creativity. Interestingly in 2007 3M report an abandonment of Six Sigma (an extreme form of BPR) in its research function (reference) and a general restriction on its use because it was reducing innovation capacity. Practice is starting to match theory.

If humans are pattern processes, then a large part of KM will involve the management of those patterns, both stimulating relevant ones to the forefront of the long term memory, disrupting established patterns to create the preconditions for innovation and increasing the number of patterns available, and their contextual relevance for decision makers. In this understanding we see the real departure of KM from information management, and the criticality of context and the creation of shared context in KM practice. It is also beholden us to pay attention to where those patterns come from. Our genetic inheritance predisposes our response to experience (nature is informed by, and enables nurture), our experiences, in particular those of tolerated failure create vivid patterns through which we filter data and make decisions but neither of these are sufficient on their own to account for human knowledge. All animals and in particular mammals have the same capacity. A major distinguishing feature of human intelligence has been our ability to manipulate our environment, to create cultures that increase familial and tribal bonds and to pass on knowledge other than through genetic evolution and experience: we are in our very essence storytellers. Not only that the greater part of our evolutionary history has been spent in an oral tradition and it is at least arguable that

the modern environment of social computing, comprising multiple fragmented conversations is a return to that.

Stories are also fractal in nature, and are linked to common work and social group experience. When a family assembles for a wedding or funeral, the family members will retell the identity stories of their family. The same is true of work groups, organisations, cultures all of which are self-similar and provide a capacity for common action. Engineers working on a long term project, create stories that define the experience and key learnings that they derive. Mentors tell stories of their own experience to those they mentor, and those mentored, in their turn modify those teaching stories and create their own. To understand what we know and how we know it, and by implication how we make decisions, we need to understand the multi-faceted narratives of our day to day discourse.

We will return to that subject when we turn to the modern practice of naturalising approaches to KM, but before doing that we need to look at another key aspect of modern science which has major implications for our understanding of human decision making, social organisation and polity.

Complex adaptive systems

From the creation myths of virtually all societies, from common conversation and from research methods we are familiar with two types of system:

- **Ordered systems**, in which there are repeating relationships between cause and effect which can be discovered by empirical observation, analysis and other investigatory techniques. Once those relationships are discovered we can use our understanding of them to predict the future behaviour of the system and to manipulate it towards a desired end state. In science we see this capability in the laws of motion and thermodynamics, experimentally derived with predictive capability. Critically in an ordered system, the nature of the system constrains agent (an individual, a group, an idea, etc.) behaviour to enable that predictability.
- **Chaotic systems** in which the agents are unconstrained, and present in large numbers. For this reason we can gain insight into the operation of such systems by the application of statistics, probability distributions and such like. The number, and the independence of the agents allows large number mathematics to come into play. In recent times we have seen some popularisation of this with varying degrees of success and intellectual integrity in the concept of The Wisdom of Crowds (Surowiecki 2004) in a range of management practice including knowledge management.

However from Chemistry (Prigogine 1984), Biology (Kauffman 1995) Economics (Arthur 1994) and increasingly Social (Jenks & Smith 2006) and Management (Axelrod & Cohen 1999) Science we are gaining an awareness of a third type of system, namely a complex adaptive system (CAS). In a CAS agent behavior is loosely constrained by the system, but in turn the agents constrain the system. As a result of this double interaction we have inherent unpredictability as the system will not return to an equilibrium state after it is disturbed, hence the phrase far from equilibrium systems which is sometimes used in preference to CAS. In such systems the agents are adapting to proximate interactions with other agents and environment. Any order or structure is emergent and only repeats, if at all by accident not by design. Of particular importance

is that due to these characteristics complex systems are highly susceptible to minor changes or weak signals. This is sometimes illustrated by the cliché that the flapping of a butterfly's wing in the Amazonian rain forest causes a hurricane in Texas, but it is a poor way of explaining the point. Douglas Adams, author of the *Hitchhikers Guide to the Galaxy* once satirically suggested that if this was the case killing the butterflies would prevent future hurricanes. There is no linear causality between the two events, the point is that multiple small perturbations within the system can produce unexpected results. There is a known condition called fundamental attribution error (reference) which describes the tendency to attribute cause where no cause exists. This is less a problem for someone brought up in, for example, a Daoist tradition which has retained an idea of a non-causal system lost to western philosophy.

The phase change between types of system is important and is more fully treated elsewhere (Kurtz & Snowden (2003)). However it is important to realise as constraints in an ordered system can easily produce the conditions under which that system shifts to being complex, or collapses into chaos. For example attempting to exert excessive control through bureaucracy may result in a slow build up of tension through frustration which leads to a collapse of control and increased levels of fraud. Because it is impossible to get anything done, people find work-arounds which enable the system to work despite itself disguising stress until the system breaks. A good example here is using receipts for legitimate expenditure to cover forbidden expenditure; banning expenditure on food for staff during late night working makes life too hard, so managers overtip a taxi driver, get a blank receipt and use that to cover the food. This is common practice in IT companies by the way, and is not considered immoral, but from that it is a short distance to more unacceptable behaviour. Enron might be considered another example of this shift.

Conceptually one of the most difficult things to grasp is that, in the context of most popular uses of the term cause, a complex system is non-causal. An additional issue of perception is that complex systems can appear as chaotic, interpreted as complete disorder or confusion or, with the benefit of hindsight as ordered. This latter case is described as retrospective coherence, or more popularly seeing things with the benefit of hindsight. Boisot (1998) illustrates this point well by taking a popular phrase used to justify knowledge management *Why didn't we join up the dots?* Figure 2 shows that the sheer number of patterns that arise from small increments in the number of dots starts to defy imagination. After something has happened it is easy to which the pattern of events was significant and to allocate blame for a failure to see that pattern, but in advance it is impossible. In the context of the tragic events of 9/11 this would include people being trained to fly aircraft, but not to take off and land. This is known as a weak signal, something that after the event is seen to have significance that is missed at the time. In our earlier referenced example, the gorilla in the basketball video is a weak signal. Retrospective coherence means that there is a natural resistance to understanding the essential unpredictability of complex systems.

Most human systems are complex. We respond and adapt to local interactions, we are constrained by systems, but we are also capable of constraining those systems and changing them. Small things lead to unintended and unforeseeable consequences. Once disturbed any human system is altered irretrievably and will not return to an equilibrium state. Returning to our two stated purposes for knowledge management, decision making and innovation we can see the impact of complexity. If we reduce the system constraints, increase agent interactivity, but increase the variety of those agents and

their environment then radically new patterns become possible. If we understand that a system is complex then our expectations of decisions are different. We are not making decisions based on forecastable outcome and best practice, as both are impossible. We cannot adopt an approach based on *fail-safe design*, but have to switch to *safe-fail experiments* and monitor for the emergence of patterns. Some we amplify, some we dampen dependent on the evolutionary direction we wish the system to take. We thus manage the evolution of the system towards an unknowable future state, we do not waste energy in trying to achieve a predefined system outcome. Experiments which fail teach us about the potential of the system.

Now this approach to management is one all too familiar to most adults. Imagine a party with some forty ten-year olds in an open space. One of the first things you would do as an adult is to establish some boundaries, which you might physically reinforce. You then introduce catalysts (a football a climbing frame etc) which may stimulate or motivate patterns of play. If a beneficial pattern of play starts to evolve you amplify it. If you spot a small group torturing an insect you disrupt the negative pattern quickly before it grows. In the language of complexity you manage the emergence of coherent beneficial patterns within boundaries, within attractors. What you would not do is create a strategic plan for the party with defined activities and milestone target achievements. Learning objectives and after action reviews let alone best practice case studies and enforcement of process do not belong. To do so would be spend much energy with a low probability of success. On the other hand managing the attractor and boundary conditions represents a low energy cost way to evolve towards a more beneficial outcome.

Another good example of this is traffic control. Figure 3 shows the layout of a roundabout near Swindon called The Magic Roundabout. This is not just a roundabout, it is a roundabout of roundabouts and traffic has multiple options. It can proceed around the outer ring, acting as if it was a single roundabout. It can also enter the centre direct with multiple options to reach the exit. The road markings act as boundaries, the traffic densities as attractors. At a very low energy cost the system copes with large volumes of traffic, the patterns adjusting to the demand. Figure 4, while a sculpture is similar to many a traffic control system at US intersections. It makes the point about the ordered alternative. Here the system (manifested in the lights) controls all traffic movements. It requires energy and programming time, if the nature of the system changes it will have to be reprogrammed. In effect a large part of the cognitive decision processing is distributed into the agents (the traffic) rather than being controlled by the centre.

This idea of distributed cognition is central to the application of complexity to human systems and has profound implications for knowledge management. It shifts the balance from centralised control in which the system controls actions (for example in predefined best practice or process) to distributed control in which the focus is on increasing interaction and connectivity between agents. A classic example of how this can work is the Grameen Bank which rejected the traditional, and ordered approach of credit scoring in favour of providing funds to self formed lending groups in which members of the same village agree to take out the same size load at the same time and guarantee each other's debts. This has produced debt repayment levels in excess of 97 per cent. In effect the lending decision is distributed to the members of the lending group: would you form a lending group with people you could not trust?

Complexity based approaches to management handle ambiguity and uncertainty by avoiding central control other than of boundary conditions, and allowing high levels

of agent interaction to create emergent patterns of meaning. Considerable progress has also been made with one aspect of complexity, namely the use of agent rules to simulate possible outcomes. One of the original, and one of the most frequently quoted examples is of Boids algorithm (<http://www.red3d.com/cwr/boids/>) which allows us to simulate the flocking behavior of birds. This type of approach has seen the development of sophisticated models for urban planning, where agents are programmed with rules and computer simulations are run over extended period to see what patterns emerge. Robot swarms (reference) have also provided insight into pattern based behaviour. It is the modern equivalent of planting an area with grass and seeing where people walk before you invest in formal paths. However there is a danger of confusing simulation with prediction, just as an earlier generation confused correlation with causation. Like Newton's Laws the approach works within boundaries, where the complexity of human decisions in the context of the model application can be rendered as simple rules we can gain considerable insight into the operation of the system. However, just as a simulation model cannot predict whether birds will flock to the left or the right of a mountain until they do it, so there are limits to the predictive power of such a model. Also in all too many cases human decision making cannot be reduced to simple rules. As we established earlier, the reality of human decision making is through the filtering of multiple fragmented patterns gained through experience of tolerated failure and through multiple sources of narrative. In addition human agency is fluid, different identities bring different patterns to the forefront of our decision making process. We are also capable of intentional and contrary action.

So, while in nature we can simulate behaviour by discrete agents operating according to rules this has only limited utility in human systems which represent a level of uncertainty over and above those found in nature; we have multiple identities, we are not single agents and we make decisions on a first fit pattern match not on the basis of discoverable rules. Accordingly understanding human complex systems will require techniques and tools that allow us to map the space in which knowledge flow is taking place. To understand the what attractor and boundary conditions are already in play, and where there is potential of change. It is to this task, and the associated management of knowledge that narrative is well suited, and the use of narrative, in particular fragmented narrative is the subject to which we will now turn.

Fragmented narrative

Czarniawska (1997) attributes the phrase *homo narrans* to Fisher (1984). Niles (1999) offers a more elaborate working of the idea that humans are fundamentally shaped by and shape the narrative structures of their existence. We know that the ability to pass knowledge between humans through story was a distinguishing feature of human evolution. No longer dependent on genetic change and imitation of parents, abstract knowledge and practical wisdom could be distributed, mutated and blended to speed learning and adaptation. Narrative remains the principle mechanism of learning and knowledge transfer within an organisation. Accordingly it is not surprising that this paper advocates that peoples narratives should be captured and interpreted as a form of research and learning in knowledge management and elsewhere. If we are *homo sapiens*, in part because we are *homo narrans* then the study of our multifaceted and fractal narratives should lead to insight and sense-making capability. Of course we are more than that; stretching my mind back to schoolboy Latin, we could also talk about *homo*

fabrilis, *homo facetus* or maybe *homo ridiculus* to reflect our toolmaking and multi-aspect forms of humour. All of these challenge the assumption of *homo economicus*, the rational actor, making decisions based on an assessment of available data on the basis of personal self interest. The introduction of pattern based intelligence, and the role of narrative in creating and forming those patterns provides new possibilities for research and knowledge management alike.

If we think about the form and manner of human exchange then it is also important to distinguish between story telling, in the sense of creating a structured narrative with a purpose, and the fragmented anecdotal exchange which characterises a human conversation. In several years of work in this field, having gathered large volumes of narrative material, the vast bulk of original material if transcribed is rarely more than a paragraph or two in length. Figure 5 introduces another key aspect of fragmented anecdotal material, namely is ability to provide context relevant material over longer periods of time. The act of abstracting original material into a structured document or case study takes place in the context of time and place. As those change the relevance of the material will reduce. Abstraction means making choices about what is relevant and the level understanding in the target audience. As a result there is inevitable loss in the process. The price of codification is abstraction, and thereby loss, the value is in rapid diffusion (Boisot 1998). Excessive levels of abstraction (for example to single words or phrases) would result in incoherence and little value, but there is a sweet spot in which the material retains value over longer periods of time. This also entails recognising that a degree of ambiguity is essential. Grandchildren will listen to their grandparents narratives when they will not read a book, engineers swap stories on the back of truck before they start work on a line repair, project managers over a drink share their experiences in narrative form; all of these are more memorable and likely to have impact that formal process.

We can take this further, in that a wider definition of narrative would include anything which tells a story: paintings, pictures, sacred objects and the modern age blogs and other URL related material, for example references to social computing sites such as Utube or Flickr. The growing popularity of social computing tools, and the voluntary participation in their use by people from all backgrounds, age and social background gives some indication that this type of fragmented narrative material is easy to create, share and represents a natural proclivity for humans. It does not require, to use some of the cliches of Knowledge Management, *investment in a knowledge sharing culture*, or the creation of *formal incentives for knowledge sharing*. People just do it. A working hypothesis of this paper is that they do so because the fragmented, socially constructed nature of the internet more accurately reflects the evolutionary patterns of human intelligence than the abstracted, information centric focus of Knowledge Management.

Current narrative based research

Lakomski (2004) provides a critique of hypothetico-deductive research (in the context of leadership research, but the points she makes are generic in nature) from a naturalising perspective as follows:

The model of the human mind has been assumed to be akin that of a symbol processor, a computer like engine that allows us to manipulate successfully a range of symbols of which language is deemed the most significant.

This view of the human mind is very limiting because it assumes that what we know, and are able to know, is expressible in symbolic form only.

... because intangibles cannot be captured in the grip of such symbolic representations as questionnaires or surveys. It might rightly be pointed out that there are qualitative means of assessing transformational leadership in terms of interpreting certain leader behaviors, or by applying leader self-reports. These are imbued with their own problems because of the inability of differentiating between competing interpretations, a core problem of interpretive social science and hermeneutics, and by the endemic unreliability of self-reports.

This directly challenges the underpinning of traditional methods of research and what might loosely be termed knowledge harvesting. The intention of this paper is not to attempt to invalidate such techniques but to recognise the limits of their applicability: the resolution of a problem in quantum mechanics cannot be achieved by more diligent attention to detail in the application of Newtonian physics. In particular we want find ways to mitigate the following issues, all in part derivative of Lakomski's issue on the implications of assuming a symbol processing model, rather than the fragmented pattern model identified earlier.

- The restriction necessarily entailed by asking a direct, or hypothesis based question and relying on the nature of the response which may be given in role, or arise from gaming behaviour. Such approaches assume too much about context, and the objectivity of the research subject.
- Inevitable cognitive bias in the interpretation of research material either through interpretation or deconstruction, in particular the pattern entrainment of the researcher reviewing primary data.
- Questions of meaning, in that numbers have a pseudo-objectivity but often lack interpretative context, and related issues of confusing correlation with causation which were referenced earlier.

The argument of this paper is that narrative techniques both provide a complementary form of what we will call pre-hypothesis research, but further that the use of narrative research techniques produces through a single intervention quantitative conclusions supported by narrative context, fragmented knowledge databases and a mechanism for measuring impact and more complex issues such as mapping ideation cultures). In doing this we create a research and knowledge harvesting method which is compatible with natural sciences understandings of the cognitive patterns of human memory and complex adaptive systems theory. This work builds on, but is substantially different from an existing body of narrative research methods which we will now review in respect of the key issue, how is that narrative, with its inherent ambiguity to be interpreted

Questions of interpretation of narrative

Czarniawska (op cit) challenges the assumptions of those who advocate *homo economicus* in organisational studies by arguing for a narrative based approach to organisational studies. Stories are at the heart of our day to day discourse and our sense-making abilities. They form a part of the common sense world in which intention, interpretation and interaction are all intermingled in any narrative. The narrator and listener assume shared context for any statement to have meaning. Stories carry with them ambiguity and their meaning can be interpreted in different ways in different contexts. So while few would disagree that narrative creates meaning, and is a meaning-making tool for humans of all levels of literacy, the question arises as to how it should be interpreted. For Czarniawska the solution to this is to examine story as genre and she focuses in particular on two drama and autobiography. Boje (1991) also takes a stand linked to drama, or more specifically, Tamara-esque drama in which the organisation is assumed to be a multiplicity of stages on which different plays are acted out by organisational members (actors) simultaneously. Boje sees narrative as occurring in fragments with fully developed stories seen as unusual. Gabriel(2000) argues that we can too easily become seduced by the story itself and lose critical faculty. He explicitly raises the tension between the expert and the voice of experience that is present in narrative work. All three are to varying degrees post-modernist in outlook. In Boje and Czarniawska (and to a lesser extent Gabriel) pose *homo narrans* as antithesis to *homo economicus*. They reject the supposed independence hypothetical-deductive approaches in favour of engagement. The researcher gathers and interprets narrative within a framework of some type such as genre. They acknowledge their engagement, but in a post modern and socially constructed world reality is at best unreachable, in the extreme irrelevant.

However, the argument of this paper is that by taking a naturalising approach to sense-making, in particular the institution of self-interpretation within a semi-structured indexing structure, we can achieve a synthesis of the two and argue that rejecting empirical science as a model for social research does not necessitate rejecting science, but that is for later. We will now proceed to summarise Czarniawska (1998) as representative of narrative research methods. The wider field of narrative and its development is summarised elsewhere. (Oliver & Snowden 2005).

Czarniawska sees narrative as entering organisational studies in four-forms:

- Research that is written as a story or tales from the field
- Collecting stories in the work place, stories of the field
- Seeing life as story making and organisational theory as story reading
- Reflection that is a form of literary critique.

She acknowledges the role of ethnographic methods in this, and the dangers of acting as a researcher in one's own culture. However this danger is dismissed in the context of organisational studies, along with the need to spend a prolonged time in the field. She distinguishes between the use of the *participant observer* in which the researcher carries out the work to learn about it or, in a weaker form *shadows* the worker. Her preference is for the *observer participant* where those who do the work are trained to gather narrative in the field. Direct involvement of the researcher is achieved through *narrative interviews* which focus on time cycles: "what happened to your unit in the last two weeks". She raises concern about the impact on the researcher that is the inevitable consequence of estrangement from the culture being studied, but concludes that it is a necessary evil or inconvenience to be shouldered in the interests of richer data. Reveal-

ingly there is little consideration of any dangers to the subject of the research from the presence of the researcher.

For Czarniawska the resultant narrative forms a series of references from which the researcher weaves a new story. Various analytical techniques are introduced. She makes a key distinction between *Conversation Analysis*, which “captures and analyzes a concrete speech situation located in a point in time and space,” and *Discourse Analysis* which “addresses many conversations that take place over time and in different locations and yet that seem to be connected”. Reading the stories and writing about them cannot be separated for long. The research report is thus a narrative orchestrated by the researcher using the original material in a form of literary collage. She points out correctly that even gluing together narrative constitutes a form of reading. Interpretation proceeds through three stages, and progress to the third is a “professional duty” for social scientists. These are:

- Explication, a reproductive translation in which the interpreter chooses to stand under the text to understand what it means.
- Explanation, in which the reader stands over the text to analyse it
- Exploration, in which the readers stand in for the author, constructing a new text from the starting point of the original text. This might involve deconstruction or reconstruction of the material.

Overall the reader/researcher should give preference to performative criteria, namely to seek descriptions associated with justification when a positive response is received by the audience. Now this could be read as pure opportunism, worthy of the worst management consultant seeking to ensure a follow through engagement. However that would be unfair. In this type of approach to narrative research the response to challenges of objectivity is not to attempt objectivity, but to argue that meaning is a social construction in which the researcher engages. They cannot avoid it, so they should not be criticised other than in the sense of justification implied by a judicial interpretation of the original material. Validation in effect is pragmatic or aesthetic with no pretense at being factual as such a position is impossible. Here we see the continuation of the literary metaphor which underpins this whole approach.

However the question of reality intrudes. Conventional field reports are designated as belonging to the field of *naive realism* the value of which is challenged by the “arrival of constructivism, relativism, and postmodernism. However forms of realism exist, appropriate to literature. These are:

- Ironic realism in which the original narratives are left, with their various paradoxes and contradictions left unresolved. (This is closer to the approach I will advocate later)
- Micro-realism in which organisational life is described at the lowest possible level in detail based on ethno-methodology
- Polyphonic realism in which multiple versions of the same event are presented in narrative form

In effect we have two contrasting techniques, hypothetical-deductive and a post modernism approach to narrative interpretation. To a degree both are defined by not being the other but both are accepting or rejecting an empirical model which is itself lim-

ited in nature. Complexity theory and understanding of pattern based intelligence provide us with an opportunity to move the agenda on, an Hegelian synthesis based on new understanding of science.

Context is key: a new way of thinking

The advantage of the questionnaire form is that it can be distributed in large numbers and the researcher need not be present. It can be completed in privacy and the results interpreted statistically. The downsides are many but in the main we can focus on the fact that the hypothesis is embodied in the question, the question covers a restricted range of options and finally the figures are subject to the “what does this mean” question: they are numbers without context.

In contrast the narrative methods described above are richer in content, but are difficult to scale given their dependency on the researcher to gather the material. They are also subject to interpretation and bias from the perspective of the researcher. This is acknowledged by Czarniawska and others who talk of the researcher engaging in a process of retelling a story. There is a related issue here of confirmation bias (Wason 1960). Once a hypothesis starts to form in the researchers mind, it will form a patterning structure in the brain which will lead to supporting data gaining more attention than contradictory data.

So both philosophies have their issues, and no research method will ever be perfect. However both theoretical work and experimentation have over the past decade, created a new set of narrative methods, based on the scientific principles outlined earlier which seek to reconcile these two positions. The working title for these approaches is Pre-hypothesis research and the method is summarised in the next section. Before proceeding to that description it is worth noting the historical origins of this approach to provide the reader with context.

- The original use of narrative was a source for mapping knowledge. The subject not from a literary tradition, nor from a communication tradition or from a formal background in research methods. Instead it arose from the deeply practical need to create a rich context from which it was possible to extract decisions and judgments to ask questions about knowledge in use (Snowden 1999). Narrative here was also shown to be a better recall mechanism for hidden knowledge than questions.
- Subsequently the work extended to the field of antiterrorism both before and after 9/11 where the approach was based on the capacity of narrative for disclosure of otherwise hard to understand factors such as intent and purpose. Also and more critically as a sensory mechanism of weak signal detection. This was based on anecdotal evidence, confirmed by subsequent experiments, that human brains are more sensitised to narrative forms of knowledge about a situation than they are to analytical processes (Lazaroff & Snowden 2006). In effect agreeing with the general criticism of sense-datum processing as a model of human intelligence.

In carrying out this work, which has taken over a decade in various forms the following conclusions were drawn. From the perspective of practice they seem commonsensical and were subsequently validated by reading in the natural sciences, but at the time were (and in some circles still are) controversial. In summary:

- That naturally occurring stories come as fragmented anecdotes. Occasionally you get a fully formed and developed story, but mostly they are anecdotal, often only a paragraph long when transcribed. Those with most meaning are often the worse constructed. In one recent case, looking at the stories of school children on leaving a secondary school in Singapore the most powerful were from the least articulate; there was less disguise. Paintings and pictures were also found often a better form of narrative expression than a pure story in textual form.
- A story is always told in a context, from a context. If you read it then it will trigger a reaction but the reaction is not necessary sympathetic to that intended or experienced by the story teller. Each reader has their own context and situation. When we also take into account that anecdotes need to be captured in their native language (try telling a story in something other than your mother tongue and you will see the problem) which adds complexity. There needs to be some common context for any translation to be effective. As will be described in the next section we determined that the best way to achieve this was for the researcher to create a tagging system of sufficient simplicity to be understood without active interpretation, and for the story teller to tag their own story. In this way the metadata represents a common context.
- That if the researcher first looks for patterns in the metadata. the way in which narrative material has been indexed or tagged (this will be clearer after the method of indexing is described later in this paper) using statistical or visual tools they are less likely to be biased by content and prematurely converge on an interpretation. Not only that but larger volumes of material could be scanned, and anomalies and clusters more easily detected. This allowed the researcher to construct and test hypotheses after data capture, using the self indexing mechanisms.
- The material so gathered formed, with simple visual and criteria based selection a valuable knowledge asset which allowed direct access by the knowledge user from an abstraction of the field, to the raw self-interpreted narrative. They reflected a natural process of knowledge recall. Faced with a difficult or intractable problem we are unlikely to look up best practice as a structured document. Instead we seek out people and other sources, for example the internet, gathering fragmented material that we select and blend with our own experience and the current context to determine how to act.

Pre-hypothesis narrative research

The issue of cognitive bias on interpretation has been identified earlier, along with the question of statistical and empirical validity. Czarniawska embraces the issue of cognitive bias, acknowledging the effective retelling of the subject's story by the researcher. Others in the action research (Argyris et al 1985) tradition acknowledge the impact of their presence but seek to mitigate. The extreme forms of post modernism argue against any objectivity. Quantitative methods are held to support a positivist position, the world is out there awaiting discovery whereas qualitative approaches represent variations of interpretivism in which the world is constructed by social agency and in consequence any research intervention will affect that reality. Much effort in recent years has seen the development of a range of formal methods derivative of the assumptions of interpretivism. Action research, Participatory Action Research, Living Theory and others see the interaction between researcher and research subject as an iterative process of enquiry that may be primarily driven by the research or the research subject. A range of

methods focus in different ways on power, from Feminism to various derivatives or Derrida and deconstruction. In the field of narrative this position is exemplified by Boje (op cit). It is not the purpose of this paper to provide a comprehensive summary or criticism of these methods other than to set the scene for what we are attempting with pre-hypothesis techniques, namely to provide a quantitative technique, which is supported by the rich context of supporting self-interpreted narrative. In its turn this provides a more objective (but not the purported objectivism of positivism) for qualitative interpretative processes by the research, and indeed the research subject that lead to sustainable action.

At the heart of this project is a view of meta-narrative as an emergent property or strange attractor arising from social interaction which is discoverable and actionable in the sense of quantum mechanics rather than the laws of motion. By taking narrative as a fragmented form of support for cognition, and using the ubiquity of the web and social computing, together with the representational and information processing capacity of computers we can considerably augment and enhance the natural pattern based intelligence that underpins human decision making, and more so to radically reduce interpretative conflict in the process. Further that research so conducted also creates a knowledge base which conforms with the naturalistic principles outlined earlier

So our goal is to attempt to utilise the rich context of narrative, but also to create objective data in which cognitive bias is minimised and where we can place some reliance on the conclusions drawn, more particularly we want to be able to move rapidly from research to action in decision making; linking back to the objectives of knowledge management stated earlier. The principle components of that approach cover, promoting questions or situations, eliciting narrative material, indexing or tagging that material by the originator at the point of origin and finally representation and discovery. These will now be outlined, before continuing to a conclusion.

Avoiding hypotheses: prompting questions and situations

An early project focused on finding ways to understand and measure the impact of museums in Liverpool on school children. Traditionally this would have been done by a questionnaire, focus group or expert interviewing. All of these have issues, expert interviewing and interpretation are subject to cognitive bias. Focus groups can be easily subject to influence by the facilitator. In this respect we conducted a series of experiments with different groups in which the facilitator was subject to peer review within sessions. In all cases peer review rapidly identified that influence was taking place in times ranging from a few minutes to a maximum of around 40 (and that was rare). Interestingly the influence was not only verbal, but could be as simple as privilege by attention; showing by a glance or a smile that a certain story was appreciated. Questionnaires in effect contain hypotheses which determine the range of answers. Did you enjoy your museum visit? Did you find the exhibits interesting? These questions prompt the adult response *It depends* as different contexts would produce different answer, in adults and children alike they often invoke the desire to give the right answer. Suggesting that an average, or overall answers is given rather defeats the point of the process in the first place. I want to know what was interesting, or which aspects appealed. An average answer is simply not good enough.

So instead a web site was provided which children could access from their school computer. They were asked two questions, designed to elicit a story. One in ef-

fect asked them to tell a story about their visit that would enable them to persuade a friend to go to the museum rather than play football, the other asked for a story about the visit that would persuade their parents not to take them to the museum, and instead to allow them to play football with their friends. Other questions have been asked in the project, but the general rules underpinning prompting questions are simple:

- They should be about the whole of the experience and should be designed to elicit narrative material rather than a simple statement.
- The question should be asked in such a way as to elicit a meaningful context in the imagination of the subject.
- The question should not privilege positive or negative experiences but should seek both.
- The subject should be allowed to answer in the third person.

Two more examples will illustrate this in Table 1, both successfully used on projects. These are shown in the table below.

Table 1: Examples of Hypothesis and Pre-hypothesis Questions

Context	Hypothesis Question	Pre-hypothesis Questions
Understanding employee attitudes	Is this company a good place to work? (Answer on a numerical scale)	Imagine you are in a bar on a Friday night and your best friend comes in and says they have been offered a job with your company. What stories from your or your colleagues experience would you tell them to encourage them to join? What stories from your or friends experience would you tell them to discourage them from joining.
Ethical auditing in the pharmaceutical sector	Do you think this company is ethical in its approach to the use of animals in drugs testing? (Answer on a scale from <i>very ethical</i> to <i>not ethical at all</i>)	Imagine you have just presented your companies work to an after school meeting of parents and children. At the end of your presentation an 11 year old stands up at the end in tears and says <i>I think you and your company are evil because you torture rabbits to produce lipstick</i> . What would you, or someone you know say in response to that question?

The purpose of a promoting question is to elicit narrative, not to gather interpretation or meaning that comes later. Situations can be used as well as questions. Examples of this include a competition with a company as to whose children could produce the best computer simulation or a painting or what it was like for their parents to work for the company concerned. A science fair format was then used to judge the competition and a video camera was taken around to record stories as parents faced the perspective of their

own and other children on work. The focus of the project was to understand issues of work-life balance. A future project in Liverpool museums will elicit response from adults as well as children as they encounter exhibits in the newly created slavery museum. Narrative material can come in a variety of forms. Oral recordings, transcribed material, url references to sites such as YouTube, pictures or paintings can all tell a story. Prompts can also be instructions: *Find a clip on Utube which summarises the attitude of management in this organisation.* The goal is to gather sense-making items (SMIs) in fragmented form which can have layers of meaning added through the process of indexing or tagging.

Addition methods to elicit narrative material

The use of a web site, or a social computing environment such as Facebook has the advantage of anonymity and is free from the presence, and consequent influence of a researcher. It allows mass, and continuous capture, and can also pick up on the volunteer aspects of social computing. Blog entries, and the results of an interesting RSS feed all provide fragmented narrative that can be used for both research and knowledge management. However there are a range of other methods that can be used, some drawn from existing research traditions. In all methods the focus is not just gathering the material, but on reducing the possibility of bias from the approach. Methods include:

- **Population sampling** takes a sample of the population and uses that sample to glean narrative material from their colleagues or associates. To take an example; five demographic bands of 50 people each, selected at random from bands created using length of service were requested to find two interview subjects (the instruction was *one representing the value of our past, one the potential of our future*) and ask them five prompting question recording the results and getting the interviewee to index the material. This method has been scaled both to bring people to web based capture and through field recording, and can be done via email or other propagation devices (such as a blog meme).
- **Participatory research** involves full participation in the role or function as an apprentice, requesting and recording narrative material in the context of day to day work. This method relies to a degree on the trusted relationship that are naturally present between master and apprentice even on short acquaintanceship. In one project students were deployed to accompany long distance lorry drivers for a week. They carried with them digital tape recorders and manual indexing sheets with the intent of gathering stories about work practices and situations. The lorry drivers were advised that the students would be responsible for transcription of their stories into a database that would be visible to management. They were further advised that should they not trust the student with the material they were under no obligation to share a single story. In practice, the act of working together and extended conversations build the necessary trust for this to be a highly successful method of gathering material.
- **Anecdote circles** are used where the stimulus of colleagues or people with similar experiences are needed to elicit material. The approach allows confidence to be built between different participants and relies on the desire to tell a better story that is the norm of human conversation. Here the question is introduced to a group who then swap stories about their common experience. When a useful anecdote emerges they

leave the group to record and index the material. This is done in preference to recording the whole session as it ensures that permission is clearly given for the anecdote itself, which is now a discrete item. The method can also create the space for other voices to emerge. This also reduces transcription bias in selection if the entire session was recorded. Issues of pattern entrainment within the group are achieved by changing the membership of groups. The dangers of facilitator bias are reduced by using three facilitators: one to lead, one to observe, one outside the event. When the observer notices influence they signal the lead who hands control to them and leaves the area to allow the third facilitator to take the role of observer.

In all of these the introduction of naivete in the story prompter can generate richer material. For example using school children to gather stories from parents and grandparents (within appropriate ethical guidelines) generates a teaching or mentoring response. Students (as described under participatory research) also assists in capturing narrative at the right level of abstraction for re-use. Engineers talking to engineering students talk at a lower level of abstraction than to other engineers.

Methods can of course be used in combinations with different types of stimulation from prompting questions to situations to pictures or videos. Overall a diverse range of interviewers, sources and methods is preferred over a limited number of interviewers who would of necessity be subject to missing the gorilla suit in a mêlée of basket ball players.

Indexing or tagging

There are two approaches generally in use to handle narrative, as well as other material. One is to adopt a classification system, frequently hierarchical assigning the material to a category. Within knowledge management the generation of a hierarchical taxonomy has been a frequent starting point, within Library Science the Dewey-Decimal. Lamb (2007) in his excellent summary of the role of Taxonomy in knowledge management points out that Taxonomy represents a form of artificial memory from the poet Simonides to modern folksonomies, enabling Clausewitz's *coup d'oeuil*, the cast ones eye to achieve awareness. In practice taxonomies have, due to the inherent limitations of card classification systems (which passed across into early computers, being hierarchical, requiring an item to be placed in a single unique category. Innovations such as facet analysis (Ranganathan 1967) allow for greater flexibility and to a large degree form an early evolutionary stage to the approach advocated in this paper. However the practicalities of such approaches had to await the development of scalable and reliable computing together with the wider awareness generated of folksonomies through social computing. In a very real sense we are now provided with two limited extremes. The rigidity of hierarchical classification and the anarchy of folksonomies

Neither present an ideal solution. Allowing people to assign whatever tag they wanted (as practiced in most social computing) would introduce massive uncertainty about the way the material was tagged. As earlier indicated there are in any even natural limits to semantic analysis by computers, which would be required to handle any reasonable volume. As over rigid classification system which attempts to remove ambiguity would be subject to the general criticism of such systems as static and non-adaptive (Weinberger 2007). In more recent years attempts have been made to create controlled vocabularies in social computing environments. This has potential within a restricted

population but is not practical for mass capture. Indeed the authors own experience is that he fails on his own blog to use even a limited vocabulary consistently! Accordingly the approach adopted, and refined over several years of experiment was to create a semi-structured tagging approach, one that could be created by the researcher if their are specific objectives, or through an emergent process using a sample of the population if enquiry is more general. The intention, and practice is to create a common interpretative grammar between subject and object.

To do this successfully we need a mixture of tag types, both ambiguous (for we are dealing with ambiguous material) as well as disambiguated material and other means by which the indexer can add meaning. Although this is a moving field, but current structure used is as follows:

- The indexer is asked to *name* their story; such names turn out to be highly significant and often contain more meaning than the content itself. They also allow for the original content to be kept private to its originator (essential in some projects) with access only granted by permission. The name on its own is enough to give the researcher or decision maker the opportunity to make sense of an overall pattern of narrative material.
- The story is then positioned on a scale or other geometric shape known as a *filter*. A simple example, using a linear scale is show in Figure 6. This offers a simple linear scale with two end labels. A more advanced version is show in Figure 7 which uses a concept known as opposing negatives. Here a desired or anticipated quality of the field is identified and the two end labels are provided as *the thing not present* and alternatively *the thing taken to excess*. The ideal would therefore be represented by marking the centre of the scale. A linear scale with two end labels provides two filters of analysis and retrieval (the left hand label represents 100 percent f itself or 0 per cent of its opposite). If a triangle is used, then for one entry six analysis filters (the distance to each corner and the vertical drop) thus minimising data entry while maximising analysis capability. Normally such filters are grouped in fives to allow scanning and assessment of the narrative across a coherent range of filters.
- The filters are designed to handle ambiguous or abstract qualities; in contrast multi-choice questions (MCQs) deal with aspects of the narrative for which there are a limited range of options. Apart from demographic and other data a range of MCQs have proved consistently useful over a range of projects. For example: Why was the story told? (to attack, to defend, to educate, to entertain, to influence, to inform, to uplift or unclear); Was the story Sacred or everyday? (this is normally elaborated in context but is important); What was the tellers relationship to the story? (Central character, reported by witness, hearsay or gossip)
- *Key words* were introduced partly to avoid arguments with more traditional approaches, but have proved useful and allow standard capabilities such as tag clouds to be used to good effect. A switch of the standard question from *What are the key words in this story* to *What are the key words you would associate with story* produced a significant reduction in the percentage of key words contained in the original content; representing another criticism of semantic analysis.
- A *free text* field is then provided to allow additional descriptions or explanations to be provided if required. This can be useful if the content is not textual in nature but is generally useful.

The original content, together with all the above is referred to as a sense-making item (SMI) to indicate that the material has been tagged, and it may in any event not be narrative in nature. The original content may not be available to analysts, or only on permission in which case it has a secure URL or reference request. Filters and MCQs may also be asked of the individual story teller regarding their general attitude or perception as well as factual issues such as demographics. The answers here are asked one time only, and are then attached to all SMI's indexed by that person (or group). In the emergent jargon of this method these are known as *stickies*, as they attach themselves to anecdotal material. The *stickies* are used both for selection and recall as well as analysis.

MCQs are generally non problematic and inherit aspects of classification systems, they can even be hierarchical in nature; but they are tagged to the item, the item is not subsumed by the classification system. Filters on the other hand deal with ambiguity. They inherit aspects of fractal analysis (Ranganathan op cit) in that they permit aspects of a separate system, or grammar of interpretation to be tagged against the item. As such they are post-coordinated, permitting the emergence of novel items not envisaged by the creators of a classification system. In the case of narrative work we also de facto permit the same item to be indexed differently by different people. This recognises the essential ambiguity of changing context needed for effective fragmented recall.

Filters can be created using samples of the population to determine emergent properties such as values and themes (Snowden 2005) and in several cases we have used large samples (50 school children in Liverpool for example) in a workshop format to define filters that have relevance to the target group, in active dialogue with those who need to carry out interpretation. Filter construction can also be more analytical. For example understanding employee attitudes may require filters in which the espoused values of the organisation are used as the centre point of a filter with opposing negatives. Figure 7 is a set of filters from a project looking on attitudes to leadership in a government agency using the opposing negatives concept. The two approaches (emergent and analytical) may also be mixed.

The function of the filters and MCQs is, as stated, to create a common grammar of interpretation between the subjects of the research, and those carrying out the research. It is a similar way between the creators of fragmented knowledge objects (anecdotes, pictures, URL references, etc.) and those who need to access them. The controlled vocabularies of both taxonomy and folksonomy alike, it is argued are in the former case too structured, and in the latter case too unstructured relying on common use of language. Both assume shared context which may not exist per se. We are also allowing the emergence of multiple micro-meta-narratives from the target population, rather than interpreting that material through the meta-narrative structures of the researcher. By allowing assumptions and hypotheses to be hidden in the filters (or combinations thereof) we also allow beliefs or meta-narratives to be tested rather than disputed. If it is not tagged, or if the tagging contradicts the assumption at the core of an analytically derived index then it is challenged.

Representation and discovery

The function of an indexing system is to allow recall, such a logic underpins the Dewey-Decimal system and the hierarchical knowledge taxonomies of many an organisation. As in any cataloguing system you have to know where to start your search, from the

broadest category, drilling down through the hierarchy. However discovery should not only be considered as the process of discovering those things that you know you need to know, but also of serendipitous encounter with the unexpected, but relevant, at the time and in the context of need. This is partly the claim and the attraction of social computing and the internet. Typing a few key words into a google search results in a long list (and increasingly long list) of items some of which are expected, some of which are unexpected. Building a set of trusted sources in an RSS feed, tracking down references to your own blog, linking of connecting through Facebook. These all represent chance or serendipitous discovery in messy only partially structured environments. Humans are comfortable with mess; look at the average desk or study and you see piles of paper, books open and marked with post-its. Every now and then a spring clean produces order before the cycle of disintegration starts up a new. There are sound reasons for this. Any structure is out of date shortly after it has been created, but s/he who lives in the mess knows where things are up to a point, and when that point is reached the spring clean happens. The author of this article had a simple process of buying a book. Knowing the name and order can be placed online, knowing the general subject area requires a visit to a major bookshop and a pleasant hour or so browsing bookshelves in the relevant section. However, when faced with an intractable or trans-disciplinary issue Foyles as was, was the place to go. Situated on the Tottenham Court Road in London this was one of the most disorganised bookshops in the world. Books were not filed by subject but by published (i.e. based on delivery) with some rough subject classification. The bookshop was staffed by enthusiastic and under paid graduates who were fascinating by books and by their subjects. A visit to Foyles was a fascinating process of unstructured discovery and conversation which never failed to turn up the expectedly relevant book. Similar processes apply to human search. The more complex the problem, the more likely a broad range of experts and acquaintanceships is likely to be brought into play.

In natural enquiry, hierarchical taxonomies are unlikely to be used in other than controlled circumstances, but in most knowledge management systems. Faced with a choice between drawing down best practice case studies or listening to the stories of half a dozen people with relevant experience, most people opt for the latter not the former and for good reason. What we want to do, reflected the fragmented, pattern based intelligence discussed earlier, is to encounter multiple fragmented anecdotes that we can blend with each other, own experience and the current context to create a contextually relevant guide for action. The semi-structured indexing approach described above alone with the mass capture of fragmented sense making items, including but not limited to narrative matches this natural process. It does so by direct query, and through visual patterns to identify relevant material, often what would be termed weak signals, in large volume data-sets. Importantly it does so without requiring knowledge of the nature or class of the knowledge objects being sought. In the pure research context, as well as in operational use, we also have statistically valid data (derived from the tags) in which any discovery or anomaly can be traced back to the originating narrative, potentially in real time.

This is probably best illustrated by example. The first will be a direct query followed by two examples (from many) of the use of visualisation. The nature of query onto a narrative database is similar to a social enquiry, ambiguous and open, narrowing through conversation. So to an illustration. An early project looked at the tricky issue of suicides in an agricultural environment using a farm chemical. In this case the client was an agrochemical company. The suicide was not pleasant, there being no cure and

death taking up to two weeks, as a result of which continued use of the chemical (despite its general utility) was under threat. Several thousand stories were gathered over a short period from a broad range of interest groups, some by proxy but in the main direct. The material was self indexed at the point of entry and a variety of scientific data was also indexed use the same structure by scientists (this could easily be isolated but was designed for serendipitous encounter. In a workshop environment client staff queries the database, attempting different combinations of filter and multi-choice questions. Motivation for this type of activity seems not to be a problem. Humans are curious, and encountering raw narrative appeals to that. One question produced a gold mine: *Show me all the stories told by a rational archetype around the theme that our product kills people from the perspective of a first witness with high emotional intensity and with the intention of attacking our company.* The query represented a range on one filter with selected options from four MCQs. It was a refinement of several queries using broader range. For this query eighteen anecdotes appeared on the screen with one clear cluster around a farmer who had videoed his brother dying having swallowed the chemical. He had then taken the resulting material around outback stations to try and discourage others doing the same, but with the opposite effect. He had acted as a Typoid Mary of suicide, everywhere he went suicides using the chemical went up. As is frequently the case with narrative discovery the pattern *once found*, made sense to experts although they had not known what they knew before the requirement for explanation. Fear of main is not a major factor in preventing suicide and a desire for revenge frequently is. A slow but irreversible death was therefore attractive to some.

In parallel with this the data was searched for metadata or index patterns. The design here is to allow the researcher to discover patterns in the metadata, and then seek understanding by looking at the supporting narrative. Looking at the narrative first would result necessarily in cognitive bias and would in any event be impossible given the volumes. There was a clear correlation in the data between stories designed to attack, with high emotional intensity and stories told by the medical profession. The correlation of filters is shown by way of another project in Figure 8 where a grey line indicates a correlation, a black line normal correlation. The blue dots are stories and can be selected and read. This was a surprise, in effect an unseen gorilla amongst basket ball players. The company had expected negativity from ecological activists (and was paying attention to their material) but assumed that the medical profession as scientists would be onside. The unexpected result resulted in a radical change of tactics.

In another, and too date the most elaborate example Figure 9 shows the use of fitness landscape as representation. This is an intelligence project and the subject matter of the narrative is Iranian intention and attitude towards the middle east. The material is open source, but was gathered and indexed by people sympathetic to the Iranian Government. The filters are then used as axes (and can be combined as either/and/or) and a landscape model. A fitness landscape is a visual representation over a range of examples in phase space where each trough represents a stable area, and the peaks an instability where the dynamic landscape can potentially reconfigure dramatically the in the next instant. By playing with the control parameters we can get an idea of where the tipping points may be. Large fluctuations are also a clue that a phase change may be about to approach and their use thus allows nuanced behaviour to be appreciated and probed. (Juarrero 1999)

In the case illustrated the model was set in one three month period and then stories were matched to the model in three subsequent periods, with narratives that did not

fit represented by yellow dots. The model shows two strong attractors, and examination of the stories that create that part of the model demonstrate that one is *anti-west*, while the other is *anti-american*. If an attempt is made to shift or change attitudes by directly addressing those issues, then only a small part of any message will be scanned and it will quickly be confirmed to the dominant and strong attractor. The opportunity lies in the flatter areas around the attractors where possibilities, for good or ill are still open.

As can be seen the model goes wildly out in period three (the second illustrated) but then stabilises again in period 4. However look at the bottom right, here we have a new pattern of stories that if we recalculated the model would appear as a new attractor. This new pattern of insight was a missed opportunity to change attitudes. This illustrates the ability of representation derived from pre-hypothesis research material to reveal weak signals that would otherwise be ignored which is one clear benefit. However another is the ability to move from an abstract representation of a total field, the originating raw material without the disintermediation of analysis or interpretation. Direct interaction of the decision maker with raw material, discovered from statistically significant patterns and anomaly detection seems more able to produce insight and willingness to act than interpreting analytical data presented by experts. Further the method allows permits representation of abstract issues such as cultural mapping and statistical analysis to be made of differences, along with accurate measures of volatility within the stabilities of the landscape. We therefore have a research instrument, a method of knowledge discovery and a means of monitoring or detecting changes, early in real time environments. The ability to see minor differences early, and explain them through the rich context of narrative seems to be more effective than the delayed response and interpretation of survey material. In the earlier case of Liverpool Museums, it allows staff to see early signs that they are succeeding or failing and respond through amplification or dampening as appropriate. Response to surveys, if fully understood would carry a higher energy cost than early intervention, motivation is also higher as the context, expressed as fragmented narrative, is carried with the numbers. Current work is investigating the use of this approach to create single measurements of impact in service provision (for example in education or health) to replace multiple outcome based targets that are all too susceptible to Goodhart's Law which can be summarised as *When a statistical instrument is used for policy it loses all value* or more colloquially as *if a measure becomes a target it ceases to be a measure*.

Conclusions

This paper has summarised Complex Adaptive Systems theory, and aspects of Cognitive Science in respect of human decision making. There are a broad range of implications for knowledge management and research that arise from an understanding of these sciences, some of which are summarised below. In looking at new methods and approaches the paper has focused on the role of narrative for research and the management of knowledge; its capture, recovery and through *conceptual blending*, creation. Knowledge management and research methods are intimately tangled and the paper has reflected this.

In large part the implications for the practice of knowledge management are implicit (and in several cases explicit) in the text above. The following list attempts a partial of those implications:

- ***The fragmented nature of human knowledge*** recall, needs to be reflected in the formal capture and distribution of knowledge. The large scale take up of social computing without obvious incentive supports the argument for the natural utility of fragmentation. Narrative approaches to knowledge capture similarly focus on fragmented sense-making material which can be captured in situ, retrospectively or on a *what if* basis.
- ***Semi-structured indexing***, as described represents a compromise between the anarchy of free tagging on social computing and the over structured and static nature of hierarchical taxonomies. It allows material to be quickly tagged and recalled in different combinations in different contexts.
- ***The confusion of correlation with causation*** plagues lessons learnt programmes in all sectors but is of particular relevance for large projects which are inherently complex in nature. Mega urban transport projects for example not only involve engineering complications, but also the multiple complexities of human interaction at all states of the project life cycle. Assuming that failures observed on one project can be avoided, or successes repeated on the basis of retrospective analysis of those projects is dangerous. Extracting multiple fragmented recollections together with speculative what if's will have greater utility than formal case studies and best practice documents.
- ***We learn more from tolerated failure***, and through spreading stories of failure that we do than from success. This implies that increasing tolerated failure will increase an attitude of anticipatory awareness in preparation for knowledge work on large projects and elsewhere. This has implications for training and health and safety as well as conventional knowledge management.
- ***Understanding context*** is as important as codifying content. This point applies from the prosaic nature of an engineering practice to the social context of place, time, culture, ideology and government policy in which large projects take place, or the organisational culture of a company knowledge management programme.
- ***Forecasting outcomes*** is not possible in a complex system and attempts to do so will produce inauthentic and/or dangerous consequences in that it will reduce anticipatory awareness, engender pattern entrained thinking and lead to false confidence. Understanding the a complex system is managed through attractors, boundaries and modulators allows a better focus and more effective use of resources, as well as increasing the capacity for contextual knowledge use and learning.
- ***Peer group pressure*** in professional environments is a particular instance of pattern entrainment which leads to a natural level of conservatism in professional communities. Path dependency and the setting of expectations through professional education and work experience may decrease the capacity of that group to sense weak signals, innovate and avoid or anticipate unexpected problems. The introduction of challenge through ritual dissent, the need to validate theory by reference to fragmented databases and the ability to describe attractor/barrier conditions all provide mechanisms to avoid some of these dangers. Forcing blending of ideas from radically different environments and disciplines can also create processes for innovation by increasing the number of patterns available to decision makers.
- ***Weak signal detection*** and the discovery of anomalies requires an ability to sense patterns in a large volume of data, with the ability, without disintermediation, to originating fragmented raw material. Various ways to achieve this were described, with the use of fitness landscapes represented one of the more advanced techniques.

- ***The nature of knowledge management*** (epistemology) will vary according to the nature of the system (ontology). If the system is ordered then it is possible to define best practice and constrain behaviour. If complex then the decision model needs to move from the fail-safe design process of order to the safe-fail experimentation through which the evolutionary possibilities of a complex system can become visible. This requires a radical shift in planning process, as well as investment strategy. understanding the boundary between order and complexity in any project is thus a vital pre-process before any formal knowledge practice is defined or practices.
- ***The universal assumption of homo economicus*** in most planning and management processes represents a major obstacle to progress in the effective management and deployment of knowledge. Many people subscribe to homo economicus while in their heart of heart realising that it is far from universally the case. Hence the myth of the *heroic leader* and ideas such as *gut feel*. The dichotomy of rationality and emotion, left and right brain processing and other myths all represent inauthentic work-arounds. Complexity gives us a new language to avoid such dichotomies, but the mind shift required should not be under estimated.
- ***Narrative methods*** by their very nature increase the level of shared context which is critical to effective information exchange, let alone knowledge. Assimilating material in narrative form is more natural for humans, and allows for greater redundancy and thereby decreases the dangers of misunderstanding and communication, itself an essential aspect of risk management.
- ***The naturalising approach*** advocated in this paper transforms the problematic concept of risk management. Using fitness landscapes to understand attractor stabilities and the level of volatility to anticipate potential phase shifts increases the probability of weak signal detection as well as increasing the overall level of anticipatory awareness. Measuring the degree of stability in attractors and barriers, and the level of understanding of predictability of modulators allows for the management of risk in complex environments. Switching to a safe-fail experimental approach to engagement with a complex system reduces risk over fail-safe design.

The nature of scientific discovery, especially when science goes through a major paradigm shift (complexity and cognitive sciences alike) is that it challenges the orthodoxies of the present. This is evident on research and knowledge management alike, This paper has attempted to outline both the science, its implications, and (in the case of narrative) new approaches that are authentic to scientific understanding but which may appear antithetical to conventional practice and academic orthodoxy. In particular it advocates replacing grand, universal narratives with small, local narratives, self created, self indexed allowing blended interpretation and recall. Recognising the radical nature of this change is the first stage to utilising the huge advantages (and cost savings) that are now possible. However it may be best to end with an cautionary note from the Bible:

“Neither do people put new wine into old wineskins, or else the skins would burst, and the wine be spilled, and the skins ruined. No, they put new wine into fresh wineskins, and both are preserved” (Matthew 9:17)

References

- Andrade, J & May, J (2004). *Cognitive Psychology*, Garland Science, Abingdon, Oxon.
- Argyris, C. Putnam, R. & Smith, D. (1985) *Action Science: Concepts, methods and skills for research and intervention*. San Francisco: Jossey-Bass.
- Axelrod, R. & Cohen (1999), M. *Harnessing Complexity: Organisational Implications of a Scientific Frontier*. New York: The Free Press
- Bennis, W.G. & O'Tool, J (2005) 'How Business Schools Lost Their Way' *Harvard Business Review* May 2005.
- Blackburn, S (1994). *Oxford Dictionary of Philosophy*, Oxford University Press, Oxford
- Boisot, M. (1998) *Knowledge Assets* Oxford.
- Boje, D.M. (1991) 'The storytelling organisation: A study of storytelling performance in an office supply firm'. *Administrative Science Quarterly*, Vol. 36: pp.106–126
- Brian Arthur W., (1994) *Increasing Returns and Path Dependence in the Economy*, University of Michigan Press, Ann Arbor,
- Christensen, C & Raynor, M (2003) 'Why Hard-Nosed Executives Should Care About Management Theory'. *Harvard Business Review* September, pp 67–74.
- Clark, A (1998) *Being There* MIT Press, Cambridge, Massachusetts
- Czarniawska, B *Narrating the organisation* (Chicago 1997).
- Czarniawska B (1998). *A Narrative Approach to Organisational Studies*, Sage, Oxford
- Davenport, T. and Prusak, L. *Working Knowledge* Harvard 1998.
- Fauconnier, G & Turner, M (2002) *The Way that we Think: Conceptual Blending and the Mind's Hidden Complexities* Basic Books.
- Fisher, W.R. (1984) 'Narration as a human communication paradigm' *Communication Monographs* 51:1–22
- Gabriel *Story Telling in Organisations* (Oxford 2000).
- Jenks, C & Smith, J (2006) *Qualitative Complexity: Ecology, Cognitive Processes and the Re-Emergence of Structures in Post-Humanist Social Theory* Routledge
- Juarrero, A (1999) *Dynamics in Action: Intentional Behaviour as a Complex System* MIT Press
- Kauffman, S *At Home in the Universe*. Oxford University Press 1995.
- Klein, G. (1998) *Sources of power: how people make decisions* MIT
- Kornblith *Naturalising Epistemology* MIT Press 1985.
- Kurtz, C & Snowden, D (2003) 'The New Dynamics of Strategy: sense making in a complex-complicated world' in *IBM Systems Journal* Volume 42 Number 3 pp 462–483
- Lakomski, G (2004). *Managing without Leadership* Elsevier
- Lambe, P (2007) *Organising Knowledge: Taxonomies, Knowledge and Organisational Effectiveness* Chandos
- Lazaroff, M & Snowden, D (2006) 'Anticipatory modes for Counter Terrorism' in Popp, R & Yen, J *Emergent Information Technologies and Enabling Policies for Counter-Terrorism* Wiley-IEEE Press
- Lucas, H. (1999) *Information Systems Concepts for Management* McGraw Hill 1979.
- Niles *Narration as a human communication paradigm* University of Pennsylvania Press

- Nonaka, I. & Takeuchi, H. (1995) *The Knowledge-Creating Company*. London: Oxford University Press.
- Nonaka, I. and Konno, N. (1998) 'The concept of "Ba": Building a Foundation for Knowledge Creation' *California Management Review*, Vol. 40, No. 3.
- Oliver, G.R./Snowden, D.J. (2005): Patterns of narrative in organisational knowledge sharing: Refolding the envelope of art-Luddism and techno-fabulism. In: Schreyögg, G./Koch, J. (Eds.): *Narratives and knowledge management. Exploring the links between organisational storytelling and knowledge management*. Berlin: Erich Schmidt Verlag, 15–72.
- Polanyi (1962) Terry Lectures at Yale University republished in 1983 by Doubleday & Company as *The Tacit Dimension* ISBN 0–8446 5999–1.
- Prigogine, I & Stengers, I (1984), *Order out of Chaos: Man's new dialogue with nature*, Flamingo
- Ranganathan, S.R. (1967) *Prolegomena to Library Classification* London: Asia Publishing House Shah, Idries *The exploits of the Incomparable Mulla Nasrudin & The subtleties of the Inimitable Mulla Nasrudin* double volume Octagon Press, London 1985.
- Snowden, D. (1999). 'Story Telling: An Old Skill in a New Context'. *Business Information Review*, 16(1), 30–37. March.
- Snowden, D. (2002) 'Complex Acts of Knowing: Paradox and Descriptive Self Awareness' in the *Journal of Knowledge Management* – Vol. 6, No. 2, (May) pp. 100–111.
- Snowden, D (2005) 'Stories from the Frontier' E:CO Issue Vol 7 noss 304 pp 155–165.
- Stewart, T. *Intellectual Capital: The New Wealth of Organisations*. New York: Doubleday 1997.
- Surowiecki, J. (2004) *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations* Little, Brown.
- Sveiby, K.E., (1997) *The New Organisational Wealth: Managing & Measuring Knowledge-Based Assets*, (Berrett-Koehler), San Francisco, CA,
- Wason, P.C. (1960) 'On the failure to eliminate hypotheses in a conceptual task'. *Quarterly Journal of Experimental Psychology*, 12, 129–140.
- Weick, K.E. & Sutcliffe, K.M. *Managing the Unexpected* Jossey-Bass 2001.
- Weinberger, D (2007) *Everything Is Miscellaneous: The Power of the New Digital Disorder* New York.
- Wexler, B.E. (2006) *Brain and Culture: Neurobiology, ideology and social change* MIT Press.

