Weak signal detection
The see-attend-act model of decision making

"What about expert searchers who have spent years honing their ability to detect small abnormalities in specific types of image? We asked 24 radiologists to perform a familiar lung nodule detection task. A gorilla, 48 times larger than the average nodule, was inserted in the last case. 83% of radiologists did not see the gorilla. Eye-tracking revealed that the majority of the those who missed the gorilla looked directly at the location of the gorilla. Even expert searchers, operating in their domain of expertise, are vulnerable to inattentional blindness."

"The invisible gorilla strikes again"
Drew, Vo & Wolfe
Psychol Sci. Sep 2013; 24(9): 1848–1853

How often are we wise after the event? With the benefit of hindsight we see that we should have paid attention to something but at the time it seemed irrelevant.   It may not even have been brought to the attention of the responsible executive as it was considered trivial. We may not have even seen the relevant data, even if we saw it then we may not have paid it much attention. Even if we or our subordinates pay attention getting people to act on something they did not expect or completely novel is problematic.

The problem of decision support is not simply about getting the right information to the right people at the right time, its a lot more complex than that and there is no single cause that we can address. This brief paper looks at the some of the issues and then at how SenseMaker® together with the Cynefin framework can help us address them.

The unavoidable realities of being human

This is by no means a complete list, but it gives a sense of the magnitude of the problem if we take a traditional information centric approach to decision support.

1. *We know best …*
   The danger of expert bias, the more we know about something, the more competent we are the less likely we are to see something that falls outside the bounds of that expertise. Sometimes known as inattention blindness this is not something that can be trained out of people, it is a part of what we are as a species

2. *Swamped by data*

   Sheer volume of data may prevent us joining up the dots, after an adverse or favourable outcome it is easy to know what is relevant and see some causal connection but at the time it is simply a matter of luck. Hindsight is a wonderful thing but it doesn't of itself lead to foresight

3. *Mediation and interpretation*

   Excessive mediation, interpretation and screening of data before it reaches the decision maker means that raw data is stripped out or summarised based on the assumptions and knowledge base of the interpreter. Middle management issues may also subconsciously eliminate material that does not match, or which implicitly and explicitly threatens their interests. The decision maker may realise the significance of data because they see the bigger picture, their sub-ordinates may not

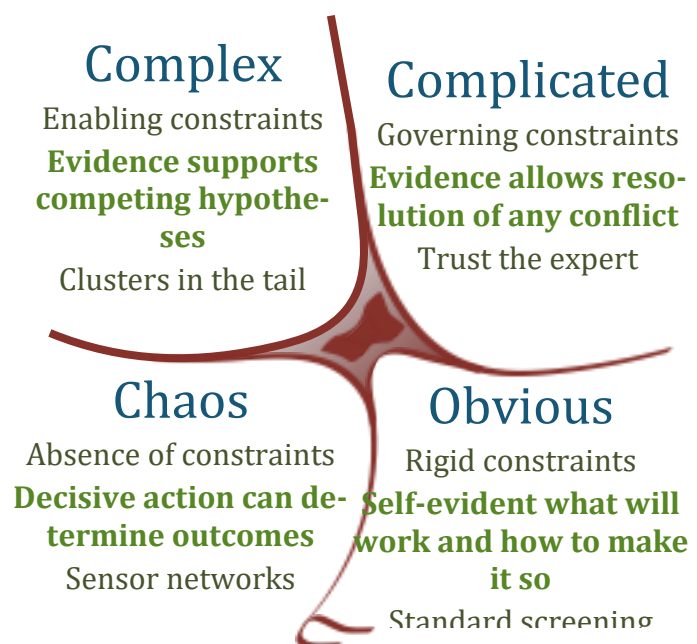4. *The courtier syndrome*

   In most organisations of any size power is based on the ability to influence the decision maker which frequently leads to bad news not being shared or access to disruptive influences being prevented. This is no fault of the decision maker per se, but it appears to emerge over time in government and industry alike.

5. *It worked for me before …*

   We all have a tendency to interpret data to match what has worked for us before and to act accordingly, without investigating in any depth. We evolved as a species to make decisions very quickly using remembered past experience and then to justify those decisions *post hoc*. In overall evolutionary terms it makes a lot of sense but the downside is that novelty may be easily missed, and the more successful we are the easier it is to miss.

The Cynefin framework

The Cynefin framework is an established decision support framework and formed the cover article of the Harvard Business Review in November 2007 (*A Leaders Guide to Decision Making* by Snowden and Boone) subsequently winning multiple awards. It is based in the science of complex adaptive systems and distinguishes five domains in which different models of analysis and decision making are appropriate. The domains are defined by the level of constraint or predictability.

Complex
Enabling constraints
Evidence supports competing hypotheses
Clusters in the tail

Complicated
Governing constraints
Evidence allows resolution of any conflict
Trust the expert

Chaos
Absence of constraints
Decisive action can determine outcomes
Sensor networks

Obvious
Rigid constraints
Self-evident what will work and how to make it so
Standard screening

It is used (amongst other things) to distinguish between different complementary approaches. The version shown distinguishes between different approaches to weak signal monitoring and consequent actions. It is being used here to provide an introduction to the various SenseMaker® capabilities which are described in the next section and which focus on the complex or chaotic domains. The lessons from this are three fold:

1. In the ordered domains (Obvious and Complicated) there are repeating relationships between cause and effect, the same thing will happen in the same way so we can expect an evidence based approach to produce the right results. Past practice can be evaluated and become best practice. The difference between Obvious and Complicated is that in the former what is happening and what needs to be done is self-evident to any reasonable person so there is no need for analysis and expert appraisal. In the Obvious domain we can apply standard operating procedures, simple screening techniques and the like. In the Complicated domain established analytic processes and or the deployment of suitable qualified experts will give rise to the correct answer in the majority of cases.

2. The main characteristic of a complex adaptive system (CAS) is that there are many emergent plausibilities and the future state of of those plausibilities will only be *knowable* in the future they cannot be *known* now. A CAS has stable elements where the propensities of those elements can be known, but at a system level we only have dispositional states not linear causality. We can know how the system might change and which vectors are more plausible than others but we do not have a predictive model. It is in the CAS space that we most frequently face strategic surprise and miss out on strategic opportunities that do not match the remembered patterns of past success. In order to understand this domain we need to avoid the pattern entrainment of retrospective coherence and hold our decisions as long as possible to avoid premature convergence on a familiar solution.

3. The chaos domain is the state of no constraints, things have broken down and we need to take decisive action fast. Entered deliberately this domain allows for innovation but it is resource intensive to use it productively - rather like nuclear fission where the energy required to keep the plasma away from the walls of the container exceeds that we can extract. Treating a CAS as if it was ordered results in catastrophic failure, which is why the bottom boundary of Cynefin is shown as a cliff. Finally the central domain of disorder is inherently undesirable as it is the state of not knowing which of the other domains you are in.

So the real domain of weak signal detection and therefore of what we can call asymmetric threat and opportunity is the complex one. Our traditional means of making decisions fall down here as there is no easy resolution through an evidence based approach, indeed for the reasons outlined early it can be plain bloody dangerous to apply those means. But it is the domain in which opportunity and threat exist and therefore one in which we need to be comfortable operating. At its simplest level we understand a CAS by acting in it, parallel safe-to-fail low cost experiments rather than planning or evaluation. It means we have to use network intelligence and dynamic real time feedback loops to see what is

emerging and respond accordingly.   We keep out options open for as long as possible and we bring as many diverse perspectives to bear as we can afford. It is this domain for which SenseMaker®, with its counter-terrorist origins was designed.

SenseMaker® and how it makes a difference

SenseMaker® is a software tool with associated methods and processes. Its features and capability can be summarised as follows:

Uses primary data without interpretation

It captures fragmented experiences, impressions and stories both reflectively in real time without requiring evaluation or interpretation. Collectively known as *fragments* or *Micro-narratives* these match the primary sense-making capability of the brain which blends together fragmented memories (both personal and those of others in narrative form) to come up with a form of action in each context. This capacity to capture anything in written, oral or pictorial form is key to weak signal detection. We can't afford to screen or restrict the material as its the chance observations that we will need to pay attention to.

One consequence of this is that SenseMaker® can be used to replace field and engineering notebooks as well as being a suggestions box, anomaly reporting device and survey instrument. Deployed on smart phones as well as the web it focuses on pervasive capture of anything that might make sense as it happens, as it is seen.

Human metadata as the primary interpretation

It allows those *fragments* to be interpreted at the point of origin into a *high abstraction* interpretative structure that prevents gaming (people knowing in advance what answer is desired or considered to the the right or low risk one). Human language evolved from abstractions, cave painting preceding the development of language itself. In consequence we are happiest with abstractions (think of the way we use metaphors and images) as uncertainty increases. Abstractions allow for *necessary ambiguity* critical for sense-making under conditions of uncertainty. It is this human metadata which allows us to scale to large networks of human respondents at little or no incremental cost.

It also allows us to sense patterns in metadata before we look at the originating material thus reducing pattern entrainment. That also overcomes issues on confidentiality of
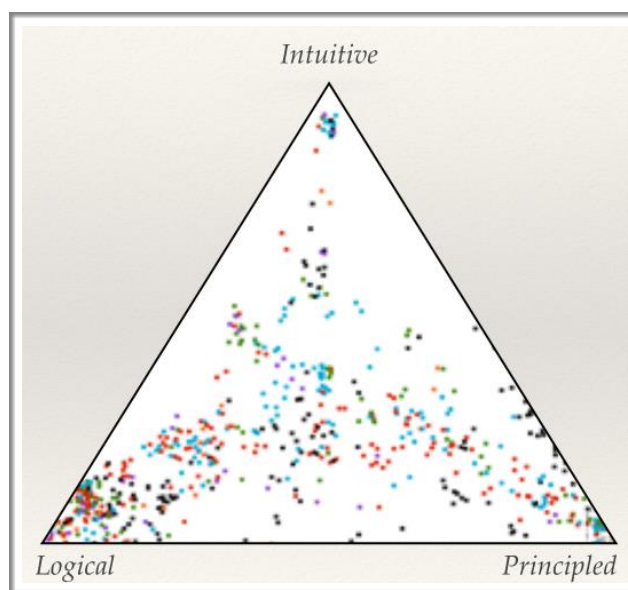
data as we only need to share the metadata, then request the original material when we understand the context and can explain why we need it. In the illustration from project mapping decision culture each dot represents an experience and the position of dot the placement by the person reporting that experience. In this case we can see the overall pattern of decision making in the organisation is balanced towards logic with little intuition. No problem under stable conditions, but indicative of an issue under conditions of high uncertainty.

Seeing the big picture, noting anomalies

SenseMaker® uses landscape and other visualisations to show both the dispositional state of the system as a whole and to focus the eye on anomalies and exceptions in real time to support decision makers with a direct link between the visualisation and the original fragments. This is know as *disintermediation* removing the mediating layers of interpretation between the decision maker and the source data. It also reduced the dangers of pattern entrainment as the fragments are only exampled to interpret a statistical pattern.

The illustration (from a live project) shows tens of thousands of self-signified fragments in a three dimensional landscape. The hollows represent stabilities, the peaks instabilities, the yellow dots outlier events. In this case the outliers are starting to cluster bottom right indicating a new opportunity or threat emerging. The decision maker can click on the model and see the data itself or request permission to access the data. The landscapes can also be presented as contour plots and linked to alert systems that advise senior decision makers when anomalies start to reach a trigger levels that require them to pay attention.

This can be used for preemptive approaches to opportunity spotting. Installed as a suggestion system or for scouts and other third party observers when fragmented ideas start to cluster it gives early indication of possible new areas where small investments would generate large returns.

Managing for serendipity

In the 1940's a Raytheon Engineer maintained the magneto of an early radar machine noticed that a chocolate bar melted in his pocket. It wasn't the first time someone had noticed this, but on this occasion he paid attention to it and was able to act on the observation to create the first micro-wave. In evolutionary biology this is known as exaptation, there a trait which evolves for one purpose enables something completely novel when the ecology shifts. Waiting and hoping that accidents will happen is one strategy, managing to make them more likely is another.

This is another use of SenseMaker® which involves scanning for opportunities not just evaluating proposals. To take one case where self-interpreted customer stories were combined with self-interpreted fragments from the various technical silos of an organisation. That clustering showed three depressions in the landscape where material from the technical silos had been interpreted in the same way as the

customer stories.   Of the five clusters identified, three then went into production; *exapting* technologies created for one purpose in a novel market.   This rapidly shortened the cycle time to market and create opportunities that were previously unknown in the organisation.

For an organisation with a University network or a large research organisation managed in silos this can be used to spot opportunities related to abstract qualities that will suggest novel products emerging from individual areas of research, but also for identifying trans-disciplinary opportunities.   This capability can also be used to associate past ideas and solutions in real time to current problems.   The high abstraction of the metadata structures allows for novel combination with few false positives than simple machine learning. Humans at the front, humans at the end, technology to scale their capability but not replace it.

Trawling for significance

One of the capabilities of SenseMaker®, originally developed in counter-terrorism, is the use of training data sets to create *classifiers*. Given that most organisations have directl or indirect access to multiple sources of past successes and failures this material can be used to create *anticipatory alerts* when something similar starts to appear. The material is signified by different groups of people (for example successful entrepreneurs, decision making executives) and the resulting micro-narratives become a training data set or *classifier*. Those classifiers are then thrown at the web, internal databases to replicate human interpretation of raw data, this can compliment raw interpretation more common in big data approaches but the human metadata creates better richer context and reduces false positives.

Critically it also shows traceability rather than being a black box which increases the possibility of acceptance. At the same time it allows for quick representation of different perspectives. For example seeing that in previous successful cases there has been disjoint between interpretation of original data between technical and financial decision makers is not significant of itself, but if that disjoint starts to repeat over rejected options then its a important decision support aid.

One possible use under discussion is the find the art movies that will be the unexpected block busters that make studios millions. Another is spotting possible
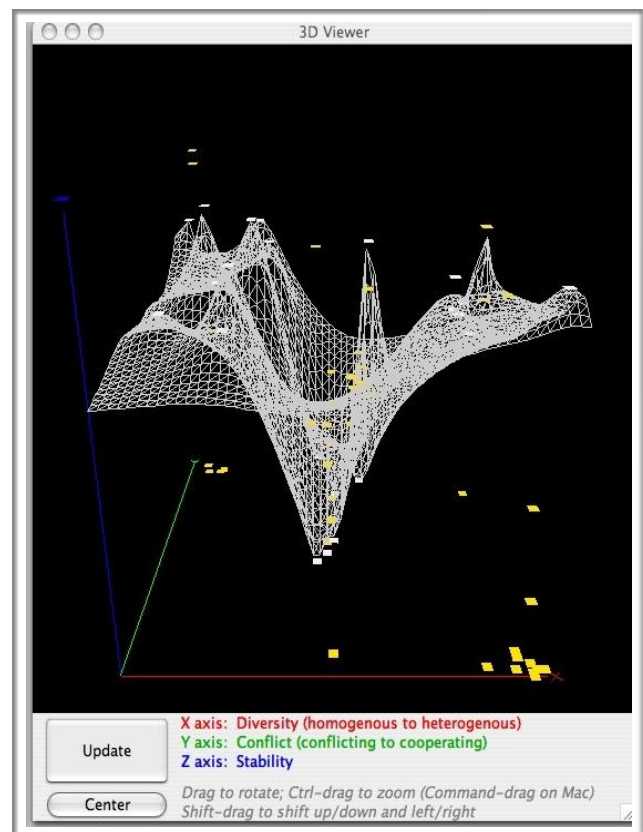
areas of premature discharge of patients from hospital something with personal costs, but also economic costs for the hospital itself in the US with high penalty costs associated with such failure.   In all these cases training data sets of naive observations prior to knowledge of outcome are used with interpretations when knowledge of outcome was present to create anticipatory alerts. Not to say *this particular thing will work* but rather to say *you need to look at this more closely*. This is a key move from anticipation and prediction to triggering humans to heightened states of alert. It also create an evidence base to persuade people to take on a novel or unusual idea.

In the area of investment appraisal it allows expertise on success and failure to be built into a sensing system without the need for the experts to be present.   Key to handling large volumes is the ability to trigger alerts where there is traceability of how the insight was gleaned. That traceability is key to acceptance to the results by decision makers

The human factor

Its all very well for a start up team to have all the right financial and technical resources in play. Its another matter all together to see if they have the right attitude. By getting start up teams to keep journals in SenseMaker® and by getting their customers to do the same we can get a real sense of the underlying attitudes and beliefs of the team.   Material that can't be gamed or interpreted to support a case. If you look at Grand Prix racing the technology can handle 98% but the final 2% is down to the driver.   The same applies in entrepreneurial and intrapreneurial teams. Attitudes are key.   Not only that the data sets obtained form this process can in turn become training data sets.

As an additional benefit the narratives create a fragmented knowledge management database that can be used by others for learning. This type of peer to peer knowledge flow is of increasing importance in international development as well as in industry; ideas need to spread and exapt in novel ways in real time to create a resilient organisation.

The wisdom of a human sensor network

The phrase wisdom of crowds is an unfortunate one, the tyranny of herds would be more appropriate. However a large network of humans with relevant experience making quantitive judgements without knowledge of how other people are assessing the data is a key new SenseMaker® enabled capability.

Once SenseMaker® is deployed within an organisation for multiple purposes (exampls include capturing material about core technologies across silos, employee satisfaction, user requirements capture, field notebooks, micro-scenario planning) then that network can be activated by a decision maker to quickly make a assessment through self-signification of a situation report or a proposal. The burden of rapidly interpreting something onto six triads is low compared with expert assessment so a

large number of people can be rapidly engaged. The resulting landscapes then inform the decision maker.

Used in international affairs for multi-agency assessment of failing nation states and rapidly changing situations this capability is now being deployed in organisations for real time decision making. Critically the decision maker can determine which networks are used and compare different results. Conclusions can then be drawn and represented as needed to get a wider perspective before the decision is made. If the overall question is confidential then the capability is simply used for aspects of the situation assessment.

In investment appraisal this can be used to engage employees, third party experts and even the applications themselves in a low cost, low time impact, ungamable approach to decision support.

The same approach has been used for micro-scenario creation with the landscapes replacing more traditional scenarios.

The above is a summary of the capabilities of SenseMaker® teamed with machine learning to provide a new approach to decision support. It's a paradigm shift in the way we think based on the three principles of managing complex adaptive systems:

1. Get the granularity right, small things combine and recombine in novel ways large things don't

2. Distribute the interpretation to a large network to make sure someone spots the Gorilla

3. Disintermediation is key, decision makers in direct contact with raw data.

See - Attend - Act

So we return to the basic decision model which separates seeing the data from paying attention to action. By distributing capture to large networks and allow them to self-interpret the material we radically increase the change that they will see things. The landscapes and other anomaly reporting means that decisions makers will pay attention to things that they would otherwise ignore. The fact that it is a quantitive technique also means that people don't read the micro-narratives before they had already realised their significance. The ability to look at the patterns and then ask *How to we get more stories like this, fewer stories like those* leads to smaller actions early, much easier that larger more resource intensive interventions later. Numbers on their own are objective but not persuasive; narrative on its own may be persuasive but not objective. Put them together and you have a powerful decision support capability.