Modelling Ignorance Levels in Knowledge-Based Decision Support

Ema Denby & John Gammack
School of Information Technology
Murdoch University
Perth, Western Australia
edenby@central.murdoch.edu.au
jgammack@murdoch.edu.au

Abstract

This paper outlines preliminary work towards a doctoral thesis which addresses the notion of *ignorance* in the context of knowledge-based decision support systems, both as an underlying source of error and of intentional choice. It is argued that a clear interpretation of the nature of knowledge that takes into account the time-dependent dynamic aspects is relevant for this analysis, and may be helpful in reaching normative principles suggested by the literature. This discussion is mainly concerned with aspects of knowledge that are typically found in organisations, and some types of ignorance that can apply. It is proposed that techniques such as fuzzy neural networks for learning domain characteristics and extracting rules, can provide viable methods for research in this area..

1. Introduction

Constantly in everyday life, and daily at all levels of organisations, people make decisions affecting themselves and others. Any decision may be evaluated on outcomes, which may be good or bad for reasons unconnected with the original decision, but to do with interactions with subsequent contextual factors. Separately from this however, any decision, purely *as a decision* may be evaluated as good or bad, and this is a distinct matter to do with the knowledge or ignorance brought to bear on the decision, and the procedures employed in making a choice. Good decisions may lead to bad outcomes, and vice versa. Presumably, the more relevant knowledge brought to bear on making a decision, the more likely a good outcome can be expected.

In this paper, the authors are concerned with a number of issues in modelling levels of knowledge and, equivalently, the ignorance involved in making decisions. It is argued that a clear interpretation of the nature of knowledge which also takes into account the dynamic or time-dependent aspects, is relevant for this analysis, and may help reach normative principles suggested by Klein & Methlie (1995). We consider the nature of knowledge as it is typically found in organisations, and some of the types of ignorance that can apply. The first part of the paper will review concepts familiar in Decision Theory with a view towards building upon those insights where applicable, and transcending their limitations where possible in a qualitative manner, leading to testable hypotheses. The view of knowledge here departs from that often used in this literature, and some typical examples of decision-making situations in organisations are considered (along with their ontological assumptions), and *ignorance* is characterised in relation to those. Primarily concerned with semi-structured and unstructured decision environments, eventually this research is intended to provide a framework for improved design and implementation of knowledge-based decision support systems (KB-DSS) in such organisational environments. In the words of the "Sage" JD Bernal[1] "The full area of ignorance is not yet mapped. We are at present only exploring its fringes."

2. Review Of Some Issues In Decision-Making Models

Much decision making has classically been characterised formally in terms of axioms embodying, rationality, utility, preference logics and optimality of various sorts (Klein and Methlie, 1995). On the view that individual's values and goals are stable, and their knowledge can be perfect in relation to some decision, models can be developed, and embodied in computable decision procedures along the

[1] See the page at: http://www.s-2000.com/quoteworld/ignorance.html (January 1999)

lines of typical decision support systems (DSS); those embodying econometric models and assumptions are another familiar class. Expert Systems or Knowledge-Based DSS are one technology particularly applicable in limited and stable domains, where decision-making can follow ordered paths; though the emphasis is on heuristic knowledge rather than purely axiomatic modelling. Fuzzy neural computing is a more recent development of knowledge modelling which allows the system to learn from exemplary data, and rules to be extracted and modified to ensure correctness.

The role of values, moral or ethical codes, and expedient pressures due to contextual factors, condition and deeply inform decision processes in many domains, and decision-making becomes subjective and pragmatic. Since Savage (1954), notions of utility, expectation and values have become subjectively referenced and axiomatised as such. These have been modelled, quantified, or otherwise designed into decision support systems in various ways. A generalisation of Savage's (1954) axioms to encompass qualitative factors has been proposed by Lehmann (1997). Elements such as deliberate neglect of some events and non-standard utilities, have been introduced as mathematical models which reduces *qualitative* decision theory to *quantitative* decision theory. On values, Keeney (1992) in particular, has argued that the focus should be upon the intended objectives that give a decision its meaning, and identify the values that lie behind a choice among alternatives.

2.1 Objective and Subjective Factors

Whether objective or subjective factors are dominant, the elements involved in making a decision must be made explicit or otherwise reified, and a typical assumption is, that the relevant knowledge *is* available or can be approximated by statistical assumptions about event probabilities. In many cases though, the really decisive presuppositions and axioms informing decision-making require to be elicited and understood, if a decision is to be appropriately assessed. These presuppositions may be deeply held beliefs of individuals, which they are unwilling to admit, or perhaps not clearly conscious about. These may reflect economic, religious, social and psychological values and characteristics of the individual upon which they may be inflexible, but which dictates their decision procedures. Along with arbitrary choices on event-neglect, and culturally conditioned values, these form potential sources of ignorance.

In decision-making involving uncertainty, risk and incomplete knowledge of the situation is much closer to that typically found in organisational settings. In group decision-making, various stakeholders each with partial and not overlapping knowledge, potentially different values and discrepant deep objectives, provide another set of familiar problems. Each member is partially ignorant, and collectively, unless all relevant knowledge is surfaced, a group's decision can be incompetent. A corollary to this is that, unless learning takes place, and is registered in the organisation beyond the temporary groups that form for localised purposes, ignorant decisions can be repeated indefinitely. The wisdom embodied in individuals who have been involved in, and experienced the consequences of typical organisational decisions, is not to be undervalued. Their knowledge is, though not necessarily infallibly, possibly the nearest to an objective model against which decisions in semi-structured and unstructured domains can be assessed for quality.

2.2 Two examples in organisations

Two examples are given here which apply in numerous organisations: (1) recruitment decisions and (2) abstracting of detailed document design for reporting and archiving purposes.

Example 1. Academic Recruitment and Selection - Choosing a new colleague is an example where deep values are likely to prevail, and chairing a recruitment panel is usually reserved for the most senior colleague, who is finally responsible for the decision. Panels may be constituted on equity, role relevance and similar grounds, but the decision makers therein are likely to be informed by their idiosyncratic human values, their role requirements, and ostensible psychological biases (cf. Kahneman & Tversky, 1973:1982). Local expediencies may also dominate a decision at any given time, particularly where knowledge of a higher level strategy is not tabled - this is another source of ignorance.

Organisational learning and memory depends on repositories of experience retained in useful forms within organisations. Often this is embodied in particular individuals with a long personal history with the organisation. Their experience with decisions made, the factors and forces identified, and their workability and consequences, is a significant knowledge asset, which DSS designers naturally have ambitions to retain. Organisations are vulnerable to the loss of such individuals, and unless that knowledge can be formalised and abstracted usefully for future generations of the organisation, wheels are likely to be reinvented. The second example addresses this familiar situation.

Example 2. Document Design Trails - Design organisations such as those in, construction and mechanical engineering, general manufacturing, producers of soft products such as documents, software and financial products, typically go through a process in which critical decisions affecting the final product are made. These are not always explicitly recorded, and their traces may only be in temporary documents such as scrap paper or notebooks. Calculations, sketches, meetings and false start activities apply, and candidate designs are accepted or rejected. Furthermore many similar designs are likely to exist in the organisation's history, which can be reused or adapted in many cases. A need is perceived for organisations to retain these creative activities for learning purposes, since early design decisions deeply affect final outcomes, and having the explicit rationale for a design facilitates future design decisions, eliminating dead ends identified in previous experience.

A system implemented by Gammack & Jenkins (1997; 1998) encourages design activity to take place on-line, in a non prescriptive electronic equivalent of a blank piece of paper. The information available to a designer, and the decisions made, are invisibly recorded via an audit trail. An extension of this work is to have experienced designers critique the decision path from conception to delivery of a product, noting key decision choice points, and the roles and sources of information at those points. Collecting behavioural protocols incidentally, then verifying those by a retrospectively constructed expert protocol is a development of standard protocol analysis techniques, likely to provide very useful and informative records for future generations of designers.

The literature on organisational learning suggest that wisdom can come from any source or level in an organisation (Senge, 1990) and is not necessarily conferred by positional seniority. Whilst this is surely true, a reconception of seniority, in terms of knowledge, experience and wisdom is required in evaluating knowledge based decision making.

In reference to the recruitment example above, *values* based on fundamentally different types of mindscape (Maruyama, 1980), which are cross cultural, pervasive, and lie beneath even value choices, may be found in the panel. Some panel members may wish to have a homogenous pool of colleagues, based on comfortable values and personal "likes": others may wish a heterogenous pool to increase diversity, regardless of personal compatibility. Some may prefer biddable subordinates, others may welcome feisty free thinkers, and so on. (cf. other examples given in Woolgar, 1998). - The distinctive patterns of reasoning or logics underlying such individual differences in decision choices has been clearly described by Maruyama (1994) with particular reference to improving business practice, and multicultural management. His work is considered by many to be an outstanding contribution to the advancement of knowledge (Caley and Sawada. 1994), and likely to provide a useful conceptual foundation for recognising patterns of knowledge and understanding typically found in organisational decision makers.

Returning once again to the Recruitment example, a member of a panel with a preference for a fit with local details on a teaching function, may be complemented by one with a preference for a specifically compatible research colleague. A departmental manager may have an understanding of the critical tradeoffs involved between these, and with respect to other functional requirements where there may be a lack. But a "senior" colleague who has experienced and since transcended all of these specific concerns, may make a decision based on a future strategic positioning requirement, which may not yet be widely known. A model of each one's knowledge and preferences is likely to show that the "wiser" colleague's is more encompassing, relatively objective, and may include the constructs of the relatively "ignorant" colleagues. All understandings of knowledge by individuals should not necessarily be given equal weight in a given situation, but ideally participants should be aware of the relative weight each individual carries. This research aims to model, in a useful way, the qualities of

knowledge that can be brought to bear on decision-making, in addition to, *domain* and *task* knowledge. The use of *inference* knowledge, and novel combinations of knowledge to solve problems, are also very important. Operationalising *ignorance* parameters as either "limitations", or as "advantages"[2], that impose upon various decision states, should assist in a more complete model for DSS.

2.3 Towards a taxonomy of ignorance

Ignorance can be construed as the state of there being corresponding knowledge 'out there' inaccessible to the decision maker. Holtzman (1989) identified several levels of ignorance by describing the equivalent knowledge paradigm within which it occurs. Table 1. presents an expanded version of the Taxonomy which identifies "Ignorance Level", "Description" and "Knowledge Required" separately for clarity. This should be useful as a starting point for mapping ignorance levels onto the decision-making situation.

Table 1. Taxonomy of Ignorance

Ignorance Level	Description	Knowledge Required
Combinatorial	Computational task too difficult, e.g. problem with 10^{40} variables.	Mathematics model available; use of supercomputers.
Watsonian	Cannot make the connection from all the clues; solution method incomplete.	Method for determining the important facts from the unimportant ones, and drawing the right conclusion.
Gordian	King Gordius tied a knot for the future king of Asia to untie. Alexander the Great was able to "untie it" by cutting the knot with his sword, thus solving the problem in an unusual way.	Lateral thinking; are there "rules" to be broken?
Ptolemaic	Attributed to the Greek mathematician and astronomer, Ptolemy, whose model of the universe centered around a stationary earth.	Evidence and observation of reality.
Magical	"No one knows how it works, but everyone knows that it works", e.g. the use of Aspirin and other similar drugs.	Trial and error.
Dark	No model is available but one is aware of the issues, e.g. "What is Life?", "Consciousness", etc.	Future of Science
Fundamental	Unaware of issue. (Ignorance is bliss!)	

From this taxonomy it is easy to discern the type of problems the decision-maker faces in situations similar to our previous examples. For instance, with the document design example, it is reasonable to suppose that in a lot of design cases designers experience *magical* ignorance when using techniques "that work" because it seems there is no explanatory basis.

A clear limitation to this taxonomy, to which Holtzman (1989) points outs, is, that it is deterministic and therefore not representative of all scenarios. One way to overcome this problem is to create possible extensions by adding the dimension of uncertainty or, disjunctive knowledge. Another

[2] Some recent research findings by Gigerenzer et al. (1999), reported by Douglas (1999), show that "ignorance pays" in stock market investments.

extension is to take into account the time-dependent characteristic of knowledge (how urgent is the issue at hand?). This is will constitute future development of this work.

2.4 A note on the position regarding human knowledge

The view of human knowledge which we adopt here is one which is systemically related to ignorance, and which, whilst constructivist and open to interpretivist processes, allows for testable formalisations of understandings in semi-structured and unstructured domains. In decision making, the uncertain is judged by comparison to an "object supposed certain" (Cusanus, 1440) and as the knowledge of anything by description is based on the knowledge of something else by acquaintance, the description is necessarily approximate and inadequate (Hawkins, 1954). Following the inspiration of Cusanus, and recent critiques in the philosophy of science (e.g. Horgan 1996) we consider formulations of knowledge to be finite human constructions, provisional and revisable, and which are buttressed by empirical reliability, historical establishment and consensual acceptance. In these terms the absolute truth of things cannot, by means of comparison, be reached by a finite intellect, but this is not to devalue the *activity* of knowledge acquisition. As Blake (1790) observed - if the fool would persist in his folly he would become wise. In Cusanus' analogy, the relationship of our intellect to the truth is like that of a polygon to a circle, infinitely precise approximations increase resemblance, but do not constitute an identity.

The authors believe that this understanding of the nature of human knowledge is helpful in modelling the constructs of individuals and groups towards a more perfected ignorance, a *learned* ignorance. Against this the relativist understandings of those with greater or lesser degrees of what is commonly understood as ignorance can be assessed. Thus whilst making no claims to an absolute truth for any particular domain, the levels of knowledge commonly accepted, agreed or validated, may be formulated for comparative assessments, and some form of differential bases for decision making can become established.

3. Research Methods and Questions

The methodology for investigating these issues derives from the theories of Kelly (1955), and in particular his repertory grid method. Kelly's method has been widely used in modelling expertise (Gammack and Young, 1984; Parsaye & Chignell, 1988; Giarratano & Riley, 1989) and in eliciting the understanding of individuals and groups in many clinical and organisational settings, but often with little regard to its constructivist host theory and its assumptions and corollaries. In this research we intend to remain close to the philosophy of the method, and to use it in establishing formulations of knowledge in organisational contexts. For clarity, we will endeavour to describe more precisely the difference between structured, semi-structured, and unstructured decision environments, and how this applies to the operational concerns described.

Based on the Repertory Grid technique, can the elicitation of ignorance be acquired in much the same way as knowledge acquisition? By selective use of features, traits, opposites, and attributes of objects of particular *cases*? A way to deal with uncertainty is to consider the use of *fuzzy qualifiers* (e.g. Somewhat True, Fairly True, etc.) and *quantifiers* (Most, Many, Usually, etc.), (Giarratano & Riley, 1989) - can these extended to include ignorance? How can the researchers deal with the problems of investigating this seemingly controversial term? More recent techniques within the connectionist paradigm such as fuzzy neural networks may be used to obtain empirical models and for testing these assumptions. By using these methods it is expected that we can construct sufficient *extensions* to the taxonomy of ignorance and how knowledge emerges. A preliminary case study (Denby & Gammack, 1999) using the evolving fuzzy neural network developed by Kasabov (1999) was applied to the domain of colour perception; this indicated that some aspects of knowledge and ignorance can be model. In addition to operationalising the taxonomy in a research setting, the authors are also concerned with three other research questions:

1) To what extent does the taxonomy of ignorance impact on the decision-making process?

2) What strategies do people use to by-pass ignorance in a given decision-making situation.

3) How can the above two questions be incorporated into a model for dealing with ignorance in a decision-making situation.

Furthermore, the criteria of *local correctness* and *extrapolability* formerly devised to deal with the validity of a formal decision method (FDM's) in a given domain, may be used to validate the correctness of the *ignorance* model.

Although this work is at an early stage, the authors have indicated the lines along which to develop an appropriate model of knowledge and ignorance in semi-structured and unstructured decision-making environments. This model will inform the design of future knowledge-based decision support systems, and facilitate intelligent management of knowledge assets.

4. Conclusion

A discussion of knowledge and ignorance was presented with two examples. It was argued that in order to better support decision-making in organizational settings we need to address not only the knowledge but also the way in which ignorance shapes decision processes. A Taxonomy of Ignorance based on the work of Holtzman (1989) was considered. Further developments will take into account *disjunctive* and *time-dependent* characteristic of knowledge; this is will constitute future developments of this work. The authors have indicated the lines along which to develop an appropriate model of knowledge and ignorance, relevant to semi-structured and unstructured decision environments. Based on a preliminary case study, it is suggested that techniques such as fuzzy neural networks can be used for testing some assumptions and model empirical data to obtain normative rules.

5. References

Denby, E. & Gammack, J. (1999) The Naming of Colours: Investigating a Psychological curiosity Using AI. 6th International Conference on Neural Information Processing. *Proc of ICONIP'99*, Perth, Western Australia, 3, 964-973.

Douglas, K. (1999) Decisions Decisions: your essential guide to making your mind up. *New Scientist*, 163 *(2202)*, 30-35.

Giarratano, J., & Riley, G. (1989). Expert Systems: Principles and Programming. Boston: PWS-Kent Publishing Company.

Hawkins, D. J. (1954) Introduction to Cusanus op cit.

Holtzman, S. (1989). **Intelligent Decision Systems**. MA: Addison-Wesley Publishing Co., Inc.

Horgan, J. (1996) *The End of Science* Abacus, London.

Kahneman, D., & Tversky, A. (1973) On the Psychology of Prediction. *Psychological Review*, 80 *(4)*, 237-251.

Kahneman, D., Slovic, P. & Tversky, A. (1973) **Judgment Under Uncertainty: Heuristics and Biases**. New York: Cambridge University Press.

Kasabov, N. (1999) Evolving Connectionist Systems for On-line, Knowledge-based Learning: Principles and Applications. *The Information Science Discussion Paper Series*, University of Otago, *2*, 1-31.

Keeney RL (1992). **Value-Focused Thinking: A Path to Creative Decisionmaking**. Harvard University Press.

Klein, M., & Methlie, L. (1995) **Knowledge-based Decision Support System with Applications in Business**. NY: John Wiley & Sons.

Lehmann D. (1997) Qualitative Preferences in Deliberation and Practical Reasoning Workshop at AAAI Spring Symposium 1997, March. http://www.medg.lcs.mit.edu/qdt/sss97/program/lehmann.ps , *Accessed January 1999*.

Maruyama M (1980) **Mindscapes and Science Theory Types.** *Current Anthropology* , *21*, 589-607

Maruyama, M (1994) *Mindscapes and Management*.UK: Dartmouth Aldershot.

Parsaye, K., & Chignell, M. (1988) **Expert Systems for Experts**. NY: John Wiley & Sons, Inc.

Savage, LJ (1954) *The Foundations of Statistics*, Dover Press: New York.

Senge, P M (1990) *The fifth Discipline* New York: Doubleday.

Woolgar S (1998) A New Theory of Innovation? *Prometheus*, 16 *(4)*, 119-136.