



Murdoch
UNIVERSITY

MURDOCH RESEARCH REPOSITORY

<http://researchrepository.murdoch.edu.au/1894/>

This is the author's final version of the work, as accepted for publication following peer review but without the publisher's layout or pagination.

Hobbs, V.J., Pigott, D. and Gammack, J. (2002) *Inaccuracy, ambiguity and irrelevance: An analysis of the nature of quality in knowledge management using the noetic prism.* In: Burstein, F., Linger, H. (eds), 2002, *Knowledge Management and Context*, Australian Scholarly Publishing, Kew Vic Australia. Pp 42-54.

Copyright © the Authors.

It is posted here for your personal use. No further distribution is permitted.

Inaccuracy, ambiguity and irrelevance: An analysis of the nature of quality in knowledge management using the noetic prism

Valerie Hobbs[†], Diarmuid Pigott[†] and John Gammack[‡]

[†]School of Information Technology
Murdoch University
Murdoch, WA 6150
{vhobbs, dpigott}@murdoch.edu.au

[‡]School of Management
Griffith University
Nathan, QLD 4111
jgammack@mailbox.gu.edu.au

Abstract: Quality, as it applies to knowledge systems, is not singular in nature but emergent from aspects of the system which themselves are subject to variation in quality. We suggest that there are three quality determinants, which we can characterise as inaccuracy, ambiguity, and irrelevance. We use the framework of the noetic prism to describe inaccuracy, ambiguity, and irrelevance formally as a result of variance in the dimensions of granularity, shape and scope in noetic material, and suggest that any unreliable system will always contain an amalgam of the three. We suggest some practical implications based on this analysis.

Keywords: knowledge management, quality, noetic prism, inaccuracy, ambiguity, irrelevance

1. INTRODUCTION

We consider that the concept of quality, as it applies to knowledge systems, is not singular in nature, but rather emergent from aspects of the system which themselves are subject to variation in quality. Overall quality, or ‘conformance to requirements’ (Crosby, 1979, p15) will be influenced by the components of the system and their quality or lack thereof. As an emergent property, quality can be assessed by reference to purposive or other performance criteria, but the intrinsic determinants that generate desired outputs require separate consideration.

Since for complex systems quality determinants may reside in, for example, the effectiveness of procedures, the integrity of digital assets, relationships inside and outside the system, and properties of the representational formalism, a blanket measure of quality is unlikely to address the areas in which each of these generative determinants may be deficient. The framework of the noetic prism (Pigott & Hobbs, 2001) allows a theoretical separation of the quality determinants applicable in knowledge systems, and ensures quality interventions are undertaken at the generative rather than the ‘symptomatic’ level of a system.

The principle of equifinality (Bertalanffy, 1968) implies that various component-based knowledge system architectures may well produce similar output standards – there are many cars displaying comparable performance, despite being based on distinctive proprietary components and designs. Although this might seem to necessitate an analysis of component quality referenced to individual cases, this would lack true generality, and we suggest that universal dimensions of quality relevant across knowledge systems can be identified, then applied in those cases.

The class of knowledge systems we consider for illustration are knowledge based systems (KBS), in particular advisory (expert) systems, whose typical architecture is well understood and for which performance measures based on output can readily be established in domains for which normative standards exist. To the extent that all knowledge management (KM) systems require organisational

knowledge to be explicitly represented in reliable forms, whether documents, procedures, data stores or any of the emergent KM technologies, issues to do with the quality of the knowledge base will apply.

Any advisory system can be seen as a system of signs that are bound together in a chain of reasoning to give contextualised, accurate and unambiguous advice. In a diagnostic expert system case, a logic, taking probabilistic or categorical input data, referenced to a body of 'facts' and heuristic rules, will rationally lead to an unambiguous or qualified output, which can be compared to benchmarks, other experts or 'gold' standards. It is worth noting that an influential quality notion, zero-defect, or 'Do it right first time' (Crosby, 1979), is not applicable in any simple way in advisory systems, since many aspects of human knowledge are only provisional constructions, representing evolving theory. More recent approaches directed at business processes, such as Six Sigma (Harry & Schroeder, 2001) emphasise continual process improvements, measurement and a customer orientation for outputs. By analogy, the internal quality of advisory systems must ensure that the processes that produce outputs can be inspected, evaluated, revised and deliver priority outputs.

When an advisory system is found to be wanting, it becomes apparent through its becoming obtrusive. Any advisory system must serve its part in the user's epistemology with complete transparency: if it obtrudes or complicates, then it may be said to have failed in its role, like dusty spectacles or broken binoculars. This is the moment when it becomes apparent that a system is no longer delivering results that match its purpose. The point of obtrusion is the point of failure: this is an *s-curve* phenomenon, with a growing number of deficiencies gradually leading up to the dramatic success/failure dichotomy.

An advisory system fails when the nature of decision support exceeds the applicability of the solution: the situation which has led to its being consulted brings with it a scale that determines the range of acceptable solutions. If the solutions presented exceed that range, it will not be applicable. However, such a system will still be used: the question is one of trust in the outcomes of the system. Even a small increase in the imprecision of an advisory system can lead to what has been termed *subcertainty*: an uneasiness or a lack of confidence in the system's advice, which is not sufficient to stop its being used, but nevertheless undermines the trust that a user must have in it to use it with comfort (Kahnemann & Tversky, 1979). And this subcertainty feeds back to compound the effect: this in turn leads to lack of use, and the system becomes marginalised from the user's lives and decision making processes.

What we must consider is how this subcertainty (as a symptom of quality deficit) can be used as an indicator of the cause of quality deficit, and as a measure of its severity.

2. INACCURACY, AMBIGUITY AND IRRELEVANCE

The individual components that comprise advisory systems will always be illuminated or adumbrated by the mechanisms whereby we instantiate them. The symbolic representation of facts has been pre-determined elsewhere, the active situation into which the outputs are interpreted cannot be fully anticipated, and the (chaining) logic makes its own assumptions. These aspects conform to semiotic reference types identified by Peirce (Peirce, 2000; Sowa, 2000) in which iconic, indexical and symbolic mechanisms are simultaneously at work. Each of these mechanisms requires a separate consideration in respect of the applicable quality criteria, addressing iconicity (representational quality), indexicality (referential accuracy) and contextualisation (applicability).

An advisory system will thus have aspects of each of the Peircean types: and, presenting different aspects to the user, will provide them with reassurance based on its role as an iconic, indexical or symbolic mechanism. Thus, an *iconic* system of signs gives us form, which we may take to be giving us statements that are unambiguous in their propositional content. An *indexical* system points clearly

and without error to the objects in the world that are under consideration, or which are the cause of the advisory system's being consulted. The *symbolic* moment of an advisory system is its ability to have sufficient of the contextual component to function as a credible basis for decision – this may be implicit, tacit, explicit, or constituted by appeals to authority.

If the system of signs can present the world in these ways, then it is also susceptible to fail in these three ways, giving errors that we can categorise as *inaccuracy*, *ambiguity*, and *irrelevance*.

Inaccuracy. The failure in *accuracy* is at first sight easily decidable: simple values when they err can be mapped one to one against the world, the discrepancy noted, and amended. The trouble is that we rarely work at the level of simples, but more often at the level of compound values. Here the possibility of one-for-one checking against the world is no longer straightforward (putting aside the various problems of unmediated auditing of values against the world!). So the mediation becomes one of checking simples, compounds and conclusions against the world. The user must then either use another system (implying an potentially infinite regress of recursive checking) or fall back on an intuitive assessment of the state of the world. An advisory system serves as part of the general epistemological mechanism: as we have seen, if it obtrudes or complicates, then it may be said to have failed in its role. And the failure becomes obvious through the system's obtrusion. This deficiency, of systematic inaccuracy, makes more, not less work for the user consulting the system, and defeats the purpose of having it in the first place.

Ambiguity. When the system fails to provide answers that have a single value (or prescribed set of values) then the user will again experience the system obtruding. This type of failure is often due to underdetermination. Such indeterminacy implies an uncertainty that again requires the user to do additional interpretive work, this time to make a personal appraisal of the most likely meaning in the range provided. Such a probabilistic interpretation renders the purpose of the system severely compromised, and this type of deficiency we call *ambiguity*.

When we look for clear answers from a set-oriented manipulation of decision-making the use of reasoning engines to combine propositions to make more complex wholes relies on a clear delineation of subject and predicate. While contingent conclusions are possible, where the delineation is not clearly defined then the simple advisory nature of the system has been replaced with one where a fair degree of consideration and arbitrary judgement will be relied upon. Ambiguity precludes the operation of that type of automatic reasoning that lies behind advisory systems. Moreover, an ambiguous subject and an ambiguous predicate may no longer be guaranteed to relate to each other – there is no guarantee that two sets of ambiguities will grow or contract in the same direction or to the same extent. Each syllogism will therefore have to be reviewed in itself and as a premise for further work. Ambiguity amounts to the inability of a sign to fulfil its iconic role, and a process of attempted disambiguation cannot be guaranteed to restore meaning.

Irrelevance. When the system gives advice that fails to be *meaningful* we can look for a failure in the contextualisation process. This may result from a failure in design, such as failure to provide results in a context, failure to record input within a context, or failure to provide the context to the mill of reason that operates within the system. All of these failures will mean that a superset of results will be delivered to the user of the system (an experience that is all too familiar to users of search engines on the Internet): when this happens it is necessary for the user to provide an additional interpretive role. The system in one way or another fails to contextualise effectively, and can be summarised as failing on the dimension of relevance. While this can be dealt with through individual attention to the cases, it is not something that can be mechanised and so a deficiency of this sort will prevent the system scaling up effectively.

A type of failure of relevance that may be impossible to predict is where the exactitudes of maintaining currency have proved too much for the administrators of the system, something we often

see in media based systems. Here, not only the keyword set, thesaurus, authority files and description guidelines, but also the data and the media gathered together, must fit the stylistic requirement of the system. If a change of audience or the passage of time results in a stylistic clash that renders the displayed results open to misinterpretation through irony, ridicule, or deconstruction, the results will no longer be meaningful. This situation is difficult to monitor as it cannot be done systematically, but only at the level of the instance. Finally, the accidental extension of a system through its repurposing can render the explicit diffuse, requiring intervention to maintain correct contextualisation (Pigott, Hobbs, & Gammack, 2001).

3. THE NOETIC PRISM AND QUALITY MODELLING

Elsewhere (Pigott & Hobbs, 2001) we have argued that existing definitions of ‘data’, ‘information’ and ‘knowledge’ derive from a flawed model of hierarchy and transformation, and that instead we can describe collectively all of the materials of computation as *noetica*. Changes in complexity of the noetica occur due to value-adding through the imposition of three different principles: granularity, or increase in aggregation; shape, or increase in set relatedness, and scope, or increase in contextualisation through the formation of networks.

We proposed the *noetic prism*, a conceptual space framed by a triangular prism that represents the location and behaviour of the noetica. The nature of any point in the noetica is characterised by the extent to which granularity, shape or scope is informing the decision-making or processing, and the complexity of the structure that it exhibits. The emphasis of operations on the noetica shifts between the vertices of the prism with each consecutive analysis and operation, because we must always look with instruments that work primarily in one of the dimensions at the expense of the others.

In Pigott & Hobbs (2001) we showed how the terms data, information and knowledge could be reconceptualised as late-binding, purpose-determined aspects of the same body of material: *data* as noetica viewed primarily along the shape vertex; *information* as noetica viewed primarily along the granularity vertex, and *knowledge* as noetica viewed primarily along the scope vertex. Particular bindings are brought out through the noetica being analysed and modelled using various tools and formalisms. Thus the prism shows how a ‘hierarchy’ of any sort (such as the traditional data-information-knowledge hierarchy and its revisions (e.g. Bellinger, 1997; Davenport & Prusak, 1998; Earl, 1994; Sveiby, 1997; Tuomi, 1999) is an inappropriate model for the disposition of the noetica, as the appearance of the noetica to any process at different points (in time or space) becomes one of focus rather than separation. In Gammack, Pigott & Hobbs (2001) we further used the noetic prism to show how it presents a useful framework for modelling context in managing knowledge.

3.1 The noetic prism: formal aspects

We briefly present some formal aspects of the noetic prism next, as described in Pigott & Hobbs (2001). The noetic prism permits a *four*-dimensional co-ordinate vector space: the vertical axis represents complexity, which we define as a measure of the intentionality stored in the noetica (as a function of time, effort and skill), and the three vertices represent the three dimensions of noetica – granularity, shape and scope (Figure 1).

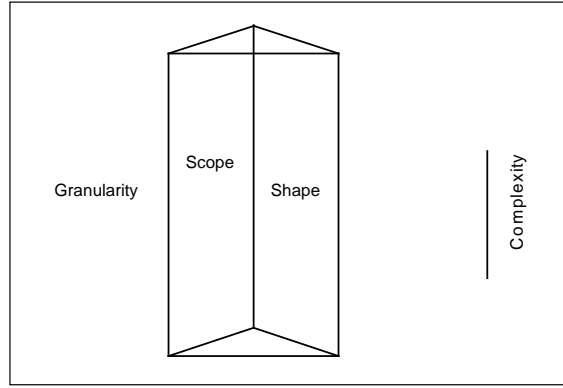


Figure 1. The noetic prism showing vertices of granularity, shape and scope.

Position in the vector space is given by the vector sum of the *importance* of each of the three dimensions. Importance is a vector of *extent* (displacement from a vertex) and *complexity* (displacement from the base of the prism). Each point \bar{N} in the noetic prism is thus determined by a 6-part value of

$$\bar{N} = (\varepsilon_G, \kappa_G, \varepsilon_S, \kappa_S, \varepsilon_C, \kappa_C)$$

where:

ε is a measure of the *extent* to which the vertex is significant at that point,

κ is a measure of *complexity* and:

ε_G	extent of granularity	κ_G	complexity of granularity
ε_S	extent of shape	κ_S	complexity of shape
ε_C	extent of scope	κ_C	complexity of scope

As this is a vector, it reduces to a 4-part value of

$$\bar{N} = (\varepsilon_G, \varepsilon_S, \varepsilon_C, \kappa_N)$$

where κ_N is the net value for complexity.

Any noetic state can thus be delineated in terms of the relative proportions of its noetica along the vertices of granularity, shape and scope. Since any point in the prism space will always have three vectors, it can be seen that is impossible to have extent along one vertex without extent along the other two.

The space enclosed by the prism is fractal in nature and can be used to represent the noetic resource for a culture, an organisation or an individual. It may also be used to represent a single item such as a database, spreadsheet, or mailing list. In each case the noetic resource has its own measure of complexity, scope, shape, and granularity.

We can use vector mathematics to represent the behaviour of noetica within the prism: for example, the effect of adding a new body of noetica to an existing body can be measured as a vector addition:

$$\begin{aligned} \bar{N}_1 + \bar{N}_2 &= \bar{N}_R \\ (\varepsilon_{G1}, \varepsilon_{S1}, \varepsilon_{C1}, \kappa_{N1}) + (\varepsilon_{G2}, \varepsilon_{S2}, \varepsilon_{C2}, \kappa_{N2}) &= (\varepsilon_{GR}, \varepsilon_{SR}, \varepsilon_{CR}, \kappa_{NR}) \end{aligned}$$

3.2 Determining location within the noetic prism

Given that we can use these vectors to place any items of noetica precisely within the prism concept-space, we must now consider how to model noetica where that placement is not precise, the effect this has on how we can make statements about such noetica, and the confidence with which we can deal with it.

The precision of location of any item in the coordinate space of the noetic prism depends on two things, the scale of the measure we are using, and the error associated with the measurement. As with any mapping system, it is the coarsest grained scale that determines the precision of the overall location. This coarseness of scale means that we have to accept some uncertainty as to the exact location of measurements below that scale. In a similar manner, known error (expressed in the prism-space as a \pm confidence value) also affects the degree of uncertainty with which we regard a value. Together, we term imprecision due to scale and error the *tolerance* of the vector space for this series of measurements. The measure of this tolerance at any point can be termed its *variance*.

We can see that any imprecision in the extent and level of complexity with respect to a vertex for a point in the prism will weaken the *certainty* with which we can refer to it. This imprecision must be expressed in terms of a *range of values* for that pair. So, for any vertex \mathbf{V} , we have:

$$((\varepsilon_V \pm y), (\kappa_V \pm x)) \Rightarrow [\varepsilon_V \dots \varepsilon_V], [\kappa_V \dots \kappa_V]$$

where x represents the variance in complexity and y the variance in extent.

So, for example, variance in granularity would be expressed as

$$([\varepsilon_G \dots \varepsilon_G'], [\kappa_G \dots \kappa_G'], \varepsilon_S, \kappa_S, \varepsilon_C, \kappa_C)$$

with the confidence limits in the extent of granularity ($\varepsilon_G \pm y$) represented by the range $[\varepsilon_V \dots \varepsilon_V]$, and the complexity of the granularity ($\kappa_V \pm x$) represented by $[\kappa_G \dots \kappa_G']$.

When we resolve the six-tuple to the four-tuple, the net complexity of the other vertices is rendered concomitantly less reliable:

$$([\varepsilon_G \dots \varepsilon_G'], [\kappa_G \dots \kappa_G'], \varepsilon_S, \kappa_S, \varepsilon_C, \kappa_C) \Rightarrow ([\varepsilon_G \dots \varepsilon_G'], \varepsilon_S, \varepsilon_C, [\kappa_N \dots \kappa_N'])$$

If we consider the same imprecision in extent and complexity for all three vertices, the certainty diminishes accordingly:

$$([\varepsilon_G \dots \varepsilon_G'], [\kappa_G \dots \kappa_G'], [\varepsilon_S \dots \varepsilon_S'], [\kappa_S \dots \kappa_S'], [\varepsilon_C \dots \varepsilon_C'], [\kappa_C \dots \kappa_C'])$$

which resolves to

$$([\varepsilon_G \dots \varepsilon_G'], [\varepsilon_S \dots \varepsilon_S'], [\varepsilon_C \dots \varepsilon_C'], [\kappa_N \dots \kappa_N'])$$

and which indicates the *net uncertainty* of any point in the prism.

This net uncertainty undermines the precision with which we can locate the point in our conceptual space, and the resultant degree of freedom increases with the distance from the origin of the imprecision. With generational increases in *complexity*, we see a matching increase in the imprecision in locating the noetica in *extent*. This is a direct consequence of the role that angles play in vector mathematics: the original ideal of point-to-point navigation within the prism increasingly becomes an indication of likelihood. (This outcome would also be expected as a result of the

transformation of the vector by obtaining the product of that vector and a scalar representing the extent.)

This variance tends to propagate upwards in a cone shape as the lateral measurement of uncertainty increases in proportion to the magnitude of the complexity. The cone of tolerance represents the worst-case for the mapping of noetica in the concept space of the prism (Figure 2).

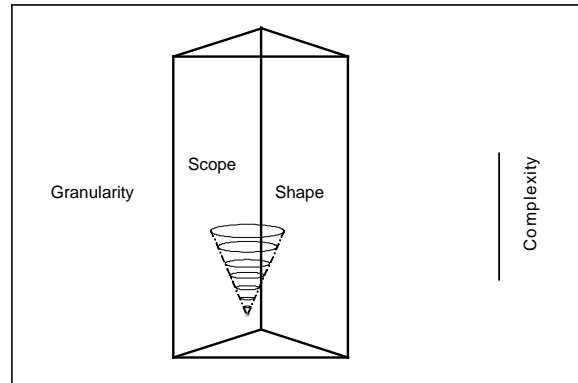


Figure 2. The propagation of variance about a point into a cone of tolerance.

We might therefore expect that the cone of tolerance would rapidly extend to cover entire regions from any single source of imprecision, rendering any placement of noetica within the prism to be so imprecise as to be useless. However, when we consider the effect of complexification on a item of noetica, we notice the *allorecurive* (Pigott & Hobbs, 2001) effect at work in the operations of complexity. Allorecurion is the principle of outward expansion in a constant pattern – the opposite direction to recursion, but with the same algorithm. With each generation of complexification, we experience the (permanent or partial) operation of occlusion with respect to the constituent noetica: manipulating a structure of a higher order of complexity allows us to effectively disregard its component lower-order structures (Figure 3).

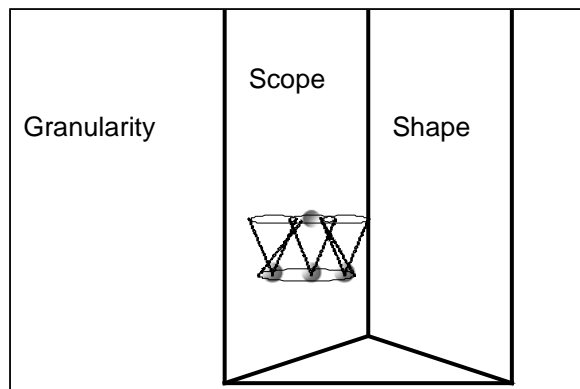


Figure 3. The mitigating effect of allorecurion on the cone of tolerance.

This effect is consistent with the predicted behaviour of systems generally: we know from systems theory (Bertalanffy, 1968), prospect theory (Kahnemann & Tversky, 1979) and attractor theory (Lorenz, 1991) that all systems tend to self-correct regardless of their content. As stated earlier, the equifinality of systems (Bertalanffy, 1968) disguises the various paths that error and scale can take to similar outcomes.

3.3 Characterising quality in the noetic prism

We can now return to our discussion of quality components, and how they can be characterised using the conceptual space of the noetic prism.

Granularity indicates instantiation (time, place, provenance, format) in the noetica: granularity is evidenced as documents, code, directories, file systems, and usage statistics. The extent of granularity is a measure of the importance of aggregated structures in the noetica. We can see that problems and sources of error associated with instantiation can arise through conversion error, inaccuracy in recording, imprecision and inaccuracy in transfer. These are direct measurements of the state of the world, and so every deficiency here distorts the picture we have of the world. This we can see results in an increased uncertainty as to the location in time or space of the noetica. We have previously described this type of uncertainty as inaccuracy. Problems with establishing the significance of granularity through variance on the granularity vertex are thus evidenced as *inaccuracy*.

Shape indicates presence of formal propositional structures in the noetica, and is seen in structures such as tables, indexes, data stores, views, object stores, databases, data warehouses, ROLAP. The extent of shape is a measure of the importance of compound structures in the noetica. Problems and sources of error that we see here can arise when a system delivers propositional results for which the subject (either singly or collectively) is unclear, resulting from such things as category errors and errors in the normalised function of data sets. We require an absolute existential import for the propositions before they can serve as function here. We have previously called this type of uncertainty ambiguity. Problems with establishing the significance of shape, through increased variance resulting from a failure of alignment or disjunction in the noetica, are thus evidenced as *ambiguity*.

Scope indicates the presence of organising principles in the noetica, and can be seen in names of fields, files and computers; classification schemes, procedure manuals, topic maps and ontologies. The extent of scope measures the importance of contextualising structures in the noetica. Problems and sources of error that we see here can arise as a function of the failure of the contexts of design, representation and inquiry to give a continuous semantic framework so that the perceived meaning of a result will have the same value as the intended meaning of the original observation. These problems can result from a poorly framed question, a badly designed expert system, or a flawed model, or through recontextualisation through relocation, migration or passage of time. This uncertainty we saw requires increased semantic mediation through the interposition of interpretive systems, and we called this type of uncertainty irrelevance. Increased variance on the scope vertex is thus evidenced as *irrelevance*.

These problems cannot exist in isolation: just as noetica will always contain components of all three vertices, so the factors that make a system unreliable will always contain an amalgam of unreliability, inaccuracy and irrelevance. As discussed earlier, this measure may be considered to be the net uncertainty: in consideration of quality determination, this can be seen to be the *subcertainty* of the system at that point.

Measuring subcertainty gives us an indication of the health of the advisory system – although not a primary measurement, it shows us where we have to begin to investigate likely causes of error. Attention to these causes gives us the minimisation in subcertainty that then delivers the overall quality improvement we are seeking.

4. PRACTICAL IMPLICATIONS

As discussed earlier, in investigating the quality components of advisory systems, we need a mechanism that will lie below the intentional aspect that the user experiences, yet which has a level

of abstraction above the components upon which it relies. We must establish a mechanism that can interpret to the system administrator the causality behind subcertainty levels. What is required is a two-fold strategy of automating the error recovery, and of preparing the material to be reflexive.

With automated monitoring the presence of such deficiencies below the s -threshold may inform us of the health of the system. The problem becomes one of finding ways of accommodating this automation, which in turn requires that many qualitative indicators be represented mathematically. Imprecise and qualitative information may be rendered quantifiable by a variety of means, for example fuzzy cognitive maps (Kosko, 1997), cumulative prospect theory (Kahneman & Tversky, 1992), and Bayesian networks (Pearl, 2000; Spirtes, Glymour, & Scheines, 2001).

The efficiency with which the source of any subcertainty can be located will also affect the system's use as a deliverer of timely advice: not only for correction as a means of improvement, but also for enabling error-checking to act as the feedback loop for an adaptive system. For this to be achieved, it is important that the nature and measure of source variance be close-coupled, and so we need to ensure that the entire system possesses a high degree of reflexivity at every level.

There are several different strategies for system reflexivity. Overall, we can stipulate the inclusion of system wide metadata, combined with standards of measurement and evaluation, to ensure that details of technical provenance, subjective content analysis and administrative requirements are retained with the indicated noetica. We may also call upon the formalisms of reflexivity in symbol processing, as established in various programming language systems, as an indication of the manner in which reflexive encapsulation of variance will not cause additional variance. There is a range of working solutions that may be considered, such as aspect-oriented programming (Kiczales, 1996), the frame representation languages of Bobrow & Winograd (1979), or Reynolds' CODIL (the Context Dependent Information Language) (Reynolds, 1971a, 1971b).

With a feedback system in place, the level of acceptable quality can be established. Given that no dynamic system (adaptive or static) can be perfect, the job of an administrator will often be one of allocating resources to attend to various sections (material, rules, categories). If we can find a mechanism for helping the administrator to determine what is acceptable quality and what are the criteria for acceptable quality, then we can establish criteria for verification, which in turn can assist in the design of future similar systems. Verification criteria must include criteria relating separately to the three prism vertices of granularity, shape and scope in order to address the potential problems of inaccuracy, ambiguity and irrelevance.

5. CONCLUSION

The framework of the noetic prism allows set-theoretic location of the semantic spaces embodied by the user and by the advisory system. It is assumed that the user is a 'larger' system in the sense of possessing an epistemic capability of bringing contextual information to constrain interpretation of outputs from the more limited, and certainly finite advisory system. The user has extra degrees of freedom than has the advisory system and so retains discretion in assigning meaning to its outputs in relation to the active world. Users have the prerogative here, since propositions embodied in an advisory system's representation are revisable. Conversely, if a trusted advisory or other similarly semantic system can bring new insights to a user, learning can occur and beliefs may be revised: these result in a development of the user's knowledge, through subsumption, assimilation and perhaps abstraction processes.

Issues concerning belief revision entail quality reference processes to ensure that established belief structures are not inappropriately changed. Kelly (1998) contrasts belief revision theory (which often emphasises retention of established beliefs when faced with new information) with learning power (the ability to arrive at true beliefs in a wide range of possible environments). In both cases the

noetic prism can help to model the evolution of epistemic states and their revision. For this the work of Gärdenfors (2000) is particularly relevant, in which a conceptual space may be built up from simpler geometrical structures that embrace the quality dimensions detailed above. Since the belief revision processes will tend to increase complexity in the overall system, by subsuming simpler structures, measures of structural complexity will be found relevant.

An advisory system is also part of a larger system that is its universal context. This situating is recognised in the move towards knowledge interchange formats and in the greater acknowledgement of the importance in developing systems within the context of broader heuristics such as the MPEG-7 program (ISO, 2000). Recognition of the role that the three prism vertices play in enframing the components of advisory systems can help both in adapting a working solution to effect interchange, and in designing a new system with a view to its immediate population with existing datasets. While our discussion has focussed on the pure explicit knowledge base such as is encoded in advisory systems, the prism is argued to have applicability to KM systems more generally.

We began by suggesting that components of quality could be identified separately by treating an advisory system as a system of signs. By examining the Peircean semiotic components of iconicity, indexicality and symbolism we defined three deficits in quality-sensitive systems: inaccuracy, ambiguity and irrelevance. We then showed how inaccuracy, ambiguity and irrelevance arise from the nature of variance within the noetic prism with respect to the three vertices of granularity, shape and scope, and how the prism conceptual space can provide a means of modelling these deficits, and of designing systems to monitor and minimise them. We believe that an analysis of the relationship between the vertices of the prism and the Peircean types may be able to clarify the nature of confidence in computing systems generally, and will return to this in a later paper.

6. REFERENCES

- Bellinger, G. (1997). *Knowledge Management — Emerging Perspectives*. Retrieved October 15, 2001, from the World Wide Web: <http://www.outsights.com/systems/kmgmt/kmgmt.htm>
- Bertalanffy, L. v. (1968). *General System Theory*: George Braziller.
- Bobrow, D. G., & Winograd, T. (1979). KRL: Another perspective. *Cognitive Science*, 3(1), 29-42.
- Crosby, P. B. (1979). *Quality is Free*. New York: Penguin.
- Davenport, T. H., & Prusak, L. (1998). *Working knowledge: how organisations manage what they know*. Boston: Harvard Business School Press.
- Earl, M. J. (1994). Knowledge as strategy: reflections on Skandia International and Shorko Films. In C. Ciborra & T. Jelassi (Eds.), *Strategic Information Systems: A European Perspective* (pp. 53-69). Chichester: John Wiley & Sons Ltd.
- Gammack, J. G., Pigott, D. J., & Hobbs, V. J. (2001). The noetic prism as a theoretical framework for contextualisation. Invited panel presentation at the Australasian Conference on Knowledge Management and Intelligent Decision Support - ACKMIDS '2001.
- Gärdenfors, P. (2000). *Conceptual Spaces: The Geometry of Thought*. Cambridge MA: MIT Press.
- Harry, M. J., & Schroeder, R. (2001). *Six Sigma, The Breakthrough Management Strategy Revolutionizing The World's Top Corporations*. New York: Doubleday.

- ISO. (2000). *Overview of the MPEG-7 Standard* (ISO/IEC JTC1/SC29/WG11 N3445). Geneva: International Organisation for Standardisation.
- Kahneman, D., & Tversky, A. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4), 297-323.
- Kahnemann, D., & Tversky, A. (1979). Prospect Theory: an analysis of decisions under risk. *Econometrica*, 47(February), 263-291.
- Kelly, K. T. (1998). The learning power of belief revision. In I. Gilboa (Ed.), *Proceedings of the Seventh Conference on Theoretical Aspects of Rationality and Knowledge (TARK-98)* (pp. 111-124). San Francisco: Morgan Kaufmann.
- Kiczales, G. (1996). Beyond the Black Box: Open Implementation. *IEEE Software*, 137-142.
- Kosko, B. (1997). *Fuzzy Engineering*: Prentice Hall, US.
- Lorenz, E. N. (1991). Dimension of weather and climate attractors. *Nature*, 353, 241-244.
- Pearl, J. (2000). *Causality: models, reasoning and inference*: Cambridge University Press.
- Peirce, C. S. (2000). What is a sign?, *The Essential Peirce* (Vol. 2): Indiana University Press.
- Pigott, D. J., & Hobbs, V. J. (2001, November 2001). *The Noetic Prism: A new perspective on the data-information-knowledge complex*. Paper presented at the Western Australia Workshop on Information Systems Research, University of Western Australia.
- Pigott, D. J., Hobbs, V. J., & Gammack, J. G. (2001). An approach to managing repurposing of digitised knowledge assets. *Australian Journal of Information Systems*(Special Issue on Knowledge Management).
- Reynolds, C. F. (1971a). CODIL: The CODIL language and its interpreter. *The Computer Journal*, 14(327-332).
- Reynolds, C. F. (1971b). CODIL: The importance of flexibility. *The Computer Journal*, 14, 217-220.
- Sowa, J. F. (2000). Ontology, Metadata, and Semiotics. In B. G. G. W. Mineau (Ed.), *Conceptual Structures: Logical, Linguistic, and Computational Issues, Lecture Notes in AI #1867* (pp. 55-81). Berlin: Springer-Verlag.
- Spirtes, P., Glymour, C., & Scheines, R. (2001). *Causation, prediction, and search*: MIT Press.
- Sveiby, K. E. (1997). *The new organisational wealth: managing and measuring knowledge-based assets*. San Francisco: Berrett-Koehler Publishers Inc.
- Tuomi, I. (1999). *Data is more than knowledge: Implications for the reversed knowledge hierarchy for knowledge management and organisational memory*. Paper presented at the 32nd Hawaii International Conference on System Sciences, Hawaii.