# Principled Development of Knowledge Based Systems - The Importance of Domain History

*Trevor Bench-Capon*[1], *Sheila Lynch*[2] *and Ray Paton*[1]

1. Department of Computer Science
The University of Liverpool
Liverpool L69 3BX UK
Tel: +151-794-3697 Fax: +151-794-3715 e-mail: tbc@csc.liv.ac.uk

2. Liverpool Hope University College
Liverpool UK

*Abstract*

*MEKAS is a methodology for knowledge analysis originally developed at the University of Liverpool. In this paper we examine the role of one important aspect of the MEKAS not found in other methods, namely the emphasis placed on the investigation of the history of the domain under investigation. Practical case studies have shown that history can play a key role in identifying the nature of the domain, and the types of theories it uses, hence determining the nature of the ontological model that the domain requires. Additionally the history gives valuable insights into the culture of the domain which facilitates elicitation, and may suggest the appropriate nature of support that can be provided by computer systems.*

## 1. Introduction

In the decade or so since Knowledge Based Systems (KBS) have become an established part of the computer science scene, there has been a movement from the initial excitement that such systems are possible to a recognition that, as with other systems, it is not enough to build them; success demands also that they be built well. This recognition has led to focus on development methodologies for KBS, perhaps the most notable of which, in Europe at least, is KADS (Wielinga et al 1992)). A major conclusion to emerge from this research is the importance of constructing an *ontology* - best thought of as an explicit conceptualisation of the domain. Thus in original versions of KADS we see the initial layer of their model as the *domain layer*, intended to describe all static and axiomatic domain knowledge in a use independent ontology. In later versions of KADS the ontology forms a crucial part of the application knowledge to be found in the domain model (Schreiber et al 1994). Interest in ontologies has been given a further boost by the work of Gruber (1993) and his development of Ontolingua, a language for writing ontologies.

What such methods do not provide, however, is much guidance for deciding on the nature of the ontology required and in constructing the ontology (Visser 1995). In KADS there is an attempt to offer some guidance in that there is a library of tasks, and it is the choice of tasks that mainly determines the form of the ontology. But even so this guidance remains somewhat abstract, because of the limited granularity of the library components (Valente and Lockenhoff, 1994). In our view there is a need for an analytic stage prior to embarking on the KADS method, so as to provide a broader and richer characterisation of the domain, which will provide a context within which questions such as what the form of the ontology will be can be answered. In order to provide such a characterisation we have developed the MEKAS methodology (Nwana et al 1994).

## 2. The MEKAS Methodology

The MEKAS methodology has been described at length elsewhere (e.g. Bench-Capon et al 1993 Nwana et al 1994). Here we will only draw attention to the salient characteristics relevant to the issues we wish to discuss in this paper.

MEKAS identifies seven characteristics common to all domains, and a MEKAS analysis will explore the domain from these seven standpoints. The result will be a characterisation of the domain by reference to these seven characteristics. The characteristics are:

- History - the way the domain has evolved over time;
- Theory - the conceptual framework of the



0-8186-7662-0/96 $05.00 © 1996 IEEE

domain;

- Meta-Theory - fundamental concepts of the domain such as time and causality;
- Global Metaphors - the language used to talk about the domain;
- Relations to Other Domains - similar sources of knowledge which may provide a source of methods, analogies, etc;
- Structure - the parts, relations, and organisation of the domain;
- Purpose - the problems which the domain addresses.

Most of these characteristics have analogues in other methodologies. Two, however, are individual to MEKAS: the emphasis placed on metaphor, and the attention paid to the history of the domain. We have discussed metaphor elsewhere, in Patoen et al (1991); in this paper we want to concentrate on the leverage provided by an understanding of the intellectual history of the domain.

Support for constructing ontologies is provided by the philosophical underpinnings of the method which have their root in the realist philosophy of science, particularly the work of Harre (1972). In the realist view, what exists within the domain is determined by the theories of the domain, and the entities which are used and postulated by the theory. Harre distinguishes three different types of theory each of which gives rise to different ontological expectations, and determine the sorts of entity which should appear in the ontology. The three types of theory are:

**Type 1**: These theories are based on heuristics and observables alone; they lack explanatory power. An example would be Newtonian kinematics in which the objects are observable. The heuristics rely on mathematical relations but have no explanatory power.

**Type 2**: These theories allow the representation of both observable and certain non-observable entities. The non-observables of these theories are, however, *in principle* observable. An example of this type of theory is Harvey's account of the circulatory system. When it was postulated the existence of capillaries was unknown, by their existence was required to support the theory. Improved microscopes, however, permitted the observation of capillaries, helping to confirm the theory.

**Type 3**: These theories extend type 2 theories by allowing for the existence of cognitive objects, which may not even be observable in principle; such as the mathematical properties of harmony and symmetry.

As well as the ontological implications, each type of theory has a particular representation associated with it. Type 1 theories consist predominantly of rules and statements, lending themselves particularly to classic "shallow" rule based expert systems. Type 2 theories extend rules and statements with more iconic or picturable representations, suggesting second generation "deep model" style KBS. Type 3 theories cannot be pictures and rely on more abstract formal languages, suggesting mathematical objects as a natural executable representation.

One central aim of the MEKAS characterisation of a domain is to establish the type of theories which are used in the domain. This classification will have important implications for the contents of the domain ontology, and for the style and feasibility of any KBS which may be constructed to solve particular tasks within the domain.

## 3. Case Studies

As part of the research described in Lynch (1996), three case studies using the MEKAS method were carried out. One of these studies explored the domain of open-textured concepts in law, focusing on the concept of "good cause for late claim" is UK Social Security Law, the other two were in the domain of colloid chemistry. The full analysis resulting from these case studies is in Lynch (1996). In this paper we focus on the contribution to the analysis from the examination of the history of these domains.

### 3.1. "Good Cause for Late Claim"

The domain of good cause was selected because it seemed representative of a number of "hard" problems in legal KBS, namely those which require the resolution of an open-textured concept. Open textured concepts are used quite often in legislation because legislators recognise that it is impossible to envisage, and explicitly legislate for, all conceivable situations in advance. Therefore they use an open-textured term indicating their general intention, but which is left to adjudicators to resolve in the light of the particular cases that present themselves. The most popular approach to such terms in AI and law has been to focus on the doctrine of *stare decicis*, by which past decisions are held to act as precedents and constrain decisions in subsequent similar cases, and to use some kind of case based reasoning to identify and enforce precedents. Issues with such systems include how cases are matched, and what happens when precedents appear to conflict. Perhaps the best known system of this type is Rissland and Ashley's HYPO (Ashley 1990)

The history of the domain showed up two important features of the domain. We will look at each of these in turn.

First it provided an important insight into the role of the adjudicator. The UK Social Security system has evolved from earlier provision by organisations such as Friendly Societies and Trade Unions. Under these schemes the adjudication of claims was carried out by an official termed "the umpire". While a game playing metaphor is not uncommon in law in general it is often couched in terms of an adversarial competition between a plaintiff and a defendant. But in this case when the claimant "wins" there is no loser: the adjudicator is an umpire, there to ensure that the interests of the claimants are properly balanced against the interests of the contributors to the fund.

This has a number of implications, which could lead to system problems if we were to come at the task from a standpoint dictated by standard adversarial law. One set of implications concern the nature of the support the system should provide: the system need not construct counter-arguments: rather its role is to assess the sufficiency of the arguments presented for the claimant. The other implications concern elicitation: it is a tenet of the MEKAS approach that knowledge engineers must harmonise their perspectives with those of the expert. So if the knowledge engineer sees the adjudicator as "opponent" when the adjudicator's self-image is that of "umpire", there will be an important attitudinal and cultural mismatch which will slow the elicitation process and distort the knowledge elicited. Of course, this kind of perspective on a role might emerge from other factors, but in this case the domain history gives a clear and immediate identification of this aspect.

The second set of implications of the domain history concern concept drift. It has long been recognised that *stare decisis* is not the whole story. Indeed, it cannot be, since conflicting decisions can be found. The doctrine can be patched up by augmenting it with principles such as: prefer the later decision; prefer the more specific decision; prefer the decision from the higher status court. This, however, provides only a mechanism for reconciling conflicting precedents: it provides no clue as to how and why decisions change. Recent work by Rissand and Friedman (1995) identifies concept drift as an important issue, and work by Berman and Hafner (1995) has also attempted to detect when concept drift is likely to occur. This, together with another important Berman and Hafner paper on the role of teleology in legal decision making (Berman and Hafner 1993), is clear evidence that the situation is much more complicated that the original ideas underlying the case based systems would support. What is interesting from the point of view of a knowledge analysis methodology is that is is these leading workers in the field who are beginning to recognise the problems, and only after more than a decade of experience in working with them. The MEKAS analysis also was able to identify these problems (although not to propose any solutions) in a brief period of analysis, from its exploration of the domain history.

Examination of the history of the "good cause" domain, shows that interpretations change, and do so in response to factors other than simply a new insight into the law resulting in a "landmark" decision. Indeed, often the decisions that set new precedents are not so much the source of change as the regularisation of a change in interpretation which has already taken place. The flexibility offered to the adjudicator is not only there to deal with the different situations that might arise across cases, but also to deal with the different environments in which the cases are decided. Particularly in the case of welfare benefits, where the law is being used as an instrument of social policy, exogenous factors are important. Among these factors are:

**Departmental Policy**: Adjudication officers are not legally qualified. Although they are legally empowered to take decisions, and take the ultimate responsibility for their decisions, in practice they are quite reliant on the guidance issued to them by the Office of the Chief Adjudicating Officer. While this guidance is the mechanism by which key judgements are made known to the adjudicators, it may also used to convey "policy", whereby nuances of interpretation may be altered to respond to current circumstances. Since the scheme has been in operation, such policy shifts have been used to liberalise and tighten interpretations without any explicit leading case or legislative change.

**Notions of Fairness**: Adjudication officers are there to balance the interests of contributors to the fund against those of claimants. Where the balance should be fairly struck is not an absolute, but changes in response to the general climate of opinion. During the 60s and 70s the prevailing climate was *benefits are rights*, and the interests of the claimant were given more weight than in more recent years where the rights of taxpayers to pay as little as possible are more emphasised. The opinions of adjudication officers not only in practice reflect such shifts in the prevailing moral climate, they are expected to do so.

**Other environmental factors**: As well as changes in the climate of opinion, there are also changes in the objective nature of society. In considering whether

good cause exists, factors such as the move to a five rather than a six day week, the widespread availability of telephones and cars, and demographic changes which have led to more people living alone and without nearby relatives, have all had their impact. Even changes in the postal system are important: whereas same day delivery was at one time quite common, today anything up to five working days needs to be allowed for the delivery of a letter posted second class.

All of these factors are quite properly taken into account when making decisions. Part of the philosophy of using lay adjudicators is that they should be responsive to these sort of factors, so that their decisions will reflect changing circumstances. These changes may, however, never show up in leading, reported cases, although if a case where such changes are presented arises, the changed situation may be made explicit in the decision. The non-adversarial nature is important here also: if the changes result in a liberalisation of the interpretation, the claimant will not appeal against award, and so the case will never appear at a level at which decisions are reported. Liberalisations will thus only be contested if they are made at the level of an appeal tribunal, where the adjudicating officer has the right to appeal against his decision being overturned. In practice, however, brakes on liberalisation, where it runs ahead of the wishes of the Department, can still be applied through policy statements.

What the consideration of the domain history showed most clearly in this case was the unsuitability of the domain as an area for the application of the existing case based techniques used in other areas of legal open texture, however successful these may have been in other domains. Any system which concentrates on cases will be inadequate, since cases do not adequately reflect shifts in the notion of what is fair, nor changes in the nature of society and expectations of its members, that are are in fact central to the decision making process. In terms of our theory types, there is little prospect of capturing a Type-2 theory, since while cases could be modelled the other factors are too complicated to capture. It would remain possible to build a shallow Type-1 model based on the guidance issued by the Chief Adjudicating Officer: this would have certain advantages in making the guidance readily available, but care would need to be taken to ensure that such a system did not trespass unduly on the discretion of the individual adjudicating officers. In matters other than in policy, such guidance will lag behind the changes that take place in the environment in which decisions are made, and over reliance on the guidance would result in an undesirable ossification of the system, removing the flexibility which is historically central.

The overall conclusion of the study was that the domain was not susceptible to treatment using current KBS technology. Whether such implications apply to other areas of legal open texture is not something we can comment on, but some of the recent work cited above suggests that there may be similar problems with them also. This could, however, only be decided by means of a similar exploration of those domains to uncover their individual histories.

### 3.2. Colloid Science

Two MEKAS analyses were carried out in the field of colloid science within an industrial research organisation. One study was with an expert with a narrowly focused domain, the production of additive packages for diesel which will improve general performance, and the other with a more broadly orientated expert, in the analytical division, responsible for providing colloid expertise to a range of different product development areas. These domains provide a sharp contrast to the previous one, being in a scientific area where there is a strong presumption that KBS techniques may be applicable. Again we will concentrate in this paper on the lessons learnt from the history of the domains.

#### 3.2.1. Additive Packages for Diesel

The Petro-Chemical industry has been interested in the possibility of adding substances to fuels to improve performance since the 1920s. By the 1940s and 1950s the use of additives had become well established. Notice that the motivation for the investigation was essentially practical: to get better performance and so produce a better, more saleable product. As one might expect, this led to a strong experimental focus in the domain. Additives to try might be suggested by some broad knowledge of their chemical structure, but confirmation was sought, and amounts gauged, by performing empirical tests and analysing the results to tune the additive package, rather than from theoretical analysis. Over the years a body of heuristic expertise has been built up, from which new packages can be proposed, evaluated and refined. This style of working continues today. From this we can see that we are dealing with a Type-1 domain, and that there is no need, or point, in probing further to uncover any deeper, explanatory theory. The history of the domain, from the standpoint of this expert, tells us that the expertise resides in knowledge of and exploitation of empirically derived heuristics, and that refinement of packages largely involves the application of the heuristics to the results of testing

particular combinations of packages.

The main lesson from domain history here is that the analysis need not be prolonged, seeking explanatory mechanisms that do not exist. Sometimes the expert may have a tacit theory which can be teased out, but here has formed no theory.

### 3.2.2. Advice on Colloid Chemistry

The second case study in this domain was carried out with an expert with a rather different perspective. Instead of being engaged in putting together particular additive packages for particular purposes, his role is to provide more general advice, across a wider range of applications. One might anticipate that this role would give a more general, reflective and perhaps less empirical perspective.

Discussion of the domain history with this expert revealed that there was indeed theoretical development in the domain, running in parallel with the commercial orientated empirical production of additive packages. The attempt to construct a theory of the domain can be seen as originating in a series of meetings (the Faraday Discussions) which began in 1954. This provided the basis for a community in which a theoretical framework could be articulated.

There has been, however, a relatively short time for a scientific theory to be developed, and the empirical results from those working in the field do not always appear to conform to the theory. This may be partly explained by the conditions in which the empirical work is carried out (and by the commercial value and hence confidentiality of much of the work), but it indicates that the theory is as yet imperfectly understood. While the development of the domain has seen many enhancements and elaborations of the theory, there is still sufficient deviation from it to mean that much caution must be exercised when applying the theory.

Also from the history of the domain we can see why different areas of the domain have been emphasised at different times. Thus in the oil crisis of the 1970s, the major drive was for increased efficiency, whereas in recent years environmental considerations have been given greater weight. The commercially driven nature of the domain, means that the experimentation is not always focused on areas most interesting from a pure science point of view, and the time constraints imposed by the commercial setting mean that the aim is always to realise a cost-effective product rather than exploring scientific niceties.

A third set of influences detectable from the domain history concern the nature of the equipment that can be used in experiments. In particular the growth of computers have enabled more sophisticated statistical analyses and the potential for simulation of aspects of the domain. Both of these aspects have accelerated the understanding of the underlying chemistry, but have also made carrying out empirical tests, allowing them to be less carefully planned.

What emerges from this is that we have a domain which can be seen primarily a scientific, which would suggest a Type-2 theory, but which is situated in a commercially driven engineering environment, which relies on a Type-1 theory. This has important implications both for knowledge elicitation and knowledge modelling. While classifying the theory type is important for the knowledge engineer, it cannot be expected to be seen as a concern of the expert. Therefore the expert will often present his knowledge in a way which switches between and intertwines material from the two theory types. In order to understand such material properly the knowledge engineer needs to be aware of the co-existence of the two theory types, and so to separate the material. The picture which needs to emerge is one of an underlying Type-2 theory, which is, however, incomplete; with the gaps in the theory covered by more heuristic Type-1 material. Such an analysis will then aid the construction of a system in the domain by identifying those areas where a "deeper" explanation can be incorporated and those areas where a "shallow" heuristic solution must suffice.

## 4. Conclusions

In this paper we have shown how the understanding of the history of a domain can play an important role in knowledge acquisition and modelling. In order to assimilate knowledge effectively the knowledge engineer must come to see the domain in the say way as the expert. This process is facilitated by classifying the type of theory underlying the domain and this in turn is facilitated by understanding the historical development of the domain. History also gives important insights into the purpose of the domain, and into the correct understanding of metaphors used by domain experts. In each of the three case studies, features of the elicitation process which might otherwise have been puzzling were explicated by a knowledge of the domain history. Finally the history reveals the culture of the expert and the context in which the tasks are performed. This in turn explains what information can be expected from the expert, and many of the presuppositions implicit in the expert's conception of the domain.

Of course, a knowledge engineer using any methodology will come to understand something of domain history through interaction with domain experts.

MEKAS, however, by making it an explicit feature of a domain characterisation, forces the knowledge engineer to think in particular about the points, and to record them so that the insights are readily available to those who subsequently use the domain analysis. Presentations of analyses from other methodologies which omit this feature often leave many of their domain choices unexplained.

## References


Ashley, K.D., (1990) *Modelling Legal Argument* MIT Press, Cambridge Mass.

Bench-Capon, T.J.M., Coenen, F.P., Nwana, H.S., Paton, R.C., and Shave, M.J.R., (1993) Two aspects of the Validation and Verification of Knowledge Based Systems, *IEEE Expert, Vol 8 No 3, pp 76-81.*

Berman, Donald H., and Hafner, Carole D., (1993). Representing Teleological Structure in Case Based Legal Reasoning. In Proceedings of the Fourth International Conference on AI and Law, ACM Press: New York, pp 50-5

Berman, Donald H., and Hafner, Carole D., (1995). *Understanding precedents in a Temporal Context of Evolving Legal Doctrine.* In Proceedings of the Fifth International Conference on AI and Law, ACM Press: New York, pp 42-51.

Gruber, T., (1993). A Translation Approach to Portable Ontology Specifications *Knowledge Acquisition* Vol 5, 199-220.

Harre, R., (1972) *The Philosophies of Science,* Oxford University Press, Oxford.

Lynch, S.C., (1996). *Some Investigations into the Extension and Refinement of the MEKAS Methodology,* Ph.D. Thesis, University of Liverpool, 1996.

Nwana, H,S., Bench-Capon, T.J.M., Paton, R.C., and Shave, M.J.R., Domain-Driven Knowledge Modelling for Knowledge Acquisition, *Knowledge Acquisition,* Vol 6, pp 243-70.

Paton, R.C., Nwana, H.S, Shave, M.J.R. & Bench-Capon, T.J.M. *From Real World Problems to Domain Characterisations,* Proceedings of 5th European Knowledge Acquisition Workshop, Crieff, May 1991.

Rissland, E.L., and Friedman, M.T., (1995) *Detecting Change in Legal Concepts* In Proceedings of the Fifth International Conference on AI and Law, ACM Press, New York.

Schreiber, A.T., Wielinga, B.J., Akkermans, J.M., and van der Velde, W., (1994) *CML: The CommonKADS Conceptual Modelling Language* in Proceedings of EKAW94, Hoegaarden, Belgium.

Valente A., and Lockenhoff, C., (1994) *Assessment* in Breuker, J., and van der Velde, W., (eds) *Common-KADS Library for Expertise Modelling,* IOS Press, Amsterdam, pp 155-74.

Visser, P.R.S., (1995). *Knowledge Specification for Multiple Legal Tasks,* Computer/Law Series, Kluwer International, The Hague.

Wielinga, B.J., Schreiber, A. T., and Breuker, J.A., (1992). KADS: A Modelling Approach to Knowledge Engineering, *Knowledge Acquisition,* Vol 4, pp 5-53.