Issues for an Ontology for Knowledge Valuation

Kieron O'Hara & Nigel Shadbolt

Intelligence, Agents and Multimedia Group Dept. of Electronics and Computer Science University of Southampton Highfield Southampton SO17 1BJ {kmo,nrs}@ecs.soton.ac.uk

Keywords: ontologies

Abstract: This paper examines the application of inferencesupporting ontologies to the issue of knowledge valuation. Knowledge, like customer goodwill or a brand, is an intangible asset for a firm; that is, a non-physical, nonfinancial claim to future wealth. Intangibles are notoriously difficult to value, but there are a number of reasons why we should try. A series of factors affecting the value of knowledge are discussed, and an ontology that could be used to express some of these is sketched. Such an ontology could be imported by e-business models, to help understand how an organisation's knowledge adds value to its operations, and therefore to enable informed management of its knowledge assets.

1 Introduction

Of all areas of business and e-commerce, the greatest effect that will be provided by developments of the semantic web idea will be on the commoditization of knowledge. The transfer of any good can be aided by developments to improve the web, but the only goods that are actually transferred by the technologies underlying the semantic web are data, information and knowledge.

The vision underlying the semantic web is a repository for *knowledge* as opposed to *information*: in other words, the semantic web will contain usable information in the form of a dynamic resource, not in static form, but represented appropriately for certain uses. Knowledge is held by a number of different types of entity (people, organisations, databases, etc), and can take a number of forms, perhaps being fuzzily held in an expert's head, or buried under an impenetrable mound of data, or explicitly stated in manuals. Unsurprisingly, a series of technologies have sprouted to help deal with the heterogeneous series of problems that such a situation creates, such as knowledge discovery techniques, information extraction technologies and knowledge modelling techniques such as problemsolving methods and ontologies.

Against this background, the Advanced Knowledge Technologies project (AKT) has been funded to extend and integrate such technologies to cover the use of knowledge over its entire lifecycle, from acquisition to deletion. As part of that effort, in this paper we will examine a number of the issues relating to the *valuation* of knowledge, and the technologies that might be used to aid this. In particular, we will set out the issues relating to ways to conceptualise the knowledge held by an organisation in order to place some kind of value upon it, both as a preliminary to managing that knowledge effectively, and as a way to support inferences about that knowledge: we will examine the idea of an *ontology* for knowledge valuation.

In Section 2, we will discuss the background to knowledge valuation, showing why it is problematic. We will then move on in Section 3 to discuss how ontologies might be used in business contexts, to provide rationales for business decisions and business structures, before sketching the form an ontology for knowledge valuation might have (Section 4), and surveying related work (Section 5) and final questions (Section 6).

2 Valuing Knowledge

We begin with the practice of knowledge valuation, which is notoriously something of a black art. We shall explain why scepticism is justified at least to some extent, and give examples of how the financial community has attempted to deal with the problems and issues (cf. [Lev, 2000]).

Knowledge in an organisation is an example of an *asset*. Assets are claims to future benefits, such as cash, which can be used to buy benefits, or a building, which can be exploited for commercial rent. Knowledge is used to extract profit benefits (e.g. from raw materials by good design). Assets can be either *tangible* or *intangible*. A tangible asset has a physical or financial embodiment, so buildings, shares or plant are tangible assets. Intangible assets do not have such embodiment, and include, as well as knowledge, such things as trademarks or organisational structures. Some intangible assets, such as patents, trademarks or copyrights, have the claim to future benefits legally secured for a period; these assets are called *intellectual property*.

Financially, the tangible assets of a firm are reported to investors, and monies spent on them can be offset against expenses. In this way, investors can draw a distinction between standard expenses (e.g. purchases of raw materials) and investments (e.g. purchases of new premises) which may be converted into future benefits to enhance future profitability. However, this is not true for intangible assets, expenditure on all of which is expensed. Hence there is no distinction, for an investor, between a firm's *investment* on knowledge acquisition and *expenditure* on paper, or between wages paid to a research scientist and the managing director's idiot brother-in-law, even though in each case one would expect the former to contribute to future profitability and not the latter [Lev, 2000].

Because there is no requirement on firms to report investment on knowledge assets, there is little methodology for knowledge valuation, and so valuing a piece of knowledge needs much guesswork; however, note that this is not the reason for the lack of reporting. Many tangible assets' value parameters are similarly based on guesswork, such as the depreciation rate of a fixed asset.

2.1 What is Value?

The first question that we can raise is that of what value is. Economics is ambiguous here, and there is a range of potential conceptualisations for the concept of 'value'. The *objective* value of a good or service is the price that it would fetch in an open, competitive bidding market. However, we must distinguish between the *value in use* and the *value in alternative use*. Suppose a firm owns a dockside warehouse. Its value as a warehouse (value in use) might be relatively small, e.g. if the firm is running down stocks. However, its value as a potential conversion into yuppie housing (alternative use) may be very different.

This leads us to an important distinction, between the value of a good as set by a market, and the value to a person, firm or organisation.

We must also distinguish between monetary and nonmonetary value; a good may not contribute to profitability, yet still make a substantial contribution to fulfilling promises in a corporate mission statement, for instance on environmental sensitivity.

2.2 Why Value Knowledge at All?

There are several reasons why some kind of reporting of knowledge assets' value is desirable.

1. There are two ways of valuing a firm. The first is to examine the value of the shares in that firm (the *market value*). The second is to add up the value of all the

tangible assets the firm owns (the *book value*). For most large firms, the market value vastly exceeds the book value. Hence a huge part of the value of a firm is typically down to the intangible assets, and one would expect the knowledge within the firm to be a chief contributor [Hall, 2000; Lev, 2000].

- 2. For a stock exchange to function efficiently, there needs to be equality of knowledge between investors. At present, if a firm invests in knowledge assets to enhance profitability, those in the firm need not report this to outsiders, and therefore will be in a better position to make investment decisions [Kyle, 1989].
- 3. Good asset management depends on information. Managers ensure they have lots of information about their tangible assets to manage them effectively. However, managers cannot get the same quality of information about their intangible assets, and hence are unlikely to manage them as effectively [Shapiro and Varian, 1999; Lev, 2000].
- 4. By being automatically expensed, investment in knowledge assets distorts a firm's profit figures. By being subtracted from the current balance, they make current profits seem lower. Then when the investment in knowledge feeds through into profit, growth figures from this lower base are exaggerated [Bushee, 1998].
- 5. Ironically, some ways of acquiring knowledge assets do appear on balance sheets, leading to an undesirable asymmetry. If a firm acquires some knowledge assets through a research programme, that is an expense. On the other hand, if it merely buys another firm that possesses those knowledge assets, the purchase, being of a tangible asset, can appear as a legitimate investment.
- 6. Many creative forms of banking, such as microlending to impoverished communities in the developing world, are based on finding alternatives to valuing tangible assets. Risk management in such areas could be improved by finding methods for valuing knowledge that entrepreneurs may possess.

2.3 Why Don't We?

So, given the reasons to value knowledge, we can infer there must be good reasons why people don't, and why information about knowledge valuation doesn't reach investors through standard financial reporting routes.

And there are.

Partial Excludability

The first reason is the partial excludability of knowledge. If someone sells a good, they contract to confer property rights on the buyer that exclude others from the benefits of the good. Knowledge is not excludable in this way – if knowledge is sold, then others can still benefit (in particular, the vendor, since short of wiping his memory he will retain the knowledge) [Lev, 2000]. Even patents can be reverse engineered; a recent study showed that patent law, despite being tightened, has not succeeded in excluding others from the benefits [Cohen et al, 2000].

Problematic Cost Structure

Knowledge has a difficult cost structure. There are relatively large initial costs (*sunk costs*) as an investment is made into knowledge acquisition, while once acquired the marginal costs of reproducing knowledge are close to zero [Shapiro and Varian, 1999]. In the long run, in a free market prices tend towards marginal costs (because competitors can always undercut a supplier above marginal cost). Hence the market value of much if not all knowledge will always fall. This is why there is so much free content on the web.

Inherent Risk

Knowledge acquisition, like any investment, is a risky process. But because of difficulties in valuing, ways of hedging this risk (e.g. by insuring against failure) are not easily available [Christensen, 1997].

Non-Marketability

For all these reasons, knowledge is very difficult to market. A price cannot be set for it objectively, because there are no large enough competitive bidding markets. A firm can buy knowledge from someone, such as a consultant, but this price is set between themselves, rather than determined independently by a series of transactions. Even intellectual property is non-marketable, although the development of Internet-based IP markets may change this [Scheutze, 1993; Griliches, 1995].

Accountancy Considerations

If there is no objective value, then the value of a knowledge asset cannot appear on a balance sheet. If a subjective value, the value to a firm of the knowledge it possesses, were to be used instead, then the auditors will be obliged to the firm itself for the knowledge valuation – which of course it might then with total freedom estimate vary highly [Lev, 2000].

Investors Respond to Information

Even though firms are not obliged to place details of their knowledge assets in their financial reports, they can of course still report investments to investors, either as footnotes to their statutory reports, or as independent press releases. There is evidence that investors do respond (i.e. by pushing up the share price) in at least some sectors of the economy, to firms' investing in knowledge assets (e.g. a programme of R&D) even though such investment will in the short run push down the profitability of the firm [Aboody and Lev, 2001]. Hence it might be argued that, although it may be desirable to produce accounts of knowledge assets, in fact the need is not so pressing because this information is already produced informally and used efficiently by markets.

2.4 What Approaches have Evolved?

In the current climate, where the reasons for knowledge asset management and reporting seem to have become more urgent (though see [Hall, 1999]), and yet accounting practices do not take investment in such assets into account, there have evolved three different approaches to the issue.

The World According to GAAP

The sceptical approach is to admit that the problems of valuing knowledge assets are too great for a response that will achieve the consensus required for investors to retain confidence in financial reporting procedures. Generally Accepted Accounting Principles (GAAP) have been around for centuries, are well-understood, have presided over historically enormous economic growth, and are not obviously responsible for serious market failure. Therefore unreliable disclosures will cause more harm than good. This position was espoused, for example, by Mike Brown, chief financial officer at Microsoft in a recent interview [beemanagement.com, 1999].

Monetary Values

The second approach is to be as clear as possible about the definitions, and then apply a number of competing valuation approaches. Dow Chemicals values its intellectual property (patents et al) this way, producing a range within which the value of the property is likely to lie [Near, 2001]. Dow uses the strategy to help managers with knowledge asset management, but in the future such information could be published in a supplementary annual report to allow investors access to this important strategic information [beemanagement.com, 1999].

Matrix of Qualitative Measures

A third approach, midway between the first two, is to set out for managers and investors a series of parameter values to aid the appraisal of a firm's knowledge asset management. The highest profile firm to adopt this strategy is Skandia [Edvinsson and Malone, 1997; Malone, 1997]. The idea here is that an intellectual capital reporting model should contain a sizeable body of indicators released alongside reports of tangible asset values. As firms that release information attract a share price premium over firms that do not release the same information about themselves [Litan and Wallison, 2000], in the end investors will force firms to release the information that is considered important. A recent authoritative survey by PricewaterhouseCoopers Netherlands for the Dutch Finance Ministry concluded that this midway approach was superior from a number of theoretical points of view [Backhuijs et al, 1999].

3 Knowledge Valuation and the Semantic Web

We have seen why the issue of knowledge valuation is difficult, and yet why it would be desirable to address it. In this section, we will examine the possibilities for the use of technology to accomplish this. In Section 3.1, we will look at ways in which the semantic web can be used to address the problem of knowledge valuation. In Section 3.2, we will focus our attention on ontologies, a particular type of knowledge-based technology, and in Section 3.3 we will give a couple of scenarios where ontologies could be of great use. Finally, Section 3.4 will consider the requirements for knowledge valuation ontologies, and list a set of desirable characteristics.

3.1 Technologies for Knowledge Valuation and Management/Investment Decisions

The semantic web has the potential to be a very potent tool for knowledge valuation. As we have seen in Section 1, the semantic web is a knowledge repository, and will be a key technology in the commoditization of knowledge. The effects of this will be widespread and, of course, largely unpredictable. However, it is possible to imagine areas where issues to do with the valuation of knowledge will be central. In this section, we will discuss three possible general areas. First, there is the possibility of using the web to create markets for knowledge. Second, the issue of cost/benefit analysis of knowledge management (KM) is introduced. Third, we will consider a more general application of KM.

Web-Based Markets

The functionality of the semantic web may provide the competitive bidding markets that will enable prices to be determined independently. There are sites which could harness the web's power in bringing together a critical mass of buyers and sellers to do this (e.g. the Patent Licensing Exchange, <u>http://www.pl</u>-x.com).

In addition, proxy agents such as Shopbots [Edwards, 2000] will also enable users to find particularly low prices; this does not in and of itself provide a competitive market, but will help standardise prices by making it easier to find the market entrants who are prepared to undercut others. This, as noted in Section 1, is an important aid to the implementation of something closer to perfect competition, while also helping prices for knowledge items to tend towards their marginal costs.

However, the gains from such approaches are in the medium to long term, for such markets need to be recognised by investors' bodies and financial regulatory organisations; they need to gain the support of knowledge asset managers, in order to play a value-determining role.

Cost/Benefit Analysis

Another way in which the KM possibilities of the semantic web can support knowledge asset management is the potential for accurate cost/benefit analyses underlying knowledge investment decisions. For example, a key KM decision is whether to produce a codified version of knowledge tacitly held in the organisation (in the terminology & Takeuchi, of [Nonaka 1995], externalisation). This is a complex decision, with a number of costs and benefits to be weighed against each other [Cowan et al, 1999]. For instance, the costs of codification will increase if new representation languages will need to be drafted; they will decrease if some or all of the knowledge already exists in explicit form. Econometric studies have

shown that particular types of codes, when used to codify particular classes of knowledge, can lead to benefits in profitability for multinationals, where knowledge sharing across the organisation is a serious problem [Schulz and Jobe, 1998]. Major costs also include those of the acquisition and maintenance of the knowledge base. Benefits will accrue by improved knowledge sharing, and for the export of the knowledge under a consultancy scheme.

One interesting trade-off is related to the excludability of the knowledge. If the organisation gains from the exclusive use of the knowledge (i.e. it is a trade secret, whether or not fully or partially patent protected), then anything that makes the knowledge more likely to leak is a relative cost. Codified knowledge is more 'leaky' than tacit knowledge, and hence codification may be too risky. However, conversely, the labour market is often used as a proxy market for tacit knowledge, and those holding such knowledge can be 'poached' for a higher salary. In such a context, the benefits of codification are higher, because then if an expert leaves the organisation, the newly-codified knowledge does not leave with him.

Technology has no exclusive answers to these difficulties, of course. However, technologies for classifying knowledge in certain ways (as tacit/codified, as appropriately/inappropriately codified, as ideally excludable or otherwise, as contested/consensual), or as highlighting potential KBs as being easy/difficult to maintain or acquire (e.g. a fast-moving domain will place time constraints on acquisition, and a heavy maintenance overhead), will clearly be of value in this process. Such possibilities point towards the use of ontologies to support cost/benefit analysis.

Knowledge Technologies for Knowledge Management

And, of course, knowledge technologies can be applied in general to the problems of KM, of which valuation is but one. Knowledge modelling techniques can be used to understand the dynamics of knowledge within an organisation. This will certainly help with knowledge asset management. It could also be a useful basis for reporting of the state of the knowledge assets within a firm to investors, perhaps in a non-statutory report. In the next subsection, we will discuss the use of ontologies in this context.

3.2 Ontologies for E-Business: Enterprise Models

Ontologies have already been exploited in the e-business context [Fox and Gruninger, 1998]. An *enterprise model* represents the structure, activities, processes, information, resources, people, behaviour, goals and constraints of an organisation. The aim is to provide the apparatus for reasoning about an enterprise and its relations with other enterprises and markets.

In particular, enterprise models can provide essential shared conceptualisations of aspects of organisations that can be very useful in distributed organisations in which there are several points of view. In other words, enterprise models can provide ontological support for an organisation.

The benefits of ontologies are well-known [Fox and Gruninger, pp.110-111], at least if applied in carefully circumscribed areas [Shadbolt et al, in press]. As well as providing shared and agreed conceptualisations, and also a language and vocabulary sufficient to express what their users wish to, they can indicate essential concepts, classes or objects – in a generic sense – that individual users may not have realised are required. Furthermore, they can save time by providing essential generic structures that individual users need only fill in with the domain-specific classes, objects etc.

In this paper, we are particularly concerned with an ontology to aid knowledge valuation. While value issues have been discussed in some depth by a number of researchers (see Section 5 below), there has been little work, for the reasons canvassed in Section 2.3, on valuing knowledge itself. Hence we shall move on to discuss the requirements for a knowledge valuation ontology. Section 3.3 will suggest some scenarios where it could be of use, while Section 3.4 suggests a few desirable characteristics a knowledge valuation ontology might have in such contexts.

3.3 Scenarios for the Use of Knowledge Valuation Ontologies

In this section, we will suggest a few possible uses of valuation ontologies, for *cost/benefit analysis*, for isolation of *communities of practice*, and for identification of *knowledge gaps*.

Cost/Benefit Analysis

As was mentioned in Section 3.1, cost/benefit analysis for KM decisions such as whether to codify or recodify a body of knowledge could be highly valuable. The scenario here would be that some sort of organisational model, perhaps along the lines of a CommonKADS organisational model [Schreiber et al, 2000], or an enterprise ontology (Section 3.2), would isolate particular uses and bodies of knowledge, and an ontology for classifying such bodies of knowledge in their organisational context along value-relevant dimensions such as those suggested in Section 3.1 would be applied, to provide an understanding of the value of the knowledge within the business plan/processes of the host organisation (i.e. an idea of the contribution it makes to the bottom line), and provide useful pointers towards a cost/benefit analysis of a proposed codification.

This would largely be a manual use of an ontology which would provide qualitative support for an essentially quantitative decision. The other scenarios in this section, by contrast, could be machine-driven.

Communities of Practice

Another use for an ontology is one that highlights particular *communities of practice* (this example is based on unpublished work performed by Harith Alani at the University of Southampton). Such communities consist of

groups of like-minded workers, possibly across companies and even technical sectors, who share a number of assumptions (tacit knowledge) about their work, maybe focusing on particular approaches or sub-disciplines. Such communities are of key importance to understanding the knowledge capabilities of an organisation [Smith and Farquhar, 2000; Cohendet et al, 2000; Collins, 1974]. This is a valuation issue in that a supportive community of knowledge users will make such knowledge more valuable for an organisation.

Traversing relationships in a valuation ontology, weighted to reflect their importance (for valuation), can show links between people and their co-authors, co-collaborators, people who publish in similar fields/journals, and can isolate groups of workers who have formed a *de facto* community, and the publications central to (understanding/membership of) that community.

There are, of course, issues here about how to draw boundaries around such communities, which relationships to track and how to weight them, and how to calculate the strength of a connection. Nevertheless ontologies, by providing the conceptual apparatus to express the relationships, hierarchies and axioms that will be of importance to defining a community, can be of help in understanding such communities, which often are difficult for management to track because of their informal nature.

Knowledge Gaps

A third scenario for a knowledge valuation ontology is for use with knowledge auditing processes, producing knowledge maps and identifying knowledge gaps [Speel et al, 1999; Smith and Farquhar, 2000]. In such a scenario, the current state of knowledge within an organisation is mapped, e.g. by creating a matrix of domain problems crossed with known solutions/best practice. Hence, such a matrix can be used to identify knowledge gaps, such as problems for which the organisation possesses either no solution, or only unreliable or expensive ones. Furthermore, the location of such problems within organisational models will provide a measure of the value of any solutions they may have, as the effect of the problem on the organisation's business processes will be shown in the model.

The compilation of such knowledge maps could be aided in a number of ways by a knowledge valuation ontology. For an organisation with a large library of knowledge sources, some automatic processing may help to link some problems with solutions. An ontology that related knowledge sources to particular people and processes within the organisation could help cut down search spaces dramatically.

3.4 Desiderata for a Knowledge Valuation Ontology

We shall suggest, then, that knowledge valuation ontologies can be of use in KM (and indeed possibly for wider purposes; the reader will have noted that much information relevant to valuation is relevant to other KM issues as well). The aim of a knowledge valuation ontology should be to allow users to express factors relevant to valuing a particular piece of knowledge. Much of this, of course, is an open question given that knowledge's non-marketability makes it very difficult to suggest an objective value. Nevertheless, our scenarios, together with approaches that suggest interesting valuation parameters (e.g. [Edvinsson and Malone, 1997]), give us a set of desirable characteristics of such an ontology. We obviously don't suggest that this list is exhaustive.

Axioms

One would expect an ontology to contain a hierarchy of classes and instances with attributes, and inheritance. However, given the vision (Section 1) of knowledge as a dynamic resource, and the necessity of combining the semantic web with inference capabilities, the question arises whether there should be inference *beyond* inheritance.

Fox and Gruninger [1999, p.111] suggest that most queries to an enterprise information system tend to be at a shallow level of processing, i.e. retrieval of information not directly stored in the database that does not require the significant amount of search characteristic of, say, expert systems and their deep-level processing. Hence much of the work of query formulation for such databases would be saved if ontologies stored the means for relatively straightforward deductions within themselves. In other words, axioms should be included in specifications of classes and objects. Some ontology-handling tools allow for the inclusion of axioms (e.g. Ontobroker [Fensel et al, 1998]), and it is clear that our second and third scenarios at least would benefit from being machine-processable.

Network Effects

Knowledge is a commodity that benefits from network effects [Shapiro and Varian, 1999]. These can be seen in any sector where the benefits of being included in a network increase with its size and interconnectedness. Examples of goods and services with network effects include programming languages or mobile phones. Network effects can be small; someone who owns a Mazda car, will be interested in how many other people drive one only to the extent to which the number is sufficient to sustain a chain of Mazda garages and suppliers of spares. Some networks effects are negative: part of the pleasure of owning an Armani suit is that you are the only kid on the block with one. The value of the suit would fall if the *hoi polloi* began to sport them.

Network effects are characteristic of advanced technology and information-based sectors of the economy. The more a piece of knowledge is used, the more valuable it is, the more likely it is to be embodied in and essential for production processes. Hence we will want a knowledge valuation ontology to enable to expression of the connectedness of a piece of knowledge or a knowledge source with a network of users or community of practice.

This would be of value, in different ways, in all three of our scenarios. In the first, the existence of a support network for a piece of knowledge affects the cost/benefit analysis of its acquisition. In the second, a community of practice is defined by the network connecting it. In the third, network connections affect which solutions are available for application to different problems.

Cumulativity

Knowledge is also cumulative. Understanding one piece of knowledge will often depend on an understanding of a wider set of knowledge sources that underpin it [Collins, 1974]. Hence the decision to acquire a particular piece of knowledge should be influenced by a strong sense of which other pieces of knowledge are related to it, both underpinning it and underpinned by it. Such a piece of knowledge may be valuable, not in itself, but in terms of which other pieces of knowledge it will provide access to. Equally, acquiring a piece of knowledge may be too costly if understanding that knowledge depends crucially on understanding a lot of more basic material. We can see that the ability to express such information would be of value in our first scenario.

Sources of Knowledge

Another important factor in knowledge valuation is the sources of knowledge available. Firms get knowledge from a number of sources, and how the firm uses knowledge will depend to a large extent on its source.

For instance, some knowledge is brought in from outside, particularly through a customer network (e.g. webbased B2B procurement can improve supply as customers themselves drive and monitor what may be a highly complex process of team-based detailed specification drafting under long-term contracts [Cairncross, 2000, p.31; Kinney, 2000]). External capital such as alliances and relationships with customers, partners, suppliers etc is important for invigorating a firm. However, it is important to have an idea of the value of these external knowledge assets as the amount of effort spent on nursing relationships may not be worth the result.

On the perspective of intellectual capital [Sveiby, 1997; Allee, 1999], there are three sources of knowledge assets.

- 1. External capital. This is capital inherited from alliances and relationships with customers, partners, suppliers, investors and the wider community.
- 2. Human capital. This is knowledge, skills, experiences and knowhow that people bring to an organisation.
- 3. Structural capital. These are systems and processes of the organisation itself that embody or encode knowledge, such as communications technologies, business models, databases, patents and other codified knowledge (what Nonaka and Takeuchi [1995] call explicit knowledge).

The intellectual capital approach to knowledge valuation sees the interplay of these three types of knowledge asset as being what generates business value, as knowledge flows around the organisation being converted to and from all these forms [Sveiby, 1997]. In different ways, all three of our scenarios would need this type of information.

Context of Knowledge

Another relevant notion is that of the *context* of knowledge within the organisation. This is set out most explicitly by knowledge modelling tools, techniques and methodologies such as CommonKADS [Schreiber et al, 2000]. CommonKADS produces a series of models of the organisational context, business processes and agents which would be concerned with the use of particular bodies of knowledge, as well as explicit models of the knowledge itself. Such models, it has been shown [O'Hara et al, 2000], can be used for quasi-valuation procedures such as certification for use in mission-critical contexts, and it is possible that the expressive power of such knowledge models could be used to produce either more formal, monetary values for bodies of knowledge or at least interesting parameters for setting in a qualitative valuation matrix. It is likely that all three of our use scenarios could exploit such information, perhaps the first in particular.

Six Challenges

A further aspect of knowledge that will affect its value is that of how manipulable it is in terms of the various challenges of knowledge management; in other words, how amenable it is to being used. There is a big difference between a piece of knowledge published on a corporate Intranet page, and that same knowledge known tacitly by one man in a corporation, who works in an obscure office.

Measuring such amenability for knowledge is not easy. However, we can use a conceptual framework developed as part of the AKT project (<u>http://www.aktors.org</u>), which postulates six fundamental challenges for KM.

- Knowledge acquisition.
- Knowledge modelling.
- Knowledge retrieval.
- Knowledge reuse.
- Knowledge publishing.
- Knowledge maintenance.

Knowledge is likely to be more valuable if it is open to approaches to these challenges (even acquisition, in the guise of making tacit knowledge explicit.

Furthermore, as these challenges form a rough sketch of a lifecycle for knowledge, being aware of what operations could be performed on knowledge is an interesting way of tracking knowledge value through its lifecycle. This value may change, and tracking such changes will itself be an importance part of knowledge maintenance. In particular, knowing when a piece of knowledge is either not worth storing, or positively harmful (negative value), is clearly important from the point of view of deleting it safely. Our first use scenario would find this sort of information useful.

4 Classes and Axioms for a Knowledge Valuation Ontology

In this section, we sketch an ontological structure for knowledge assets that may sit alongside existing ecommerce ontologies, such as the ENTERPRISE ontology [Uschold et al, 1998], or those underlying e-catalogs [Keller, 1995], or the project ontology currently under development in the AKT project. The intention is to allow the various aspects of knowledge valuation discourse discussed in Section 3.4 to be expressed within the ontology, to provide a way of understanding the properties of knowledge held by an organisation that pertain to its value.

The main class structure revolves around the distinction between knowledge sources, widely described as a fundamental distinction in that leveraging value from knowledge involves understanding the flow of knowledge between all these sources ([Sveiby, 1997], and implicit in [Nonaka and Takeuchi, 1995]). Hence the main classes would look something like: Knowledge asset

External knowledge asset

Human knowledge asset

Structural knowledge asset

The hierarchy underlying all these would of course have to be extended further. For instance, the hierarchy for external assets might look like the following.

External knowledge asset Collaboration Collaboration with individual Collaboration with organisation Informing paper Informing project Partnership User relationship

Under such a conceptualisation, the knowledge assets – the instances of the ontology – are the relationships between the organisation and its partners, collaborators, etc, together with such items as papers and reports. Knowledge assets would be assets *of* something. Hence we have an obvious axiom:

```
(all x y)(knowledge-asset(x y) => organisation(x)
    v individual(x))
```

For simplicity, we will assume that we are dealing with the knowledge assets of an organisation, but of course an audit could be carried out for an individual as well. Note also that the axioms themselves which we state in this paper are not strictly important – the main point is that certain predicates are introduced by them, and are likely to feature in a valuation ontology. The examples we give are intended to be indicative; the reader need not trouble to plough through the formalese to establish the exact content.

Many, if not all, knowledge asset instances will have base indicators of value, which would be expressed as attributes of the instances. For example:

- Publications might include an attribute quantifying its citations, or the citation rates of the authors, and the impact factor of the journal in which it was published.
- Readership figures may be given for books.
- Individuals may have as attributes the ranking of their institution, the number of publications or patents for which they are responsible, the amount of research funding for which they are responsible, or simple indicators of rank (e.g. Professor, CEO).

- Software can be evaluated in a number of creative ways: [Shadbolt et al 1999] discusses the empirical evaluation of knowledge acquisition techniques and tools, with basic but interesting ways of measuring their effectiveness.
- An organisation model could help pinpoint business processes' contributions to the bottom line.

These base values themselves are enhanced by being placed in a fruitful context. So, for example, reasoning about collaborative work is accomplished in an intuitive way. Consider, for example, the AKT project, which is a collaboration between five British universities. The hierarchy gives us an inheritance axiom.

```
(all x y)(collaboration-with-organisation(x y) =>
    knowledge-asset(x y))
```

Hence we would have:

(knowledge-asset (southampton, akt))

And so on for the other four partners. Then the extent of the AKT collaboration can be deduced from a simple query for all instances of

(collaboration-with-organisation (?x akt))

Information, the entities that inform the organisation in question, concerns papers, projects of which the organisation is not a member, web pages, adverts, etc. These will be entities to which the organisation has access (maybe via specific individuals within it). Partnerships and user relationships (being less formal than collaborations, and generally bilateral) may be expressed directly as a relation between a pair of organisations or individuals.

The result of characterising the space in these terms is that it can be read off immediately what are the knowledge assets of an organisation, and furthermore which type of intellectual capital they represent. There are of course implications for an organisation depending on which of the types of intellectual capital an asset is. Structural knowledge assets are pretty well embedded in the firm, whereas human knowledge assets can quit after giving notice, and external knowledge assets can disappear at any stage.

With public domain knowledge, access may depend on external interfaces (e.g. subscriptions to journals, web address books, knowledge that requires special expertise or software to interpret – all of which may be provided for the organisation by customers, suppliers, partners etc). The value to an organisation of such public domain knowledge may well depend on the existence within it of (a representative of) an appropriate community of practice, as discussed in the second of our three scenarios above.

Another relevant aspect of knowledge assets that we wish to capture with our knowledge valuation ontology is that of the networks with which they are associated. There are different ways of capturing this information. For example, an organisation doing its own knowledge audit can claim a knowledge asset of another organisation with which it has some relationship (on the ground that if the asset is known about it has almost certainly been shared).

(all x y a)(knowledge-asset (y a) & (exists z)(collaboration-with-organisation (y z) & collaboration-with-organisation (x z)) => knowledge-asset (x a))

Other network effects will require alternative characterisations. For instance, suppose a process (e.g. a ticketing process) is a knowledge asset of an organisation (i.e. a structural knowledge asset). Then the network effect is expressed by the number of users of such a process. In that case, the class of user relationships becomes important for delineating the organisations which may benefit from the process.

```
(all x y p)(structural-knowledge-asset (x p) &
```

process (p) & (exists y)(user-relationship (x

y) & user-of (y p)) => knowledge-asset (y p))

Then a quantification of the network of users of p can be found by querying

(knowledge-asset (?x p))

Further effects can be given a straightforward monetary characterisation, particularly with respect to intellectual property. For instance, if an organisation has a patent, it may have licensing income, which we may express as follows. (all x p) (structural-knowledge-asset (x p) &

patent (p) & (exists y)(user-relationship (x y)
& licensee-of (y p)) => (exists a)(integer (a)
& price (y p a)))

Examination of network effects can also be aided by exploitation of transitivity of many of these relations, as in our second scenario.

To express context- and content-based aspects of the value of a knowledge asset, one approach would be to use structures supplied by knowledge models such as those of CommonKADS. For instance, various aspects of an organisation may be expressed using predicates such as:

(agent-of (x y))
(business-process-of (x y))
(generic-task (x))
(theory (x))

These could be linked with appropriate axioms, as well as organisation-specific terms. It is an empirical question which information is required for knowledge valuation issues. For instance, if a theory is being used as part of a business process, much will depend on whether alternatives to the theory exist.

Note also that a knowledge model will typically suggest a decomposition of an organisation into 'bite-sized chunks' which can be used to produce a finer-grained analysis of its knowledge assets. For instance, certain papers, projects, websites, etc, will only be of interest to certain areas of the firm. These divisions will generally be connected to the content (e.g. research topic) of the asset. Of course content is very difficult to pin down in the abstract, but using such structures, it should be possible to produce a useful list topics of interest to an organisation based on the contents of a knowledge model of it.

5 Related Work

There has been a great deal of work performed on the application of the semantic web to e-commerce. Many focus on ontology-based technologies, or use ontologies to structure the interactions between the two. Within the scope of this paper, we will focus on some interesting and recent approaches to the issue: for a fuller survey, see [Fox and Gruninger, 1998, pp.112-115].

5.1 **TOVE**

TOVE is the Toronto Virtual Enterprise model [Fox and Gruninger, 1998], which is intended to provide a shared enterprise terminology for an organisation, in which all the terms have as unambiguous a meaning as possible in first-order logic, with an implementation in Prolog and a graphic representation. The TOVE project makes a distinction between a GEM (Generic Enterprise Model), which contains a taxonomy of object classes, relations between the classes and a set of attributes plus their definitions (i.e. an ontology), and a DEM (Deductive Enterprise Model), which is a GEM plus axioms and an inference engine. TOVE is a DEM, and includes deductive capabilities beyond inheritance for much the same reasons that we gave in Section 3.4.

The TOVE approach sees enterprise ontologies as integrating multiple conceptualisations and ontologies within the enterprise. Hence a knowledge valuation ontology would be integrated into the organisation by the TOVE ontology. TOVE itself, being concerned with internal communication, does not address the valuation issues that we have discussed in this paper.

5.2 AIAI ENTERPRISE Ontology

The Edinburgh ENTERPRISE project [Uschold et al, 1998] is an environment for integrating methods for capturing important aspects of an organisation based on a semiformal collection of terms and definitions which are required for business-related discourse.

There are four main subdivisions of the ontology. *Activities* are actions or processes, which are performed by actors, consuming resources. The *organisation* itself is modelled as either a legal entity or an organisational unit, the former type having various legal rights and responsibilities. The *strategy* of an organisation is a plan to achieve a purpose. Finally, *marketing* covers the important factors relating to sales, brands, customer needs, etc.

The ontology is then used to underlie the application of a toolkit for enterprise modelling, providing the communication between a procedure builder for capturing process models, an agent development toolkit and a task manager for visualisation, integration and support for process enactment. The role of the ontology is seen pretty much in the same way as TOVE (though it is less formal).

It is interesting that many of the requirements for a valuation ontology are distributed across the ENTERPRISE ontology. For example, a knowledge asset is a resource that is used by various activities; networks would be represented as part of the organisation section of ENTERPRISE; much of importance would also be subsumed under strategy and marketing. This is an intriguing comment on the ubiquity of KM across a modern business organisation.

The ENTERPRISE ontology is concerned with general activities and potential sales, though knowledge valuation is less of a sales issue than a management issue. It unbundles sales and general process- or strategy-related considerations, whereas the requirement for knowledge valuation is to bundle them up again.

5.3 Electronic Business Models

Electronic Business Models (EBMs) [Gordijn et al, 2000] are closer to the concerns of this paper by being centrally concerned with the issue of value, in particular the exchange of valued objects, value addition and the mechanisms of value exchange. The aim is to express at a high level of abstraction how value is created, interpreted and exchanged within a "multi-party stakeholder network", i.e. a market (described in European rather than Anglo-Saxon terms!).

Electronic Business Models relate value, quite properly, to markets. They are designed to cope with areas where value is established by exchange. However, we have seen that the marketability of *knowledge* assets is quite limited. Furthermore, we have also seen that some [Edvinsson and Malone, 1997; Norton and Kaplan, 1996] argue that the best measures of knowledge value are qualitative rather than quantitative.

Furthermore, the model focuses on notions of value that are internal to the good itself. A value object is valued using a multi-dimensional utility function. However, the value of knowledge is strongly connected with consumption of itself in relation to other objects or knowledge assets; its value strongly depends on the context of exploitation. It is possible that these complex valuation effects could be expressed using the notions of value interfaces [Gordijn et al, 2000, pp.261-2], market segments [p.264], and composite actors and value objects [pp.265-267]. The result, though, would be complex at this high level of abstraction, and it may be that a special purpose ontology such as we propose will have the advantage of relative simplicity.

Gordijn et al are sometimes ambiguous about value (as economists are themselves), allowing actors to value objects subjectively [p.260]. As we have seen (Section 2.3), this cannot contribute to the understanding of the value of knowledge (because of the resulting accountancy loophole), though of course consumers will have a maximum price they are prepared to pay.

But the valuation of knowledge is much less dependent on market transactions anyway, partly because it is only partially excludable, and hence not a proper *exchange*, and partly because much of the purpose of valuation is to facilitate management (e.g. is this knowledge worth making explicit and sharing with another department?), rather than to extract the most advantageous deal.

5.4 Knowledge Mapping and Auditing

As we saw in our third scenario in Section 3.3, the idea of knowledge mapping is that high-level knowledge models can be created in graphical formats, to be used as management tools. Knowledge auditing can then take the form of examining the knowledge maps to see where weaknesses and strengths lie [Speel et al, 1999].

Such conceptual maps are interesting from the point of view of a valuation ontology, in that they can provide base cases for attributes of instances of the ontology. Particular papers, for example, can receive valuations by making explicit how much the mapping matrices they cover, and showing which problems they address (and hence how urgently they are required).

More interestingly, such techniques also give a handle on how to value knowledge that has not at present been acquired. For example, if it is known that a product that would meet a series of requirements would be worth x_x , while the current product is worth $y < x_x$, and a matrix identifies areas where processes to achieve the requirements are unknown, then the value of the missing processes will be related to (x - y). Unknown knowledge would naturally appear as a fourth major class of knowledge asset in terms of the ontology sketched in Section 4.

5.5 Agent-Based Approaches

One way for a value to be conferred on an item such as knowledge will be for an individual negotiation to be made for the license. This is, in fact, currently the standard way in the absence of a competitive bidding market; consultants and clients enter a negotiation, and a piece of research is done, or a license to use a patented process or service is conferred. The value of the exchange is not determined by a market independently of the transaction.

Such negotiations can be performed by software agents trading off constraints in the context of uncertainty and resource limitations in open systems to produce an outcome that is as mutually beneficial as possible [Faratin et al, 2000].

An ontology for knowledge valuation in this context might be of use in giving an extra framework for valuation beyond the constraints given to the agents. For instance, understanding the relationships between knowledge to be purchased and other related knowledge in the field may lead the firm to alter its estimates of the value of the knowledge by realising, e.g. the network effects on the knowledge, or its proximity to related bodies of knowledge that it may itself possess (perhaps explicitly modelled). In other words, an ontology of knowledge valuation may help a firm have a better idea of the ultimate value of some knowledge to it beyond the relationship between the costs of acquisition and the immediate contribution it would make to profitability.

6 Discussion

Knowledge valuation is a thorny issue. As we have seen, there are many reasons to want to do it, and many reasons why it cannot easily be done, and a few reasons at least for thinking that technology can help a little. We have presented a sketch of an ontology to allow some of the relevant relations to be expressed, which could be an interesting addition to wider e-business models.

There are two chief issues to be addressed in order to turn the sketch into reality. First, clearly the ontology needs to be fleshed out, and extended if possible to include as many of the desiderata of Section 3.4 as possible. Second, the accountancy debates need to be followed, in order to see exactly how the reporting of investment in intangible assets is to be specified, and therefore how to integrate any knowledge valuation technology into the accountancy practice. For example, if the general move is towards a Skandia-type matrix, then interesting and informative parameters, and ways of establishing their values, must be invented. This is probably the simplest solution, as the various parameters of the matrix are independent of each other. A balanced scorecard approach will require more in the way of integration, in order to understand when balance is achieved (cf. [Norton and Kaplan, 1996]). A fully-fledged system of applying monetary values will be the most complex, although there is little agreement currently as to how this can be done to intangibles.

Using inference-supporting ontologies to understand the value of knowledge has the potential to be an important tool for the management of knowledge assets. We have seen what a large part of a company's value is down to the intangible assets, and knowledge is one of the most important of those. Knowing more about what it is worth is a key factor in using it properly.

Acknowledgments

This work is supported under the Advanced Knowledge Technologies (AKT) Interdisciplinary Research Collaboration (IRC), which is sponsored by the UK Engineering and Physical Sciences Research Council under grant number GR/N15764/01. The AKT IRC comprises the Universities of Aberdeen, Edinburgh, Sheffield, Southampton and the Open University.

References

[Aboody and Lev, 2001] David Aboody and Baruch Lev. *The Productivity of Chemical Research and Development*, Working paper, Stern School of Business, New York.

[Allee, 1999] Verna Allee. New tools for a new economy. *Perpectives on Business and Global Change*, 13(4), 1999.

[Backhuijs et al, 1999] J.B. Backhuijs, W.G.M. Holterman, R.S. Oudman, R.P.M. Overgoor and S.M. Zijlstra. *Measuring and Reporting Intellectual Capital: Experiences, Issues and Prospects*, PricewaterhouseCoopers Netherlands, www.oecd.org/dsti/sti/industry/indcomp/act/Ams-conf/ <u>Technical-meeting/netherlands-pwc.pdf</u>

[beemanagement.com, 1999] beemanagement.com. Understanding Intangibles, http://www.beemanagement.com

[Bushee, 1998] B. Bushee. The influence of institutional investors on myopic R&D investment behavior, *Accounting Review*, 73: 305-333, 1998.

[Cairncross, 2000] Frances Cairncross. Inside the machine, *The Economist*, November 18th, 2000.

[Christensen, 1997] Clayton Christensen. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*, Harvard Business School Press, Boston.

[Cohen et al, 2000] Wesley Cohen, Richard Nelson and John Walsh. *Protecting Their Intellectual Assets: Appropriability Conditions and why US Manufacturing Firms Patent (or Not)*, NBER working paper 7552, Cambridge, Mass., <u>http://papers.nber.org/papers/w7552</u>

[Cohendet et al, 2000] Patrick Cohendet, Frederic Creplet and Olivier Dupouet. Interactions between epistemic communities and communities of practice as a mechanism of creation and diffusion of knowledge. In Domenico Delli Gatti, Mauro Gallegati and Alan Kirman (eds.) *Interaction and Market Structure: Essays on Heterogeneity in Economics*, Springer-Verlag, Berlin, 2000.

[Collins, 1974] H.M. Collins. The TEA set: tacit knowledge in scientific networks, *Science Studies*, 4:165-186, 1974.

[Cowan et al, 1999] Robin Cowan, Paul A. David and Dominique Foray. The explicit economics of knowledge codification and tacitness, presented at the 3rd TIPIK Workshop, BETA, University of Louis Pasteur, Strasbourg, 24th April, 1999, <u>http://www-econ.stanford.edu/faculty/</u> workp/swp99027.pdf

[Edvinsson and Malone, 1997] Leif Edvinsson and Michael S. Malone. *Intellectual Capital: Realizing Your Company's True Value by Finding its Hidden Brainpower*, Harper Business, New York.

[Edwards, 2000] John Edwards. Is that your best offer? *CIO Magazine*, 1st Nov, 2000.

[Faratin et al, 2000] P. Faratin, C. Sierra and N.R. Jennings. Using similarity criteria to make negotiation trade-offs. In *Proceedings of the Fourth International Conference on Multi-Agent Systems*, Boston, pp.119-126.

[Fensel et al, 1998] Dieter Fensel, Stefan Decker, M. Erdmann and Rudi Studer. Ontobroker: the very high idea. In D.J. Cook (ed.) *Proceedings of the 11th Annual Florida Artificial Intelligence Research Symposium*, AAAI Press, Menlo Park, Calif., 1998.

[Fox and Gruninger, 1998] Mark S. Fox and Michael Gruninger. Enterprise modeling. *AI Magazine*, Fall 1998, 109-121.

[Gordijn et al, 2000] Jaap Gordijn, Hans Akkermans and Hans Van Vliet. What's in an Electronic Business Model? In Rose Dieng and Olivier Corby (eds.) *Knowledge Engineering and Knowledge Management: Methods, Models and Tools*, Springer-Verlag, Berlin, 2000, pp.257-273. [Griliches, 1995] Zvi Griliches. R&D and productivity: econometric results and measurement issues. In Paul Stoneman (ed.) *Handbook of the Economics of Innovation and Technological Change*, Blackwells, Oxford.

[Hall, 1999] Robert Hall. *The Stock Market and Capital Accumulation*, NBER working paper 7180, Cambridge, Mass., <u>http://papers.nber.org/papers/W7180</u>

[Hall, 2000] Robert Hall. *E-Capital: The Link Between the Stock Market and the Labor Market in the 1990s*, Working paper, Hoover Institution and Stanford University, http://www.stanford.edu/~rehall/SMCA-d%205-12-00.pdf.

[Keller, 1995] Arthur C. Keller. *Smart Catalogs and Virtual Catalogs*, <u>http://www-db.stanford.edu/pub/keller/1995/</u>virtual-catalogs.ps

[Kinney, 2000] Sam Kinney. An Overview of B2B and Purchasing Technology: Response to Call for Submissions, Federal Trade Commission, <u>http://mba.tuck.dartmouth.edu/</u> anderson/kinney_white_paper.htm

[Kyle, 1989] Albert Kyle. Informed speculation with imperfect competition, *Review of Economic Studies*, 56: 317-356, 1989.

[Lev, 2000] Baruch Lev. Intangibles: Management, Measurement and Reporting, Stern School of Business, NYU, <u>http://www.stern.nyu.edu/~blev/intangibles.doc</u>

[Litan and Wallison, 2000] Robert E. Litan and Peter J. Wallison. *The GAAP Gap: Corporate Disclosure in the Internet Age*, AEI/Brookings Joint Center for Regulatory Studies, http://www.aei.brookings.org

[Malone, 1997] Michael S. Malone. New metrics for a new age. *ASAP: Forbes Supplement on the Information Age*. http://www.forbes.com/asap/97/0407/040.htm

[Near, 2001] David Near. Specific applications, to be presented at the 4th Intangibles Conference: Advances in the Measurement of Intangible (Intellectual) Capital, Stern School of Business, New York, May 2001.

[Nonaka and Takeuchi, 1995] Ikujiro Nonaka and Hirotaka Takeuchi. *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*, OUP, New York.

[Norton and Kaplan, 1996] David Norton and Robert Kaplan. *The Balanced Scorecard: Translating Strategy into Action*, Harvard Business School Press, Boston.

[O'Hara et al, 2000] Kieron O'Hara, Nigel Shadbolt and Jeni Tennison. Certifying KBSs: using CommonKADS to provide supporting evidence for fitness for purpose of KBSs. In Rose Dieng and Olivier Corby (eds.) *Knowledge Engineering and Knowledge Management: Methods, Models and Tools*, Springer-Verlag, Berlin, 2000, pp.419-434.

[Scheutze, 1993] Walter Scheutze. What is an asset? *Accounting Horizons*, 7: 66-70, 1993.

[Schreiber et al, 2000] Guus Schreiber, Hans Akkermans, Anjo Anjewierden, Robert de Hoog, Nigel Shadbolt, Walter Van de Velde, and Bob Wielinga. *Knowledge Engineering and Management: The CommonKADS Methodology*, M.I.T. Press, Cambridge, Mass, 2000.

[Schulz and Jobe, 1998] Martin Schulz and Lloyd A. Jobe. Codification and Tacitness as Knowledge Management Strategies: an Empirical Exploration, Carnegie Bosch Foundation and Center for International Business Education and Research, <u>http://cbi.gsia.cmu.edu/newweb/</u> 1998WorkingPapers/ct/ct.pdf

[Shadbolt et al, in press] Nigel Shadbolt, Kieron O'Hara and Hugh Cottam. The use of ontologies for knowledge acquisition. In Jose Cuena, Yves Demazeau, Ana Garcia and Jan Treur (eds.) *Knowledge Engineering and Agent Technology* IOS Press, Amsterdam, in press.

[Shadbolt et al, 1999] Nigel Shadbolt, Kieron O'Hara and Louise Crow. The experimental evaluation of knowledge acquisition techniques and methods: history, problems and new directions, *International Journal of Human-Computer Studies*, 51: 729-755, 1999.

[Shapiro and Varian, 1999] Carl Shapiro and Hal R. Varian. Information Rules: A Strategic Guide to the Network Economy. Harvard Business School Press, Boston, Mass., 1999.

[Smith and Farquhar, 2000] Reid G. Smith and Adam Farquhar. The road ahead for knowledge management: an AI perspective, *AI Magazine*, Winter 2000, 17-40.

[Speel et al, 1999] Piet-Hein Speel, Nigel Shadbolt, Wouter de Vries, Piet Hein van Dam and Kieron O'Hara. Knowledge mapping for industrial purposes. In *Proceedings* of the Knowledge Acquisition Workshop (KAW 99), University of Calgary, Banff, 1999.

[Sveiby, 1997] Karl-Erik Sveiby. *The New Organizational Wealth: Managing and Measuring Knowledge-Based Assets.* Berrett-Koehler, San Francisco, 1997.

[Uschold et al, 1998] Mike Uschold, Martin King, Stuart Moralee and Yannis Zorgios. The Enterprise Ontology, Knowledge Engineering Review, 13:71-88, 1998.