

Ontology definition and construction, and epistemological adequacy for systems interoperability: A practitioner analysis

Journal of Information Science
37(3) 246–253
© The Author(s) 2011
Reprints and permission: sagepub.
co.uk/journalsPermissions.nav
DOI: 10.1177/0165551511401804
jis.sagepub.com



Jim Lumsden

Fife Council, UK

Hazel Hall and Peter Cruickshank

Edinburgh Napier University, UK

Abstract

Ontology development is considered to be a useful approach to the design and implementation of interoperable systems. This literature review and commentary examines the current state of knowledge in this field with particular reference to processes involved in assuring epistemological adequacy. It takes the perspective of the information systems practitioner keen to adopt a systematic approach to in-house ontology design, taking into consideration previously published work. The study arises from author involvement in an integration/interoperability project on systems that support Scottish Common Housing Registers in which, ultimately, ontological modelling was not deployed. Issues concerning the agreement of meaning, and the implications for the creation of interoperable systems, are discussed. The extent to which those theories, methods and frameworks provide practitioners with a usable set of tools is explored, and examples of practical applications of ontological modelling are noted. The findings from the review of the literature demonstrate a number of difficulties faced by information systems practitioners keen to develop and deploy domain ontologies. A major problem is deciding which broad approach to take: to rely on automatic ontology construction techniques, or to rely on keywords and domain experts to develop ontologies.

Keywords

ontology; systems interoperability

1. Introduction

There exists a considerable body of literature that concerns the development and use of ontological modelling. This is rich in theoretical concepts, as well as associated frameworks and models, all of which aim to increase the understanding of the field. Among the potential benefits of ontological modelling is the creation of a development space in which common concepts, tools and techniques can be used to come to an agreed understanding of a domain, and thus facilitate the creation of interoperable systems. The concept of the semantic web develops this theme, attempting to provide a paradigm for data standards and interaction to allow machine-based reasoning and system interoperability. This level of agent-based autonomy is underpinned by an unambiguous codified description of the entities that inhabit a given domain, and the relationships they can have with each other: the domain ontology.

This article provides a critical evaluation of the published literature on methodologies for the development of ontological solutions to problems of information systems interoperability. A particular area of interest is the value of this published material to information systems practitioners. The review reveals that practitioners interested in developing and deploying domain ontologies face a number of challenges. A major problem is deciding which broad approach to take: to rely on automatic ontology construction techniques, or to rely on keywords and domain experts to develop ontologies. It is also recognized that ontologies support the related functions of search and information retrieval, as well as forming part of the

Corresponding author:

Jim Lumsden, Fife Council, Fife House North Street, Glenrothes, Fife, KY7 5LT, UK.

Email: jim.lumsden@fife.gov.uk

infrastructure of the semantic web. These areas of research are extensive. Therefore, in order to keep the scope of this paper reasonable, the focus here is the issue of interoperability, with particular reference to the question of the agreement of meaning.

The analysis derives from the first author's participation in a project designed to explore the practicality of developing a system to improve the interoperability of information used by 32 Scottish local authorities in the management of Scottish Common Housing Registers. The purpose of Common Housing Registers is to provide a straightforward means by which applications for social housing are managed and processed. Ideally they should support systems where landlords select tenants from a single pool of applicants, and those who seek housing interact with just one source of information and advice. At the time that the work discussed here was undertaken the systems that support Scottish Common Housing Registers were not subject to ontological modelling.¹ It was hoped, however, that the end goal of interfacing, integrating and allowing individual Common Housing Register systems to work together might be built up from a shared understanding of the vocabulary of the domain. Thus the question of how meaning is best agreed was regarded as a key stage in working towards system interoperability. For this reason a practitioner analysis of the relevant literature on ontological modelling was deemed an appropriate starting point for this work.

2. Ontologies, epistemological adequacy and systems interoperability

Where the term 'ontology' refers to 'an explicit specification of a conceptualisation' [1, p. 200], ontology *development* allows a model of a domain of knowledge to be constructed in which an agreed taxonomy, or system of entity classification as a repository of meaning, is overlaid with explicit descriptions of the relationships that those entities share. This definition of ontology development can be further refined by considering the steps taken in analysing domain knowledge, making domain assumptions explicit, clarifying the difference between domain knowledge and operational knowledge, determining the agreement of meaning, and making it possible for domain knowledge to be reused [2, p. 1]. The development of an ontology ultimately offers a 'shared and common understanding of a domain that can be communicated between people and application systems' [3, p. 75]. In the case of systems design an ontology can be used as an *interlingua*, translating data and queries from one (computer) system to another [4, p. 17], removing the need for a priori understanding. Thus when a remote system presents a query in a form which it understands, it receives in return a response that makes sense.

To assess the value of an ontology, three main criteria have been proposed [5]. These are epistemological adequacy, operationality and reusability [5, p. 41]. In combination with sub-criteria, an examination of an ontology along these axes will reveal how well it communicates ontological commitments between actors, and the extent to which it supports knowledge acquisition.

The focus of this paper is epistemological adequacy. Epistemological adequacy presents a number of challenges, not least due to the shifting nature of meaning from context to context. For example, earlier work on ontology development identifies that even commonly used terms project multiple meanings when used informally. Taking the context of housing, for instance, in different parts of Scotland the term 'cottage' may refer to a semi-detached two-storey property of any size, or to a more traditional small single storey house. Equally, a number of apparently different terms can actually share meaning: '4-in-a-block', 'maisonette', 'flat' and 'walk-up flat' are all terms that may be treated interchangeably in a discussion of one particular type of dwelling. Added to this, domain experts regularly agree terms that are, in fact, highly ambiguous and often inhibit shared understanding [4, p. 22]. More recent studies confirm such findings. For example, research completed in 2009 reveals how domain experts and knowledge engineers present a range of interpretations of the same data [6].

Such 'inaccurate' expert responses have been identified as a major issue in ontology development. A particular case where this is problematic occurs when the ontologist is unaware of the ontological implications of contributions from subject specialists. A typical outcome here is that sub-classes are confused with classes [7, p. 542]. For example, when housing administrators are asked about the errors that commonly occur in form filling, they are expected to refer to sub-classes such as omissions, inaccuracies or lack of supporting documentation. Instead they may identify examples of these sub-classes, such as 'missing signature' rather than 'omission'. Further complications arise in scoping a domain where there are multiple and diverse interests: the greater the number of interested parties, the greater the risk of lengthy arguments over meaning. For example, in the case of Common Housing Registers, the expert housing staff of 32 local authorities may be interested in debating the definition of the term 'bedroom'.² Where knowledge is to be integrated across domain boundaries (for example, landlords and tenants may have opinions on what constitutes a 'bedroom' from their own lay perspective, and these differ from the interpretation of a housing officer), shared understanding becomes a greater challenge.

Given the difficulties in collecting and analysing common terms from subject experts, systems practitioners – and particularly those with a background in information *systems*, as opposed to information *science* – may be tempted to devise an automatic tool for this purpose, especially in cases where there exists a plethora of electronic sources in the subject

domain of interest. However, although search tools may easily identify key terms, they struggle to identify the subtleties of meaning, such as the nature of the relationships between terms identified. Automated systems also miss cases where context can reverse the apparent meaning of a term. In the provision of social housing, for example, ‘applicants’ apply for accommodation independently, or on behalf of themselves and their ‘dependents’. The applicant–dependent relationship may initially appear to be straightforward: one applicant may have zero, one, or more dependents. However, entitlement for housing, and the type of property to be provided, takes into account a number of features of the applicant–dependent relationship, such as their real-life connection and the age of the parties. Furthermore, the existence of more than one ‘dependent’ modifies the case further.

Related to the issue of meaning in the broad sense is the question of the *level of detail* to which it is necessary to seek agreement, and how much detail is required for the ontology to be considered valid. Again, earlier research illustrates approaches. For example, four ontologies in the legal field were examined to determine their relative merits [5, p. 30]. They all differed in the level of detail, each with its own impact on expressive power and reusability. This analysis confirmed that while expressive power increases as an ontology becomes more detailed, its potential for reusability decreases [5, p. 54]. Others have warned that efforts dedicated to agreeing fine levels of semantic detail can be counterproductive (for example, [7, p. 542]).

An implication of this discussion of epistemological adequacy is that regardless of how many people are involved in exercises to agree meanings for a formal ontology, ultimately there is a need for a final arbitrator of classification to confirm the vital common terms. Depending on the context, this role is often invested with a large degree of influence as to how the subject domain is represented in the system. Indeed, a powerful ‘owner’ of classification may be positioned to alter reality through decisions on the meaning of terms, where ‘acting on the model changes the world’ [8, p. 11].

The process of agreeing meaning, and not least decisions on who is charged with such a task, is key to the development of an ontology. Indeed, even if other conditions for ontology development are ideal – for example the domain in question is closed and entities are few and stable – if the ontologist is not a domain expert and/or does not have access to subject specialists, then there is a strong risk that any output from such an exercise will be unusable as a representation of shared meaning. A lack of commonly agreed understanding when developing systems can lead to poor communication between people and organizations, and difficulties in defining specifications, as well as limited potential for the re-use and sharing of application code, and the interoperability of systems [4, p. 3].

3. ‘Amateur’ collaborative efforts as alternatives to ontology development: applicability in the context of local government administration

It has been established that key to the process of agreeing meaning is the constituency involved in the task: ontological modelling is ideally executed within a community of practice. If those who participate in the work do not offer a formal requisite level of domain expertise, then the potential value of the outcome of the effort is called into question. It has also been noted that domain experts find it difficult to switch into an ontology frame of mind, so even if they are available to contribute, they may not contribute effectively [9, p.7]. Added to this, there is some evidence to suggest that the amount of effort involved in developing a formal taxonomy from such exercises does not actually merit the investment. For example, a study of systems integration/interoperability in the Norwegian healthcare sector found that increasing efforts to integrate systems was actually counterproductive [10, p. 276]. A conclusion of this work was that disorder is an inherent feature of integrated systems and that, in effect, the more time and resources devoted to joining systems together, the worse the result for the end-users [10]. These findings build on earlier research which suggests, for example, that unless database administrators have a priori knowledge of every context in which their systems will be queried, then the effort of ontological modelling does not match the anticipated reward of system interoperability [11, p. 9]. This implies that boundaries should be set as to levels of investment made in formal efforts to agree shared meaning for systems development.

As an alternative to formal ontology development it is worth considering whether an informal collaborative effort on the part of interested amateur parties may offer a more worthwhile and valid means of determining shared meaning of a domain. A folksonomy approach [12, p. 1] recognizes the apparent lack of logic and chaotic nature of information content, while rejecting the need to take heed of database structures, or the requirement for ontological commitments. Shirky [8] is a strong advocate of systems where users create their own links between information artefacts, adding value by forging new meanings as they do so. Here the power in determining shared meaning is removed from the recognized domain experts and ontologists, and instead passes over to the mass of users. It is argued that such a large corpus of contributors, working collaboratively in a similar way to Wikipedians, may provide a route to agreed meaning within a domain [13, p. 109]. Key here is the ease with which content on a wiki-like platform can be overwritten and then restored, so that only content that enjoys a critical mass of support will survive. The approach allows a process of iteration to take place whereby entities, relationships and meanings are proposed and modified until an acceptable consensus is reached. As a result, the

eventual domain model will be acceptable to all. As an illustration of this, Shirky cites the range of Wikipedia articles that cover controversial issues. These have survived the ruthless editing and potential for vandalism that are an inherent part of the system [13, p 116].

It should be recognized, however, that such an approach only has the potential to succeed under certain conditions:

1. There should be a sufficient number of people concerned with the development of the ontology to represent the full range of views on meaning in order to build the domain description.
2. Those involved need to care enough about the subject to engage in the process of iteration and reiteration until consensus is achieved.
3. The number of individuals who actively participate in the process should be large enough to guard against domination by a particular cohort (for example, those who simply have most time to devote effort to the enterprise).

These conditions may be easily met among a community of enthusiasts or hobbyists, such as recreational code breakers, political activists or sports fans. However, it is less obvious that this would be the case in the quest to agree meaning for an ontology focused upon a sphere of local government administration, as is the case of the Common Housing Registers of Scotland, as specified for the larger study to which this analysis of the literature contributes.

A further major hurdle for the agreement of meaning based upon the contribution of amateurs (as opposed to domain experts working in a community of practice) is the issue of the validity of agreed terms. The risk of a wiki-like approach is that it leads to agreement at the level of the lowest common denominator, modelling facets that are universally agreed, but ignoring aspects of the domain deemed 'too difficult'. An example from housing is the issue as to whether or not someone with an established connection to the area in which accommodation is sought should be treated as a priority case. Some – but not all – landlords award 'local connection' points as a means of improving the opportunity of local people to access housing stock in a particular area. Those who follow the practice see this as an essential housing management tool. Others consider this poor practice of questionable legality. Views are polarized, and debate often generates more heat than light.

Similarly, 'truthiness' [14], i.e. a desired perception of truth, could take precedence over truth per se. So, for example, if enough contributors think a poodle is a bulldog, a bulldog it is. Less amusing examples emerge in the context of the provision of social housing. It is possible that a wiki-style approach could lead to sub-optimal solutions where the subject matter is prone to populism, such as the housing of sex offenders. The definition of data relationships in these circumstances is probably best left to thought leaders rather than to the mass of opinion, even among professionals.

4. Theory and practice of ontological development as published in the extant literature and its value to the information systems practitioner

If information systems professionals reject the alternatives to formal ontology development as inadequate, then it is worth considering the literature that can point to practical solutions related to the question of how to seek agreement of meaning. This can then be applied to the specifics of the system under construction. In the literature of ontology development a number of methodologies are proposed. As can be seen from the examples outlined below, they may depend solely upon automatic term generation, upon contributions from subject experts, or a combination of both.

A framework for 'ontology learning' [15, p. 196] demonstrates the drawbacks of automatic creation or redesign of ontologies. This framework tested natural language processing techniques as a means of discovering domain terminology. The goal was then to use these to map out a concept hierarchy and to specify ontological relationships. The output provided by the automatic tools for an ontology on the subject of football was compared with another created by domain experts. It was found that the two ontologies differed to a large degree. The automated ontology could only approach a 47% match with that of the experts, and that was after some massaging. Most of the other comparative indicators in the study were considerably worse than this. A further perspective offered is one that depends more heavily on human input and is based on software engineering principles, in particular the Rational Unified Process [16, p.259].³ Here domain experts and knowledge engineers work together on multiple iterations of the ontology to such an extent that they learn together over the duration of the exercise, with the advantage that the domain expert grows an awareness of the consequences of providing inappropriate subject terms. A drawback of this approach, which will be recognized by information systems practitioners seeking scope for systems interoperability, however, is that the progressive refinement of the model results in an ontology that is highly specific to a particular context. Other frameworks propose the modelling of the domain by taking into account both (a) the knowledge of experts derived from interviews and workshop alongside (b) efforts to build a hierarchy of meaning automatically by interrogating, for example, databases and electronic document management systems, as well as existing ontologies (for example, [17, p. 4] and [18, p. 2]). Others propose that the need for multiple perspectives for building a final ontology depends on the construction of several other ontologies. For example, with

reference to a theoretical framework for ontology development for an e-commerce system, four separate ontologies were considered [19, p. 3175].

An analysis of the literature surveyed above on frameworks for ontological modelling would point the information systems practitioner to the conclusion that a relatively straightforward approach for the agreement of meaning would be a keyword analysis of the domain in question, in combination with consultations with experts to refine the final output. There appears to be no obligation, for instance, to carry out extensive document searches with lexical taggers as part of this process. To determine how such an approach might be implemented in practice, the information systems practitioner may turn to the literature related to the application of ontological modelling in real-world environments. Unfortunately, however, although there is a relatively sound theoretical basis for the deployment of ontological modelling in systems development, concrete examples of successful systems implementations underpinned by ontology development, are somewhat scarce.

Much of the practical work in this area that *can* be identified has been undertaken under the auspices of European-funded projects in response to problems of poor data integration across individual countries and Europe as a whole.⁴ In response to the issues of this nature raised by the European Commission, the main conclusion of a study into the requirements for a full e-government ontology that would comply with electronic government targets in Greece, for example, found that central and local government had been unable to ensure data interoperability [20, p. 187]. While publication of this type of work may have value for other purposes, it does not offer much to an information systems practitioner who seeks a method for the agreement of meaning to improve the interoperability of local government systems. Indeed, of the paucity of published work available, any lessons of earlier implementations take the form of practices to avoid, rather than practices to emulate, as can be seen elsewhere beyond the domain of government. For example, a case study of the development of an ontology to support the interoperability of electronic product catalogues takes the reader through the stages of domain analysis, modelling entities in the domain, and establishing relationships between entities [21]. A close reading of the work reveals, however, that the building of the ontology appears to have been a means of arriving at a relational database schema, not least because the goal of the exercise was to build an operational procurement catalogue. There are a number of drawbacks of the system as implemented, for example, an attempt to deploy richer relational terms (substitute – for, complement – to) leads to minor, but potentially inappropriate outcomes, such as the substitution of a pencil for a pen in a stationery order.⁵ A further weakness of the work is the poor distinction between classes and instances, a result of which is that the ontology holds 881,000 classes. If these are easy mistakes to make, the lesson here is that ontologists need to pause during the development process to confirm that the product that they are trying to create actually has real-world value.

Another set of case studies [5] emphasises how important it is to appreciate that it is not possible to find a neutral point of view from which everything can be described. This is illustrated with reference to the design of a series of ontologies in the legal field where, although each ostensibly had very similar aims, there was clear divergence in the agreement of terms at the very highest level. A key point that an information systems practitioner needs to appreciate from a reading of this work is that the starting perspective of the ontologist strongly influences final domain representation. So, for example, an ontology developed for functional decomposition and system analysis will be substantially different from one developed for the construction of a body of knowledge, even when they purport to describe an identical domain. The assessment of the relationship between the initial standpoint of the researcher, as well as the use to which the completed ontology will be put, is a vital component of project controls and risk. In the context of housing information systems, for example, it is important to take into account that how the data is handled may have an impact on housing applicants' life chances. Therefore landlords would wish to ensure that the underlying ontology is constructed from an appropriate perspective. It should not be forgotten, however, that the existence of more than one ontology describing a particular domain is not unusual (perhaps for the reasons given above), as demonstrated by the existence of ontology repositories, and tools such as vocabulary crosswalks and mappings are available to establish relationships between ontologies.

The analysis above shows that, although there is a body of literature on theoretical considerations in determining the agreement of meaning, there is little that points the information systems practitioner to best practice in general, and even less that addresses the specifics of the starting point of this study (identifying agreed terms for systems operability in Common Housing Registers), nor its scale (for information systems deployed by 32 Scottish local authorities). While many of the publications keenly promote a particular framework or modelling approach (often one created by the authors of the papers in question) few provide conclusions of value to those attempting to create systems in the 'real' world. Such papers (for example, [22–26]) frequently detail an application developed as a 'proof of concept'. However, from an information systems practitioner perspective, 'proof of concept' papers expend much effort in proposing solutions to the easy part of a problem, while glossing over the hard part. A particularly instructive example of this is provided in the discussion of an ontology for the animal domain of 'canine' [27]. Through a complex series of algorithms it is determined that a fox is a furry member of the canine class. Most of the conclusions reached here are based on the prior work of taxonomists,

who actually achieved the same end results without the benefit of ontological modelling. While it is clearly the case that automatic tools can be applied to domains where taxonomies have not been prepared in advance, examples are not prominent in the literature. Given the difficulties of identifying earlier studies that might demonstrate the obvious benefits of ontological modelling over a conventional systems development approach, the information systems practitioner may then question whether this is actually worth pursuit.

Related to ontology design methodology is the ontology representation language. Many such languages exist, including Ontolingua, LOOM, OCML, OIL, DAML + OIL and OWL, with vary degrees of formality and expressive power [28]. OWL is the latest of these, and was developed by the W3C Web Ontology Working Group, drawing on the features of preceding languages to create a sophisticated language specifically for the semantic web [29]. From the practitioner's perspective, OWL has the advantage of being supported by generic tools including editors, reasoning systems, query languages and storage technologies [30].

The considerable body of robust languages and tools provides a strong platform, in that they facilitate the creation of an ontology i.e. the work of the ontologist. The benefit of these is less obviously applicable to the domain expert or information systems practitioner. Workbench style environments such as Protégé go some way to addressing this issue. The possibility of a disconnection between tools useful to the ontologist and tools for the practitioner is exemplified in [31], where the field of ontology development languages is reviewed and the conclusion is reached that the steps in the ontology lifecycle are gathering, extraction, organization, merging, refinement and finally retrieval. Retrieval is seen as the ultimate goal, but 'it will be a while yet before we see smart software applications'. For those whose main interest is the creation of smart software applications, this leaves a need unaddressed.

5. Conclusion

A practitioner analysis of the published literature on methodologies for the development of ontological solutions to problems of information systems interoperability uncovers the difficulties related to representing knowledge in a structured format. Many approaches are proposed. Indeed, it is interesting to note how an activity designed to promote understanding has clearly led to a proliferation of standards from which developers may choose. A limitation of the proposed approaches, however, is that they derive largely from the work of academics whose theoretical output does not offer working examples for 'real world' application. As such, the value of this published research to a practitioner attempting a systematic approach to systems development or integration supported by technology is limited. Major problems faced by an information systems practitioner keen to develop and use ontologies relate to decisions as to which broad approach to take: to rely on automatic ontology construction techniques, or to rely on keywords and domain experts. This was experienced first-hand through involvement in the case study of the Scottish Common Housing Registers as described above.

A benefit of exploring these issues, however, and one that is especially useful to practitioners who work in information systems, is that in addressing this body of literature the real complexities related to the processes of agreement of meaning are made obvious. This raises awareness of the importance of such work, and possible approaches. It is clear, for example, that some broad messages from this review could be applied in cases similar to that of the Scottish Common Housing Registers, where ontological modelling is actually achieved in practice (which was not ultimately the case in the project discussed in this paper). For example, keyword analysis combined with input from subject experts could offer a straightforward approach to ontology design; practitioners need to bear in mind that the level of detail to which the work is completed needs to take into account that high expressive power sacrifices interoperability (and vice versa); ontology design is not a 'neutral' activity, and any arbitrator in the final decisions on subject terms holds a degree of power.

The ontology created in such a process has obvious applications to explore the mapping of terms from folksonomies onto a clear ontology, and to use this mapping for retrieval of information from the internet. It could be particularly productive to work in conjunction with other projects that explore local government domains in the context of semantic web technologies.⁶

This review highlights the scope for further work in this area. For example, this can be appreciated at the level of broad issues of interest to the domain of information science, such as the power of context in information sharing [32] – whether this be through the use of ontologies as is the case of the content of this paper, or with respect to other areas of research such as studies of information behaviour. Particular concerns of the practitioner community have also been highlighted, not least the accessibility of fully documented case studies from which future practical implementations may draw lessons on good practice.

Acknowledgement

The authors thank Valeri Wiegel for his assistance in the preparation of the original manuscript of this paper. They are also grateful for the generosity of the two anonymous referees for their detailed and constructive comments in its preparation for publication.

Notes

1. It should be noted that in the event, an ontology was *not* developed as a part of the project.
2. The definition of a bedroom, while intuitively obvious, is important in this context because the number of rooms in a house that can be defined as bedrooms affects not only decisions around present housing circumstances (for example, a current indicator of overcrowding), as well as the size of property to which an applicant might be entitled. Typically, disputes arise over downstairs rooms intended upon construction as bedrooms, but subsequently put to another use by the occupant. Some systems might define bedrooms according to their intended use, others by the de-facto use. In everyday parlance a bedroom may be ‘in the eye of the beholder’, but for an interoperable system an agreed definition is vital.
3. The Rational Unified Process (RUP) is an iterative software development framework owned by IBM. The framework is designed to be customized and configured by the development organizations and software project teams that will select the elements of the process that are appropriate for their needs. The process consists of four project lifecycle phases, six software engineering disciplines and three supporting disciplines.
4. It is tempting to put this down to language barriers. However, of course, the practice of describing a domain in terms of logical relationships, controlled vocabularies and taxonomies should actually help to eliminate problems caused by variations in language. This applies whether the issues arise from differing domain knowledge or differing mother tongues – assuming the underlying legal terms can be aligned, which is often the case in domains that are governed under EU Directives. The METALEX project (www.metalex.eu) explored this area further. This may form an interesting contrast to the case discussed by Visser and Bench-Capon [5].
5. ‘Is – a’ overloading is a common misuse of the subclass–superclass relationship, which results in oversimplification of relationships between the entities in the domain and a focus on ‘one to many’ style simple relationships [7, p. 542].
6. In this context, it is interesting to note that the ESD-Toolkit team is currently adapting the Integrated Public Sector Vocabulary (IPSV www.esd.org.uk/standards/ipsv/) to incorporate Scottish terms, and are exploring its extension into other European countries with the Smart Cities project (www.smartcities.info/).

References

- [1] T.R. Gruber, A translation approach to portable ontology specifications, *Knowledge Acquisition* 5(2) (1993) 199–220.
- [2] N. Noy and D. McGuinness. *Ontology development 101: a guide to creating your first ontology* (2001). Available at: www-ksl.stanford.edu/people/dlm/papers/ontology-tutorial-noy-mcguinness-abstract.html (accessed 30 August 2010).
- [3] J. Davies, D. Fensel and F. van Harmelen. *Towards the Semantic Web: Ontology Driven Knowledge Management* (John Wiley & Sons, London, 2002).
- [4] M. Uschold and M. Gruninger, Ontologies: principles, methods and applications, *Knowledge Engineering Review* 11(2) (1996) 93–136.
- [5] P.R.S. Visser and T.J.M. Bench-Capon, A comparison of four ontologies for the design of legal knowledge systems, *Artificial Intelligence and Law* 6(1) (1998) 27–57.
- [6] J. Evermann, Theories of meaning in schema matching: an exploratory study, *Information Systems* 34(1) (2009) 28–44.
- [7] J. Kingston, Multi-perspective ontologies: resolving common ontology development problems, *Expert Systems with Applications* 34(1) (2008) 541–550.
- [8] C. Shirky, Ontology is over-rated: categories, links and tags (2005). Available at: www.shirky.com/writings/ontology_overrated.html (accessed 30 August 2010).
- [9] J. Fraser, N. Adams, A. MacIntosh, A. McKay-Hubbard, T.P. Lobo, P.F. Pardo, R.C. Martinez and J. S. Vallecillo, Knowledge management applied to e-government services: the use of an ontology, *Proceedings of the 4th IFIP international working conference on Knowledge management in electronic government* (Springer-Verlag, Berlin, 2003).
- [10] G. Ellingsen and E. Monteiro, The slight surprise of integration. In: C. Sørensen, Y. Yoo, K. Lyytinen and J.I. deGross (eds.) *Designing Ubiquitous Information Environments: Socio-technical Issues and Challenges* (Springer, New York, 2005) 261–274.
- [11] A.M. Ouksel and I. Ahmed, Ontologies are not the panacea in data integration: a flexible co-ordinator to mediate context construction, *Distributed & Parallel Databases* 7(1) (1999) 7–35.
- [12] T. Vander Wal, Folksonomy (2007). Available at: <http://vanderwal.net/folksonomy.html> (accessed 30 August 2010).
- [13] C. Shirky, *Here Comes Everybody: The Power of Organising without Organizations* (Penguin Books, London, 2008).
- [14] K. Maxwell, *Macmillan English Dictionary: Word of the week archive* (2006). Available at: www.macmillandictionaries.com/wordoftheweek/archive/060213-truthiness.htm (accessed 30 August 2010).
- [15] R. Gacitua, P. Sawyer and P. Rayson, A flexible framework to experiment with ontology learning techniques, *Knowledge Based Systems* 21(3) (2008) 192–199.
- [16] A. De Nicola, M. Missikoff and R. Navigli, A software engineering approach to ontology building, *Information Systems* 34(2) (2009) 258–275.
- [17] E. Blomqvist and A. Ohgren, Constructing an enterprise ontology for an automotive supplier, *Engineering Applications of Artificial Intelligence* 21(3) (2008) 386–397.
- [18] N. Adams, J. Fraser, A. Mackintosh and A. McKay-Hubbard, Towards an ontology for electronic transaction services, *Intelligent Systems in Accounting, Finance & Management* 11(3) (2002) 173–181.

- [19] F. Garcia-Sanchez, R. Valencia-Garcia, R. Martinez-Bejar and J.T. Fernandez-Breis, An ontology, intelligent agent-based framework for the provision of semantic web services, *Expert Systems with Applications* 36(2) 3167–3187.
- [20] G. Prokopiadou, C. Papatheodorou and C. Moschopoulos, Integrating knowledge management tools for government information, *Government Information Quarterly* 21(2) (2004) 170–198.
- [21] T. Lee, I. Lee, S. Lee, S. Lee, D. Kim, J. Chun, H. Lee and J. Shim, Building an operational product ontology system. *Electronic Commerce Research & Applications* 5 (2005) 16–28.
- [22] Y.J. Chen, Y.M. Chen and H.C. Chu, Development of a mechanism for ontology-based product lifecycle knowledge integration, *Expert Systems with Applications* 36(2) (2009) 2759–2779
- [23] H. Cheng, Y. Lu and C. Sheu, An ontology-based business intelligence application in a financial knowledge management system, *Expert Systems with Applications* 36(2) (2009) 3614–3622.
- [24] A. Naumenko, S. Nikitin, V. Terziyan and A. Zharko, Strategy alliances in the paper industry – XML vs ontology-based integration platforms, *The Learning Organisation* 12(5) (2005) 492–514.
- [25] A. Saito, K. Umemoto and M. Ikeda, A strategy-based ontology of knowledge management technologies, *Journal of Knowledge Management* 11(1) (2007) 97–114.
- [26] Q.N.N. Tran and G. Low, MOBMAS – a methodology for ontology-based multi-agent systems development, *Information & Software Technology* 50(7–8) (2008) 697–722.
- [27] H. Wang, S. Liu and L.T. Chia, Image retrieval with a multi-modal ontology, *Multimedia Systems* 13(5-6) (2008) 379–390.
- [28] O. Corcho, M. Fernandez-Lopez and A. Gomez-Perez, Methodologies, tools and languages for building ontologies. Where is their meeting point? *Data & Knowledge Engineering* 46 (2003) 41–64
- [29] I. Horrocks, P. Patel-Schneider and F. van Harmelen, From SHIQ and RDF to OWL: the making of a web ontology language. *Journal of Web Semantics* 1 (2003) 7–26
- [30] S. Bowers, J. Madin, M. Schildhauer, Owlifier: Creating OWL-DL ontologies from simple spreadsheet-based knowledge descriptions. *Ecological Informatics* 5 (2010) 19–25
- [31] J. Pulido, M. Ruiz, R. Herrera, E. Cabello, S. Legrand and D. Elliman, Ontology languages for the semantic web: a never completely updated review. *Knowledge Based Systems* 19 (2006) 489–497.
- [32] H. Hall and M. Goody, KM culture and compromise – interventions to promote knowledge sharing supported by technology in corporate environments, *Journal of Information Science* 33(2) (2007) 181–188.