

A Program for Building a State of the Art Enterprise Ontology: Report on Progress

Chris Partridge¹, Milena Stefanova²

¹*The BORO Program*, partridge@BOROProgram.org.

²*LADSEB CNR*, Italy, stefanova@ladseb.pd.cnr.it.

Abstract: This paper is a report on progress of the CEO project whose goal is to build a state of the art enterprise ontology. The project is currently at the stage of harvesting insights from the best existing enterprise ontologies. The goal of this stage is to synthesising from these a Base Enterprise Ontology. This will then be used as the foundation for the construction of the ‘industrial strength’ Core Enterprise Ontology (CEO). The synthesis is intended to build upon the strengths and eliminating – as far as possible – the weaknesses from the selected ontologies. Among other things, this paper describes one of the main achievements of this work to date: the development of the notion of a person (entities that can acquire rights and obligations) enabling the integration of a number of lower level concepts. In addition, it identifies some of the common ‘mistakes’ in current enterprise ontologies – and proposes solutions.

1. INTRODUCTION

This paper reports on the initial results from an ongoing project – the CEO (Core Enterprise Ontology) project that is part of the BORO Program [BORO]. The CEO project aims to build an ‘industrial strength’ enterprise ontology, that will be suitable as a basis for facilitating, among other things, the semantic interoperability of enterprises’ operational systems.

This initial stage of the project involves a collaboration with the European IKF (Intelligent Knowledge Fusion) [EUREKA] project. This European IKF project will ultimately develop a Distributed Infrastructure and Services System (IKF Framework) with appropriate toolkits and techniques for supporting

knowledge management activities. The project is planned to take 3.5 years (it started in April 2000) and has participants in Italy, UK, Portugal, Spain, Hungary and Rumania.

The IKF project aims to develop two kinds of ontology:

- a *Reference Ontology* composed of a *Top Level Ontology* and several *Core Ontologies* [B1997]. The top level ontology contains primitive general concepts to be extended by lower-level ontologies. The core ontologies span the gap between various application domains and the top level ontology.
- *Domain Ontologies*. The vertical applications will build ontologies for their specific domains. For example, there is an IKF/IF-LEX project that is building a domain ontology for bank supervision regulations, focusing on money laundering.

2. SYNTHESIS WORK PLAN

The first stage of the CEO project is a synthesis of the best existing enterprise ontologies. The scope of this stage is large – and so the work has been divided into more manageable chunks.

As [B1997] states, a core ontology contains “the categories that define what a field is about.” A first rough intuitive guess of what these categories might be has proved not only a useful basis for segmenting the work but also helped to clarify the focus on the important aspects for the CEO.

Our guess of what the categories might be is:

- *parties* (persons) which may enter in
- *transactions* (composed of agreements and their associated activities), involving
- *assets*.

A review of existing ontologies was made and a sample of the best selected on the basis of:

- the relevance of their content to the Core Enterprise categories, and
- the clarity of their characterisation of the intended interpretations of this content [G1997], [G1993] & [P1996].

This gave us the following list:

- Toronto Virtual Enterprise - TOVE [F1993], [F1996] & [TOVE],
- AIAI’s Enterprise Ontology - EO [EO], [U1997] & [U1998],
- Cycorp’s Cyc® Knowledge Base - CYC [CYC],
- W.H. Inmon’s Data Model Resource Book - DMRB [I1997] & [H1997].

The work proceeds by analysing one category in one ontology at a time, and then re-interpreting the previous results in the light of any new insights.

Initially, the work focuses on individual ontologies but as it proceeds there is enough information to start undertaking comparisons between ontologies. The final analysis will encompass analyses of both the individual ontologies and comparisons between them.

In each of the ontologies, the concepts and relations relating to the category being considered are examined for the clearness and uniformity of their descriptions and formalisations. Further, each concept is analysed for its coverage and extendibility in cases where the coverage is not complete. Relations between concepts that are not explicitly described, but clearly exist, are identified as well. In addition, for the sake of a clear interpretation, we have found it necessary to consider the top concepts (whether or not they are explicitly described).

An important part of the analysis is testing each concept and its relations against a number of standard examples and more specialized concepts. Further, a check is made against a number of standard difficult cases. Both these checks help to identify weaknesses in the coverage of the ontologies.

A key concern in the analysis is to understand how the various concepts interlink with one another, to better understand the unifying structure of the enterprise field.

At various stages during the analysis an interim ontology is synthesised from the strengths found in the analysis, in such a way as to eliminate the known weaknesses – and itself analysed. In the final synthesis, all the categories in all the ontologies are combined into a base CEO ontology.

At this time, our work is concluding the analysis of the Parties (Persons) category for the EO and TOVE ontologies – and early drafts of synthesised ontologies are being reviewed. There is still substantial work that needs to be done in determining the precise relations between concepts, such as *LEGAL ENTITY* and *OWNERSHIP* within the EO.

3. INITIAL FINDINGS

Both the ontologies have many important insights and provide much useful material. However, our most general findings, at this stage, are that none of the ontologies:

- adequately meet our criteria of clear characterisation, or
- really share a common view of what an organisation is.

These findings, taken together mean that the creation of the synthesised base CEO ontology cannot just be a simple merging of the common elements of the selected ontologies.

We now illustrate these findings with examples and show how we synthesised a resolution to some of them.

3.1 Clear Characterisation

When the characterisation is unclear, it is difficult to work out the intended interpretation – in the worst case, impossible to decide between competing interpretations. There are many different ways in which the characterisation

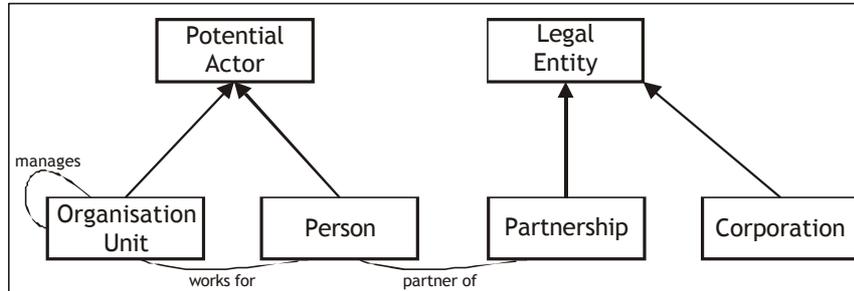


Figure 1. Simplified EO overview

can be unclear – as we describe below.

A clear overview of the structure helps understanding. Neither TOVE nor EO had one, so we developed graphical representations based upon ER diagrams. Figures 1 & 2 provide simplified versions of these.

Both TOVE and EO make use of a number of top concepts. A top ontology – or top concepts – can provide a useful structure that segments the enterprise and other domains into general categories that help in defining and the use of domain concepts and relations. However, if this is not done properly it can have the opposite effect.

Some of the problems we encountered with the top concepts and the domain analysis in the selected ontologies are:

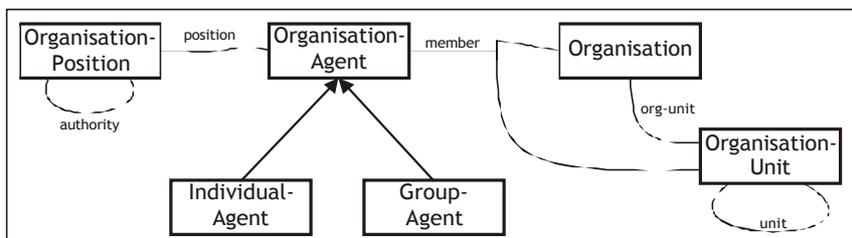


Figure 2. Simplified TOVE overview

- *Insufficient characterisation of the disjointness of top concepts.* For example, in the informal EO the relationship between the top concepts *ENTITY*, and *ROLE* is not clear – in particular, whether *ROLES*

can be *ENTITIES* or not, and so whether they can enter into *RELATIONSHIPS*.

- The same lack of care in characterising disjointness (and overlapping) exists at the domain level in both TOVE and EO. We found this can make it impossible to definitively determine the intended interpretation. For example, in TOVE the formalisation allows an *ORGANISATION-UNIT* to be an *ORGANISATION* – though this seems counter-intuitive, and probably not what the authors intended.
- *Not applying top concepts*. TOVE states that a fluent is "a [type of] predicate or function whose value may change with time". But it does not identify which predicates in its ontology are fluents – leaving this to the readers, who have to make their own judgements. Supplying such information would have helped not only the users of the ontology but also its creators and designers. For example, the TOVE's creators end up (probably unintentionally) having to regard *ORGANISATION* as a fluent – when in the normal (commonsense) use of the concept it is not.
- *Messy formalization trajectories*. EO formalizes its concepts in logical systems (Ontolingua and KIF), which rely on their own (different) top concepts. An attempt for a clear formalisation trajectory has been made [U1997], but unfortunately this does not match very well with the informal specification. For example, in the informal EO it is stated that each *RELATIONSHIP* is also an *ENTITY*, but is not defined as such in the formalization. Furthermore some *RELATIONSHIPS* are defined in the formalization as *classes* and others are defined as *relations* without explaining what the motivations for these choices are (e.g., *SALE* is a *RELATIONSHIP* formalized as a *class*, *HAVE-CAPABILITY* is a *RELATIONSHIP* formalized as a *relation*). This becomes a more serious problem if the formalisation is meant to be taken as the more accurate version.
- *Failing to use general concepts to achieve uniformity*. Both TOVE and EO fail to use top concepts to describe in a uniform way core relations and concepts. This hampers understanding. Typical examples are the *part-of* relation, used in describing the decomposition of organizations into smaller units, and the relation, which shows the different ways for *participation* in organizations. For example, TOVE introduces two kinds of *part-of* relations: *org-unit* (between *ORGANISATION* and *ORGANISATION-UNIT*), and *unit* (between two *ORGANISATION-UNITs*). These relations express *ORGANISATION* and *ORGANISATION-UNIT* decompositions, but are not explicitly unified under a common relation. In the EO several ways of participating in a company are considered, as a partner (*partner_of* relation

between *PERSON* and *PARTNERSHIP*), as an employee (*works_for* relation between *PERSON* and *OU*), as a shareholder in a corporation (only in the informal EO specification, Uschold, et al, 1997). These ways of participation are not unified in the EO.

- *Insufficient analysis.* As an example consider the EO concepts of *OWNERSHIP* and *SHAREHOLDING* (Uschold, et al, 1997), which are formally unrelated, while *SHAREHOLDING* as evident from its informal and formal, definitions represents the ownership relation between a *CORPORATION* and its owners.

3.2 Common view of an organization

TOVE and EO do not share a common view of an organisation. Even a cursory glance at Figures 1 & 2 (which give a broad picture) reveals there are significant differences.

There are many examples in both TOVE and EO of how a better analysis would have led to more similar views:

- *Insufficient analysis.* In TOVE, for example, it seems that an *ORGANISATION* is not an *AGENT*, but has *AGENTS* as members. Yet as EO recognises there are many examples of organisations (such as the EU or NATO), which have other organisations as members.
- *Missing Links.* In the EO, the relation between the concepts *OU* and *LEGAL ENTITY* is unclear. All that we are told is that a *LEGAL ENTITY* "may correspond to a single *OU*" (Uschold, et al, 1997). No further analysis (informal or formal) of the link between these two concepts is given.
- *Implicit context dependencies.* In the EO, the concept *LEGAL ENTITY*, is not well thought out – having several (informally inconsistent) descriptions. It seems that the intended meaning actually depends on a particular jurisdiction (in this case on the current UK jurisdiction) – though it is not clear that the authors recognise this. This dependence is inappropriate in the modern global economy – and it raises potential problems should the UK jurisdiction change. For example, the *LEGAL ENTITY* concept would no longer be the "union of *PERSON*, *CORPORATION*, and *PARTNERSHIP*".

3.3 Unifying the Core Concepts: Person

A vital missing element from both the ontologies is a unifying core category. Part of the synthesis work is to try and incorporate these in preparation for a synthesised common view.

So we introduced the unifying category **PERSON** (*PARTY*), which can be a **NATURAL PERSON** or **SOCIALLY CONSTRUCTED PERSON** (*SOCIAL PERSON* in short). This acts as the catalyst for transforming the ontologies into ones with similar characteristics. The natural next step (which we plan to undertake soon) is to merge them into a single synthesised ontology.

The result of introducing **PERSON** into the EO ontology is shown in Figure. 3. A comparison of this with Figure 1 shows how **PERSON** has unified the taxonomy.

To give the reader some idea of how this transformation was effected, we describe the steps we went through. The EO concepts **PERSON** (human being) and **OU** are renamed **NATURAL PERSON** and **SOCIAL PERSON**. These are then generalised into the concept **PERSON**. **LEGAL ENTITY** is taken completely out and substituted with the context independent notion of **LEGALLY CONSTRUCTED PERSON** (*LEGAL PERSON* in short) – a subtype of **SOCIAL PERSON**.

Note that **LEGAL PERSON** is not the same concept as the EO **LEGAL**

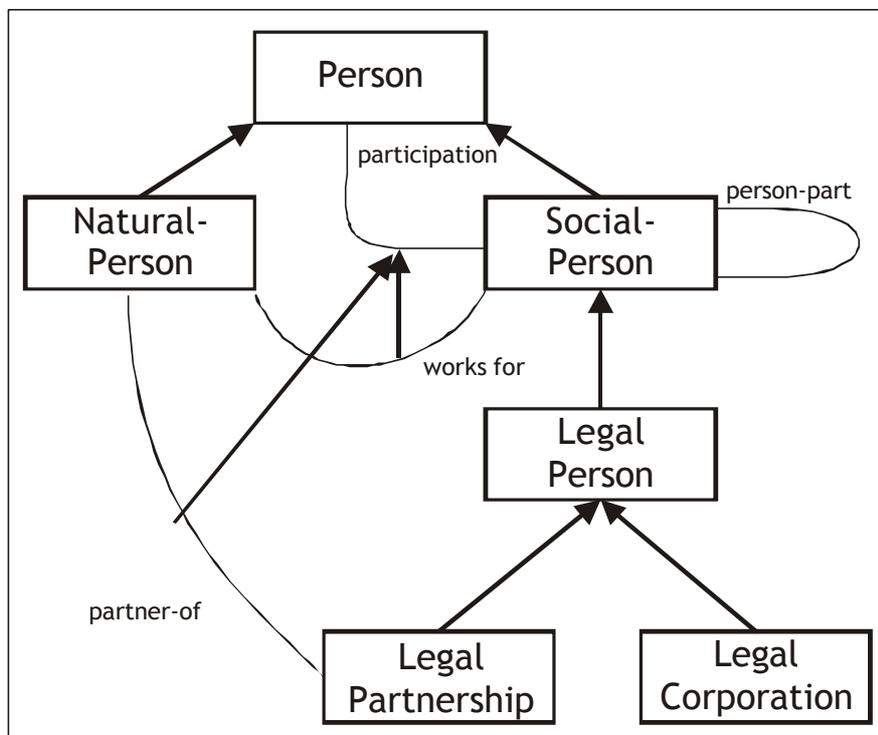


Figure 3. EO transformation

ENTITY, since it is intended to represent parties, which are constructed according to a legal jurisdiction, but not necessarily recognised by it as legal persons (in EO terms, **LEGAL ENTITY**s). For example, in the UK a partnership is

not legally recognized as a person (it cannot sign contracts in its name) but it is a *LEGALLY CONSTRUCTED PERSON*, because there are legal constitution rules for partnerships. Finally the two participation relations, *partner_of* and *works_for* are consolidated under a general *participation* relation, and the relation *manages* is renamed into *person-part* (which is a particular kind of *part_of* relation).

The result of introducing *PERSON* into the TOVE ontology is shown in Figure 4. As before, a comparison of this with Figure 1 shows how *PERSON* has unified the taxonomy. The transformation steps between Figure 2 and Figure 4 are similar in many respects to those between Figures 1 and 2.

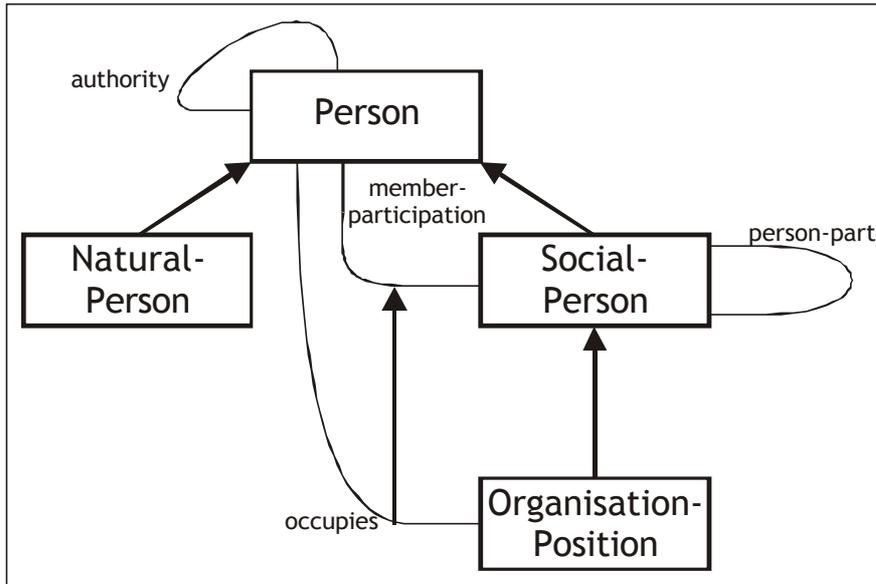


Figure 4. TOVE transformation

4. CONCLUSION

Even at this early stage our work has revealed both the need for a substantial improvement in enterprise ontologies to bring them up to ‘industrial strength’ and a route for doing so. Hopefully, our future work will go some way towards realising this across the whole enterprise ontology.

5. ACKNOWLEDGEMENTS

We would like to thank to the BORO Program and IKF Project for helping to make this research possible. Furthermore we would like to thank to Alessandro Oltramale, Claudio Masolo, and Nicola Guarino for the numerous fruitful discussions we had on topics related to ontologies and organisations.

6. REFERENCES

- [BORO], <http://www.BOROProgram.org>
- [B1997] Breuker J, Valente A., Winkels R. 1997, Legal Ontologies: A Functional View in P.R.S. Visser and R.G.F. Winkels, Proceedings of the First International Workshop on Legal Ontologies,
- [CYC] <http://www.cyc.com/publications.html>
- [EO] <http://www.aiai.ed.ac.uk/project/enterprise/enterprise/ontology.html>
- [EUREKA] http://www3.eureka.be/Home/projectdb/PrjFormFrame.asp?pr_id=2235
- [F1993] Fox, M.S., Chionglo, J., Fadel, F., 1993, *A Common-Sense Model of the Enterprise*, Proceedings of the Industrial Engineering Research Conference.
- [F1996] Fox, M.S., Barbuceanu, M., and Gruninger, M., 1996, *An Organisation Ontology for Enterprise Modelling: Preliminary Concepts for Linking Structure and Behaviour*, Computers in Industry, Vol. 29, pp. 123-134.
- [G1993] Gruber, T., 1993, *Toward Principles for the Design of Ontologies Used for Knowledge Sharing*, in *Formal Ontology in Conceptual Analysis and Knowledge Representation*, Nicola Guarino and Roberto Poli Eds.
- [G1997] Guarino, N., 1997, *Semantic Matching: Formal Ontological Distinctions for Information Organization, Extraction, and Integration*. In M.T. Pazienza (Ed.) *Information Extraction: A Multidisciplinary Approach to an Emerging Information Technology*.
- [H1997] Hay David C., 1997, *Data Model Patterns: Conventions of Thought*, Dorset House.
- Inmon W.H. 1997, *The Data Model Resource Book: A Library of Logical Data and Data Warehouse Models*, John Wiley and Sons.
- [P1996] Partridge, C. 1996, *Business Objects: Re-Engineering for Re-Use*, Butterworth-Heinemann.
- [TOVE] <http://www.eil.utoronto.ca/tove/>
- [U1997] Uschold, M., King, M., Moralee, S., and Zorgios, Y. 1997, *The Enterprise Ontology*, AIAI, The University of Edinburgh,
- [U1998] Uschold, M., King, M., Moralee, S., and Zorgios, Y. 1998, *The Enterprise Ontology*, in *The Knowledge Engineering Review*, Vol. 13, Eds. M. Uschold and A. Tate,