

# Task Neutral Ontologies, Common Sense Ontologies and Legal Information Systems

*T.J.M. Bench-Capon  
Department of Computer Science  
The University of Liverpool  
Liverpool  
UK*

**Abstract:** In this paper I consider the notion of a task neutral ontology intended to represent the common sense needed by any application. If such a thing existed, it could serve as the starting point for the construction of any application ontology. If such a thing does exist, Wordnet is a plausible and often mentioned candidate. I show, however, that Wordnet is not in fact suited to this role. I then argue that nothing could be suited to this role: that any ontology requires a task both to determine its structure and to determine the knowledge that it should contain. In sum, I contend that task neutrality is not a realistic aspiration for ontologies.

## 1. Introduction

It has been suggested that ontologies can provide benefits for legal information systems. Among these suggested benefits are the ability to reuse a knowledge base developed for one task in another task, or, what is perhaps a special case of this, to provide a representation of common sense knowledge which can be incorporated in any legal information system. Against this it has been argued that to build an ontology which is not so abstract as to be of little practical use for an information system, it is necessary to have a clear understanding of the task to be carried out with the represented knowledge and this will have considerable impact on the ontology which results. Certainly many of the methodologies for constructing ontologies take task as a crucial initial starting point (e.g. Jones et al 1988). In this paper I want to consider the argument, so as to address the question: can a suitable task neutral common sense ontology be built for use in legal information systems?

I shall begin by thinking about countries and supposing we are trying to develop an ontology to support a legal information system to address issues relating to movement across borders. The system is supposed to handle both movement of people and movement of currency. Such a system will need to include the concept of a country in its ontology.

## 2. Wordnet: A Task Neutral Ontology?

A popular starting point for constructing ontologies is Wordnet (Fellbaum, 1998). As is well known, Wordnet is a "lexical database of the English language", and is available on-line (<http://www.cogsci.princeton.edu/cgi-bin/webwn>). Essentially it offers a number of synonym sets, organised into a hierarchy so that the hyponyms and hypernyms of each synonym set can be found. Each synonym set is also associated with a brief natural language description. Wordnet is intended to be task neutral and has good coverage, so that one is tempted to think that it should provide a useful starting point for constructing an ontology.

Suppose therefore we turn to Wordnet when we wish to start to construct our ontology of countries. If we find country, we find that the word exists in five senses of which two seem relevant to us:

1. the territory occupied by a nation
2. a politically organised body of people under a single government

We might first think that sense two is what we require, and so retrieve its hyponyms. We receive the following:

Reich, suzerain, sea power, world power, city state, ally

This will probably convince us that we wanted sense 1. For this we get the following hyponyms:

fatherland, buffer state, kingdom, European country, African country, Asian Country, South American country, North American country, and then a list of other countries (such as Cuba, Haiti, Israel, Australia, Seychelles, and Vanuatu) not included in any of the other hyponyms.

This is certainly more promising: it does give us access to the individual countries that we probably expected to form the leaves of our ontology. But is the structure of use to us? Here we have some problems.

- First we find some simple errors of fact. "United Kingdom" is included as a hyponym of "kingdom" which described in the gloss as "a country with a king as head of state". This is an (understandable) mistake, since the current head of state is a queen, We also find that "Great Britain" is given as a synonym of "United Kingdom" whereas in fact this term does not include Northern Ireland: the full title of the UK is "the United Kingdom of Great Britain and Northern Ireland". Such errors are perhaps unimportant: we wanted to acquire a structure from Wordnet, rather than factual information
- But there are structural problems as well. "United Kingdom" is included as a hyponym of "kingdom" , while its constituent parts (England, Scotland, Wales and Northern Ireland) are included as hyponyms of European country. Thus UK is not a sibling of Italy, France and the other countries, which we probably need for our purposes, since it is the UK which controls immigration and currency, rather than its member countries.
- Is the division by continent appropriate? Those of us who regularly travel in Europe know that the broad division applied at immigration points is between European Union and non-European Union countries. (In passing, European Union does not appear in Wordnet, although some older terms such as "European Economic Community" and "Common Market" are given as synonyms and without hyponyms). On entering the US, on the other hand, the division is into the US and rest of the world.
- The divisions of European country are also interesting, but unhelpful. We find the following hyponyms: Scandinavian country, Balkan Country, and then individual European countries. Unfortunately this means that EU countries such as Sweden and France are not siblings, and non EU countries such as France and Switzerland are. Worse, no longer existing countries appear at this level: Flanders appears (described as "a medieval country in northern Europe") as a sibling of Belgium and the Netherlands. East Germany still appears as a republic in north central Europe. Other past countries such as Burgundy do not appear at all. Sometimes former names for countries are treated as synonyms (Ghana and Gold Coast), but sometimes as hyponyms (China, Cathay). One slang term for a country appears ("Blighty" as a hyponym of "England"), whereas others do not (no "Emerald Isle" for "Ireland". Fictional countries such as Ruritania (rightly?) do not appear.

So an initial hierarchy derived from Wordnet might look like Figure 1. Obviously I have not included a full set of hyponyms for each term.

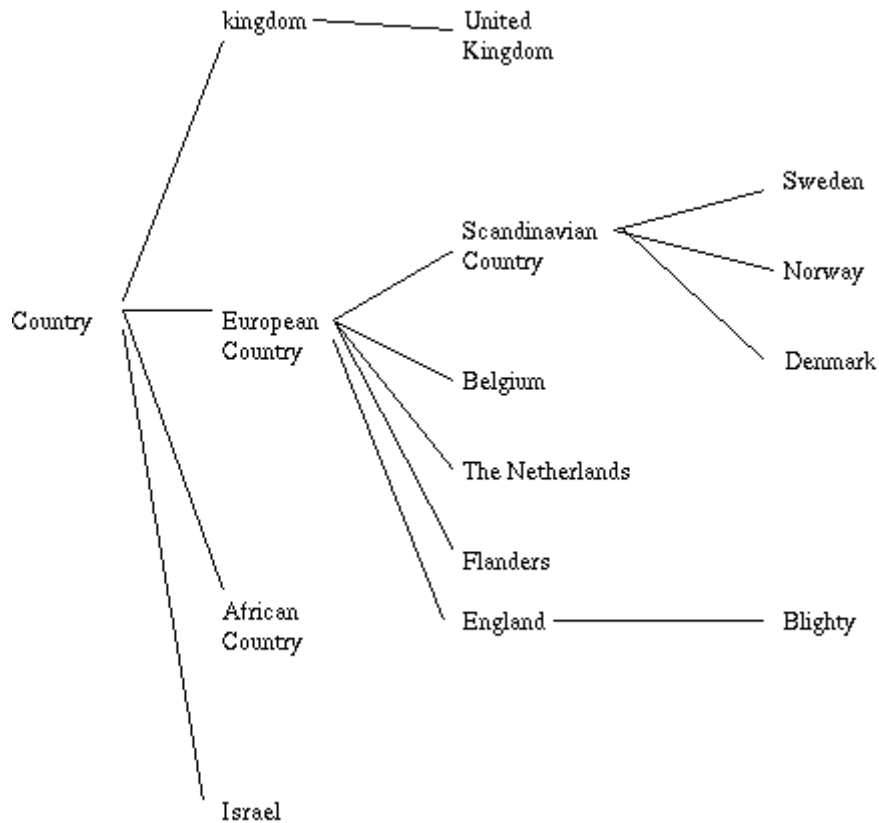


Figure 1: Ontology from Wordnet

### 3. An Ontology for Immigration

Would this assist us in building our ontology for immigration? I think not. What we would need to do is first extract the names of the countries from the hierarchy. This could not be done automatically, since terms we need appear at all three levels below country, and the levels all contain terms we do not want, and require our skill and judgement (aided by the glosses) to select. We would then need to introduce our own intermediate terms, not to be found in Wordnet, to get something like Figure 2.

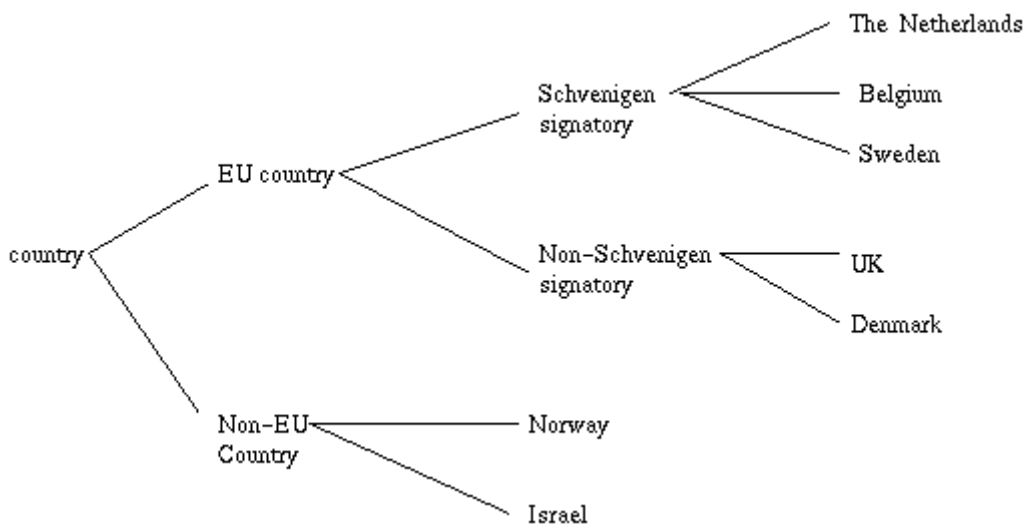


Figure 2: Ontology for Immigration

Note, however, that Figure 2 is directed toward European immigration: other countries would need different intermediate terms. Would Wordnet be the best starting point for this exercise? I think not: a

list of countries from an encyclopaedia would give use want we want more directly. Wordnet contains information we do not need and makes distinctions that are not helpful for our purpose.

#### 4. An Ontology for Currency Control

We now consider this hierarchy with respect to currency. Here we want to classify countries according the currency they use: the Euro countries, the sterling block (should this still exist), the dollar block (if this exists) and so on. Wordnet is again no help for this: it contains many interesting facts on US and UK coinage past and present, and European coins past (the sou and the real are there, but the franc and the peseta are not). But it provides no link from coinage to countries.

But the hierarchy in Figure 2 is of no help either. Probably a non-EU country such as Monaco could (does?) use the Euro, and sterling may be valid in non-EU countries as well. Similarly the Schvenigen distinction is orthogonal to membership of the Euro: the UK may eventually join the Euro, but is unlikely to give up its border controls. So the currency ontology might look like Figure 3. (Obviously the information in Figure 3 does not represent the current situation: but it does represent a possible, if highly improbable, future organisation).

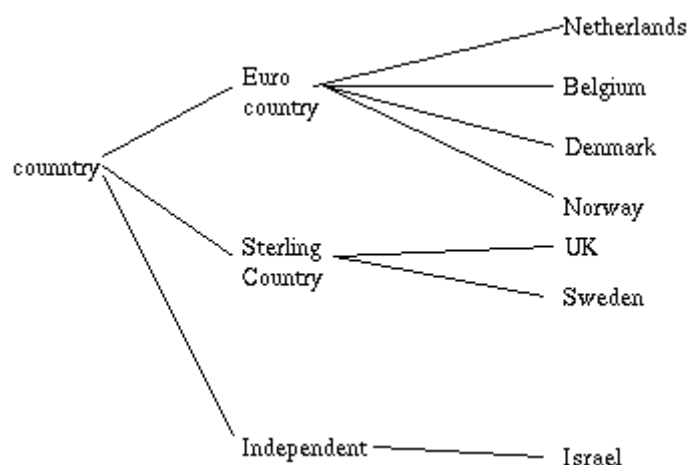


Figure 3: Ontology for Currency Control

#### 5. A Shared Ontology

Figures 2 and 3 represent ontologies customised to two different tasks. Can they be sensibly merged into one task neutral hierarchy? I would say not. We could some scheme involving multiple inheritance that might do the trick theoretically, but would not result in a very agreeable structure for use by an information system. A better way of combining the two might be to flatten the hierarchy into two levels, and use attributes to represent the distinctions that we need for the two tasks. Thus we define a country as having three attributes: *EU-member*; *Schvenigen-Signatory* and *Currency*, and then assign the appropriate values to the individual countries which represent its children.

This ontology would hold the information as the ontologies Figures 2 and 3. But it has two drawbacks if regarded as a task neutral ontology for an information system.

- Part of the attractions of ontologies in the information system context is their similarities with other systems involving inheritance hierarchies. Gains from using inheritance hierarchies are well known (e.g. Bench-Capon 1991). These gains, however, are best realised when the hierarchy is *deep* and has a small *fan-out*. In the modest context of the ontologies in figures 2 and 3, there was some effort to realise these properties. Even Wordnet, which allows considerable fan-out at some stages, attempts to provide this: as for example when it groups countries by continent. The shared ontology above, however, suggest that we move towards a shallow, high fan-out ontology to capture multiple tasks. My contention is that it is the *task* that motivates the distinctions that deepen the hierarchy and control the fan-out.
- The two level ontology can be regarded as an ontology *common* to both the ontology of figure 2 and the ontology of figure 3. The attributes capture the distinctions between the leaf entities which

could be used to derive either of these ontologies. It is not, however, *task neutral*. It may be indifferent to the two tasks which motivated the ontologies of figures 2 and 3, but it is determined by those tasks. A huge number of attributes could have been ascribed to country. For example: wine growing, side of the road driven on, cricket playing, English speaking. The list is, endless. Is there a point in selecting a large, but arbitrary set of possible attributes? I would say not, since when confronted with a task, we may still lack what we need. Since we cannot aspire to completeness, why complicate matters by representing things that just might become useful? And to determine what might become useful is to *anticipate a new task*. I contend that task neutrality is a chimera: the most that can be achieved is an ontology which is even-handed between several, *identified* tasks.

## 6. There is no Such Thing as Task Neutrality

None of this is intended to denigrate Wordnet. Wordnet is a useful resource for what it was intended; namely support for natural language understanding and information retrieval. It is probably more suited for novels originally written in English than for scientific papers, and more suited for English newspaper reports than legal texts. None the less this is not a problem: if I am right that task neutrality is impossible, the distinctions it manifests must inevitably be drawn from some real or imagined task.

What I am criticising is the claim that Wordnet can be used as a suitable seed or starting point for an information system intended to support some particular reasoning activity, such as the application of a certain domain of law. It was never intended as such and cannot be easily used as such. Moreover, if my view on task neutrality is correct, no improvement on Wordnet, nor any other effort is likely to provide this. In knowledge modelling there are no short cuts, few prospects for automation, and limited scope for finding the task already done.

## 7. Concluding Remarks

First I should state that I do believe in ontologies. I think no principled information system can be without one. Only so can the conceptualisation of the domain used in the system be made explicit, which is essential for the confrontation and explication of design choices, and for verification and validation of the resulting product. Also they open the possibility for re-use of knowledge for another particular task, with its own ontology. Such reuse, however, requires the production of a third, common, ontology, and some translation work. What ontologies cannot do is make the knowledge available for effort-free reuse. Nor in my view, can ontologies be sensibly developed except in the context of a particular task. Moreover such tasks need to be rather precisely specified: for example, "assessing claims for immigration", rather than "legal reasoning".

So, let us build ontologies for particular systems as part of our development methodology. But let us not build ontologies not driven by any task, in the hope that they may one day prove to be the answer to every problem that may in a domain.

## References

- Bench-Capon, T.J.M., (1991). Knowledge Representation: An Approach to Artificial Intelligence: Academic Press: London.
- Fellbaum, C., (ed) (1998). WordNet, an electronic lexical database. MIT Press, Cambridge, Mass.
- Jones, D.M., Bench-Capon, T.J.M., and Visser, P.R.S., (1998). Methodologies for Ontology Development. In Cuena, J., (ed) Proceedings of XV IFIP World Computer Congress, IT and KNOWS, pp 62-75.