

A Note on Ontology Localization

Philipp Cimiano^{a,*}, Elena Montiel-Ponsoda^b, Paul Buitelaar^c, Mauricio Espinoza^b and Asunción Gómez-Pérez^b

^a *Semantic Computing Group, Cognitive Interaction Technology Excellence Center (CITEC), Universität Bielefeld, Germany*

E-mail: cimiano@cit-ec.uni-bielefeld.de

^b *Ontology Engineering Group, Universidad Politécnica de Madrid, Spain*

E-mail: {emontiel, jespinoza, asun}@fi.upm.es

^c *Unit for Natural Language Processing, Digital Enterprise Research Institute, National University of Ireland, Galway*

E-mail: paul.buitelaar@deri.org

Abstract. We revisit the notion of ontology localization, propose a new definition and clearly specify the layers of an ontology that can be affected by the process of localizing it. We also work out a number of dimensions that allow to characterize the type of ontology localization performed and to predict the layers that will be affected. Overall our aim is to contribute to a better understanding of the task of localizing an ontology.

Keywords: ontology localization, multilinguality, ontology engineering

1. Introduction

Ontology localization has been defined by Suárez-Figueroa and Gómez-Pérez (2008) as “*the adaptation of an ontology to a particular language and culture*”. In this short note, we propose a more general definition which de-emphasizes the adaptation to a particular language as follows: “*ontology localization is the process of adapting a given ontology to the needs of a certain community, which can be characterized by a common language, a common culture or a certain geo-political environment.*” The adaptation to the language spoken by the target community is thus one possible aspect of ontology localization. In this sense, ontology localization is clearly a specific type of ontology reengineering activity where the requirements (for the reengineering) are provided by the needs of the community to which the ontology is adapted. Ontology localization is thus an activity with very pragmatic goals, i.e. fostering reuse of ontologies already available for the domain in question instead of building them from scratch. Therefore, there is also an economic aspect of ontology localization as it has the potential to reduce costs compared to building a completely new ontology for the target community.

Ontology localization is a transformation process that takes an ontology as input and produces an (adapted) ontology as output, whereby the output can be the same ontology, extended with labels in additional languages, or a new ontology. Ontology localization can affect two different layers: the surface – or *lexical layer* – of an ontology, or the *conceptualization* itself. We consider the lexical layer of an ontology to include all the labels, definitions and accompanying documentation in natural language that make the ontology human understandable. It can already be foreseen that the lexical layer will always undergo modifications, regardless of whether the target community speaks a different language or not. The underlying reason for this is that any discrepancies at the conceptualization layer – due to differences in the cultural or geo-political background – will have an impact on the lexical layer as language is the means by which we experience reality.

*Corresponding author: Philipp Cimiano, Semantic Computing Group, Cognitive Interaction Technology Center of Excellence (CITEC), Universität Bielefeld, 33615 Bielefeld

In this contribution we aim at clarifying the notion of ontology localization as well as the different layers of an ontology it affects. We describe how the localization of the different layers (lexical and conceptual) interact and introduce different dimensions that characterize the localization process and determine by a large extent the ontological layers that are affected by the localization activity. Overall, our goal is to contribute to a better understanding of the ontology localization process.

This short note is structured as follows: in Section 2 we discuss in more detail the different ontology layers (lexical vs. conceptualization) that can be distinguished and how these different layers interact; in Section 3 we then present an overview of different dimensions that can be identified in determining the type of localization that is to be performed and the layers that are affected. Section 4 discusses several real-life cases in which ontology localization has been performed for different purposes. Finally, before concluding, in Section 5 we discuss how ontology localization relates to other areas.

2. Layers in Ontology Localization and their Interaction

2.1. Lexical Layer

The lexical layer of an ontology comprises: i) the labels of the concepts, properties and individuals defined in the ontology, ii) natural language definitions of these entities, as well as iii) the documentation accompanying the ontology, which describes its scope, purpose, usage etc. The inclusion of definitions of ontology entities in natural language is in fact regarded as good practice and even explicitly mentioned as part of the “clarity” criterion that ontologies should fulfill according to Gruber (1995).

Obviously, the lexical layer is language-specific and is thus clearly affected by any ontology localization process, even when the adaptation is done within the same linguistic system. This means that the changes motivated by the cultural environment in which the ontology is to be used – be it within the same linguistic system or not – will be reflected at the lexical layer.

A straightforward way to localize the lexical layer is to provide a 1:1 translation for each label, definition and the accompanying documentation. However, as we will see below, whether a 1:1 translation is suitable ultimately depends on the purpose for which an ontology is localized.

2.2. Conceptualization Layer

While the translation of labels is an important aspect of the ontology localization process, the conceptualization may also need to be adapted if so required for example by a different cultural or geo-political context. Consider an ontology about political functions and charges. Most democratic systems distinguish for example the role of *head of government* in the sense of head of the executive power vs. the role of *head of state* with mainly representative function. An ontology designed to model political functions and charges in Germany would further distinguish between the *Bundeskanzler* (*chancellor*) playing the role of the *head of government* and the *Bundespräsident* playing the role of the *head of state*.

If we want to use an ontology about political charges engineered for the German geo-political and cultural environment in other geo-political environments, e.g., the UK or Spain, we will need to adapt the conceptualization expressed by the ontology. In the case of the UK, we would introduce the class of *prime minister* as *head of government* and the *queen* as *head of state*. In the case of Spain, we would introduce the class of *presidente* (*president*) as *head of government* and the *monarca* (*monarch*) as *head of state*. While one could argue that this adaptation can also be achieved at the lexical level, e.g., by adding additional labels (*prime minister*, *presidente*) for the class *Bundeskanzler* or (*queen*, *monarca*) for the class *Bundespräsident*, this is clearly insufficient as these concepts have different extensions and even intensions. In this case, the adaptation to a different geo-political and cultural reality may require more than a 1:1 translation, i.e. a change to the underlying conceptualization itself.

It is important to emphasize that the adaptation of the conceptualization layer will be primarily driven by the inexistence of conceptual equivalents (or concepts with the same granularity level) in the target

community whenever the final purpose of the ontology is to be equally valid in the source and target communities. If the concept of *Bundeskanzler* serves the function of head of government in the German culture and we aim at reusing the ontology in the UK, we should not translate '*Bundeskanzler*' literally as '*federal chancellor*' merely because the word exists in the English language unless the purpose of the localization is to "paraphrase" in English how the German political structure is organized. We will come back to this in section 3.

2.3. Interaction between Layers

The different layers that we have outlined above do certainly interact in the sense that changes to one layer can not be performed completely independently of changes to the other layers. This means that changes to the conceptualization will be inevitably reflected at the lexical layer, and changes to the lexical layer may also end up influencing the conceptual layer.

- First of all, **changes in the conceptualization will also require the adaptation of the lexical layer** (see the example of *Bundeskanzler* and *prime minister* or *presidente* above). The dependency of changes to the lexicon on changes in the conceptualization is clearly unavoidable if the target ontology should have appropriate labels in the language spoken by the target community.
- Secondly, in some cases where only the lexical layer of the ontology is changed to document the ontology for another language, **unintended shifts in meaning may occur** in case the term chosen as a 1:1 translation has different connotations in the target community than in the source community. An example for this is the English term '*apartment*' understood as a self-contained housing unit that occupies part of a building. The direct translation into Spanish, i.e. '*apartamento*', however, mainly refers to a housing located at holiday resorts which is typically occupied only for a limited period of time. Thus, sticking to literal translations in ontology localization can indeed lead to a phenomenon that has been termed *concept misalignment* (Ajani et al., 2007).

A further example for a meaning shift is the translation of '*Ehe*' in German into '*matrimonio*' in Spanish for instance. While '*matrimonio*' also embraces homosexual marriages in Spain, this is not the case of the German counterpart. To some extent it could be argued that the meaning of the term *matrimonio* has evolved over time or that it has different connotations when used in the legal context vs. common language where it is possibly still mainly used to refer to heterosexual marriages. In any case, this is a further example showing that a straightforward 1:1 translation can indeed have unintended connotations, thus requiring either that the conceptualization is adapted or the meaning of the concept is made explicit by stating whether it embraces both heterosexual and homosexual marriages or not.

- Thirdly, even the **adaptation of the lexical layer might require changes in the conceptualization**. Imagine a geographical ontology designed by speakers of the French language. Speakers of the French language might be more inclined to include the distinction between rivers flowing into the sea ('*fleuve*') and rivers flowing into other rivers ('*rivière*') into their ontology. The fact that this distinction is directly lexicalized in the French language (in contrast to other languages such as English, Spanish or German), makes a French ontology designer prone to include this distinction into the ontology. When localizing this ontology into a different language (say English, Spanish or German), an ontology engineer has two basic choices:
 1. **keeping the distinction** between rivers flowing into the sea and rivers flowing into other rivers in the conceptualization. This means that there will be no direct lexicalization in terms of one designation that can be used as label for each concept, but a paraphrase in the target languages.
 2. **remove this distinction** and keep only the concept of a *river* without distinguishing further between rivers flowing into the sea and rivers flowing into other rivers. In this sense the ontology engineer is de-constructing the original ontology by removing distinctions that result from granularity layers that are not completely shared by the cultures involved. Even in this case it would be possible to keep the terms '*rivière*' and '*fleuve*' - possibly together with a definition by way

of a gloss - in the lexical layer of the ontology. The LIR model for example makes this possible as it was developed for the purpose of capturing cultural differences in a lexical layer which is maintained separately from the ontology (Peters et al., 2007).

The decision will be taken by considering whether the real-world distinction between rivers flowing into the sea and rivers flowing into other rivers is a relevant one considering the applications that the target ontology is supposed to support.

The fact that decisions at different layers are clearly dependent on each other makes ontology localization and in particular the development of semi-automatic techniques supporting it a challenging and non-trivial endeavour.

3. Dimensions of Ontology Localization

Localization implies the existence of an input ontology that is ‘adapted’ to serve the purposes of a different linguistic and/or cultural community. From our viewpoint, localization has an eminently practical importance as it fosters the reuse of already conceptualized knowledge in different linguistic and cultural settings. As already outlined in Section 1, this adaptation may have different implications, that is, different layers of the ontology will be affected by the localization to different extents.

We have identified three crucial dimensions that - to some extent - determine the type of localization to be performed. These dimensions, which have been already outlined by Espinoza et al. (2009b), rely on some of the most well-established and acknowledged translation theories. For more on this we refer the interested reader to Nord (1997). In the following we spell out each of these dimensions:

- **international (standardized) domain vs. culturally influenced domain:** Some domains are clearly “internationalized” or “standardized” as a byproduct of globalization activities driven by the need to exchange data on a global level. This is often the case in very technical domains, e.g., in engineering and finance where standards for processes (e.g. the ISO standards) or for reporting (e.g. XBRL¹ in the financial domain) exist. Other domains are more culturally influenced, e.g. in the public administration of various countries on issues such as taxation, law, political charges etc. The resulting models of the same domain in different communities will necessarily diverge.
- **functional vs. documental localization:** Inspired by Functionalist theories to translation (Nord, 1997) we state that an ontology might be localized with different goals in mind. We might require for example that the target ontology has the **same function** in the target community as the original ontology in the source community. In this case we talk about **functional localization**. Take again the example of ontologies in use within public administration, e.g., ontologies modelling immigration procedures. If we want to port these models to a different geo-political reality, we will need to change the conceptualization to fit the requirements of the target community and to make sure that the ontology can have the same function in applications that the original model had in the source community. Functional localization thus typically implies the creation of a new ontology on the basis of the old one, adapted to the requirements of the target community. In the simplest case the ontology can be reused as is but from a practical point of view this will rarely be the case. In **documental localization** the purpose is only to support the use of the original ontology by members of another (linguistic) community. Let us take again the example of the ontology modelling immigration procedures. In order to make these procedures accessible to an immigrant (a member of a different cultural and linguistic community), we need to document the meaning of these procedures in their language. This does not involve the creation of a completely new ontology, but requires that the meaning of the ontology is documented in a different language.

¹<http://www.xbrl.org>

- **interoperable vs. independent ontology**: One important aspect when reengineering an ontology to meet the needs of a certain target community is how interoperable with respect to each other the new ontology and the original ontology should be. There is clearly a trade-off here between meeting the special needs of the target community and maintaining a certain level of interoperability. If the target ontology should still be used to exchange data between the source and target systems, the changes to the conceptualization should be restricted to those strictly needed to accommodate both cultures in order to ensure a certain degree of **interoperability**. If the target ontology will be used as an **independent** ontology in a functionally equivalent manner (see point above on the functional dimension), then significant changes to the conceptualization are acceptable to meet the needs and capture the specificities of the target community.

Given these different dimensions of the ontology localization task, we can now define which ontology layers will be affected depending on the type of localization that is to be carried out. The following table summarizes this for the international vs. culturally influenced domain and functional vs. documental localization dimensions:

Purpose / Type of domain	International	Culturally-influenced
Functional	n.a.	conceptualization, lexical layer
Documental	lexical layer	lexical layer

According to the above table, the configuration corresponding to an internationalized domain and functional localization is not applicable. The reason for this is that according to our definition the functional localization implies creating a new ontology. In the case of an internationalized domain it is however preferable to share one ontology across national borders instead of developing a different one.

We might however document the meaning of the classes and relations defined in the ontology modelling an internationalized domain in different languages so that it is accessible to speakers of various languages. The localization will thus affect only the lexical layer of the ontology in this case.

In the case of a culturally influenced domain, the main distinguishing criterion is whether the ontology is supposed to be used in a different geo-political and cultural environment, in which case the conceptualization needs to be adapted, or the goal is to allow people with a different cultural and linguistic background to access and use the ontology.

The degree of interoperability affects the case of the functional localization of a culturally influenced domain only. Depending on the degree of interoperability desired (in particular the granularity at which the ontologies need to be interoperable), the conceptualization can change more or less. Imagine that each country has a different system for capturing census information. If we want to be able for example to count the number of citizens in the EU belonging to a certain confession or religion, then the corresponding concepts have to be reused across countries or at least be aligned with each other.

4. Examples of localization projects

Domain type and function of the localized ontology can be combined to result in different scenarios. In the following, we describe real-life use cases of localization projects, illustrating the interplay between the different layers and dimensions discussed.

Use Case 1 (GenomaKB): In the GenomaKB project², terminology experts of the Institute of Applied Linguistics at the Universitat Pompeu Fabra in Barcelona (Spain) created a biomedical knowledge base of the human genome in three languages (Spanish, English and Catalan) to assist terminologists, translators and scientific journalists working in this domain. The starting point was an ontology that models the domain with links to three further models on terminological, textual and factographic information. Domain experts from the three linguistic communities worked together to come up with a common and consensual conceptualization of the domain. Once the ontology was stable, its concepts were linked to the terms in

²<http://genoma.iula.upf.edu:8080/genoma/index.jsp>

English, Spanish and Catalan stored in the terminological module. Here the conceptualization is a good example of what we understand by an *internationalized domain*, reflecting the common view of all the cultures represented in the project. This is a clear example of a localization project on an internationalized domain, affecting only the lexical layer in the ontology. Note that *interoperability* is not an issue here as there is only one ontology shared across countries.

Use Case 2 (New to Holland): The New to Holland project website³ provides access to an ontology-based information portal developed in the Netherlands by the company BeInformed⁴ for the Dutch government with the goal of informing immigrants on their rights and duties, e.g. on the process of applying for an immigration permit. The underlying conceptualization of the New to Holland ontology reflects certain specific characteristics of Dutch immigration procedures that need to be localized in other languages. In this scenario, the ontology is modeling what we have called a *culturally-influenced domain* and the purpose of localization is to *document* specifics of Dutch administration services in several other languages. This is therefore clearly a case of localization for documental purposes, i.e., for the purpose of explaining the meaning of concepts and procedures in the language of target users of applications that build on the adapted ontology.

Use Case 3 (WordNet related projects, e.g. EuroWordNet⁵, Meaning⁶, GlobalWordNet⁷, Kyoto⁸): In the different projects that have been underway since the beginnings of the EuroWordNet project for linking WordNets in different languages to the Princeton English WordNet (Miller et al., 1990), we can find different strategies for the construction of the multilingual WordNets. Although WordNet can not be considered an ontology in a strict sense, we believe that these projects reflect the difficulties of having to perform a *functional localization* of a general lexicon to different target languages. The objective of each lexicon is to capture the specificities and particularities of each language, while maintaining a considerable degree of *interoperability* with the remaining WordNets. Guaranteeing interoperability among WordNets representing *culturally-influenced domains* is not a trivial task. In most cases the creation of the local WordNets consisted of reusing the English WordNet and adapting it to the specific needs of each culture (the so-called Expand Model described by Vossen and Boersma (1999)). This method was followed in the case of the Spanish WordNet in the EuroWordNet project or the Japanese WordNet in the current Kyoto project (Bond et al., 2009; Vossen et al., 2008). As each WordNet is used for the same purposes in each culture (e.g. for NLP tasks), the localization is assumed to be functional. As a consequence, not only the lexical layer of the English WordNet, but also the underlying conceptual structure required modifications to accommodate the specifics of the target cultures. While each WordNet is in principle *independent*, interoperability is achieved by creating a mediator ontology (the so-called Interlingua) with mappings between the different WordNets.

Use Case 4 (European Union Directives): Ajani et al. (2007) present a project aimed at supporting both the translation and drafting of European Union Directives (EUD). They assume that there is one EU ontology representing legal concepts relevant at the EU level as well as ontologies modelling the legal concepts in each of the member countries. This is an interesting use case where we have both an *internationalized domain* as well as national ontologies modelling a *culturally influenced domain* that all need to be synchronized. Lesmo et al. clearly argue that 1:1 translation of terms is not always possible as they might have a different meaning in the context of EU law compared to national law. They present a model - the Legal Taxonomy Syllabus - that allows to i) align the different concepts across the EU and national ontologies, ii) model the ambiguity in interpretation of terms with respect to the different ontologies and iii) capture how a given legal concept in the EU ontology can be expressed in the languages of the different member states.

³<http://www.newtoholland.nl>

⁴see <http://www.beinformed.nl>

⁵<http://www.ilc.uva.nl/EuroWordNet/>

⁶<http://www.lsi.upc.edu/nlp/meaning/>

⁷<http://www.globalwordnet.org/>

⁸<http://www.kyoto-project.eu/>

5. Related Areas

Ontology localization is related to a number of areas. We discuss the relation to the following fields: software localization, thesauri translation and machine translation.

5.1. Software Localization

Localization is by now a core issue for the software industry where it is heavily related to the notion of internationalization, i.e. of developing products which can be commercialized world-wide. Esselink (2003) for example states that “(...) *localization revolves around combining language and technology to produce a product that can cross cultural and language barriers. No more, no less.*” In fact, the analogy to localization in software engineering supports the understanding of ontology localization that we have put forth in this note. As much as the localization of an ontology can affect two layers, i.e. the lexical layer or the conceptualization itself, localization in software engineering can affect the ‘surface’ of a software product or the actual functionality and behavior of the software. Selling software in a certain country (say Germany) requires that the documentation, online help as well as graphical user interfaces are translated into German. This is what we call ‘*surface localization*’ as it does not affect the core of the software in terms of functionality or behavior. This is similar to the “label translation” aspect of the ontology localization process. However, it might well be the case that the functionality and behavior of the software itself has to be changed to comply with the different processes and rules in place in another country. In this case we do actually change the functional core of the software in the same way we change the conceptualization to meet the requirements of a given geo-political and cultural environment. The most obvious difference between software localization and ontology localization is the fact that software and hardware products to be localized have previously undergone a process of “internationalization”. This basically means that those features of products considered “specific to a certain locale” are adapted to support changes or additions already at production time (e.g. the support of international natural language character sets or the addition of functionalities specific to foreign markets). The main reason for this is that software developers saw early on the need to localize their products for international markets, a need that is now arising in the Semantic Web.

5.2. Thesauri Translation

The issue of localization is also crucial in the field of thesauri development. Much as in ontology localization, the goal here is to reuse existing thesauri or create multilingual systems for the purpose of indexing documents across languages. The issue of developing automatic approaches that can reduce the costs in translating thesauri into other languages is also crucial (Lin et al., 2009). In contrast to ontology localization, the goal in thesauri translation is to find a reasonable translation such that the thesaurus can support the same applications (e.g. indexing of documents) across languages.

A crucial difference to ontology localization is the fact that the issue of changing the conceptualization does not arise as severely in thesaurus translation. In thesauri, the semantics of terms is typically neither formally defined in terms of axioms nor are the hierarchical relations between terms interpreted as strictly as in ontologies, where the subclass-of relation is formally interpreted in terms of (extensional) subsumption. Thesauri are thus typically loosely defined structures in comparison to ontologies, so that the question whether a certain translation has unintended meaning shifts is not as relevant.

5.3. Machine Translation

Machine translation deals with the translation of surface strings (sentences, documents etc.) from a source language into a target language. Different approaches to machine translation exist, including i) statistical, ii) transfer-based and iii) interlingua-based approaches. Interlingua-based approaches are based on a language-independent universal representation to which texts in the source language are mapped to and from which a translation in the target language is generated (Mitamura et al., 1991). Interlinguas share with ontologies the fact that they are language-independent representations. However, ontologies are typically domain-specific and do not aim to represent a universal language that can be used to translate textual input from a source language into a target language. Prominent interlinguas include for example the Universal Networking Language (UNL)⁹ or the Mikrokosmos ontology (Mahesh and Nirenburg, 1995).

Different machine translation techniques are applicable to the problem of translating the lexical layer of an ontology into a target language. However, as argued above, changes in the lexical layer and in the conceptual layer depend on each other. Thus, off-the-shelf MT techniques have a restricted use here. The task of translating ontology labels actually needs to take into account the whole context and purpose of the ontology. So far, first tools tackling this challenge have been developed. An example is LabelTranslator (Espinoza et al., 2008, 2009a), a system created with the aim of supporting a semi-automatic localization of ontologies by providing a functional translation of ontology labels. It is released as a plug-in of the ontology editor NeOn Toolkit¹⁰. In its current version it supports the translation from and into English, Spanish and German. The ontological context of each of the labels to be translated is used to find the most appropriate translation candidate in the target language. In a first step, candidate translations are obtained from multilingual lexical resources and/or machine translation web services. Then, translation candidate senses are retrieved from different ontology pools. For this purpose, the system accesses Semantic Web search engines such as Watson to retrieve the senses that correspond to the candidate ontology concepts from different ontologies. In doing this, it not only obtains natural language descriptions of the concept or synonyms, if available, but also its local context. This permits to compare the ontological contexts of candidate translations with the one of the original label and perform a ranking to offer the most appropriate translation. Further, there has been some research on systems to semi-automatically enrich ontologies with natural language information extracted from multilingual linguistic resources (Pazienza and Stellato, 2005) as well as work on integrating ontologies with lexical resources (Oltamari and Stellato, 2008).

6. Conclusion

In this note we have proposed a more general definition of ontology localization as *“the process of adapting a given ontology to the needs of a certain community, which can be characterized by a common language, a common culture or a certain geo-political environment”*. This definition is more general than previous ones in the sense that it emphasizes that adaptation to a specific language is not the only goal and purpose of the localization activity. We have further characterized the task of ontology localization along three dimensions: degree of internationalization, purpose (functional vs. documental) and degree of interoperability. From these dimensions we have derived different types of localization activity which affect the two layers (conceptualization and lexicon) in different ways, having different inputs and outputs. We have also discussed several real-life scenarios corresponding to different configurations along the above mentioned dimensions. Finally, we have argued that changes to the conceptualization and to the lexicon are clearly not independent from each other but interact in a number of unforeseeable ways that need to be balanced by the person, agent or algorithm performing or supporting the localization.

Ontology localization is per se an important and practical activity with high economic impact as it allows to reuse ontologies engineered for a specific linguistic and cultural community to fit the needs of

⁹<http://www.undl.org/>

¹⁰<http://neon-toolkit.org>

a different community. In this sense localization is a special type of re-engineering activity inheriting all the known difficulties involved in the task of engineering an ontology. Specific methodologies to adapt an ontology to a different community characterized by a common language, a common culture or a certain geo-political environment are certainly needed if localization is expected to be performed at a reasonable cost and with high-quality output.

From a general point of view, the methodological support of multiple languages and the localization process in particular is crucial for multinational companies managing knowledge across national borders. This is the case, for instance, for public organisations such as the Food and Agriculture Organization of the United Nations (FAO) or the World Health Organisation (WHO). Therefore, it seems crucial to provide methodological and technological support for the process of localizing an ontology. This is in particular an important goal of the NeOn project, where ontology localization has been integrated into the NeOn methodology as one of the activities to be performed in ontology engineering whenever multilingual ontologies are needed and methodological guidelines have been provided for this purpose (Espinoza et al., 2009b). However, the solutions proposed in NeOn focus on the specific needs of certain use cases and there still remains the challenge of extending not only the methodological but also the technological support to cover a wider range of localization scenarios and use cases.

Overall, the aim of this short note has been to shed light on the notion of ontology localization. Future work will have to be devoted to developing new or adapting existing ontology engineering methodologies to the specifics of the ontology localization task as well as to develop tools to support users in this task. A further challenge to address is to abstract from the specific use cases described in this note in order to yield general interaction patterns involving the different layers. A further question is whether the identified dimensions are rich and exhaustive enough in order to capture the complex matrix of requirements and features needed in ontology localization. Finally, an interesting question is to understand in how far ontology localization might profit from more foundational levels of knowledge organization as aimed for in DOLCE (Masolo et al., 2003) or in upper level ontologies that have been designed with the goal of simplifying the language-ontology interface such as the Penman model (Bateman, 1990).

Acknowledgements: This work is supported in part by the European Union under Grant No. 248458 for the Monnet project¹¹ as well as by the Science Foundation Ireland under Grant No. SFI/08/CE/I1380 (Lion-2). We thank all anonymous reviewers as well as Alessandro Oltramari in particular for relevant comments which have helped to improve the paper.

References

- Ajani, A., Lesmo, L., Boella, G., Mazzei, A., and Rossi, P. (2007). Terminological and ontological analysis of european directives: the multilingualism in law. In *Proceedings of the 11th International Conference on Artificial Intelligence and Law*, pages 43–48.
- Bateman, J. (1990). Upper modeling: A general organization of knowledge for natural language processing. In *Proceedings of the Workshop on Standards for Knowledge Representation Systems*.
- Bond, F., Isahara, Fujita, S., Uchimoto, K., Kuribayashi, T., and Kanzaki, K. (2009). Enhancing the japanese WordNet. In *Proceedings of the 7th Workshop on Asian Language Resources (ALR-7)*, pages 1–9.
- Espinoza, M., Gómez-Pérez, A., and Mena, E. (2008). Enriching an ontology with multilingual information. In *Proceedings of the European Semantic Web Conference (ESWC 2008)*, pages 333–347.
- Espinoza, M., Gómez-Pérez, A., and Montiel-Ponsoda, E. (2009a). Multilingual and localization support for ontologies. In *Proceedings of the European Semantic Web Conference (ESWC 2009)*, pages 821–825.
- Espinoza, M., Montiel-Ponsoda, E., and Gómez-Pérez, A. (2009b). Ontology localization. In *Proceedings of the 5th International Conference on Knowledge Capture (KCAP)*, pages 33–40.
- Esselink, B. (2003). The evolution of localization. Guide to Localization edited by Multilingual Computing and Technology.
- Gruber, T. (1995). Toward principles for the design of ontologies used for knowledge sharing. *International Journal of Human-Computer Studies*, 43(5-6).
- Lin, J., Murray, C., Dorr, B. J., Haji, J., and Pecina, P. (2009). A cost-effective lexical acquisition process for large-scale thesaurus translation. *Language Resources and Evaluation*, 43(1).
- Mahesh, K. and Nirenburg, S. (1995). A situated ontology for practical NLP. In *Proc. Workshop on Basic Ontological Issues in Knowledge Sharing, collocated with the International Joint Conference on Artificial Intelligence (IJCAI-95)*.

¹¹<http://www.monnet-project.eu/>

- Masolo, C., Borgo, S., Gangemi, A., Guarino, N., and Oltramari, A. (2003). Ontology library (final). WonderWeb deliverable D18.
- Miller, G., Beckwith, R., Fellbaum, C., Gross, D., and Miller, K. (1990). Introduction to WordNet: An on-line lexical database. *International Journal of Lexicography*, 3(4):235–244.
- Mitamura, T., Nyberg, E. H., and Carbonell, J. G. (1991). An efficient interlingua translation system for multi-lingual document production. In *Proceedings of the 3rd Machine Translation Summit*.
- Nord, C. (1997). *Translating as a Purposeful Activity. Functionalist Approaches Explained*. Manchester: St. Jerome.
- Oltramari, A. and Stellato, A. (2008). Enriching ontologies with linguistic content: An evaluation framework. In *Proceedings of OntoLex 2008, collocated with LREC*.
- Pazienza, M. and Stellato, A. (2005). The protégé Ontoling plugin - linguistic enrichment of ontologies in the Semantic Web. In *Poster proceedings of the 4th International Semantic Web Conference (ISWC)*.
- Peters, W., Montiel-Ponsoda, E., and de Cea, G. A. (2007). Localizing ontologies in OWL. In *Proceedings of OntoLex 2007 Workshop, collocated with ISWC*.
- Suárez-Figueroa, M. and Gómez-Pérez, A. (2008). First attempt towards a standard glossary of ontology engineering terminology. In *Proceedings of the 8th International Conference on Terminology and Knowledge Engineering (TKE2008)*.
- Vossen, P. Bloksma, L. and Boersma, P. (1999). The dutch WordNet. Technical report, University of Amsterdam.
- Vossen, P., Agirre, E., Calzolari, N., Fellbaum, C., Hsieh, S., Huang, C., Isahara, H., Kanzaki, K., Marchetti, A., Monachini, M., Neri, F., Raffaelli, R., Rigau, G., Tescon, M., and VanGent, J. (2008). Kyoto: a system for mining, structuring and distributing knowledge across languages and cultures. In *Proceedings of the Sixth International Language Resources and Evaluation Conference (LREC'08)*.