ESSOT: an expert supporting system for ontology translation

Title	ESSOT: an expert supporting system for ontology translation
Author(s)	Arcan, Mihael;Dragoni, Mauro;Buitelaar, Paul
Publication Date	2016
Publisher	Springer Verlag
Repository DOI	10.1007/978-3-319-41754-7_6

ESSOT: an Expert Supporting System for Ontology Translation

Mihael Arcan[1] Mauro Dragoni[2] Paul Buitelaar[1]

[1] Insight Centre for Data Analytics, National University of Ireland, Galway
[firstname.lastname]@insight-centre.org
[2] FBK- Fondazione Bruno Kessler, Via Sommarive 18, 38123 Trento, Italy
dragoni@fbk.eu

Abstract. To enable knowledge access across languages, ontologies, mostly represented only in English, need to be translated into different languages. The main challenge in translating ontologies with machine translation is to disambiguate an ontology label with respect to the domain modelled by the ontology itself; however, a crucial requirement is to have translations validated by experts before the ontologies are deployed. Real-world applications have to implement a support system addressing this task to help experts in validating automatically generated translations. In this paper, we present ESSOT, an Expert Supporting System for Ontology Translation. The peculiarity of this system is to exploit the semantic information of the label's context to improve the quality of label translations. The system has been tested within the Organic.Lingua project by translating the modelled ontology in three languages, whereby the results are compared with translations provided by the Microsoft Translator API. The provided results demonstrate the viability of our proposed approach.

1 Introduction

Nowadays, semantically structured data, i.e. ontologies or taxonomies, typically have labels stored in English only. Although the increasing number of ontologies offers an excellent opportunity to link this knowledge together, non-English users may encounter difficulties when using the ontological knowledge represented in English only [1]. Furthermore, applications in information retrieval or knowledge management, using monolingual ontologies are limited to the language in which the ontology labels are stored. Therefore, to make ontological knowledge accessible beyond language borders, these monolingual resources need to be enhanced with multilingual information [2].

Since manual multilingual enhancement of domain-specific ontologies is very time consuming and expensive, we engage a domain-aware statistical machine translation (SMT) system, called OTTO, embedded within the ESSOT system to automatically translate the ontology labels. As ontologies may change over time, having in place an SMT system adaptable to an ontology can therefore be very beneficial. Nevertheless, the quality of the SMT generated translations relies strongly on the translation model learned from the information stored in parallel corpora. In most cases, the inference of translation candidates cannot always be learned accurately when domain-specific vocabulary, like ontology labels, appears infrequent in a parallel corpus. Additionally, ambiguous labels built out of only a few words do not express enough semantic information to guide the SMT system in translating a label correctly in the targeted domain. This can be observed in domain-independent systems, e.g. Microsoft Translator,[3]

[3] http://www.bing.com/translator/

where an ambiguous expression, like *vessel* stored in a medical ontology, is translated as *Schiff*[4] (en. *ship*) in German, but not into the targeted medical domain as *Gefäß*.

In this paper, we present ESSOT with the domain-aware SMT system, called OTTO, integrated into a collaborative knowledge management platform for supporting language experts in the task of translating ontologies. The benefits of such a platform are *(i)* the possibility of having an all-in-one solution, containing both an environment for modelling ontologies which enables the collaboration between different type of experts and *(ii)* a pluggable domain-adaptable service for supporting ontology translations. The proposed solution has been validated in a real-world context, namely Organic.Lingua,[5] from quantitative and qualitative points of view by demonstrating the effort decrease required by the language experts for completing the translation of an entire ontology.

2 Related Work

In this section, we summarize approaches related to ontology translation and present a brief review of the available ontology management tools with a particular emphasis on their capabilities in supporting language experts for translating ontologies.

The task of ontology translation involves generating an appropriate translation for the lexical layer, i.e. labels stored in the ontology. Most of the previous related work focused on accessing existing multilingual lexical resources, like EuroWordNet or IATE [3,4]. This work focused on the identification of the lexical overlap between the ontology and the multilingual resources, which guarantees a high precision but a low recall. Consequently, external translation services like BabelFish, SDL FreeTranslation tool or Google Translate were used to overcome this issue [5,6]. Additionally, [5] and [7] performed ontology label disambiguation, where the ontology structure is used to annotate the labels with their semantic senses. Similarly, [8] show positive effect of different domain adaptation techniques, i.e., using web resources as additional bilingual knowledge, re-scoring translations with Explicit Semantic Analysis, language model adaptation) for automatic ontology translation. Differently to the aforementioned approaches, which rely on external knowledge or services, the machinery implemented in ESSOT is supported by a domain-aware SMT system, which provides adequate translations using the ontology hierarchy and the contextual information of labels in domain-relevant background text data.

Concerning the multilingual ontology management tools, we identified three that may be compared with the capabilities provided by MoKi: *Neon* [9], *VocBench* [10], and *Protégé* [11].

The main features of the *The NeOn toolkit*[6] include the management and the evolution of ontologies in an open, networked environment; the support for collaborative development of networked ontologies; the possibility of using contexts for developing, sharing, adapting and maintaining networked ontologies and an improved human-ontology interaction (i.e. making it easier for users with different levels of expertise and experience to browse and make sense of ontologies).

VocBench[7] is a web-based, multilingual, editing and workflow tool that manages thesauri, authority lists and glossaries using SKOS-XL. Designed to meet the needs of

[4] Translation performed on 2.3.2016
[5] `http://www.organic-lingua.eu`
[6] `http://neon-toolkit.org/wiki/Main_Page`
[7] `http://vocbench.uniroma2.it/`

semantic web and linked data environments, VocBench provides tools and functionalities that facilitate both collaborative editing and multilingual terminology. It also includes administration and group management features that permit flexible roles for maintenance, validation and publication.

Protégé[8] is a free, open source visual ontology editor and knowledge-base framework. The Protégé platform supports two main ways of modelling ontologies via the Protégé-Frames and Protégé-OWL editors. Protégé ontologies can be exported into a variety of formats including RDF(S), OWL, and XML Schema.

While the first two, *Neon* and *VocBench*, are the ones more oriented for supporting the management of multilinguality in ontologies by including dedicated mechanisms for modelling the multilingual fashion of each concept; the support for multilinguality provided by *Protégé* is restricted to the sole description of the labels. However, differently from MoKi, none of them implements the capability of connecting the tool to an external machine translation system for suggesting translations automatically.

3 The Organic.Lingua Project

Organic.Lingua is an EU-funded project that aims at providing automated multilingual services and tools facilitating the discovery, retrieval, exploitation and extension of digital educational content related to Organic Agriculture and AgroEcology. More concretely, the project aims at providing, on top of a web portal, cross-lingual facility services enabling users to *(i)* find resources in languages different from the ones in which the query has been formulated and/or the resource described (e.g., providing services for cross-lingual retrieval); *(ii)* manage meta-data information for resources in different languages (e.g., offering automated meta-data translation services); and *(iii)* contribute to evolving content (e.g., providing services supporting the users in content generation).

These objectives are reached in the Organic.Lingua project by means of two components: on the one hand, a web portal offering software components and linguistic resources able to provide multilingual services and, on the other hand, a conceptual model (formalized in the "Organic.Lingua ontology") used for managing information associated with the resources provided to the final users and shared with other components deployed on the Organic.Lingua platform. In a nutshell, the usage of the Organic.Lingua ontology is twofold: *(i)* resource annotation (each time a content provider inserts a resource in the repository, the resource is annotated with one or more concepts extracted from the ontology) and *(ii)* resource retrieval (when web users perform queries on the system, the ontology is used, by the back-end information retrieval system, to perform advanced searches based on semantic techniques). Due to this intensive use of the ontology in the entire Organic.Lingua portal, the accuracy of the linguistic layer, represented by the set of translated labels, is crucial for supporting the annotation and retrieval functionalities.

4 Machine Translation for Ontology Translation

Due to the shortness of ontology labels, there is a lack of contextual information, which can otherwise help disambiguating expressions. Therefore, our goal is to translate the identified ontology labels within the textual context of the targeted domain, rather than

[8] http://protege.stanford.edu/

in isolation. To identify the most domain-specific source sentences containing the label to be translated we engage the OnTology TranslatiOn System, called OTTO[9] [12]. With this approach, we aim to retain relevant sentences, where the English label *vessel* belongs to the medical domain, but not to the technical domain, which would cause a wrong, out-of-domain translation. This process reduces the semantic noise in the translation process, since we try to avoid contextual information that does not belong to the domain of the targeted ontology.

Statistical Machine Translation Our approach is based on statistical machine translation, where we wish to find the best translation $\mathbf{e}$, of a string $\mathbf{f}$, given by a log-linear model combining a set of features. The translation that maximizes the score of the log-linear model is obtained by searching all possible translations candidates. The decoder, which is essentially a search procedure, provides the most probable translation based on a statistical translation model learned from the training data.

For a broader domain coverage of an SMT system, we merged several parallel corpora necessary to train an SMT system, e.g. JRC-Acquis [13], Europarl [14], DGT (translation memories generated by the *Directorate-General for Translation*) [15], MultiUN corpus [16] and TED talks [17] among others, into one parallel dataset. For the translation approach, the OTTO System engages the widely used Moses toolkit [18]. Word alignments were built with GIZA++ [19] and a 5-gram language model was built with KenLM [20].

Relevant Sentence Selection In order to translate an ontology label in the closest domain-specific contextual environment, we identify within the concatenated corpus only those source sentences, which are most relevant to the labels to be translated. Nevertheless, due to the specificity of the ontology labels, just an *n-gram overlap* approach is not sufficient to select all the useful sentences. For this reason, we follow the idea of [21], where the authors extend the semantic information of ontology labels using Word2Vec [22] for computing distributed representations of words. The technique is based on a neural network that analyses the textual data provided as input and outputs a list of semantically related words. Each input string, in our experiment ontology labels or source sentences, is vectorized using the surrounding context and compared to other vectorized sets of words in a multi-dimensional vector space.Word relatedness is measured through the cosine similarity between two word vectors.

The usage of the ontology hierarchy allows us to further improve the disambiguation of short labels, i.e., the related words of a label are concatenated with the related words of its direct parent. Given a label and a source sentence from the concatenated corpus, related words and their weights are extracted from both of them, and used as entries of the vectors to calculate the cosine similarity. Finally, the most similar source sentence and the label should share the largest number of related words.

OTTO Service in Action The OTTO service[10] works as a pipeline of tasks. When a user invokes the translation service, all labels contained in the ontology context (where as "context" we mean the set of concepts that are connected directly or with a maximum distance of N arcs with the concept that has to be translated) are extracted from the ontology and stored in the message that is sent to the OTTO service (left part in Figure 1).

[9] http://server1.nlp.insight-centre.org/otto/

[10] for more information how to invoke the service, see also: http://server1.nlp.insight-centre.org/otto/rest_service.html

When the service receives the translation request, the service looks into the model for the best candidate translations by considering the contextual information accompanying the ontology label to translate. A ranked list based on log probabilities of candidate translations within the JSON output format (Figure 1) is generated from the OTTO service and sent back to the user that will select, among the proposed translations, the one to save in the ontology.


```
{
   "label2translate":"vessels",
   "concept_context":[
      "blood",
      "medical",
      "disease",
      "biomedical",
   ],
   "translate2":"de"
}
```

```
{
   "possible_translations":{
      "blutgefäßen":-15.8438,
      "gefäßen":-2.4100,
      "halsgefäße":-2.6682
   },
   "time":"24 wallclock secs",
   "source_label":"vessels",
   "best":"gefäßen"
}
```


Fig. 1: JSON representations provided to and from the OTTO system.

5 Supporting The Ontology Translation Activity With MoKi

The translation component described in the previous section has been integrated in a collaborative knowledge management tool called MoKi[11] [23]. It is a collaborative MediaWiki-based[12] tool for modelling ontological and procedural knowledge in an integrated manner[13] and is grounded on three main pillars:

- each basic entity of the ontology (i.e., concepts, object and datatype properties, and individuals) is associated with a wiki page;
- each wiki page describes an entity by means of both unstructured (e.g., free text, images) and structured (e.g. OWL axioms) content;
- multi-mode access to the page content is provided to support easy usage by users with different skills and competencies.

In order to meet the needs of the specific ontology translation task within the Organic.Lingua project, MoKi has been customized with additional facilities: *(i)* connection with the OTTO service that is in charge of providing the translations of labels and descriptions associated with the ontology entities; and *(ii)* user-friendly collaborative features specifically targeting linguistic issues. Translating domain-specific ontologies, in fact, demands that experts discuss and reach an agreement not only with respect to modelling choices, but also to (automated) ontology label translations.

Below, we present the list of the implemented facilities specifically designed for supporting the management of the multilingual layer of the Organic.Lingua ontology.[14]

[11] http://moki.fbk.eu

[12] Wikimedia Foundation/Mediawiki: http://www.mediawiki.org

[13] Though MoKi allows to model both ontological and procedural knowledge, here we will limit our description only to the features for building multilingual ontologies.

[14] A read-only version, but with all functionalities available, of the MoKi instance described in this paper is available at https://dkmtools.fbk.eu/moki/3_5/essot/

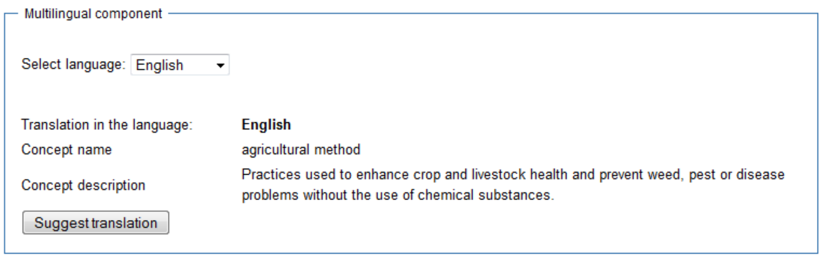

Fig. 2: Multilingual box for facilitating the entity translation

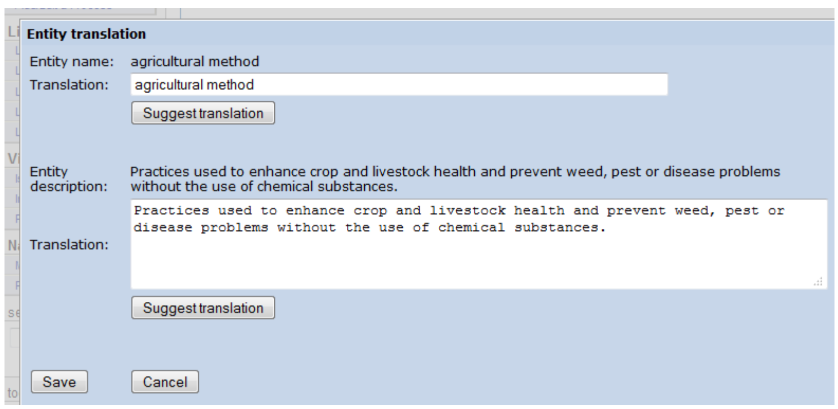

Fig. 3: Quick translation box for editing entities translations

Domain And Language Experts View The semi-structured access mode, dedicated to the Domain and Language Experts, has been equipped with functionalities that permits revisions of the linguistic layer. This set of functionalities permits to revise the translation of names and descriptions of each entity (concepts, individuals, and properties).

For browsing and editing of the translations, a quick view box has been inserted into the mask (as shown in Figure 2); in this way, language experts are able to navigate through the available translations and, eventually, invoke the translation service for retrieving a suggestion or, alternatively, to edit the translation by themselves (Figure 3).

Approval And Discussion Facilities Given the complexity of translating domain specific ontologies, translations often need to be checked and agreed upon by a community of experts. This is especially true when ontologies are used to represent terminological standards which need to be carefully discussed and evaluated. To support this collaborative activity we foresee the usage of the wiki-style features of MoKi, expanded with the possibility of assigning specific translations of ontology labels to specific experts who need to monitor, check, and approve the suggested translations. This customization promotes the management of the changes carried out on the ontology (in both layers) by providing the facilities necessary to manage the life-cycle of each change.

These facilities may be split in two different sets of features. The first group may be considered as a monitor of the activities performed on each entity page. When changes are committed, approval requests are created. They contain the identification of the expert in charge of approving the change, the date on which the change has been performed, and a natural language description of the change. Moreover, a mechanism for

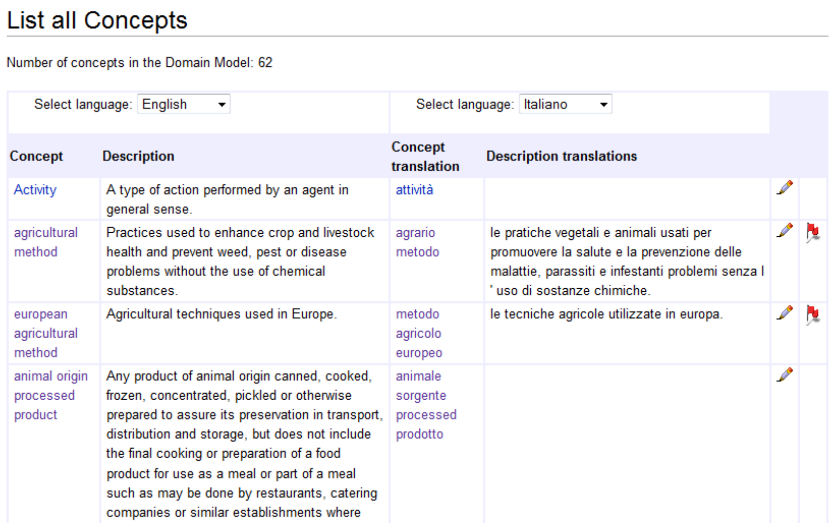

List all Concepts

Number of concepts in the Domain Model: 62

Select language: English | Select language: Italiano

Concept	Description	Concept translation	Description translations
Activity	A type of action performed by an agent in general sense.	attività	
agricultural method	Practices used to enhance crop and livestock health and prevent weed, pest or disease problems without the use of chemical substances.	agrario metodo	le pratiche vegetali e animali usati per promuovere la salute e la prevenzione delle malattie, parassiti e infestanti problemi senza l ' uso di sostanze chimiche.
european agricultural method	Agricultural techniques used in Europe.	metodo agricolo europeo	le tecniche agricole utilizzate in europa.
animal origin processed product	Any product of animal origin canned, cooked, frozen, concentrated, pickled or otherwise prepared to assure its preservation in transport, distribution and storage, but does not include the final cooking or preparation of a food product for use as a meal or part of a meal such as may be done by restaurants, catering companies or similar establishments where	animale sorgente processed prodotto	

Fig. 4: View for comparing entities translations

managing the approvals and for maintaining the history of all approval requests for each entity is provided. Instead, the second set contains the facilities for managing the discussions associated with each entity page. A user interface for creating the discussions has been implemented together with a notification procedure that alerts users when new topics/replies, related to the discussions that they are following, have been posted.

"Quick" Translation Feature For facilitating the work of language experts, we have implemented the possibility of comparing side-by-side two lists of translations. This way, the language expert in charge of revising the translations, avoiding to navigate among the entity pages, is able to speed-up the revision process.

Figure 4 shows such a view, by presenting the list of English concepts with their translations into Italian. At the right of each element of the table, a link is placed allowing to invoke a quick translation box (as shown in Figure 3) that gives the opportunity to quickly modify information without opening the entity page. Finally, in the last column, a flag is placed indicating that changes have been performed on that concept, and a revision/approval is requested.

Interface And Ontology Multilingual Facilities In order to complete the set of features available for managing the multilingual aspects of the Organic.Lingua project, MoKi has been equipped with two further components that permit to switch between the languages available for the tool interface: to add a new language to the ontology and to select the language used for showing the ontology in different views.

Through these facilities, it is also possible to add a new language to the MoKi interface and to manage the translation of its labels. This module has been implemented on top of the multilingual features of MediaWiki.

Concerning the ontology, when new labels are added to the ontology, the OTTO service described in Section 4 is invoked for retrieving the translations related to its labels and descriptions. Finally, the Ontology Export functionality has been revisited by adding the possibility to choose the export languages, among the available ones.

6 Evaluation

Our goal is to evaluate the usage and the usefulness of the MoKi tools and of the underlying service for suggesting domain-adapted translations.

In detail, we are interested in answering two main research questions:

RQ1 Does the proposed system provide an *effective* support, in terms of the quality of suggested translations, to the management of multilingual ontologies?

RQ2 Do the MoKi functionalities provide an *effective* support to the collaborative management of a multilingual ontology?

In order to answer these questions, we performed two types of analysis:

1. Quantitative: we collected objective measures concerning the effectiveness of the translations suggested by the embedded machine translation service. This information allows to have an estimation of the effort needed for adapting all translations by the language experts.
2. Qualitative: we collected subjective judgements from the language experts involved in the evaluation of the tool on general usability of the components and to provide feedback for future improvements.

Six language experts have been involved in the evaluation of the proposed platform for translating the Organic.Lingua ontology in three different languages: German, Spanish, and Italian. Most of the experts had no previous knowledge of the tool, hence an initial phase of training was necessary.

After the initial training, experts were asked to translate the ontology in the three languages mentioned above. Experts used MoKi facilities for completing the translation task and, at the end, they provided feedback about tool support for accomplishing the task. A summary of these findings and lessons learned are presented in Section 6.2.

6.1 Quantitative Evaluation Results

The automatic evaluation on label translations provided by OTTO is based on the correspondence between the SMT output of OTTO and reference translations (gold standard), provided by domain and language experts. For the automatic evaluation we used the BLEU [24], METEOR [25] and TER [26] algorithms.

BLEU is calculated for individual translated segments (n-grams) by comparing them with reference translations. Those scores, between 0 and 1 (perfect translation), are then averaged over the whole *evaluation dataset* to reach an estimate the automatically generated translation's overall quality. METEOR is based on the harmonic mean of precision and recall, whereby recall is weighted higher than precision. Along with standard exact word (or phrase) matching it has additional features, i.e. stemming, paraphrasing and synonymy matching. Differently to BLEU, the metric produces good correlation with human judgement at the sentence or segment level. TER is an error metric (lowers scores are better) for machine translation measuring the number of edits required to change a system output into one of the references.

We evaluate the automatically generated translations into German, Italian and Spanish provided by OTTO and the Microsoft Translator API. Since reference translations are needed to evaluate automatically generated translations, we use the translated labels provided by the domain experts. The Organic.Lingua ontology provides 274 German,

	English → German			English → Italian			English → Spanish		
System	BLEU	METEOR	TER	BLEU	METEOR	TER	BLEU	METEOR	TER
Microsoft	0.037	0.196	**0.951**	**0.135**	0.286	0.871	0.210	0.369	0.733
OTTO	**0.074**	**0.310**	0.991	0.130	**0.342**	**0.788**	**0.257**	**0.444**	**0.667**
Microsoft n-best	**0.076**	0.279	**0.872**	0.145	0.328	0.829	0.274	0.402	0.657
OTTO n-best	0.074	**0.310**	1.000	**0.150**	**0.408**	**0.719**	**0.333**	**0.523**	**0.566**

Table 1: Automatic translation evaluation of the Organic.Lingua ontology by the Microsoft Translator API and OTTO System (bold results = best performance)

354 Italian and 355 Spanish existing translations out of 404 English labels. As seen in Table 1, with the help of contextual information OTTO significantly outperforms (p-value < 0.05) Microsoft Translator API when translating English labels into German (51.3% averaged improvement over all metrics) or Spanish (51.7%) and produces comparable results when translating into Italian (10.5%).

Since both translation systems can provide additional, less probable translations of an English label, we identify with METEOR the best translation (due the gold standard) out of a set of possible translations. In this setting, the Microsoft Translator API provided on average 1.2 translations per label for German and Italian, and 1.6 for the Spanish, respectively. The OTTO system provided 9.5, 10.4 and 8.9 possible translations for German, Italian and Spanish. This additional information, seen as OTTO n-best and Microsoft n-best in the last part in Table 1 allows us to provide better translation candidates of labels. Compared to the first-best translation evaluation (upper part in Table 1), the translation quality improves for both systems. In the n-best scenario, OTTO demonstrates a 13.7% averaged improvement in terms of the evaluation metrics over Microsoft and 14.4% over OTTO first-best scenario when translating into Italian. For Spanish the improvements are 21.8% and 20.8%, respectively. Only for German (-2.1%; -4.9%) Microsoft Translator performs better.

6.2 Qualitative Evaluation Results

To investigate the subjective perception of the six experts about the support provided for translating ontologies, we analysed the subjective data collected through a questionnaire. For each functionality described in Section 5, we provide the information how often each aspect has been raised by the language experts.

Language Experts View
Pros: Easy to use for managing translations (3)
Usable interface for showing concept translations (2)

Approval And Discussion
Pros: Pending approvals give a clear situation about concept status (4)
Cons: Discussion masks are not very useful (5)

Quick Translation Feature
Pros: Best facility for translating concepts (5)
Cons: Interface design improvable (2)

The results show, in general, a good perception of the implemented functionalities, in particular concerning the procedure of translating a concept by exploiting the quick translation feature. Indeed, 5 out of 6 experts reported advantages on using this capability. Similar opinions have been collected about the language expert view, where the users perceived such a facility as a usable reference for having the big picture about the status of concept translations.

Controversial results are reported concerning the approach and discussion facility. On the one hand, the experts perceived positively the solution of listing approval requests on top of each concept page. On the other hand, we received negative opinions by almost all experts (5 out of 6) about the usability of discussion forms. This result shows us to focus future effort in improving this aspect of the tool.

Finally, concerning the "quick" translation facility, 5 out of 6 experts judged this facility as the most usable way for translating a concept. The main characteristic that has been highlighted is the possibility of performing a "mass-translation" activity without opening the page of each concept, with the positive consequence of saving a lot of time.

6.3 Findings and Lessons Learned

The quantitative and qualitative results demonstrate the viability of the proposed platform in real-world scenarios and, in particular, its effectiveness in the proposed use case. Therefore, we can positively answer to both research questions, **RQ1**: the back-end component provides helpful suggestions for performing the ontology translation task, and **RQ2**: the provided interfaces are usable and useful for supporting the language experts in the translation activity. Besides these, there were other insights, either positive and negative, that emerged during the subjective evaluation that we conducted.

The main positive aspect highlighted by the experts was related to the easy and quick way of translating a concept with respect to other available knowledge management tools (see details in Section 2), which do not enable specific support for translation. The suggestion-based service allowed effective suggestions and reduced effort required for finalizing the translation of the ontology. However, even if on one hand, the experts perceived such a service very helpful from the point of view of domain experts (i.e. experts that are generally in charge of modelling ontologies, but that might not have enough linguistic expertise for translating label properly with respect to the domain), the facilities supporting a direct interaction with language experts (i.e. discussion form) should be more intuitive, for instance as the approval one.

The criticism concerning the interface design was reported also about the quick translation feature, where some of the experts commented that the comparative view might be improved from the graphical point of view. In particular, they suggested (i) to highlight translations that have to be revised, instead of using a flag, and (ii) to publish only the concept label instead of putting also the full description in order to avoid misalignments in the visualization of information.

Connected to the quick translation facility, experts judged it as the easiest way for executing a first round of translations. Indeed, by using the provided translation box, experts are able to translate concept information without navigating to the concept page and by avoiding a reload of the concepts list after the storing of each change carried out by the concept translation.

Finally, we can judge the proposed platform as a useful service for supporting the ontology translation task, especially in a collaborative environment when the multilingual ontology is created by two different types of experts: domain experts and language experts. Future work in this direction will focus on the usability aspects of the tool and on the improvement of the semantic model used for suggesting translations in order to further reduce the effort of the language experts. We plan also to extend the evaluation on other use cases.

7 Conclusions

This paper presents ESSOT, an Expert Supporting System for Ontology Translation implementing an automatic translation approach based on the enrichment of the text to translate with semantically structured data, i.e. ontologies or taxonomies. The ESSOT system integrates the OTTO domain-adaptable semantic translation service and the MoKi collaborative knowledge management tool for supporting language experts in the ontology translation activity. The platform has been concretely used in the context of the Organic.Lingua EU project by demonstrating the effectiveness in the quality of the suggested translations and in the usefulness from the language experts point of view.

Acknowledgement

This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289 and the European Unions Horizon 2020 programme MixedEmotions (Grant Number 644632).

References

1. Gómez-Pérez, A., Vila-Suero, D., Montiel-Ponsoda, E., Gracia, J., Aguado-de Cea, G.: Guidelines for multilingual linked data. In: Proceedings of the 3rd International Conference on Web Intelligence, Mining and Semantics, ACM (2013)
2. Gracia, J., Montiel-Ponsoda, E., Cimiano, P., Gómez-Pérez, A., Buitelaar, P., McCrae, J.: Challenges for the multilingual web of data. Web Semantics: Science, Services and Agents on the World Wide Web **11** (2012) 63–71
3. Cimiano, P., Montiel-Ponsoda, E., Buitelaar, P., Espinoza, M., Gómez-Pérez, A.: A note on ontology localization. Appl. Ontol. **5**(2) (2010) 127–137
4. Declerck, T., Pérez, A.G., Vela, O., Gantner, Z., Manzano, D.: Multilingual lexical semantic resources for ontology translation. In: Proceedings of the 5th International Conference on Language Resources and Evaluation, Saarbrücken, Germany (2006)
5. Espinoza, M., Montiel-Ponsoda, E., Gómez-Pérez, A.: Ontology localization. In: Proceedings of the Fifth International Conference on Knowledge Capture, NY, USA (2009)
6. Fu, B., Brennan, R., O'Sullivan, D.: Cross-lingual ontology mapping - an investigation of the impact of machine translation. In Gómez-Pérez, A., Yu, Y., Ding, Y., eds.: ASWC. Volume 5926 of Lecture Notes in Computer Science., Springer (2009)
7. McCrae, J., Espinoza, M., Montiel-Ponsoda, E., Aguado-de Cea, G., Cimiano, P.: Combining statistical and semantic approaches to the translation of ontologies and taxonomies. In: Fifth workshop on Syntax, Structure and Semantics in Statistical Translation (SSST-5). (2011)
8. McCrae, J.P., Arcan, M., Asooja, K., Gracia, J., Buitelaar, P., Cimiano, P.: Domain adaptation for ontology localization. Web Semantics: Science, Services and Agents on the World Wide Web **36** (2015) 23–31
9. Espinoza, M., Gómez-Pérez, A., Mena, E.: Enriching an ontology with multilingual information. In: Proceedings of the 5th European semantic web conference (ESWC'08). ESWC'08, Springer (2008) 333–347

10. Stellato, A., Rajbhandari, S., Turbati, A., Fiorelli, M., Caracciolo, C., Lorenzetti, T., Keizer, J., Pazienza, M.T.: Vocbench: A web application for collaborative development of multilingual thesauri. In Gandon, F., Sabou, M., Sack, H., d'Amato, C., Cudré-Mauroux, P., Zimmermann, A., eds.: The Semantic Web. Latest Advances and New Domains - 12th European Semantic Web Conference, ESWC 2015, Portoroz, Slovenia, May 31 - June 4, 2015. Proceedings. Volume 9088 of Lecture Notes in Computer Science., Springer (2015) 38–53
11. Gennari, J., Musen, M., Fergerson, R., Grosso, W., Crubézy, M., Eriksson, H., Noy, N., Tu, S.: The evolution of protégé: an environment for knowledge-based systems development. Int. J. Hum.-Comput. Stud. **58**(1) (2003) 89–123
12. Arcan, M., Asooja, K., Ziad, H., Buitelaar, P.: Otto – ontology translation system. In: ISWC 2015 Posters & Demonstrations Track. Volume 1486., Bethlehem (PA), USA (2015)
13. Steinberger, R., Pouliquen, B., Widiger, A., Ignat, C., Erjavec, T., Tufis, D., Varga, D.: The JRC-Acquis: A multilingual aligned parallel corpus with 20+ languages. In: Proceedings of the 5th International Conference on Language Resources and Evaluation (LREC). (2006)
14. Koehn, P.: Europarl: A Parallel Corpus for Statistical Machine Translation. In: Conference Proceedings: the tenth Machine Translation Summit, AAMT (2005)
15. Steinberger, R., Ebrahim, M., Poulis, A., Carrasco-Benitez, M., Schlüter, P., Przybyszewski, M., Gilbro, S.: An overview of the european union's highly multilingual parallel corpora. Language Resources and Evaluation **48**(4) (2014) 679–707
16. Eisele, A., Chen, Y.: Multiun: A multilingual corpus from united nation documents. In Tapias, D., Rosner, M., Piperidis, S., Odjik, J., Mariani, J., Maegaard, B., Choukri, K., Chair), N.C.C., eds.: Proceedings of the Seventh conference on International Language Resources and Evaluation, European Language Resources Association (ELRA) (5 2010) 2868–2872
17. Cettolo, M., Girardi, C., Federico, M.: Wit^3: Web inventory of transcribed and translated talks. In: Proceedings of the 16^{th} Conference of the European Association for Machine Translation (EAMT), Trento, Italy (May 2012) 261–268
18. Koehn, P., Hoang, H., Birch, A., Callison-Burch, C., Federico, M., Bertoldi, N., Cowan, B., Shen, W., Moran, C., Zens, R., et al.: Moses: Open source toolkit for statistical machine translation. In: Proceedings of the 45th Annual Meeting of the ACL on Interactive Poster and Demonstration Sessions, Association for Computational Linguistics (2007) 177–180
19. Och, F.J., Ney, H.: A systematic comparison of various statistical alignment models. Computational Linguistics **29** (2003) 19–51
20. Heafield, K.: KenLM: faster and smaller language model queries. In: Proceedings of the EMNLP 2011 Sixth Workshop on Statistical Machine Translation, Edinburgh, Scotland, United Kingdom (July 2011) 187–197
21. Arcan, M., Turchi, M., Buitelaar, P.: Knowledge portability with semantic expansion of ontology labels. In: Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics, Beijing, China (July 2015)
22. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. ICLR Workshop (2013)
23. Dragoni, M., Bosca, A., Casu, M., Rexha, A.: Modeling, managing, exposing, and linking ontologies with a wiki-based tool. In Calzolari, N., Choukri, K., Declerck, T., Loftsson, H., Maegaard, B., Mariani, J., Moreno, A., Odijk, J., Piperidis, S., eds.: Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC-2014), Reykjavik, Iceland, May 26-31, 2014., European Language Resources Association (ELRA) (2014) 1668–1675
24. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: BLEU: a method for automatic evaluation of machine translation. In: Proceedings of the 40th Annual Meeting on Association for Computational Linguistics. ACL '02 (2002) 311–318
25. Denkowski, M., Lavie, A.: Meteor universal: Language specific translation evaluation for any target language. In: Proceedings of the EACL 2014 Workshop on Statistical Machine Translation. (2014)
26. Snover, M., Dorr, B., Schwartz, R., Micciulla, L., Makhoul, J.: A study of translation edit rate with targeted human annotation. In: Proceedings of Association for Machine Translation in the Americas. (2006)