

Translating Ontologies in Real-World Settings

Mihael Arcan¹, Mauro Dragoni²(✉), and Paul Buitelaar¹

¹ Insight Centre for Data Analytics, National University of Ireland, Galway, Ireland

{Mihael.Arcan,Paul.Buitelaar}@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 that are often represented only in English, need to be translated into different languages. The main challenge in translating ontologies is to disambiguate an ontology label with respect to the domain modelled by ontology itself. Machine translation services may help in this task; however, a crucial requirement is to have translations validated by experts before the ontologies are deployed. For this reason, real-world applications must implement a support system addressing this task to relieve experts in validating all translations. In this paper we present the Expert Supporting System for Ontology Translation, called ESSOT, which exploits the semantic information of the label's context for improving the quality of label translations. The system has been tested within the Organic.Lingua project by translating the ontology labels in three languages. In order to evaluate further the effectiveness of the system on handling different domains, additional ontologies were translated and evaluated. The results have been compared with translations provided by the Microsoft Translator API and the improvements demonstrate a better performance of the proposed approach for automatic ontology translation.

1 Introduction

Nowadays, most of the semantically structured data, i.e. ontologies or taxonomies, have labels stored in English only. Although the increasing amount 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].

Another important reason to translate ontologies is that they may already exist in different languages, but without aligning the concepts across languages we are not able to combine, compare or extend them. Furthermore, government institutions may be obliged to publish their ontologies or other structured data in their native language, e.g. financial reports need to be written in the language in which the financial institution operates. Therefore, performing ontology based data analytics would fail on providing reports in German, Spanish,

or other official European languages, if we would use only existing ontologies in English. Additionally, medical ontologies, e.g. the ICD Ontology¹ can be used to standardize medical reports, but physicians, researchers or patient organizations will access these reports in their native language; therefore, only a cross-lingual aligned ontology may give an appropriate overview. Another example is the Europeana project,² in the heritage domain, where preservation of the cultural treasure of Europe shows the need of cross-lingual alignment of different resources.

Since manual multilingual enhancement of domain-specific ontologies is very time consuming and expensive, we engage a domain-aware statistical machine translation (SMT) 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 specific vocabulary, like ontology labels, appears infrequent in a parallel corpus. Additionally, ambiguous labels built out of only a few words do not always express enough semantic information to guide the SMT system in translating a label correctly in regards to the targeted domain. This can be observed in domain-unadapted SMT systems, e.g. Microsoft Translator,³ 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, a collaborative knowledge management platform with a domain-aware SMT system 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 two different settings: (i) in a real-world context, namely Organic.Lingua,⁵ from quantitative and qualitative points of view, and, quantitatively only, (ii) on a set of ontologies aiming to evaluate the effectiveness of the SMT service in different domains.

The paper is structured as follow. Section 2 introduces the context of our main use case, the Organic.Lingua project. Section 3 provides an overview of the ESSOT architecture; while, Sect. 4 focuses on user facilities implemented for supporting the ontology translation task. In Sect. 5 we reported the evaluation conducted on both the translations suggested by the service and the usability of the platform. Finally, Sect. 6 provides a general overview about ontology translation and knowledge management tool; while, Sect. 7 concludes the paper.

¹ <http://www.who.int/classifications/icd/en/>.

² <http://www.europeana.eu/portal/>.

³ <http://www.bing.com/translator/>.

⁴ Translation performed on 16.04.2016.

⁵ <http://www.organic-lingua.eu>.

2 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:

- 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. The list of available concepts is retrieved by using an ontology service deployed in the ontology management component. Then, this list is exploited for annotating the learning resources published on the Web portal.
- 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. Moreover, the Web portal is equipped with a graphical semantic tree that exploits the content of the ontology for facilitating the browsing of the resource repository classification. Finally, the ontology is used also by the Cross-Language Information Retrieval component for query expansion purposes.

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. The maintenance of such an accuracy requires a precise methodology, and dedicated tools, for avoiding the loss of effectiveness of the components deployed on the platform.

3 Platform Architecture

In this Section, we present a general overview of the platform for managing the life-cycle of translating ontological entities. Figure 1 shows the architecture diagram, where we distinguish two main blocks:

- the service-side: containing the components for the machine translation models used for suggesting translations when requests are performed by users; and,

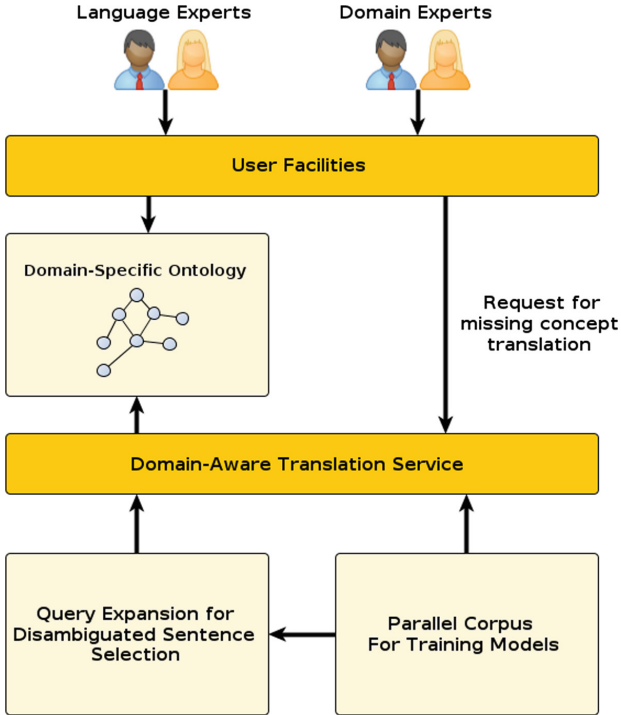


Fig. 1. Diagram of the overall architecture of the ESSOT platform

- the user-side: containing the facilities implemented for supporting experts in managing the multilingual layer of ontologies.

The service side contains the components used for creating and updating the model used by the domain-aware machine translation service. Such components are described below, while in Sect. 4 we provide a description of the facilities implemented for supporting users in managing ontologies.

Statistical Machine Translation. Our approach is based on statistical machine translation, where we wish to find the best translation \mathbf{e} of a source 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 functions as a search procedure, provides the most probable translation based on a statistical translation model learned from sentence aligned corpora.

For a broader domain coverage of datasets necessary to train an SMT system, we merged several parallel corpora, e.g. JRC-Acquis [3], Europarl [4], DGT (translation memories generated by the *Directorate-General for Translation*) [5], MultiUN corpus [6] and TED talks [7] among others, into one parallel dataset. For the translation approach, we engage the widely used Moses toolkit [8]. Word

alignments were built with GIZA++ [9] and a 5-gram language model was built with KenLM [10].

Query Expansion for Sentence Selection. Due to the shortness of ontology labels, there is a lack of contextual information, which can otherwise help disambiguating short or ambiguous expressions. Therefore, our goal is to translate the identified ontology labels within the textual context of the targeted domain, rather than in isolation. With this selection approach, we aim to retain relevant sentences, where the English label *vessel* or *injection* belongs to the medical domain, but not to the technical domain. 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.

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 [11], where the authors extend the semantic information of ontology labels using Word2Vec⁶ 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 [12]. 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. A score of 1 would represent a perfect word similarity; e.g. *cholera* equals *cholera*, while the medical expression *medicine* has a cosine distance of 0.678 to *cholera*. Since words, which occur in similar contexts tend to have similar meanings [13], this approach enables to group related words together.

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 used 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.

4 Supporting Users in the Ontology Translation Activity

The ESSOT system has been equipped with facilities supporting the collaborative translation of domain-specific ontologies in order to satisfy the requirements of the ontology translation task from a user perspective.

Concerning users, we identified two distinct groups: the Domain Experts and the Language Experts. Domain Experts are in charge of the modelling aspect of ontologies (i.e. creation of concepts, individuals, properties, and the relationships between them); while Language Experts are responsible of managing the labels associated with each entity by evaluating their correctness and, eventually, by

⁶ <https://code.google.com/p/word2vec/>.

providing a more fine-grained adaptation with respect to the domain described by the ontology.

Below, we present the list of the implemented facilities specifically designed for supporting the management of the multilingual layer of ontologies. Here, we focused on the interface that have been specifically implemented for managing the *Organic.Lingua* ontology.⁷ However, such facilities can be adopted, in general, for managing the multilingual aspect of any ontology.

Domain And Language Experts View. The page dedicated to the management of an ontology label, specifically designed for the Domain and Language Experts, has been equipped with functionalities that permits revisions of the linguistic layer. This set of functionalities permits to revise translations of names and descriptions of each entity (concepts, individuals, and properties).

For facilitating the browsing and the editing of the translations, a quick view box has been inserted into the mask (as shown in Fig. 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 (Fig. 3).

Multilingual component

Select language: English

Translation in the language:	English
Concept name	agricultural method
Concept description	Practices used to enhance crop and livestock health and prevent weed, pest or disease problems without the use of chemical substances.

Suggest translation

Fig. 2. Multilingual box for facilitating the entity translation

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 a wiki-style paradigm [14], 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.

⁷ 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. 3. Quick translation box for editing label translations

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 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 Fig. 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.

5 Evaluation

Our goal is evaluating the usage and the usefulness of the ESSOT user facilities and 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?

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.
europaean 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 label translations

RQ2 Do the ESSOT 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. Qualitative: the tool has been validated in the context of the *Organic.Lingua* project where we collected subjective judgements from the language experts. They have been involved in the evaluation of the tool on the general usability of the components and by providing feedback for future improvements.
2. Quantitative: beside the user evaluation, 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.

5.1 User Evaluation Context

Eleven 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. They were all experts of the agricultural domain, therefore, labels used by them have to be considered as a gold standard from the domain point of view. From the mother tongue perspective, the evaluation was performed by three German, four Spanish and four Italian native speaking experts. Most of them had no previous knowledge of the tool, hence an initial phase of training was necessary. The training was organized according to the following steps:

- A one-day overall introduction to the tool.
- A few short, on-line, training sessions with the ESSOT tool guided by ontology and tool experts, targeted to help domain experts to better understand the capabilities of the tool.
- Hands-on usage of the tool: language experts were left to “play” with ESSOT in order to become familiar with the functionalities that they would use during the revision process. This exercise also had the secondary objective to collect doubts and problems encountered by experts.

After the initial training, experts were asked to translate the ontology in the three languages mentioned above. The experts used ESSOT facilities for completing the translation task and, at the end, they provided feedback on the tool support for accomplishing the task. A summary of these findings and lessons learned are presented in Sect. 5.2.

5.2 Qualitative Evaluation Results

To investigate the subjective perception of the eleven experts about the support provided for translating ontologies, we analysed the data collected through a questionnaire. For each functionality described in Sect. 4, we provided the information how often each aspect has been raised by the language experts.

Language Experts View

Pros: Easy to use for managing translations (9)

Usable interface for showing concept translations (3)

Approval And Discussion

Pros: Pending approvals give a clear situation about concept status (4)

Cons: Discussion masks are not very useful (8)

Quick Translation Feature

Pros: Best facility for translating concepts (8)

Cons: Interface design improvable (3)

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, 9 out of 11 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.

Results concerning the approach and discussion facility are inconclusive. On the one hand, the experts perceived positively the solution of listing approval requests on top of each concept page. This fact is connected with a personalization that we embed into the ESSOT home page. Indeed, after the login, users are able to see the list of pending approvals require their action. This way, it is more easy for them to locate the translations that have to be evaluated and, eventually, to approve or to modify them. On the other hand, we received negative opinions

by almost all experts (8 out of 11) 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, 8 out of 11 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.

5.3 Quantitative Evaluation Results

The automatic evaluation on label translations provided by ESSOT is based on the correspondence between the automatically generated output and reference translations (gold standard), provided by domain and language experts. For the automatic evaluation we used the BLEU [15], METEOR [16] and TER [17] algorithms.

BLEU is calculated for individual translated segments (n-grams) by comparing them with reference translations. Those scores, between 0 and 100 (perfect translation), are then averaged over the whole *evaluation dataset* to reach an estimate of 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.

Datasets. To demonstrate the performance of the proposed framework, we use several ontologies coming from different domains:

- The **Organic.Lingua** ontology contains 291 concepts in the agricultural domain. All concepts within the ontology have been translated into 16 languages. In addition, mappings to Agrovoc and Eurovoc have also been defined.
- The DOAP (Description of a Project) Ontology⁸ defines the vocabulary to describe software projects. It was created to convey semantic information associated with free and open source software projects. It holds translations of labels into 6 languages,⁹ whereby we use German and Spanish translations as the gold standard, which is compared with the automatically generated labels.
- The Geoskills ontology¹⁰ holds the competencies, topics and educational contexts in five different languages, i.e., English, German, French, Spanish and Dutch.

⁸ <https://github.com/edumbill/doap>.

⁹ French, Spanish, German, Czech, Portuguese, Japanese.

¹⁰ <https://github.com/i2geo/GeoSkills>.

- The STW Thesaurus for Economics [18] provides the vocabulary of more than 6,000 standardized subject headings (in English and German) and 20,000 additional entry terms (keywords) belonging to the economical domain. In addition to that, the entries are richly interconnected by 16,000 broader/narrower and 10,000 related relations.
- The Thesaurus for the Social Sciences (TheSoz) [19] enables indexing documents and research information in the social sciences. In overall it stores about 8,000 standardized subject headings in English, German and French.

Table 1. Automatic translation evaluation of the targeted ontologies by the Microsoft Translator API and our proposed system (bold results = best performance)

	English → German			English → Italian			English → Spanish		
<u>Organic.Lingua</u>	BLEU	METEOR	TER	BLEU	METEOR	TER	BLEU	METEOR	TER
Microsoft	3.7	19.6	95.1	13.5	28.6	87.1	21.0	36.9	73.3
ESSOT	7.4	31.0	99.1	13.0	34.2	78.8	25.7	44.4	66.7
<u>DOAPOntology</u>	BLEU	METEOR	TER	BLEU	METEOR	TER	BLEU	METEOR	TER
Microsoft	6.0	24.3	93.3	/	/	/	20.2	32.7	81.6
ESSOT	6.4	31.3	91.8	/	/	/	20.1	35.3	72.5
<u>GeoSkills</u>	BLEU	METEOR	TER	BLEU	METEOR	TER	BLEU	METEOR	TER
Microsoft	11.8	22.3	97.5	/	/	/	16.3	30.1	88.9
ESSOT	14.0	30.7	94.6	/	/	/	16.7	35.9	83.0
<u>STW</u>	BLEU	METEOR	TER	BLEU	METEOR	TER	BLEU	METEOR	TER
Microsoft	6.5	15.1	89.9	/	/	/	/	/	/
ESSOT	6.8	21.7	98.8	/	/	/	/	/	/
<u>TheSoz</u>	BLEU	METEOR	TER	BLEU	METEOR	TER	BLEU	METEOR	TER
Microsoft	12.1	37.0	88.1	/	/	/	/	/	/
ESSOT	20.4	39.8	103.5	/	/	/	/	/	/

We evaluate the automatically generated translations into German, Italian and Spanish provided within ESSOT 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 as the gold standard.

The Organic.Lingua ontology provides 274 German, 354 Italian and 355 Spanish existing translations out of 404 English labels. As seen in Table 1, the contextual information for label translation used in ESSOT, significantly outperforms¹¹ the Microsoft Translator API. When translating English labels into German we gain a 51.3% averaged improvement over the commercial system and 51.7% for Spanish. In addition to that, it produces comparable results when translating into Italian (10.5% improvement).

¹¹ The approximate randomization approach in MultEval [20] is used to test whether differences among system performances are statistically significant with a p-value < 0.05.

Besides the evaluation on the *Organic.Lingua* ontology labels, the multilingual gold standard within the DOAP and GeoSkills ontology enables the evaluation of automatically translated ontology labels into German and Spanish. In detail, the results for the DOAP ontology show similar performance between both translation systems. On the other hand, the results of the GeoSkills ontology labels show statistically significant (p-value < 0.05) improvements over the Microsoft Translator API. For the STW and TheSoz ontology, which enables automatic evaluation of German language, only translations of the TheSoz labels show significant improvements of our system. Although the evaluation metrics show slight improvements when translating the STW ontology, the improvements are not significant.

As a final evaluation, we manually analysed the TheSoz translated labels regarding the most frequent errors of the both translation systems. The first observation is related to compound words, a frequent error class when translating into German. We observed that Microsoft often provided a non-compound translation in German. As an example, labels like *company takeover*, *working week* or *crime fighting* were translated word by word into German, i.e., *Übernahme der Firma*, *wöchentliche Arbeitszeit* or *Bekämpfung von Kriminalität*. Although these translations can be seen as correct translations, the provided gold standard in the ontology preferred German compounds *Unternehmensübernahme*, *Wochenarbeitszeit*, *Verbrechensbekämpfung*. Besides a small amount of wrong translations (*partnership* into *marriage*, *translator* into *translator*, *young worker* into *Junge Arbeitnehmer*), Microsoft's system showed expected problems in disambiguating short expressions. Due to the shortness of the labels, the ontology label *driver* was translated as *Treiber*, which is correct in the IT domain (as *hardware driver*), but not in the targeted domain. Similarly, *stroke*, without contextual information, was translated as *Strich* (en. *line*, *dash*), although *Schlaganfall* would be the correct translation into German. For these ambiguous labels, our proposed system, which used a disambiguated contextual information, provided correct translations, i.e. *Fahrer* from *driver* or *Schlaganfall* from the English label *stroke*. On the other hand also our system did not always perform best. The largest observed error class were out-of-vocabulary issues, i.e. alignments between source and target language, which were not learned during the SMT training step. For example, TheSoz labels *bonapartism*, *shamanism*, *patriciate* or *praxeology*, which are not stored in our translations models, were provided as untranslated words on the target side.

5.4 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 Sect. 6), which do not enable specific support for translation. The suggestion-based service allowed effective suggestions and reduced the effort required for finalizing the translation of the ontology. As example, we may consider the Organic.Lingua specific use case, where the time for translating the ontology was reduced from 3.5 h (completely manual translation) to 2.1 h (translation performed with ESSOT). This point confirms the capability of the domain-aware translation service of providing translations adapted to the specific topic of the ontology experts are going to model. 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 modeling ontologies but that might not have enough linguistic expertise for translating label properly with respect to the domain), facilities supporting the 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.

6 Related Work

In this section, we want to summarize approaches related to the pure ontology translation task and to present a brief review of the most known ontology management tools current available by emphasizing 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 [21, 22]. Their 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 ser-

vices like BabelFish, SDL FreeTranslation tool or Google Translate were used to overcome this issue [23, 24]. Additionally, [23, 25] performed ontology label disambiguation, where the ontology structure is used to annotate the labels with their semantic senses. Similarly, [26] 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 text data. Current frameworks for ontology label translation are accessing directly commercial systems, such as Google Translate or Microsoft Translate, whereby both systems are unable to detect the domain when translating short ambiguous expression, e.g. *vessel*, *injection*, *track*, *head*, *equity*. In this paper, we demonstrate a platform supporting a machine translation system to translate ontology labels in a domain-specific context.

If we perform a “skimming” of the systems available for ontology management, we identified four of them that may be compared with the capabilities provided by ESSOT : *Neon* [27], *VocBench* [28], *Protégé* [29], and *Knoodl*¹². However, they do not fully support experts in the specific task of translating ontologies. 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* and *Knoodl* is restricted to the sole description of the labels. Finally, none of them implements the capability of connecting the tool to an external machine translation system for suggesting translations automatically.

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. ESSOT system integrates a domain-adaptable semantic translation component and a collaborative knowledge management facilities for supporting language experts in the ontology translation activity. The platform has been concretely used in the context of the *Organic.Lingua* EU project and on a set of multilingual ontologies coming from different domains by demonstrating the effectiveness in the quality of the suggested translations and in the usefulness from the language experts point of view.

Acknowledgments. This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289.

¹² <http://www.knoodl.com>.

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)
3. 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)
4. Koehn, P.: Europarl: a parallel corpus for statistical machine translation. In: Conference Proceedings: The Tenth Machine Translation Summit, AAMT (2005)
5. Steinberger, R., Ebrahim, M., Poulis, A., Carrasco-Benitez, M., Schlüter, P., Przybyaszewski, M., Gilbro, S.: An overview of the European union’s highly multilingual parallel corpora. *Lang. Res. Eval.* **48**(4), 679–707 (2014)
6. 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), Calzolari, N.C., (eds.) Proceedings of the Seventh conference on International Language Resources and Evaluation, European Language Resources Association (ELRA), vol. 5, pp. 2868–2872 (2010)
7. Cettolo, M., Girardi, C., Federico, M.: Wit³: Web inventory of transcribed and translated talks. In: Proceedings of the 16th Conference of the European Association for Machine Translation (EAMT), Trento, Italy, pp. 261–268, May 2012
8. 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, pp. 177–180 (2007)
9. Och, F.J., Ney, H.: A systematic comparison of various statistical alignment models. *Comput. Linguist.* **29**, 19–51 (2003)
10. Heafield, K.: KenLM: faster and smaller language model queries. In: Proceedings of the EMNLP 2011 Sixth Workshop on Statistical Machine Translation, Edinburgh, pp. 187–197, July 2011
11. 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
12. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. ICLR Workshop (2013)
13. Harris, Z.: Distributional structure. *Word* **10**(23), 146–162 (1954)
14. 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, 26–31 May 2014, European Language Resources Association (ELRA), pp. 1668–1675 (2014)
15. 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 2002, pp. 311–318 (2002)

16. 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)
17. 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)
18. Borst, T., Neubert, J.: Case study: Publishing stw thesaurus for economics as linked open data. Case study (2009)
19. Zapolko, B., Schaible, J., Mayr, P., Mathiak, B.: TheSoz: a SKOS representation of the thesaurus for the social sciences. *Semantic Web* 4(3), 257–263 (2013)
20. Clark, J., Dyer, C., Lavie, A., Smith, N.: Better hypothesis testing for statistical machine translation: controlling for optimizer instability. In: Proceedings of the Association for Computational Linguistics (2011)
21. Cimiano, P., Montiel-Ponsoda, E., Buitelaar, P., Espinoza, M., Gómez-Pérez, A.: A note on ontology localization. *Appl. Ontol.* 5(2), 127–137 (2010)
22. 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)
23. Espinoza, M., Montiel-Ponsoda, E., Gómez-Pérez, A.: Ontology localization. In: Proceedings of the Fifth International Conference on Knowledge Capture, New York (2009)
24. 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 2009. LNCS, vol. 5926, pp. 1–15. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-10871-6_1](https://doi.org/10.1007/978-3-642-10871-6_1)
25. 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)
26. McCrae, J.P., Arcan, M., Asooja, K., Gracia, J., Buitelaar, P., Cimiano, P.: Domain adaptation for ontology localization. *Science, Services and Agents on the World Wide Web, Web Semantics* (2015)
27. Espinoza, M., Gómez-Pérez, A., Mena, E.: Enriching an ontology with multilingual information. In: Bechhofer, S., Hauswirth, M., Hoffmann, J., Koubarakis, M. (eds.) ESWC 2008. LNCS, vol. 5021, pp. 333–347. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-68234-9_26](https://doi.org/10.1007/978-3-540-68234-9_26)
28. Stellato, A., Rajbhandari, S., Turbati, A., Fiorelli, M., Caracciolo, C., Lorenzetti, T., Keizer, J., Paziienza, 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.) ESWC 2015. LNCS, vol. 9088, pp. 38–53. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-18818-8_3](https://doi.org/10.1007/978-3-319-18818-8_3)
29. 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), 89–123 (2003)