

Content Translation: Computer-assisted translation tool for Wikipedia articles

Niklas Laxström
University of Helsinki
Dept. of Modern Languages
and
Wikimedia Foundation
nlaxstrom@wikimedia.org

Pau Giner
Wikimedia Foundation
pginer@wikimedia.org

Santhosh Thottingal
Wikimedia Foundation
sthotttingal@wikimedia.org

Abstract

The quality and quantity of articles in each Wikipedia language varies greatly. Translating from another Wikipedia is a natural way to add more content, but the translation process is not properly supported in the software used by Wikipedia. Past computer-assisted translation tools built for Wikipedia are not commonly used. We created a tool that adapts to the specific needs of an open community and to the kind of content in Wikipedia. Qualitative and quantitative data indicates that the new tool helps users translate articles easier and faster.

1 Introduction

Wikipedia is the most multilingual encyclopedic knowledge archive, with over 280 languages with varying amount of content. Knowledge available for a user is limited by the languages used to access it. Translation has been a common way to expand knowledge across languages in Wikipedia. The editing activity of the top 46 language editions of Wikipedia shows that 25% of edits by multilingual users are for the same article in different languages (Hale, 2013).

It is not necessary to use any tool to translate Wikipedia articles. However, it is a complicated process and mainly done by experienced Wikipedia editors.

There were many attempts to build tools to support translation of articles. None has seen widespread use: in our research only few users reported using those tools when translating Wikipedia articles.

In this paper we present a new approach to support translation which has been designed taking into account the unique needs of Wikipedia content and their community. Content Translation (CX) is a new tool that automates many steps of the translation process and validates the approach in practice. It was first enabled on 8 Wikipedias as an opt-in feature to create new articles in January 2015. Selected language pairs have machine translation (MT) support.

2 Previous work

MediaWiki, the software powering Wikipedia, is translated to hundreds of languages using the Translate extension. No such solution was available for translating Wikipedia articles, leaving a gap in the translation support.

There were at least ten instances of translation tools built for Wikipedia¹. Those tools can be divided into two groups based on whether the tool creators already possessed MT software. The first group is composed of companies such as Google and Microsoft, but also smaller companies and researchers. The other group of tools has been created just for Wikipedia article translation, mostly by volunteers.

Among the earliest tools were GTT by Google and WikiBhasha by Microsoft, using their own MT services (García and Stevenson, 2009; Kumarana et al., 2011). Later, Casmacat for professionals and researchers (Alabau et al., 2013) and CoSyne for multilingual MediaWikis (Bronner et al., 2012), unlike Wikipedias which are monolingual.

Common to all such tools is that they are not integrated into Wikipedia. To use them one needs to go to another website or install software. CX

© 2015 The authors. This article is licensed under a Creative Commons 3.0 licence, attribution, CC BY.

¹Details collected at https://meta.wikimedia.org/wiki/Machine_translation

is integrated into Wikipedia and provides a WYSIWYG editor (what you see is what you get).

3 Designing the translation experience

The design of CX was aimed at improving the existing process users followed when translating. Following the principles of User-Centered Design (Norman and Draper, 1986), we organised periodic user research sessions to (a) better understand the user needs during the existing translation process, and (b) validate new ideas on how to improve this process.

3.1 User research

We recruited 106 participants using a survey². From their responses we identified dictionaries (76% of participants used them), and Wikipedia (60%) as their most used tools when translating. MT (53%), spell checkers (48%) and glossaries (42%) were also common. Less than 6% of the participants mentioned tools specifically aimed at Wikipedia article translation, such as those described in Section 2, and no tool was mentioned by more than one participant.

We organised 16 research sessions. Sessions were organised in two parts. In the first part, contextual inquiry techniques (Beyer and Holtzblatt, 1998) were applied to observe user behaviour while translating, and identify their needs. The second part was a usability testing study (Nielsen, 1994) to evaluate different design ideas in the form of prototypes.

3.2 Design principles

The research sessions were instrumental to guide the design of the translation experience³. The following design principles summarise the approach we followed when designing the tool.

Freedom of translation

There is a significant diversity in Wikipedia content across languages. On average, two articles from different languages on the same topic have just 41% of common content (Hecht and Gergle, 2010). In contrast to other kinds of content, such as software user interface strings or documentation, Wikipedia articles are not intended to be exact translations that are always kept in sync. In order to support that content diversity, CX does not force

users to translate the full article. As illustrated in Figure 1, users add content to the translation one paragraph at a time. When a paragraph is added, an initial translation based on MT is provided for the user to edit. MT is used if available, but the user can also start with the source text or an empty paragraph if that is preferred.

Unlike other tools that define a strong boundary to translate on a per sentence basis, working at a paragraph level allows users to reorganise sentences and accommodate different editing patterns.

Provide context information

In CX the original article and the translation are shown side-by-side. Each paragraph is dynamically aligned vertically with the corresponding translated paragraph, regardless of the difference in length. This allows users to quickly have an overview of what has already been translated and what has not.

Contextual information reduces the need for the user to navigate and reorient. When translating a sentence, the corresponding sentence in the original document is highlighted. In addition, when manipulating the content, options are provided anticipating the user's next steps. In Figure 1, based on the user's text selection, the user can explore the article related to the selected text (in the source or target languages), or turn the selected text into a link. Dictionary can be accessed inside the tool by selecting a word or by using the search box in the tools column.

Focus on the translation

We identified many steps in the translation process that could be automated. Users spend time making sure each link they translated points to the correct article in the target Wikipedia, and recreating the text formatting that was lost when using an external translation service. They also look for categories available in the target Wikipedia to classify the translated article, and save constantly during the process to avoid losing their work.

CX deals with those aspects automatically. When adding a paragraph, the initial translation preserves the text format. Modifications to the translated content are saved automatically. Links point to the right articles if they exist and existing categories are added to the article thanks to Wikidata⁴, a structured data knowledge repository, that maps corresponding concepts across lan-

²<https://goo.gl/iKQIDh>

³A detailed design specification is available at <https://www.mediawiki.org/wiki/CX>

⁴<https://www.wikidata.org>

The screenshot displays a Wikipedia article titled 'Cupcake' in two languages: Spanish (español) and Catalan (català). The Spanish version is on the left, and the Catalan version is on the right. Both versions feature an image of chocolate cupcakes and a descriptive paragraph. The Spanish text describes a 'cupcake' as a small portion of 'tarta' (cake) for one person, baked in a mold with paper liners called 'cápsulas'. The Catalan text describes a 'cupcake' as a small portion of 'pastís' (cake) for one person, prepared in a mold with paper liners called 'càpsules'. A yellow highlight in the Spanish text notes that it is often confused with muffins and magdalenas. A '+ Add translation' button is visible between the two columns. On the right, a sidebar contains a search bar with 'pastís' entered, a WYSIWYG editor with bold, italic, and list icons, and a link list with 'Pastís' highlighted. A 'Publish translation' button is at the top right, and a 'Provide feedback' button is at the bottom right.

Figure 1: The source and translated content side-by-side and additional tools on the right.

guages. As those aspects are automated, users can focus on adapting content for the initial version of the article rather than on technical and formatting tasks.

Quality is key

One of the concerns raised early by the participants was about MT quality. Users were concerned about the potential proliferation of low quality content in Wikipedia articles.

In order to respond to that concern, CX keeps track of the amount of text that is added from MT without further modification by the users. When a given threshold is exceeded, a warning is shown to users encouraging them to focus on quality more than quantity.

4 WYSIWYG implementation

MediaWiki’s wikitext is not standardised. For a long time, the only way to use wikitext was to render it to HTML with MediaWiki. Parsoid⁵ is a Wikimedia project that implements a second parser for wikitext. To follow the principle *focus on translation* principle we only provide limited editing and formatting options and side-step a lot of complexity of Wikipedia article structure

⁵<https://www.mediawiki.org/wiki/Parsoid>

without negatively affecting the translation process. CX is the first translation tool that provides a WYSIWYG editor using the annotated HTML provided by Parsoid.

Some MT services neither support HTML input nor provide reordering information. Preserving markup is an essential requirement for CX because wikitext adaptation and WYSIWYG editing are based on the markup. We devised an algorithm that can reconstruct the reordering information by making the MT service do some additional work.

5 MT evaluation

We use the subjective evaluations of MT quality for a given language pair to decide whether we will include a MT service for a language pair in the tool. To evaluate a MT for a given language pair, we ask the potential future users of the tool to translate articles using it and tell whether it was useful for them or not.

We run a MT service on our servers using the open-source Apertium project (Forcada et al., 2011), but we support other MT providers as well.

6 Evaluation

Currently the tool is only available to self-selected users (most of them experienced editors), hence the results cannot be generalised to the whole community. Further studies on the resulting quality over a long term will help.

The low deletion ratio for articles created using CX suggests that there are no major problems in terms of quality. In three months of exposing the tool as an opt-in feature, 900 articles were published using CX with an overall deletion ratio lower than 1% across all languages, which is lower than the deletion rate for all new articles.

We noticed that there is a significant difference between the number of created articles in different target language Wikipedias, which cannot be explained by the number of active users, number of available articles to translate nor the availability of MT. For example in three months the Catalan Wikipedia saw 455 articles created by translating from Spanish with CX, but in the Portuguese Wikipedia only 25. Both language pairs have MT provided by Apertium. Statistics about the tool are collected publicly⁶.

We have not yet made precise measurements on translation time saving, but we got positive reports from our users. In a roundtable⁷ organised with editors of the Catalan Wikipedia, an experienced editor reported a 70% time saving.

We found that English is the most used source language, consistent with Hale's findings on multilingual user behaviour (2013).

7 Conclusions

We developed a tool that addresses the specific needs of an open community and the specifics of the kind of content in Wikipedia. CX is a computer-assisted translation tool with a WYSIWYG editor and automatic link adaptation. CX supports multiple different MT providers, but by integrating the open source Apertium project we were able to quickly provide MT for multiple language pair We developed MT education and tracking features to address community concerns about proliferation of poor quality translations.

User feedback for CX is supportive and data

also shows that quality of the published translations is good, alleviating the community concerns. The low translation activity in multiple languages where the tool is already available needs further research. Close integration in Wikipedia allows CX to recruit users and suggest articles to translate in ways not possible with the previous tools.

References

- Alabau, Vicent, Ragnar Bonk, Christian Buck, Michael Carl, Francisco Casacuberta, Mercedes García-Martínez, Jesús González, Philipp Koehn, Luis Leiva, Bartolomé Mesa-Lao, et al. 2013. Casmacat: An open source workbench for advanced computer aided translation. *The Prague Bulletin of Mathematical Linguistics*, 100:101–112.
- Beyer, H. and K. Holtzblatt. 1998. *Contextual Design: Defining Customer-centered Systems*. Interactive Technologies Series. Morgan Kaufmann.
- Bronner, Amit, Matteo Negri, Yashar Mehdad, Angela Fahrni, and Christof Monz. 2012. Cosyne: Synchronizing multilingual wiki content. In *Proceedings of the Eighth Annual International Symposium on Wikis and Open Collaboration*, WikiSym '12, pages 33:1–33:4, New York, NY, USA. ACM.
- Forcada, Mikel L, Mireia Ginestí-Rosell, Jacob Nordfalk, Jim O'Regan, Sergio Ortiz-Rojas, Juan Antonio Pérez-Ortiz, Felipe Sánchez-Martínez, Gema Ramírez-Sánchez, and Francis M Tyers. 2011. Apertium: a free/open-source platform for rule-based machine translation. *Machine translation*, 25(2):127–144.
- García, Ignacio and Vivian Stevenson. 2009. Reviews-google translator toolkit. *Multilingual computing & technology*, 20(6):16.
- Hale, Scott A. 2013. Multilinguals and wikipedia editing. *CoRR*, abs/1312.0976.
- Hecht, Brent and Darren Gergle. 2010. The tower of babel meets web 2.0: User-generated content and its applications in a multilingual context. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 291–300, New York, NY, USA. ACM.
- Kumarana, Narend, S Ashwani, and D Vikram. 2011. Wikibhasha: Our experiences with multilingual content creation tool for wikipedia. In *Proceedings of Wikipedia Conference India*, Wikimedia Foundation.
- Nielsen, Jakob. 1994. *Usability Engineering*. Interactive technologies. AP Professional.
- Norman, D.A. and S.W. Draper. 1986. *User centered system design: new perspectives on human-computer interaction*. New Perspectives on Human-Computer Interaction Series. Lawrence Erlbaum Associates.

⁶https://www.mediawiki.org/wiki/Content_translation/analytics

⁷<https://blog.wikimedia.org/2014/09/29/round-table-with-editors-from-the-catalan-wikipedia/>