

Towards an effective toolkit for translators

Andreas Eisele, Christian Federmann, James Hodson

Abstract

As advances in the field of machine translation (MT) continue to allow for greater distribution of multi-lingual information, an increasing number of browser-based toolkits are being developed in an attempt to empower translators, and decrease the gap between speed, and quality of output. However, these systems so far have failed to satisfy the translation community due to a fundamental misunderstanding of the way translators work. There has been an essential divergence in the study and implementation of machine translation systems—away from the original goal of providing perfect target-language texts, towards an attempt to provide output that can be largely understood for quick dissemination. The realisation that current methodology in the field does not allow for such high quality to be attained has turned researchers' attentions to the development of systems that act as supporting processes during several of the traditional stages of translation. Nonetheless, such efforts have continued on the assumption that high quality translation from the get-go is the most important factor, rather than assessing how translators' workflows react to poor quality target language text generated by such systems. Our aim has been to identify techniques that reduce the time that translators spend on the whole transfer process - by realising that a SMT core can provide much more than a raw translation, even at this relatively early stage in its development.

Our research evaluates the shortcomings of existing systems, and proposes a new kind of integrated online toolkit that addresses these limitations based on an understanding of the most significant and unnecessary struggles between translators and MT systems. Consequently we believe that with just a handful of changes it is possible to envisage an arrangement that will allow translators to regularly outperform non-MT-based working strategies, both in terms of speed and quality.

1 Introduction

Since Martin Kay’s seminal work in [4], the observation that MT and human translators are good at different aspects of translation has motivated the search for better ways to combine the relative strengths of both and to make MT useful as an auxiliary tool in the professional translation workflow.

However, the quest for tools that are truly effective in a professional setting has proven to be a difficult task, and up to now, the impact of MT technology on translation has been rather limited. This stands in striking contrast to translation memory (TM) technology, which has been remarkably successful on the market, although it is addressing much simpler tasks than MT proper. Statistical MT, especially the phrase-based variant of it, can be seen as a generalisation and extension of TM to the case where the units for which translations are to be re-used are not complete sentences, but smaller groups of words, the aim being to generate a coherent and equivalent sentence, not necessarily a direct replacement. Hence it appears plausible to assume that recent progress in SMT may also provide new ways to achieve improvement in TM related tasks.

One work package within the research project EuroMatrixPlus aims at bringing the “Wiki” idea of collaborative content development to translation, and to use this as a basis for the development of new MT-based translation tools. These tools build on ideas such as post-editing machine translation output, translation memories, phrase glossing, and interactive machine translation, in a way that is more sensitive to traditional translation workflows. Furthermore, through this model we want to allow users to easily create translations that can then be crowd-sourced and used as training data for more powerful and tailored translation models. Integration of the different approaches explored within the project into one common product will allow us to showcase the relative strengths and weaknesses of the competing technologies to a general audience, and it will help us to highlight to what extent the integration of different knowledge sources (*linguistic rules, statistical evidence from corpora, user feedback etc.*) leads to perceptible improvements in translation quality and system usability for end users of the technology.

This paper is structured into three sections. The first section will motivate research work on web-based translation toolkits, highlighting the difficulties faced by different user groups, and identifying useful (*symbiotic*) relationships that can result. In the second section, we showcase the techniques that are implemented within existing toolkits and research, together with their shortcomings with respect to real-world translators. We will concentrate our efforts on Google’s Translator Toolkit, and recent research prototypes that attempt to apply more predictive elements to aid the translator. Finally, we will present our proposals for tightening the gap between these tools and their users within our own prototype-implementation: WikiTrans. We suggest methods to promote the sharing of more context-specific information by the user, as well as supplementary crowd-sourcing techniques for enhancing the amount and quality of training data that can be extracted from any given user translation. However, these are techniques that inevitably make the robustness and quality of the resulting machine translation system highly dependent upon having a threshold of dedicated users, which is ultimately difficult to achieve in a non-commercial setting.

2 Motivation

So, why ought we to care about developing MT-based toolkits for translators at all? And, if we should care, then what is the best way to begin thinking about this goal? We start, rather contrary to the dominant thought-process in the field, by thinking about translation without computers, or, rather, without machines. The reason we approach the problem in this manner is that translators without technology are actually surprisingly good at their job, and have developed many “tricks-of-the-trade” to ensure accuracy, efficiency, and consistency. Furthermore, in developing tools to help translators, our aim is to do just that, and not to inadvertently create bottlenecks of (*unnecessary*) inefficiencies. Thus, it seems clear to us that without a solid knowledge of translation in reality, developments in the field will only serve to distance it from those users it relies on the most – translators. Avoiding this alienation should be a top priority in order to ensure that a sustained dialogue exists between developers and experts for years to come, and that this dialogue leads to increased quality and usability.

2.1 The basic translation model

The dominant translation model can be summarised as follows:

- Integration of Source text
- Identification of Challenges
- Source language resolution
- Skeleton translation
- Target language resolution
- Fleshing
- Polishing

That is to say, most translators logically follow a process similar to this one when translating a source text. Initially it is necessary to understand the nature of the piece to be translated, the message it is attempting to convey, and the audience for which it is intended. Then, it is a case of identifying those areas of a translation that will cause the most difficulty in terms of the cultural perspectives offered, referential level (*e.g.*, *self-referential sentences*), and problematic constructions or vocabulary. As an answer to some of these challenges raised, a translator may even choose to modify the original piece in order to give a similar effect that is more conducive to translation into the target language. Such resolution of problematic source language elements can often be useful for ensuring that the meaning and purpose of the text to be translated is retained as unmodified as possible during transfer. Next the translator will actually begin to translate words, phrases, sentences and groups of sentences into the target language without necessarily solving the remaining problems—the idea being to build a coherent basis for the work in its new language, without finalising the manner in which those challenges will be solved. Once a firm grounding has been achieved, then the resolution of outstanding issues must occur. In many cases the answers to these will be pure judgement calls on the part of the translator (*how much of the original are we comfortable throwing*

overboard in order to keep the work navigable?). Related to this resolution is the fleshing out of the piece, adding in those details that were not central to the original, and marking the piece for style. As a final step in the translation process the work ought to be checked for consistency in all areas, and for fidelity to the original. Polishing a translation can make the difference between a text that reads like a dish-washer manual, and a Pulitzer prize winning novel.

2.2 Adding machine translation

That's the model, but where does MT come into all of this? Isn't the goal of MT to provide a translation that is as close to the end of this process as possible? The goal of MT as it stands is to provide output that covers all steps mentioned. Except that MT does not cover any of the steps mentioned in a way that is analogous to the way a translator would. Your average SMT engine takes a sentence, computes the matching sentence with highest probability, and spits it out the other end. If there is any resolution going on it's implicit, and most likely only an accidental by-product of the corpus being used. However, when developing a toolkit that involves both MT and real translators there is an opportunity to use MT in a more ambitious way at various stages of the translation workflow: performing duties that could dramatically increase the efficiency of translation. The main task when building a toolkit for translation is to find an appropriate balance between MT, information derived from MT, and the human translator.

In order to formulate an idea of the best division of labour we must first consider the different scenarios in which such a toolkit would be used, the kinds of translators that would be using it, and so the most obvious areas in which help is needed. One of the main advantages of the web is that it facilitates collaboration and the sharing of information across a variety of boundaries. Such freedom has created a whole new group of translators—the non-professional who is interested in making information from the internet accessible to more people across language barriers. In contrast to the professional translator, this new group may not even be fluent in the source language, and so may be incapable of performing a decent job on the first three aspects of our translation workflow (*above*). Furthermore, most of the translation tasks being completed by this group would reflect the most habitually accessed information - news, factual articles (*e.g.*, *Wikipedia*), and popular culture. A restricted range of input types may give us hope of being able to achieve satisfactorily high quality translation with MT alone, but this is severely off-set by the reality that much of the source text will be written in a variant of the source language that has no well-defined parallel corpus.¹ As such, we can be assured that the register of language in the translation (*however coherent*) will in all likelihood not be concomitant with the original. Thus, our toolkit should help our translator to make target language decisions about structure and lexicon based upon alternative analysis of

¹A parallel corpus must be large enough, and of a high enough quality, in order to be able to successfully train a MT system. Currently such corpora only exist in areas where there has been large-scale need for free public multi-lingual access to a text. Usually this occurs with political documents, which contain standardised language, not conforming to the lexicon and structure of popular culture or encyclopedia articles.

the text – and not assume that its job is done once the initial translation is provided. Of course, the tools that could be useful to a professional translator will have a wider remit than those that would be useful to the hobbyist, and any toolkit ought to be aware of this fact, and cater to both groups.

2.3 Towards a usable translation toolkit

With the building of an online toolkit for translation come several related opportunities that can be pursued – adding both to the accuracy of the final output, and the value of the product. The most obvious of these opportunities is collaboration between groups of translators/editors on the same project, with everyone having access to the part of the project that is assigned to them to translate/correct. By providing an efficient platform for versioning control, permission management, and project-workflow-management, it is possible to facilitate the completion of large and complicated tasks by individuals (*volunteers or otherwise*) with diverse skills, and in diverse geographic locations. This is an area that is starting to be realised with the advent of “social translation”, through which individuals find others interested in the same kinds of projects by means of a forum or marketplace. Once a group of translators has been formed, various versioning tools are used in order to manage the translations that they create, but no further tools are usually provided in order to increase the efficiency of translators. Most importantly, there currently exists no prominent tool that caters to the social translation market by means of Machine Translation.

The translation toolkit paradigms that we have mentioned thus far bring tangible benefits both to translators themselves, and to the end consumers of the works. However, the benefits of online toolkits do not stop there. So far, we have covered the benefits that come from the system, but not those benefits that relate to the system itself. With a large number of translators creating a large number of translations between a large number of language pairs, it is feasible to expect that future MT systems will take advantage of these high-quality, readily available parallel corpora, and use them to improve accuracy and performance recursively. This kind of crowd-sourcing works on the current corpus in two ways, adding to and supplementing it, and providing the ability to correct the mistakes present therein. Continuous expansion and enhancement of the corpus being used allows for the creation of sub-corpora for specific domains, and for the identification of different language registers within the source—therefore allowing the system to tailor the part of the corpus being used towards the type of translation being performed. By harnessing the large volume of translation that occurs by means of the toolkit it becomes possible to reduce the update-lifecycle of SMT systems, and tighten the feedback loop between corrected translations and training data. There are, of course, pitfalls to be avoided when designing such systems in order to ensure the continuing improvement in quality, and adaptability to different users needs, but we hope that we have managed to successfully highlight the need to bring MT closer to the translators themselves – it is beneficial for translators, and even beneficial for the future of non-supervised MT.

3 Perspectives on current technology

Recently, Google released a toolkit designed to facilitate the post-correction of results generated by their SMT engine in a web-based setting. This toolkit offers a set of additional functions, such as the ability to upload TM files, share the translation task with co-workers, and then upload the finished translation as a new article to Wikipedia in the target language. One of the main motivations for making this toolkit available free of charge to potential users seems to be the possibility of harvesting substantial corrections to SMT output, allowing its owners to use the collected material in order to incrementally increase the quality of the translation system. Our attempts to use the toolkit on a number of original articles from Wikipedia have lead us to the conclusion that this resource, in its current form, does not reach the level of usability desired by its users. It is clear even from preliminary usability studies and discussion with potential users of such systems, that although the quality of SMT output is increasing, the collateral tools that are often packaged with these do little more than collect useful data from translators, rather than helping them directly with the task at hand.

In this section we intend to provide an overview of the state-of-the-art in MT post-editing research, and the toolkits that are freely available for translators to use. Our purpose is to critique the effectiveness of current toolkits in response to the translation model identified in the previous section, and to use our analysis as a basis for sharing suggestions on best-practice across the industry (*see the following section*). Our research naturally concentrates on the comparative advantage that a translator can expect to receive from any given toolkit, over manual post-editing of MT output.

3.1 Current research efforts

Current research in this area has largely focused upon novel methods of presenting MT output to the translator. This comes in the form of “top x” lists of alternative translations for the same source language sentence (*essentially a list of decreasingly coherent sentences that the system did not think were appropriate*); or sentences that are presented to the user based on active constraints imposed by the translator’s typing (*searching the MT-translation for the best output sentence that matches the words typed thus far, and the original source sentence, see [5]*). This research is promising in that it proposes a way of unobtrusively providing translation help through active use of MT (*rather than simply providing a user with the raw output*), and at the same time offers the possibility to adjust the SMT model based upon the choices that a translator makes—therefore making a more accurate choice more probable the next time a similar sentence is translated (*integration of the source text*). Furthermore, it offers the possibility of augmenting the training corpus with each sentence that a translator completes that is not already present. However, such tools, because they rely directly and solely on the output of MT (*with little further processing*), are unable to provide the translator with further information that may be useful during the post-editing process. Also, they have the drawback of forcing the user to think about a translation linearly (*i.e., from start to finish*) rather than being able to break it down into constituent parts (*as our model attempts to do*). Other tools for aiding

post-editing of MT include the ability to share translations among users in a versioning-style system, which allows collaboration between users on both sides of the language barrier, as well as translators with different technical backgrounds (*which is especially good for fleshing and polishing of the final piece*).

Aside from tools to aid translators, research has also been focusing upon the methods used to leverage the power of crowd-sourcing and social translation to improve existing MT systems. Such research envisions the “correct” human-post-edited sentences forming a feedback loop to amend and augment the MT system’s training corpus, without needing to re-train the entire system every time [1]. Obviously, if implemented successfully such a system would allow for continuous improvement of MT output (*by reducing the weights of bad translations in a given context*). However, avoiding the pitfalls of over-training and correct identification of context-specific information is currently especially difficult to achieve without severely reducing the usefulness of post-edited (*accurate*) translations (*i.e., only using the final sentences as a means to correct the training data*). Incorporating the idea of social translation into the above, there is some hope that more accurate context-specific information could be gathered from multiple people working on the same (*or similar*) translation projects. Furthermore, social translation is seen as the holy-grail for attracting a necessary threshold of users in this internet-age.

3.2 Available online toolkits

Slowly, a number of online toolkits are appearing in order to facilitate the translation process. These tools vary from purely social translation communities (*such as dotSUB*), to raw MT such as AltaVista’s Babelfish, and Microsoft’s Bing Translator. However, the tools that are of particular interest and promise lie somewhere between. Tools such as Google’s Translator Toolkit and SYSTRANet’s Translator offer a MT system for performing translation between several language pairs, along with the ability to post-edit the output, make use of TM files, and collaborate online with other users. These systems are promising because the other kinds of tools available already perform well-defined tasks². The target of MT-based toolkits is much broader in scope. These toolkits should aim to achieve high quality translations by providing the kind of environment in which even the casual translator becomes able to make informed choices about the raw MT output—leading to an efficient use of a translator’s time in areas where they are most needed. In an ideal situation, such a toolkit would reduce translation time to that of heavily proof-reading a document. However, the current reality is rather different. In many cases, using a translation toolkit actually increases the amount of time that it takes to achieve a high quality, readable translation. Why? Because the raw output from the system is of such unsupported poor quality that it stifles the translator’s creativity rather than offering inspiration. As these toolkits offer the raw output along with separate tools for performing reference and editing, they are essentially the same as using a MT engine,

²Pure MT allows for quick and easy jist translations for casual consumption, whilst social translation communities allow volunteers to organise themselves around projects of particular interest to them personally – especially in areas such as video subtitling where web-based MT is currently unable to compete with human translation.

and then printing the output and working off-line with traditional textual resources. These toolkits do have the advantage that they allow the management and sharing of translation memory files, but translation memory files are nothing new, and should not be claimed as a key feature of these systems – more ambition is required in order to extract useful information from the processing of the text, the experience gained from previous translations, and collateral reference tools, and providing this seamlessly to the translator in an easily digestible manner. Currently such flexibility is not being pursued by these systems, leaving the field wide open to innovation.

4 A vision of the future

In this section, we present some of the design goals of the WikiTrans system currently under development as part of the EuroMatrixPlus project. Some of the features we will present have already been implemented, whereas others will require very sophisticated techniques. Although we do not promise that all of these features will be incorporated, as EuroMatrixPlus is part of a larger community effort (*one that has brought about very powerful tools like the Moses decoder*), we hope that by putting even very ambitious goals up for discussion we can help to propagate and develop the vision of a platform on which human translators and translation algorithms could complement each other effectively.

4.1 Machine Translation quality

In order to maximize the usefulness of Machine Translation technology for “outbound” translation (*i.e., the production of documents that can be published*), the quality of raw MT output should be as high as possible. This goes far beyond the properties that are usually measured by MT evaluation campaigns, where fluency and adequacy are appraised by human judges, or where the textual similarity between MT output and reference translations is measured with simplistic means such as the BLEU score and related metrics. If the ultimate goal is to minimize the effort to turn MT results into high quality translations, then there will (*in the long run*) be no alternative but to measure exactly what counts, which is the effort of human post-editors making use of MT systems in real life. As previously mentioned, recent work in the project [5] has gone a long way in providing detailed measurements of the areas where significant time is spent in post-editing MT results, and how different ways to use MT technology can make translation more effective. This line of work needs to be continued and broadened, and the measurements that are so far only part of carefully designed studies ought to be an integral part of all MT systems. Some very important properties to allow MT tools to achieve higher quality output will be the ability to automatically determine the domain and style of the source documents, and to adapt their internal process accordingly. An awareness of domain-specific terminology and translation memories, along with the ability to incorporate such resources into the workings of the MT engine will be vital components that enable MT systems to evolve with language, as well as increase their accuracy for any given source text.

We currently see a keen competition between rule-based and statistical approaches, where the former are on average better in the generation of linguistically well-formed results, whereas the latter are typically better in the integration of heterogeneous types of knowledge, including domain-specific terminology and world knowledge. It is a highly interesting question how these different qualities will affect tools for human post-editing. Ultimately the goal would be to integrate as many relevant knowledge sources as possible, and we believe that the hybrid approaches that are currently under development, once implemented and optimized for a certain application domain, have the potential to minimize the time human post-editors spend revising MT output.

4.2 Post-editing user interfaces

Despite recent progress in MT quality and the possibility of continuous improvement, it would be misguided to expect fully automatic translations of a publishable quality except for in very rare circumstances. Most applications of outbound translation will continue to require human effort in order to amend the output of MT. However, it is not implausible to assume that today’s predominant setup, in which MT and human post-editing are strictly sequential phases, could be significantly reconsidered in the near future. We could imagine a much closer integration of MT functionality with user interfaces, whereby human translators have a significant impact on the computation as it is occurring. Work in this direction needs to be guided by empirical studies, involving prototypes, in order to find out which of the many possible functionalities are actually beneficial, and how they should be presented to the user. For instance, as alluded to earlier, it appears beneficial to present the user with several potential translations in cases where there are several high-probability alternatives for a given source sentence. Nonetheless, if we were to push this idea to an extreme and always display several choices, then this would clearly strain the attention of the user, diminishing productivity.

Several recent studies towards better user interfaces (*such as work within the TransType projects [2, 7, 8] and in EuroMatrixPlus [6]*) indicate a high potential for novel interfaces, and have started to collect empirical evidence on the time-saving effects of certain kinds of MT post-editing interactions. In what follows, we would like to indicate some possible extensions to these approaches, aiming at a translator’s workbench that will combine existing technologies in novel ways to optimally support translators—both casual and professional.

Layout of texts

Source and target texts should be presented in a way that makes the correspondence between the texts easy to spot, e.g., in a side-by-side view. This will minimize the effort to search for corresponding phrases on screen. Scrolling through source and target text should be automatically synchronized with scrolling of the corresponding text. There ought to be a clearly visible distinction between raw MT output and those parts that have already been proof-read/edited. For instance one could display text generated by the MT system in gray, whereas text entered during postediting could be black. This

distinction should not only be provided on the textual level, i.e., to mark paragraphs and sentences, but also on the intra-sentence level. Such a feature would allow the translator to concentrate his/her efforts on the areas requiring the most work.

Auxiliary tools

The user interface needs to include a number of auxiliary tools that help to accelerate the search for optimal expressions in the target language. Starting with conventional utilities for spotting problems in text quality via spelling-, grammar- and style checkers, but also including tools for context-sensitive lexical look-up that display a number of likely expressions for a given source language term (*e.g., activated by hovering the mouse over the source term*). It would be very easy to pick the right choice from such a list, and insert it into the target text.

Editing functionalities

The interface must offer effective methods for correcting the target text. A large subset of the following operations could be offered:

- mark a word or an expression as undesirable
- select a better alternative for a word or an expression
- move a part of the sentence around
- change morphological property (*such as number, case*) of a complex phrase in a single operation, i.e., propagation of grammatical changes throughout the sentence and document

Automatic update

After any number of edit steps of these types, the user would have the option to request a re-computation of the translation. The new translation takes the changes entered by the user into account, but is allowed to adapt the parts not yet affected by user input to the new constraints. In many cases these constraints will lead to a higher quality of MT output, and even offer the opportunity to use the translator's feedback as further training data – adjusting the automatic update process accordingly.

Automatic alignment

Input to the system may come in a variety of different formats, from web sites, to various file types, to manually-entered text. Robust filtering is needed in order to chunk the original text into sections that will achieve the largest number of properly aligned target sentences. In situations where the system is unable to find an alignment (*or defaulted to a different translation algorithm*), the sentence should be flagged to the user as requiring immediate attention, and the toolkit should guide the user through correction of the errors. Furthermore, in order for the system to take advantage of the user corrections

that will be entered, it will be necessary to compute alignments between the edited text and the original on-the-fly. Although a discussion of efficient ways to achieve this is beyond the scope of the paper, it is clear that such an alignment could also be used to keep track of the parts of the original sentence that the user has translated, and to highlight any portions that have been missed. It is this continuous push and pull between user and system that will generate the highest quality translation most efficiently. In a similar vein, the toolkit could contain functionality for the automatic alignment of source and target text, thereby allowing the user to provide previously translated texts from a similar domain, and have these texts be used in order to improve translation quality in this and contextually-similar situations. Once again the user would have the option to correct alignment errors, thereby teaching the system new translations that may be useful, and will usually be of a high quality (*though obviously this should be verified separately*).

Feedback into MT engine

The toolkit should allow quick updates of the translation engine's underlying models based on post-editing, and automatic re-translation of those parts of the current document that have not been reviewed. The quicker the system can learn, the more attractive it will be for the user to help the system to learn the right things.

4.3 Collaborative infrastructure

An important goal of the infrastructure we intend to build for EuroMatrixPlus is the ability for a group of users to collaborate on the post-editing of larger documents, such as translations of long Wikipedia articles or other open-content material, such as, e.g., documentation of open-source software. Users should be made aware of the progress by others working on related material, and the improvements in the translation models used by a group of users should be shared. This is most easily realized when the translation engines are not installed independently on computers of different users, but run on a jointly used server (*an approval process would need to exist in order to ensure the translation quality was not being purposefully damaged*).

Another challenge is the need to keep track of changes in all different language versions of a multilingual set of corresponding documents. Assume a document such as a Wiki page has been translated into several languages at a certain point in time. Now any of the resulting documents may be changed by users in that language, and as each such change may add information that would also be relevant for the renderings in the other languages, one will at least consider corresponding updates of pages in all languages. A simplified scenario of this type has been described in [3], but as this scheme does not make any use of machine translation, all that can be done in their case is the management of information about which part of a document in some language is still in synch with versions in other languages and would need to be updated in order to re-establish cross-lingual correspondence. Generalisation of this scheme to the WikiTrans environment will allow us to go a step further: In cases in which updates happen to a version in some

language, one could automatically propagate them to the other languages via domain-optimized MT, and offer the readers in the other languages the choice between a human written (*but partially outdated*) version or a version that is up to date (*in which parts have been automatically translated*). The fact that this functionality would be embedded in a Wiki context would facilitate the resolution of such divergences via post-editing—thus synchronising all versions.

5 Conclusions

We hope that by offering a brief overview of current advances in the field of MT-assisted human translation toolkits, we will provide a framework for future developments, as well as motivating a best-practice approach to the empirical studies necessary in this area. It is clear that there exists a need for more effective tools that leverage the power of MT, with the flexibility and efficiency of social translation and crowd-sourcing methods. Achieving close integration without losing sight of the human is essential in all “social” applications, but even more so for translation, where the user is not just a consumer with some editing capacities, but actually moulds the underlying system being used.

The suggestions that we have offered for future development are by no means perfect, nor are they exhaustive. Our suggestions represent a class of improvements that must be incorporated in order to achieve the integration between MT-system and translation model that we discussed in the first section. Such tools will also serve to mitigate some of the disadvantages of over-training, as well as keeping the system current in a way that is otherwise currently unachievable.

Although we strongly believe that our recommendations would increase translation quality and efficiency in these situations, it is important to remember that more research and empirical studies are required in order to understand and find the most effective implementation methodologies, and that these are (*for the most part*) techniques that make the robustness and quality of the resulting MT system highly dependent upon having a threshold of dedicated users—which has been historically difficult to achieve in an open-source translation setting.

References

- [1] Chris Callison-Burch, Colin Bannard, and Josh Schroeder. Improved statistical translation through editing. In *Proceedings of the 2004 meeting of the European Association for Machine Translation*, Malta, 2004.
- [2] George Foster, Philippe Langlais, and Guy Lapalme. User-friendly text prediction for translators. In *2002 Conference on Empirical Methods in Natural Language Processing (EMNLP 2002)*, Philadelphia DC, 2002.
- [3] Louis-Philippe Huberdeau. Architecture for collaborative translation synchronization. Unpublished Draft, available from <http://blog.lphuberdeau.com/wordpress/wp-content/uploads/2008/02/architecture.pdf>, 2008.

- [4] Martin Kay. The proper place of men and machines in language translation. *Machine Translation*, 12:3–23, 1980/1997. First appeared as a Xerox PARC Working paper in 1980.
- [5] Philipp Koehn. A process study of computed aided translation. *Machine Translation*, 23(?), 2009. Accepted for publication.
- [6] Philipp Koehn and Barry Haddow. Interactive assistance to human translators using statistical machine translation methods. In *MT Summit XII*, 2009.
- [7] Elliott Macklovitch. The contribution of end-users to the transtype2 project. In *Proceedings of the 6th Conference of the Association for Machine Translation in the Americas (AMTA-2004)*, Washington DC, 2004.
- [8] Elliott Macklovitch. Transtype2: The last word. In *Proceedings of 5th International Conference on Language Resources and Evaluation (LREC-2006)*, Genoa, Italy, 2006.