Interactive Poster: The Chinese Room – Understanding and Correcting Machine Translation

Joshua S. Albrecht* Rebecca Hwa† G. Elisabeta Marai‡

Department of Computer Science
University of Pittsburgh

Abstract

We present an interface for visualizing information from a diverse type of linguistic resources to facilitate the correction of machine translated sentences. The goal of this project is to aid users with little or no fluency in the language of the original text to gain a greater comprehension of machine translations. This paper describes our interface and the linguistic resources it exploits to help users to detect and correct translation errors.

Index Terms: H.5.2 [Information Interfaces and Presentation]: User Interfaces—Graphical user interfacesI.2.7 [Artificial Intelligence]: Natural Language Processing—Machine translation

1 Introduction

The field of Machine Translation (MT) is concerned with developing methods for automating the task of translating between two natural human languages, such as Chinese and English. Because this task is challenging even for skilled human translators, and requires a considerable amount of world knowledge that cannot be easily encoded in straightforward algorithms, the sentences produced by current MT systems are often difficult or impossible to understand. Consider the following example output:

> "He utter eyes and not the slightest attention As leakage."

The resulting output is more accurately described as a jumble of words than an English sentence, even though a very high quality MT system produced this output. The translation for the original sentence should have been:

> "His eyes were wide apart; nothing in their field of vision escaped."

Many researchers are working hard to improve MT directly by creating better algorithms and fully automatic systems. We pursue an alternative solution – allow human users access to the significant amount of information available to an MT system, and let the user correct the MT output. This idea has been proposed as far back as 1980 by Martin Kay [2], but that work and subsequent approaches have focused on improving the performance of professional translators. Our goal is to allow even users who are not bilingual to gain the information contained in the original sentence.

To this effect, we have begun creating an interface for understanding and correcting MT outputs. To do so, we must present users with considerable information beyond machine translation, such as the original source language text, dictionary translations for each source language word, parse trees and parts of speech for the source language words, the ability to retranslate selected segments, the ability to see likely alternative translations, and a search function to find source phrases in other contexts. All of this information could easily overwhelm the user, so it is essential that we design an interface that presents it in an understandable manner.

*e-mail: jsa8@cs.pitt.edu
†e-mail:hwa@cs.pitt.edu
‡e-mail: marai@cs.pitt.edu

2 Methods

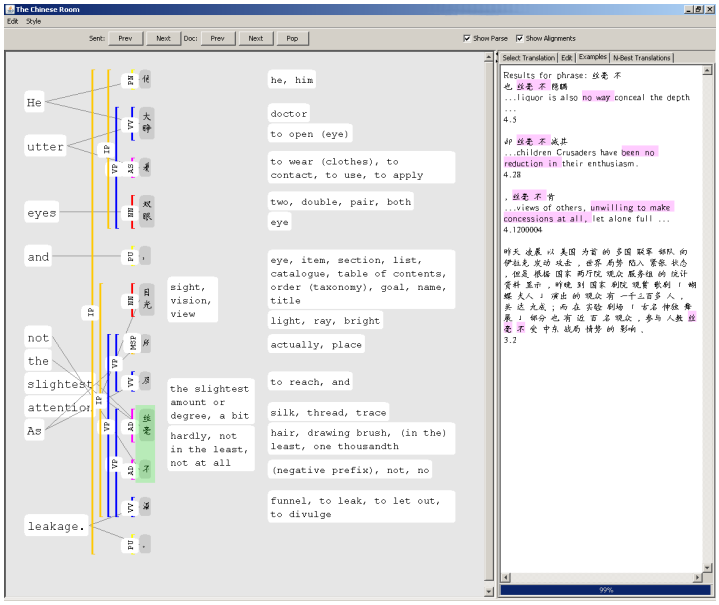

Figure 1: The left pane is the main sentence view, showing machine translation (white words on left), source sentence (column of grey words), alignments (grey lines), parse tree (colored lines, shown in more depth in Figure 2), and dictionary definitions (both columns of white boxes on the right). The first column of definitions is for whole words, the second is for definitions of single characters. The right pane is the extra information pane, currently showing the results of a search for the selected phrase (highlighted in the main view in green.) Matches in the search results are highlighted in pink.

In order for a user to understand a source sentence even in the presence of many MT errors, we have attempted to provide access to as much information as possible. First, we display the original sentence, the machine translation, and the alignments (mappings) between the words and phrases in each. These alignments, like much of the information we can display, are not always correct (since they are automatically generated). Nonetheless, they are useful for showing why the machine translation created a certain English word or phrase for some Chinese words. We also use the alignments to guide the layout of the translation by clustering phrases together, but users are free to rearrange the words and definitions on the screen. They can also interactively edit the words

and definitions as they attempt to understand the sentence.

Another major source of information is a bilingual dictionary (to provide definitions in English for Chinese words). This is challenging because there are often alternative segmentations of the Chinese sentence into words that would lead to different definitions. We make clear the definition association by highlighting alternative segmentations when the user mouses over a definition.

We display automatically generated phrase structure annotations for the Chinese sentence. The user can expand and collapse the phrases to keep the workspace visually uncluttered. The role (noun, verb, adjective, etc.) that a word fulfills, as well as the structure of the phrases both offer clues to the meaning of original sentence.

Finally, we allow a user to interactively select Chinese phrases and search for similar phrases in large monolingual and bilingual corpora. The user can then inspect the translation of these similar sentences to better understand how a given phrase should be translated. Alternative translations can also be requested from the original MT system, which provides a number of possible translations.

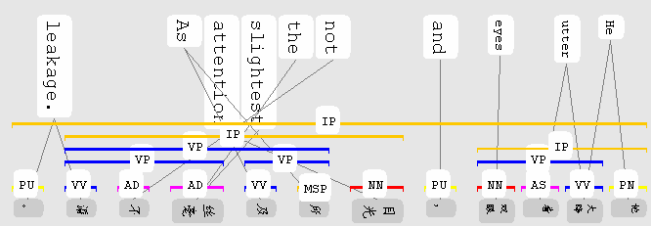

Figure 2: View of entire parse tree. Colored lines correspond to syntactic groups (constituents), which can be of various types, as indicated by the labels and doubly encoded with color (red is for nouns, blue is for verbs, etc). The phrase "he utter eyes" has been collapsed by the user to take up less display space.

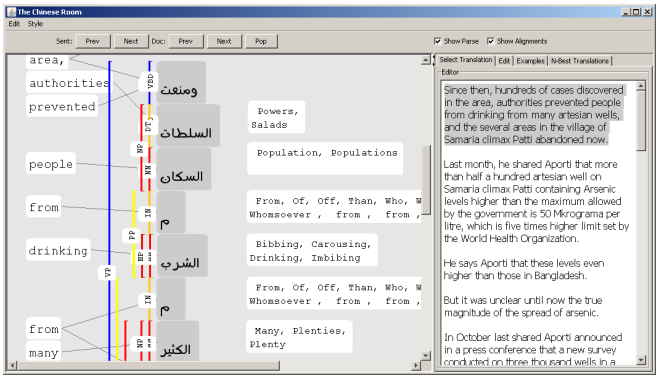

Figure 3: Our system also supports Arabic. Note that there is only one column of definitions because word segmentation is not an issue in Arabic. The right pane shows the document overview; the sentence currently selected by the user is highlighted.

3 Results

Figure 1 shows a screen shot of the entire application. The screen is split into two panes. The right pane contains 4 different views that can be accessed by the tabs at the top. It is currently displaying results for a search for the selected phrase (highlighted in the main view in green.) Matches in the search results are highlighted in pink. The user can verify from these search results that the correct translation for the current phrase is roughly "not at all", since that is how similar phrases were translated in the examples. The program searches a large monolingual corpus as well, and provides automatic translations of these matches at the request of the user. Other tabs on the right panel can be used to switch between search results, alternative translations for the selected phrase, a document overview (showing all sentences at once) and an editing pane (where the current translation can be edited). Figure 2 shows the entire sentence, with a better view of the whole parse tree. Figure 3 shows an example of machine translation from Arabic.

4 Discussion

The current design is the result of multiple iterations based on the feedback from expert users, who now find the tool quite useful for correcting broken machine translations. We are currently conducting experiments with novice users to determine how useful the interface is for a wider range of users.

In addition to evaluation, there are a number of interesting remaining challenges. First, the many types of information displayed (alignments, segmentation, parse trees, definitions, etc) often contain mistakes. Finding methods to convey the certainty of these resources may help users discover and correct areas of poor translation. Another challenge is simplifying the parse tree so that it can be more easily understood by novice users. Finally, we have already created basic support for Arabic as well as Chinese, and would like to extend it to other languages.

5 Related Work

While there is a considerable body of knowledge on how to create MT systems (see [3] for an overview), there has been surprisingly little work in the area of visualizing machine translation. Recently, a tool called DerivTool [1] was created for the purposes of interacting with the core of an MT system. However, their focus was on directly improving a specific MT system, not attempting to empower users with MT technology.

There are a number of established methods for visualizing trees, such as the work of [4], which has inspired our parse tree designs. Additional relevant work includes the work of Tufte, whose Principles of Design have inspired our graphical design [5].

6 Conclusion

We have presented our current research in the area of visualizing Machine Translation, and very briefly demonstrated the capabilities of our existing system. Our interface is designed to present a large number of linguistic resources traditionally used for MT to a user, with the hope that a human can help solve some of the large issues facing MT currently. Future work includes resolving several remaining interface issues with respect to the amount of data available to the user and the reliability of that data, as well as continuing with the evaluation of the interface for its intended task. A formal user study will provide valuable feedback about the effectiveness of this approach.

References

[1] S. DeNeefe, K. Knight, and H. H. Chan. Interactively exploring a machine translation model. In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 97–100, Ann Arbor, Michigan, June 2005. Association for Computational Linguistics.

[2] M. Kay. The proper place of men and machines in language translation. *Machine Translation*, 12(1-2):3–23, 1997. Originally circulated as Xerox Tech Report CSL-80-11.

[3] A. Lopez. Statistical machine translation. *ACM Computing Surveys*, 40(3), September 2008.

[4] T. Munzner, F. Guimbretière, S. Tasiran, L. Zhang, and Y. Zhou. Treejuxtaposer: scalable tree comparison using focus+context with guaranteed visibility. *ACM Trans. Graph.*, 22(3):453–462, 2003.

[5] E. R. Tufte. *Envisioning Information*. Graphics Press, Cheshire, Connecticut, USA, 1990.