

CASMACAT: A Computer-assisted Translation Workbench

V. Alabau*, C. Buck[‡], M. Carl[†], F. Casacuberta*, M. García-Martínez[†]
U. Germann[‡], J. González-Rubio*, R. Hill[‡], P. Koehn[‡], L. A. Leiva*
B. Mesa-Lao[†], D. Ortiz*, H. Saint-Amand[‡], G. Sanchis*, C. Tsoukala[‡]

*PRHLT Research Center, Universitat Politècnica de València

{valabau, fcn, jegonzalez, luileito, dortiz, gsanchis}@dsic.upv.es

[†]Copenhagen Business School, Department of International Business Communication

{ragnar.bonk, mc.isv, mgarcia, bm.ibc}@cbs.dk

[‡]School of Informatics, University of Edinburgh

{cbuck, ugermann, rhill2, pkoehn, hsamand, ctsoukal}@inf.ed.ac.uk

Abstract

CASMACAT is a modular, web-based translation workbench that offers advanced functionalities for computer-aided translation and the scientific study of human translation: automatic interaction with machine translation (MT) engines and translation memories (TM) to obtain raw translations or close TM matches for conventional post-editing; interactive translation prediction based on an MT engine's search graph, detailed recording and replay of edit actions and translator's gaze (the latter via eye-tracking), and the support of e-pen as an alternative input device. The system is open source software and interfaces with multiple MT systems.

1 Introduction

CASMACAT¹ (Cognitive Analysis and Statistical Methods for Advanced Computer Aided Translation) is a three-year project to develop an advanced, interactive workbench for computer-assisted translation (CAT). Currently, at the end of the second year, the tool includes an array of innovative features that combine to offer a rich, user-focused working environment not available in any other CAT tool.

CASMACAT works in close collaboration with the MATECAT project², another open-source web-based CAT tool. However, while MATECAT is concerned with conventional CAT, CASMACAT is focused on enhancing user interaction and facilitating the real-time involvement of human translators. In particular, CASMACAT provides highly interactive editing and logging features.

¹<http://www.casmacat.eu>

²<http://www.matecat.com>

Through this combined effort, we hope to foster further research in the area of CAT tools that improve the translation workflow while appealing to both professional and amateur translators without advanced technical skills.




Figure 1: Modular design of the workbench: Web-based components (GUI and web server), CAT server and MT server can be swapped out.

2 Design and components

The overall design of the CASMACAT workbench is modular. The system consists of four components. (1) a front-end GUI implemented in HTML5 and JavaScript; (2) a back-end implemented in PHP; (3) a CAT server that manages the editing process and communicates with the GUI through web sockets; (4) a machine translation (MT) server that provides raw translation of source text as well as additional information, such as a search graph that efficiently encodes alternative translation options. Figure 1 illustrates how these components interact with each other. The CAT and MT servers are written in Python and interact with a number of software components implemented in C++. All recorded information (source, translations, edit logs) is permanently stored in a MySQL database.

These components communicate through a well-defined API, so that alternative implementations can be used. This modular architecture al-

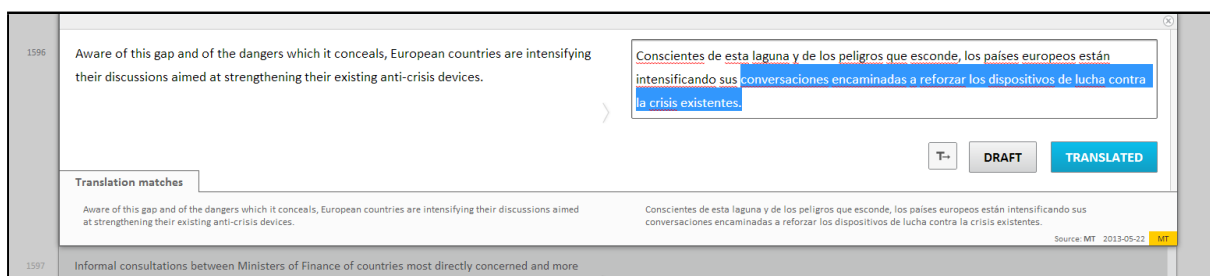


Figure 2: Translation view for an interactive post-editing task.

allows the system to be used partially. For instance, the CAT and MT servers can be used separately as part of a larger translation workflow, or only as a front-end when an existing MT solution is already in place.

2.1 CAT server

Some of the interactive features of CSMACAT require real-time interaction, such as interactive text-prediction (ITP), so establishing an HTTP connection every time would cause a significant network overhead. Instead, the CAT server relies on web sockets, by means of Python’s Tornado.

When interactive translation prediction is enabled, the CAT server first requests a translation together with the search graph of the current segment from the MT server. It keeps a copy of the search graph and constantly updates and visualizes the translation prediction based on the edit actions of the human translator.

2.2 MT server

Many of the functions of the CAT server require information from an MT server. This information includes not only the translation of the input sentence, but also n-best lists, search graphs, word alignments, and so on. Currently, the CSMACAT workbench supports two different MT servers: Moses (Koehn et al., 2007) and Thot (Ortiz-Martínez et al., 2005).

The main call to the MT server is a request for a translation. The request includes the source sentence, source and target language, and optionally a user ID. The MT server returns an JSON object, following an API based on *Google Translate*.

3 Graphical User Interface

Different *views*, based on the MATECAT GUI, perform different tasks. The translation view is the primary one, used when translating or post-editing, including logging functions about the

translation/post-editing process. Other views implement interfaces to upload new documents or to manage the documents that are already in the system. Additionally, a replay view can visualize all edit actions for a particular user session, including eye tracking information, if available.

3.1 Post-Editing

In the translation view (Figure 2), the document is presented in segments and the assistance features provided by CSMACAT work at the segment level. If working in a post-editing task without ITP, up to three MT or TM suggestions are provided for the user to choose. Keyboard shortcuts are available for performing routine tasks, for instance, loading the next segment or copying source text into the edit box. The user can assign different status to each segment, for instance, “*translated*” for finished ones or “*draft*” for segments that still need to be reviewed. Once finished, the translated document can be downloaded in XLIFF format.³

In the translation view, all user actions related to the translation task (e.g. typing activity, mouse moves, selection of TM proposals, etc.) are recorded by the logging module, collecting valuable information for off-line analyses.

3.2 Interactive Translation Prediction

Here we briefly describe the main advanced CAT features implemented in the workbench so far.

Intelligent Autocompletion: ITP takes place every time a keystroke is detected by the system (Barrachina et al., 2009). In such event, the system produces a prediction for the rest of the sentence according to the text that the user has already entered. This prediction is placed at the right of the text cursor.

Confidence Measures: Confidence measures (CMs) have two main applications in

³XLIFF is a popular format in the translation industry.

MT (González-Rubio et al., 2010). Firstly, CMs allow the user to clearly spot wrong translations (e.g., by rendering in red those translations with very low confidence according to the MT module). Secondly, CMs can also inform the user about the translated words that are dubious, but still have a chance of being correct (e.g., rendered in orange). Figure 3 illustrates this.

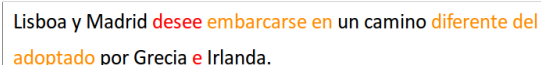


Figure 3: Visualisation of Confidence Measures

Prediction Length Control: Providing the user with a new prediction whenever a key is pressed has been proved to be cognitively demanding (Alabau et al., 2012). Therefore, the GUI just displays the prediction up to the first wrong word according to the CMs provided by the system (Figure 4).

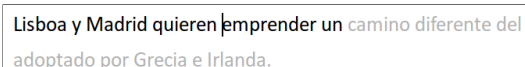


Figure 4: Prediction Length Control

Search and Replace: Most of CAT tools provide the user with intelligent search and replace functions for fast text revision. CASMACAT features a straightforward function to run search and replacement rules on the fly.

Word Alignment Information: Alignment of source and target words is an important part of the translation process (Brown et al., 1993). To display their correspondence, they are highlighted every time the user places the mouse or the text cursor on a word; see Figure 5.

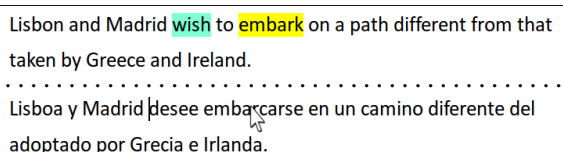


Figure 5: Visualisation of Word Alignment

Prediction Rejection: With the purpose of easing user interaction, CASMACAT also supports a one-click rejection feature (Sanchis-Trilles et al., 2008). This feature invalidates the current prediction made for the sentence that is being translated, and provides the user with an alternate one.

3.3 Replay mode and logging functions

The CASMACAT workbench implements detailed logging of user activity data, which enables both automatic analysis of translator behaviour and retrospective replay of a user session. Replay takes place in the translation view of the GUI and it displays the screen status of the recorded translation/post-editing process. The workbench also features a plugin to enrich the replay mode with gaze data coming from an eye-tracker. This eye-tracking integration is possible through a project-developed web browser extension which, at the moment, has only been fully tested with SR-Research EyeLinks⁴.

4 E-pen Interaction

E-pen interaction is intended to be a complementary input rather than a substitution of the keyboard. The GUI features the minimum components necessary for e-pen interaction; see Figure 6. When the e-pen is enabled, the display of the current segment is changed so that the source segment is shown above the target segment. Then the drawing area is maximised horizontally, facilitating handwriting, particularly in tablet devices. An HTML canvas is also added over the target segment, where the user's drawings are handled. This is achieved by means of MINGESTURES (Leiva et al., 2013), a highly accurate, high-performance gesture set for interactive text editing that can distinguish between gestures and handwriting. Gestures are recognised on the client side so the response is almost immediate. Conversely, when handwritten text is detected, the pen strokes are sent to the server. The hand-written text recognition (HTR) server is based on iAtrous, an open source HMM decoder.



Figure 6: Word substitution with e-pen interaction

5 Evaluation

The CASMACAT workbench was recently evaluated in a field trial at Celer Soluciones SL, a language service provider based in Spain. The trial involved nine professional translators working with the workbench to complete different post-editing tasks from English into Spanish. The pur-

⁴<http://www.sr-research.com>

pose of this evaluation was to establish which of the workbench features are most useful to professional translators. Three different configurations were tested:

- PE: The CASMACAT workbench was used only for conventional post-editing, without any additional features.
- IA: Only the Intelligent Autocompletion feature was enabled. This feature was tested separately because it was observed that human translators substantially change the way they interact with the system.
- ITP: All features described in Section 3.2 were included in this configuration, excepting CMs, which were deemed to be not accurate enough for use in a human evaluation.

For each configuration, we measured the average time taken by the translator to produce the final translation (on a segment basis), and the average number of edits required to produce the final translation. The results are shown in Table 1.

Setup	Avg. time (s)	Avg. # edits
PE	92.2 ± 4.82	141.39 ± 7.66
IA	86.07 ± 4.92	124.29 ± 7.28
ITP	123.3 ± 29.72	137.22 ± 13.67

Table 1: Evaluation of the different configurations of the CASMACAT workbench. Edits are measured in keystrokes, i.e., insertions and deletions.

While differences between these numbers are not statistically significant, the apparent slowdown in translation with ITP is due to the fact that all translators had experience in post-editing but none of them had ever used a workbench featuring intelligent autocompletion before. Therefore, these were somewhat unsurprising results.

In a post-trial survey, translators indicated that, on average, they liked the ITP system the best. They were not fully satisfied with the freedom of interactivity provided by the IA system. The lack of any visual aid to control the intelligent auto-completions provided by the system made translators think that they had to double-check any of the proposals made by the system when making only a few edits.

6 Conclusions

We have introduced the current CASMACAT workbench, a next-generation tool for computer assisted translation. Each of the features available in the most recent prototype of the workbench has been explained. Additionally, we have presented an executive report of a field trial that evaluated genuine users’ performance while using the workbench. Although E-pen interaction has not yet been evaluated outside of the laboratory, it will be the subject of future field trials, and a working demonstration is available.

Acknowledgements

Work supported by the European Union 7th Framework Program (FP7/2007-2013) under the CASMACAT project (grant agreement n° 287576).

References

- Vicent Alabau, Luis A. Leiva, Daniel Ortiz-Martínez, and Francisco Casacuberta. 2012. User evaluation of interactive machine translation systems. In *Proc. EAMT*, pages 20–23.
- Sergio Barrachina et al. 2009. Statistical approaches to computer-assisted translation. *Computational Linguistics*, 35(1):3–28.
- Peter Brown et al. 1993. The mathematics of statistical machine translation: Parameter estimation. *Computational linguistics*, 19(2):263–311.
- Jesús González-Rubio, Daniel Ortiz-Martínez, and Francisco Casacuberta. 2010. On the use of confidence measures within an interactive-predictive machine translation system. In *Proc. of EAMT*.
- Philipp Koehn et al. 2007. Moses: Open source toolkit for statistical machine translation. In *Proc. of ACL*, pages 177–180.
- Luis A. Leiva, Vicent Alabau, and Enrique Vidal. 2013. Error-proof, high-performance, and context-aware gestures for interactive text edition. In *Proc. of CHI*, pages 1227–1232.
- Daniel Ortiz-Martínez, Ismael García-Varea, and Francisco Casacuberta. 2005. Thot: a toolkit to train phrase-based statistical translation models. In *Proc. of MT Summit X*, pages 141–148.
- G. Sanchis-Trilles et al. 2008. Improving interactive machine translation via mouse actions. In *Proc. of EMNLP*, pages 485–494.