

10 The relevance of metadata during the localisation process: An experiment

In localisation, there is a constant need to automate processes in an attempt to reduce the cost and time associated with those processes. One of the main ways of achieving this objective is to reuse previously localised data and metadata by using standardised translation memory formats – such as the LISA Translation Memory eXchange (TMX) format or the OASIS XML Localisation Interchange File Format (XLIFF). Although XLIFF has emerged as the most desirable format for the exchange of localisation data, much work still needs to be done before it can be generalised to most localisation activities and can be implemented in a straightforward way. Our research aims to study the importance of the localisation metadata associated with the translation suggestions provided by computer-assisted translation (CAT) tools. We analysed the ways in which localisation data and metadata (D/M) can be represented in the current specification of XLIFF (i.e. XLIFF 1.2). We designed a new format called the Localisation Memory Container (LMC) to organise previously localised XLIFF files in a single container. Finally, we developed a prototype (XLIFF Phoenix) to leverage the D/M from the LMC to improve the translator's task by helping CAT tools not only to produce more translation suggestions easily (similar to bi-text approaches), but also to enrich those suggestions with relevant metadata. In order to test whether the 'CAT-orientated' enriched metadata are helpful to the translator, we designed an experimental translation task using Swordfish II. Different levels of D/M were introduced in the same document for the different participants, which enabled us to study the influence of the D/M on the translation decision-making process and on the machine–human interaction. In this chapter, we present our research objectives, the methodology and procedures, and an analysis of the results of the experiment. Finally, we extract some preliminary conclusions about the importance of metadata during the localisation process, both as displayed in our experiment and as an integrated function of the tool's automated capabilities.

Keywords: Computer-assisted translation (CAT), localisation, metadata, XLIFF

Introduction

In this chapter, we present the experimental work that we carried out to elucidate the relevance of metadata in the localisation process. By metadata, we mean the type of information that usually accompanies computer data (e.g. a text file, a JPEG image file, a whole CD) and that tells us more about the actual data. This information is used either with a descriptive function for the ‘reader’ of the data (i.e. a computer user interested in the properties of a file; a car CD player showing titles, names of artists) or with a performative value, causing the performing software to ‘make decisions’ based on such metadata (the simplest example being file headers, which help software read or play a file correctly).

The focus of our research is on the decision-making process carried out by a human translator in the process of localising software documentation with the help of computer-assisted translation (CAT) or localisation tools. In that context, metadata become relevant if they improve the process by making translations better, faster and/or cheaper.

The structure of this chapter is as follows: we start by providing an overview of the field being examined, that is, a brief description of the localisation standard that forms the basis for our research, and, in particular, the metadata that were used. In the following section, we present the tools that we used for the processing of data, of metadata and of documents for our experiment. Subsequently, a full account of the ways in which we designed and implemented the experiment is presented. This is followed by the preliminary results of the experiment and by an analysis of the limitations of our approach, which, in turn, serves as a basis for our future work. We conclude by providing recommendations and conclusions.

The object: XLIFF, translation matches and metadata

XLIFF as a container for localisation

XLIFF (XML Localisation Interchange File Format) is a data container that carries localisation content from one localisation process to another. XLIFF files, as containers of data, accommodate chunks of previously created information or groups of objects that are central to a localisation process. Such information or groups of objects consist of, for example, different objects or files in a localisation product (software, web, etc.), workflow information, supplementary material (translation memories, terminologies, etc.) and other customised data (XLIFF Technical Committee, 2008a).

The main feature of XLIFF is that it enables interoperability between different tools. For example, a localisation object might be converted into XLIFF, translated and commented on with one tool, before being exported into a different tool for further processing of the translatable elements, of memories or of any additional information. Finally, the localisation object can be returned to a translated version of its original file format. The XLIFF Technical Committee (2008a) provides a clear representation of the ways in which XLIFF is handled in this process (Figure 10.1). The original file is split into a Localisation Data section, whose translation units (TUs) can be processed by CAT tools, and a so-called skeleton section with non-localisable file structure data. These two parts are finally merged back into the translated material, generally in its original format.

XLIFF, in its basic form, is an XML language that may contain one or more files to be translated. As any other XML representation, it is structured by means of tags, consisting of elements that can have zero, one or more attributes. The root element `<xliff>` can contain one or more `<file>` elements. The `<file>` element [outer rectangle inside Figure 10.2] consists of both a `<header>` element, which contains metadata about the `<file>`, and a `<body>` element, which contains the extracted translatable data from the `<file>`. The translatable data within `<trans-unit>` elements [inner rectangle inside Figure 10.2] are organised into `<source>` and `<target>` paired elements (XLIFF Technical Committee, 2008a).

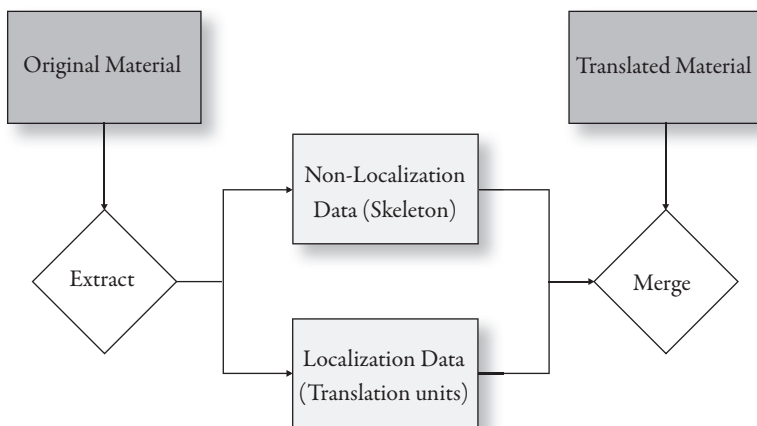


Figure 10.1: Extraction/merge principle of XLIFF
(XLIFF Technical Committee 2003: 10)

```

<?xml version="1.0" encoding="UTF-8"?>
<xliff version="1.2" xmlns="urn:oasis:names:tc:xliff:document:1.2">
  <file original="HelloTradumatica.txt" source-language="en" target-language="es"
    datatype="plaintext">
    <body>
      <trans-unit id="1" translate="yes" state="needs-review-translation">
        <source>Hello, Tradumatica!</source>
        <target>¡Hola, Tradumatica!</target>
      </trans-unit>
    </body>
  </file>
</xliff>
  
```

Figure 10.2: Basic XLIFF file

An XLIFF file can also accommodate data from previous localisation processes, such as translation matches for the source segments (either in the target language of the current project or any other languages that may be of use). This is achieved through the `<alt-trans>` element, which, among other possible attributes, includes information on how well-matched each source segment (i.e. the translation unit (TU) to be translated) and each translation suggestion is (Figure 10.3). This is how XLIFF is able to integrate both the files to be translated and the translation memories, which would be leveraged by means of the translation or localisation tools of choice.

```
<trans-unit id="1" translate="yes" state="needs-review-translation">
  <source>Hello, Tradumàtica!</source>
  <target>¡Hola, Tradumàtica!</target>
  <alt-trans alttranstype="proposal" match-quality="71.0">
    <source xml:lang="en">Goodbye, Tradumàtica!</source>
    <target xml:lang="es">¡Adiós, Tradumàtica!</target>
  </alt-trans>
</trans-unit>
```

Figure 10.3: alt-trans element in XLIFF

In Figure 10.4, we can see all the elements from the previous two examples displayed on the tool that we used in our experiments.

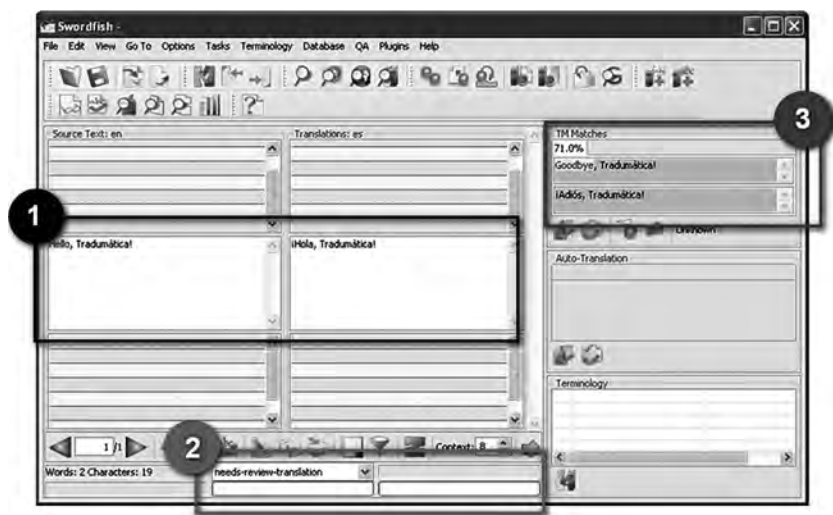


Figure 10.4: XLIFF data and metadata as shown by Swordfish II

Choice of XLIFF metadata

Metadata can be extremely important in various phases of the localisation process (Anastasiou and Morado, 2010), particularly when data are shared among people, in similar or different roles, along the globalisation cycle. However, not all metadata have the same relevance for localisation

professionals, and this depends very much on the CAT tools and the types of content being localised. XLIFF provides many mechanisms to add metadata related to the product, to the process, to the supplementary material or to any other localisation factors. However, little use is made of this functionality to improve the localisers' task by enriching the source data that they need to process. What is more, are the potential metadata available (in the current specification) embedded in the appropriate regions of the core structure so that the tools can make the most of them? Are localisers trained to perceive the importance of interacting with localisation metadata through their tools? Are tool makers aware of the need to improve the availability and the usability of such metadata?

In order to identify potentially relevant localisation metadata for our research, we first carried out a complete analysis of the current XLIFF 1.2 specification and we looked for attributes that could describe the localisation process by answering the traditional Zachman Framework (Warren, 2007) WH questions: when, who, where, when, why and how. See Table 10.1 for a schematic overview of our analysis.

Table 10.1: Zachman Framework applied to XLIFF 1.2 attributes

What	is it? is being localised?	original source-language target-language category version
	does it do?	datatype
Who	is localising/translating?	contact-name
	localised similar files earlier? (matches)	contact-email
Where	is it localised [company/country]?	company-name target-language
	is it localised [tool]?	tool-id tool-name
	is it going to be sold to/used by?	<i>skopos</i>
	did the translation matches come from?	state-qualifier

When	were matches obtained? does the translation need to be done by?	date <i>skopos</i>
Why	is it being localised? am I the translator?	job-id
	[purpose]	<i>skopos</i>
How	is it localised [tool, process, auxiliary material]?	tool-id tool-name
	is it being localised?	state state-qualifier

Our next step was to check whether these attributes were already present in the alt-trans element, which is the place where translation suggestions (i.e. matches) can be introduced in an XLIFF 1.2 file, as mentioned above.

Table 10.2: Attributes in alt-trans

3. Attribute	4. Allowed in the element	Included in our analysis
Alttranstype	Yes	No
Approved	No	Yes
Category	No	Yes
Company-name	No	Yes
Contact-email	No	Yes
Contact-name	No	Yes
Coord	Yes	No
Crc	Yes	No
Css-style	Yes	No
Datatype	Yes	Yes
Date	No	Yes
Exstyle	Yes	No
Extradata	Yes	No
Extype	Yes	No
Font	Yes	No
Format	No	Yes

3. Attribute	4. Allowed in the element	Included in our analysis
Help-id	Yes	No
Job-id	No	Yes
Match-quality	Yes	No
Menu	Yes	No
Menu-name	Yes	No
Menu-option	Yes	No
Mid	Yes	No
Origin	Yes	No
Original	No	Yes
Phase-name	Yes	No
Rename	Yes	No
Restype	Yes	No
Source-language	No	Yes
State	No	Yes
State-qualifier	No	Yes
Style	Yes	No
Target-language	No	Yes
Tool	Yes	No
Tool-company	No	Yes
Tool-id	Yes	Yes
Tool-name	No	Yes
Translate	No	Yes
Ts	Yes	No
Version	No	Yes
Xml:lang	Yes	No
Xml:space	Yes	No

In Table 10.2, we have summarised the results from the metadata obtained. In the current XLIFF version, only a set of predefined attributes are allowed to be introduced in the <alt-trans> element without breaking the validity

of the file. To avoid creating invalid files, we decided to include the items identified in our analysis through processing instructions (W₃C, 2008), which also allowed us to keep our ‘attribute-value’ structure intact for later reuse.

Pre-processing tools

Localisation memory container (LMC)

The information that is kept in the XLIFF file format can be reused for future projects. However, XLIFF was not originally designed to act as a translation memory (TM) in itself. The use of XLIFF as a TM was first associated to its conversion to the LISA Translation Memory Exchange (TMX) standard (Raya, 2004, 2005; XLIFF, 2008b), but the conversion from XLIFF into TMX tags, by mapping the former items with the equivalent ones in the latter format, results in the removal of many pieces of information that cannot be matched in TMX. This is, among other reasons, because both standards were created with different goals and they fulfil different necessities during the localisation process (XLIFF, 2008b). Instead of using that approach, we wanted to maintain the whole data set that an XLIFF document contains by storing the document itself inside a bigger container.

We worked on the design and development of a self-describing XML language that could contain XLIFF files and we called it the Localisation Memory Container. As its name suggests (and in contrast with traditional TMs), this language enables us to encapsulate much more than simple translation units. It also allows us to encapsulate workflow information, agents implied, tools used and, of course, translation data. An additional benefit of this approach is that in-context matches can be obtained, as the whole relevant document could be found and consulted, in contrast with TMX, where translation units are not necessarily found in a particular order.

In other words, the Localisation Memory Container was developed as a data descriptor to facilitate the storage of previous XLIFF documents in a single file. The LMC document has a <header> and a <body> as its main structural elements. In the header, the following four attributes can be found: (1) author (of the LMC), (2) last-mod (latest modification of the LMC), (3) total-files (contained) and (4) started (creation date). The body can include any XML data (e.g. TMX) that are correctly specified with a namespace. However, for the purposes of our research, we used only XLIFF files.

Along with the LMC language, we developed a tool that allowed us to bring the different XLIFF documents together in the Container: the LMC Builder.

XLIFF Phoenix

XLIFF Phoenix¹ is a CAT tool that allows for the reuse of previously localised XLIFF documents included in an LMC. This is achieved by filtering the information included in the documents and matching the documents with a new set of XLIFF documents imported into the system. The resulting document is an enriched file that contains translation recommendations (in the <alt-trans> element) and embedded metadata.

Two sets of files are loaded into the program (the XLIFF file to be translated and the LMC). The user can process the files and look for concordances among the translation units. If a translation match is found, it will be leveraged in the <alt-trans> element of the corresponding translation unit in the final enriched file along with the metadata associated with it.

1 The technical implementation of the tool was carried out between June 2010 and August 2010 by CSIS second-year intern Seán Mooney, who worked as an intern for eight weeks with a grant from the CNGL. He was supervised mainly by Lucía Morado Vázquez, Chris Exton and Dimitra Anastasiou, from the Localisation Research Centre (LRC), University of Limerick. The tool was developed using the Netbeans platform. The programming language chosen for this implementation was Java.

Finally, the enriched XLIFF file can be saved for later use in a CAT tool with XLIFF support. Figure 10.5 shows how our file is displayed in Swordfish II. On the ‘match properties’ window, you can see the metadata that XLIFF Phoenix recovered from the LMC.

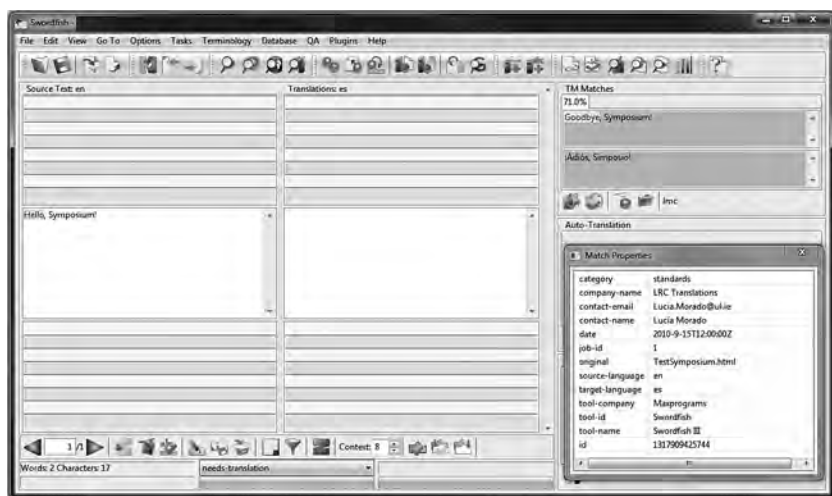


Figure 10.5: Enriched XLIFF file in Swordfish II

Design of the experiment

Few experimental studies into the use of translation memories have been conducted over the last years (Bowker, 2005; Christensen and Schjoldager, 2010). There are various reasons for this lack of research, ranging from the difficulty in obtaining valid participants, useful research objects and supporting data (e.g. documents, TMs), and the different rates of development between academia and the industry to the difficulty in controlling all the different variables, such as participant profiles, workflow, time pressure, text types, translation instructions, TM/MT programs and language combinations chosen (Pym, 2011).

Mechanics of the experiment

The experiment presented in this chapter was divided into three phases. In phase 1, the participants were asked to complete a questionnaire so as participant-related background information could be obtained. In phase 2, the focus of the experiment, the participants were given instructions for the translation assignment and they were asked to complete the translation of a document with the aid of our CAT tool. In phase 3, the participants were asked to fill in a task-related questionnaire so as the participants' impressions and opinions about the translation process could be obtained.

Date and place

The experiments were carried out on 16 December 2010 in a computer room in the Faculty of Translation and Documentation at the University of Salamanca. In order to avoid contamination or any leaking of the results, the whole process was carried out in a single session in the same room. Although all participants worked simultaneously in the same room, they were not allowed to communicate with each other.

Participants

The participants in this experiment were 10 students (9 women and 1 man) in their final year of their undergraduate programme in translation and interpreting. All participants had enrolled in an elective localisation module and had previous experience with CAT tools. In addition, all participants had selected the language combination 'English into Spanish'. Seven participants declared that, on average, they worked between 2 and 4 hours a day on translation activities, with 2 participants declaring that they spent 5 to 7 hours, and one participant claiming to spend at least 8 hours a day on translation-related tasks. None of the participants received any monetary compensation for participating in the experiment.

Translation assignment and data

In order to ensure a high degree of ecological validity, we designed a translation task that was as similar to a professional situation as possible.

Text type is one of the variables that, according to Pym (2011), is bound to affect the validity of an experiment involving TMs. It has always been suggested that certain text types are more inclined to work better with TM systems than others (Bowker, 2005; Christensen, 2003; Christensen and Schjoldager, 2010). Other studies also appear to establish a relation between the translation of technical texts and an increased use of TM systems by translators (Fulford and Granell-Zafra, 2005; Lagoudaki, 2006). This relation could be explained by the nature of the technical texts themselves, which generally use a limited range of terms and normally contain lexical and phraseological repetitions, all of which make this the most suitable text type for TM usage (Bowker, 2005; Christensen, 2010). On account of these findings, we decided to use a technical text in our experiment. This decision was also influenced by the access that we had to a wide range of technical texts, donated to us by Microsoft, one of the CNGL's industrial partners. More precisely, we were able to use the original Microsoft Office help documentation and its official Spanish translation.

One of the Microsoft's documents was selected for the translation task. Because of the estimated duration of the experiment (i.e. 1 hour), we extracted only the first 30 translation units and converted the document to an XLIFF file, which was handed out to the translators on the day of the experiment. Along with the XLIFF file to be translated, we also provided the students with the document in HTML format to allow them to see the text in its original context.

In order to enrich the XLIFF file to be translated with metadata, we carried out a series of processes. First, we aligned the original English file with the official translation file by using Maxprograms's Stingray Document Aligner. Subsequently, we manually modified some of the original translation units and their target segments to artificially produce translation matches with a match percentage lower than 100 per cent, or we introduced a local variant into the target segment. We also added more metadata items than the number of metadata present in the aligned document. After that,

we merged the XLIFF file into an LMC. Finally, we introduced the original English file in XLIFF format into XLIFF Phoenix and we processed it against the LMC. This resulted in an enriched file with translation suggestions and their corresponding metadata.

After obtaining our master version of the enriched file, we created three versions of it by taking into account different levels of data and metadata. In each version, the 30 segments were divided into three sections, and they arranged as follows: one section had no suggestions at all (referred to hereafter as ‘ \emptyset ’); another section had translation suggestions with only one metadata item – the translation match percentage or TM; the third section contained translation suggestions with ten items of metadata associated to it (referred to hereafter as ‘TM+MD’).

We split the participants into three groups (A, B and C). Each group received a different version of the master version of the enriched file (containing different data and metadata levels) as explained above. Figure 10.6 represents the six possible data/metadata distribution combinations, which would allow for a maximum of six groups of participants with different text types.

Since there were 10 participants, we decided to have only three groups and text types, as explained above. We chose three that did not start with the \emptyset data (no leveraged data or metadata), so each group of participants could enjoy some contact with the TM right from the beginning.

TM+MD	TM	TM
TM	TM+MD	\emptyset
\emptyset	\emptyset	TM+MD
TM+MD	\emptyset	\emptyset
\emptyset	TM+MD	TM
TM	TM	TM+MD

Figure 10.6: The six possible data/metadata distribution solutions

Since there were 10 participants, we decided to have only three groups and text types, as explained above. We chose three that did not start with the \emptyset data (no leveraged data or metadata), so each group of participants could enjoy some contact with the TM right from the beginning.

GROUP A	GROUP B	GROUP C
TM+MD	TM	TM
TM	TM+MD	\emptyset
\emptyset	\emptyset	TM+MD

Figure 10.7: Chosen distribution of data

Nature of the data and metadata introduced into the translation suggestions

The translation suggestions were introduced following the same match pattern for the three different text types and groups of participants. This means that they always had the same match level in the same data/metadata part. For instance, if in group A the TM+MD section had two exact matches in all the segments, group B also had two exact matches in all the segments of the TM+MD section. One problem that we encountered was that even if we actually introduced two suggestions per segment, our tool fully prioritised the segment with the highest match and, in order to visualise the second suggestion completely, the participants were required to click on the appropriate tab (see Figure 10.8). Although the participants had had prior training in using the tool itself and they were aware of the possibility to use the tool, the recordings of the experiment showed us that they hardly ever consulted the second option, let alone accepted it. We could anticipate here that participants would prioritise the translation match percentage over any other information that they might get.

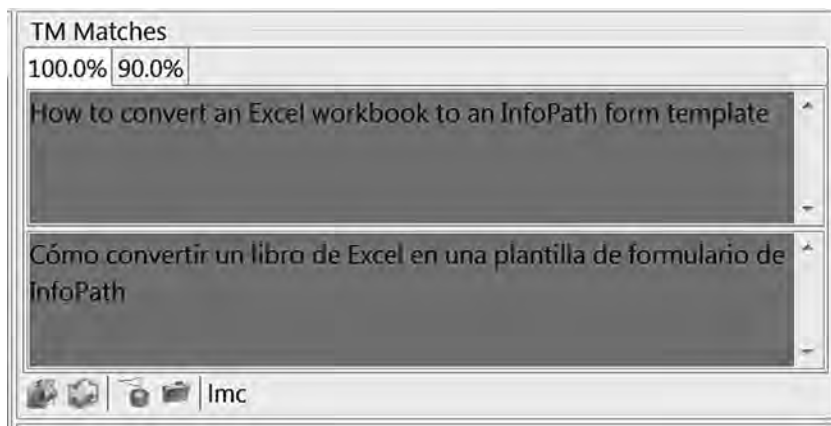


Figure 10.8: Several translation suggestions shown in Swordfish II

Translation assignment and instructions

Before starting with the translation task, the participants were given general instructions for their assignment. They were asked to complete the translation using Swordfish II and any other useful resource that they were able to access via a browser and an Internet connection. They were told to behave as they would in a real translation assignment scenario, with the only restriction being that the participants were not allowed to communicate with each other.

The following information about the translation suggestions that they might encounter was offered to the participants: the participants were given the names of the official translators of Word and Excel and they were told that four different sources were in use for the translation suggestions: the Spanish-Spain locale, the Spanish-Mexico locale, the Galician-Spain locale or a machine translation system. They were also told that the most recent official translations were produced after 2008. The last instruction given to the participants was the recommendation that they should use their common sense in case of doubt and that blindly following the TM would

not the most advisable course of action (Bowker, 2005). We concede that these instructions could condition the participants' answers to our questions and could have led the participants to concentrate on a particular set of metadata items. On the other hand, our intention was to put the participants to work in an authentic/real scenario, in which the translation suggestions could carry metadata that would somehow be relevant and meaningful for them, and to study if this information could influence the participants' decisions in any way.

Data-collection process

In order to obtain the maximum amount of information (i.e. data) from the experiment and to enhance its validity (Oates, 2006: 37), we decided to use three different data-collection methods: (1) questionnaires, (2) translated documents and (3) recordings of screen movements. With respect to the two questionnaires, the first questionnaire was conducted at the beginning of the experiment and sought to gather participants' background data (e.g. age, translation experience, experience with CAT tools). The second questionnaire was conducted just after the translation task and was related to the tasks that the participants had just completed. As far as the translated documents were concerned, we analysed the actual translations that the participants produced. For the recording of screen movements, we recorded the participants' on-screen behaviour (using CamStudio) to analyse their actions (e.g. use of external resources, time spent on each translation unit, use of the translation memory provided).

Tool used

The task was carried out using a customised version of Swordfish II. Its main developer, Rodolfo Raya, kindly agreed to modify the tool for us to have a window displaying metadata for the current translation suggestion (see Figure 10.5 above).

Preliminary results

The first lesson that we learnt after completing the experiment is that our approach of presenting two translation suggestions for the same segment and expecting the participants to choose between these two suggestions was unsuccessful, mainly for practical reasons. It was extremely time-consuming and the participants opted to follow only the first suggestion (i.e. the suggestion with the highest translation match rate). This last finding could lead us to think that the participants based their first judgements on the translation match metadata information.

Time

In terms of completion time, the average time for all the participants was 55.7 minutes. There was no significant difference between groups (Figure 10.9). However, if we take into account the three sections with different grades of D/MD, some differences emerge (Figure 10.10). The \emptyset section consistently took more time to complete than the other two sections. In two cases (groups A and C), the TM+MD section required less completion time than the TM and \emptyset sections.




Figure 10.9: Task completion time




Figure 10.10: Task completion analysed in sections

Quality

Wright (2006: 260–261) identifies three existing translation quality metrics: (1) SAE J 2450 (‘for the evaluation of translation quality for service documentation in the automotive industry’), (2) the ATA Framework for Standard Error Marking (‘a far more comprehensive evaluative tool than the SAE metric, but this higher level of complexity reflects the fact that it is designed for assessing a full range of text types and subject matters’) and (3) the LISA QA Model, which is intended ‘for measuring not only translation quality, but also all aspects of the localisation process as well’.

For the evaluation of our tasks, we chose the LISA QA Model, and more specifically, its 1.0 Printed and Online Documentation template (LISA, 1995), adapted to the requirements of version 3.1. The initial template starts with a certain number of points, each error found in the translation is annotated and it deducts different points depending on the category of the error and its severity.

In the LISA QA Model 3.1 (LISA, 2007), the quality analysis is divided into seven tasks: (1) Documentation Language, (2) Documentation Formatting, (3) Help Formatting, (4) Help Formatting-Asian, (5) Software Formatting, (6) Software Functionality Testing and (7) Documentation Formatting-Asian. Because of the nature of our translation task, we focused only on the first task (i.e. Documentation Language). This task has seven possible error categories:

1. Mistranslation;
2. Accuracy (sub-categories: omissions, additions, cross-references, and headers/footers);
3. Terminology (sub-categories: glossary adherence, context);
4. Language (sub-categories: grammar, semantics, punctuation and spelling);
5. Style (sub-categories: general style, register/tone and language variants/slang);
6. Country (sub-categories: country standards, local suitability and company standards) and
7. Consistency.

The analysis of the participants' translation task was carried out by a lecturer in localisation with extensive experience in the assessment of translation and localisation. The results of the analysis (Figure 10.11) showed us that the Ø section consistently obtained lower marks than the other two sections. The TM/MD section obtained a better assessment than the TM and Ø sections for Groups A and C (the same pattern was observed in the time analysis). You can see the results of the QA analysis in more detail in Figure 10.12 (Group A), Figure 10.13 (Group B) and Figure 10.14 (Group C).




Figure 10.11: LISA QA Analysis, all groups




Figure 10.12: LISA QA Analysis, Group A

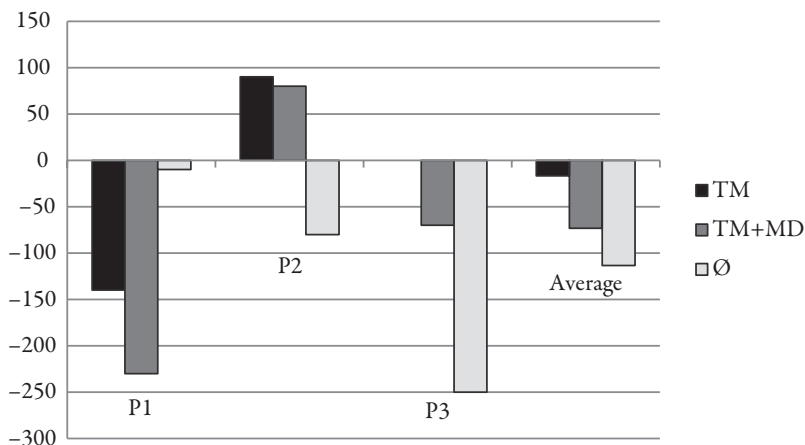


Figure 10.13: LISA QA Analysis, Group B




Figure 10.14: LISA QA Analysis, Group C

Relevance

In the task-related questionnaire, we obtained information about the participants' relation with the metadata presented during the translation task.

First, we asked the participants which of the three sections, if any, made their work easier (Ø, TM/MD or TM), with 9 participants preferring TM/MD and 1 participant choosing TM. A similar question [*If you had to choose between any of the three systems for your daily work, which one would you choose?*] was presented at the end of the questionnaire and the result was exactly the same. In relation to this finding, we asked the participants whether the MD had had a positive influence on their work, with 8 participants agreed and 2 disagreed.

Subsequently, we asked the participants about the level of distraction that the metadata produced in their work. We did this using a 5-point Likert scale, with 1 being ‘very confused’ and 5 ‘not distracted at all’. The average was 4.4.

Finally, we asked the participants which were the most useful/consulted metadata for them (see Figure 10.15) and they answered the following: Fuzzy Match, Contact-name, Date, Source language, Target Language and Category. Asked about the opposite, that is, the least useful items, the answer was the following: Origin, Company-name, Contact-email, Job-id, Tool-company, Tool-id, Tool-name and Tool-version (see Figure 10.16). The answers to these questions can be seen side by side in Figure 10.17.




Figure 10.15: Useful and consulted MD comparison




Figure 10.16: Less useful or not consulted MD




Figure 10.17: Useful and consulted vs less useful and not consulted MD

Limitations of the study and future work

The first aspect to be considered for a future experiment is the nature and the number of the raw materials and the representativeness of the sampled participants. Professional localisers and translators, as well as students, should be used in experiments to obtain a broader picture of the impact and relevance of metadata. An increase in the number of subjects would also improve the validity of results. Only one language combination was used. However, this should not affect the actual mechanisms at play, which have to do with a combination of translator and localiser competence, localiser-friendliness of the tool and the representative and performative potential of XLIFF.

If we look at the type of document used for translation, one piece of criticism that could be levelled at the experiment presented in this chapter is that it may be too static for the whole process to be regarded as an instance of 'localisation'. Although it can be considered localisation of software documentation (in html format), even this kind of product has some interactive features (in our particular document, hyperlinks, alternative definitions, etc.) that may need special attention. The process undertaken by the subjects was mostly seen as plain translation of technical content, which also limited (or was limited by) the scope of the questions that the metadata were able to answer (by confining them mainly to the area of translation suggestions). However, knowledge about the product itself, its functionality and interactivity, and the *skopos* of the localisation assignment and other process-related details can also be highly relevant for the localiser.

It would also be extremely important to analyse the cultural and socio-linguistic (e.g. text, discourse, grammar, lexis) conditions of the document, as well as the characteristics of the translation suggestions and associated metadata that can be injected through the LMC and the XLIFF Phoenix tools. This is particularly relevant if an error-based assessment of the subjects' production is to be made. In so doing, we can anticipate what may influence translation decisions concerning comprehension, reliance on the TM, and the participants' own function as localisers. Surely, a larger

corpus of related documents and products to derive translation memory from would be advisable too, which could also help to fine-tune the capabilities of both the LMC and the XLIFF Phoenix pre-processing tools.

As for the use of one particular localisation tool, familiarity with it on the part of the participants is important. However, the main functionalities of CAT tools are well known and fairly stable by now, regardless of the actual interface and the workflow design. What concerns us most is the ability of our tool to present the information – or perform – on the basis of relevant metadata provided in the XLIFF file to the participants. This, in itself, that is, the rethinking of translation and localisation tool design, together with the enhancement of metadata in the localisation process, is the goal of our research. What is more, as the tool is not the best-known of translation and localisation tools, having our participants become acquainted with it might be an advantage, since they would not be conditioned by their own pre-conceptions and usage habits. Nonetheless, one thing that we should take into account is the (in)ability of the tool to display the product in a WYSIWYG mode, or any other visual and functional aids that have to do with the information that our metadata can provide.

Finally, further preparation of the ways in which the LISA QA model or any other qualitative or quantitative analysis methods are used would also be advisable. In general, the relevance (within the field of research) of the performance of the localisers should be re-evaluated, and some kind of comprehensive yet flexible benchmark performance criteria should be established. As far as translation decisions and actions are concerned, eye-tracking software could also be a valuable addition to any future experiments.

Recommendations and conclusions

Because of the limitations of this initial experiment, we cannot draw definitive conclusions about the relevance of enhancing metadata information for the localisation process. However, there are some clear indicators that point to the validity of our hypothesis, and we presented some findings that

could modestly be presented as recommendations to the XLIFF Technical Committee.

Our experiment has indicated that some metadata elements are taken into consideration more often than others. Although these elements seemed more relevant to our participants, this may also have to do with the training of students or with the experience of the translators with the translation process or the use of CAT tools. There is a second important factor that needs to be taken into account and that we cannot control a priori: the usability of the metadata provided has a lot to do with the ways in which the tool presents the data or influences the data. We hope that this caveat can serve as a reminder to translator training institutions and to tool makers that it is in everyone's interest to provide translators with the cognitive baggage, the skills and tools for a critical approach to the translation and localisation task. Awareness of the relevance of process-, product- and function-related information should be combined with its usable integration into both representation languages (e.g. XLIFF) and professional software.

As for the increase in the quality and productivity that enriched translation suggestions seem to offer, it could be argued that the participants may have been 'blindly' confirming high translation matches, which would speed up the process (particularly when compared with those parts where no suggestions are provided and the participants (translators) are left to their own devices to create new translations). Also, since matches were mainly extracted from the official Microsoft translation, this source could also explain the better performance. However, there are at least four elements to bear in mind before accepting these arguments offhandedly. First, we must not underestimate the whole revision process carried out (as well as the time spent) by the subjects on both their own production and the matches that they incorporated into their work. Second, as Gow (2003) points out in her master's thesis, more time is often spent on validating, amending or improving the textual cohesion of translation suggestions than on producing translations from scratch. Third, the localisation lecturer who assessed the participants' performances did not regard all the segments from the official Microsoft translation as the best possible solutions. Sometimes, the alternative solutions provided by the participants presented fewer QA problems than the official solutions. And fourth, the

metadata included in the translation suggestions can give the translator (and the tool) useful cumulative information for the decision-making process, which seems to affect the translator's learning curve with respect to the document or product to be translated.

In sum, we can say that the main lesson learnt from this experiment is that all the different variables that may have compromised the validity of our experiment have also become for us the main assets for a fresh new approach to investigating the relevance of metadata in localisation. In the next steps that we will take, we will concentrate on the specific information that translators and localisers need to deal with to improve their performance. One of the features that we will look into is translator and localiser competence, and how this assembly of knowledge and of skills can be represented and acted upon by the forthcoming XLIFF specifications and the tools that support this standard. This will help us to enhance our future studies into the effects of these resources on translators' work, ultimately with a view to influencing developments in the industry as well as in the academic field. Hopefully, this will empower translators and localisers in the critical task of mediating texts, products and services between cultures and/or markets.

Acknowledgements

We would like to thank Rodolfo Raya for his advice and time and for allowing us to use his tool in the experiment. In addition, we would like to express a heartfelt thank you to Dag Schmidkte and to Microsoft for providing us with original Microsoft material and with TMs. We would also like to thank the XLIFF TC for its support and advice. Finally, we would like to express our gratitude to the localisation students from the University of Salamanca who participated in the experiment and to the participants in the pilot study presented in this chapter.

This research has been supported by the Science Foundation Ireland (Grant 07/CE/I1142) as part of the Centre for Next Generation Localisation at the University of Limerick.

Bibliography

- Anastasiou, D., and Morado, L. (2010). Localisation standards and metadata. In *Metadata and Semantic Research, 4th International Conference, MTSR 2010 Proceedings, Communications in Computer and Information Science* (pp. 255–276). Berlin: Springer.
- Bowker, L. (2005). Productivity vs quality?: A pilot study on the impact of translation memory systems. *Localisation Focus*, 4(1), 13–20.
- Christensen, T. P., and Schjoldager, A. (2010). Translation-memory (TM) research: What do we know and how do we know it? *Hermes – Journal of Language and Communication Studies*, 44, 89–101.
- Fulford, H., and Granell-Zafra, J. (2005). Translation and technology: A study of UK freelance translators. *The Journal of Specialised Translation*, 4, 2–17.
- Gow, F. (2003) *Metrics for Evaluating Translation Memory Software* (Master's thesis), Ottawa: University of Ottawa. <<http://www.localisation.ie/resources/Awards/Theses/Metrics%20for%20Evaluating%20Translation%20Memory%20Software.doc>>.
- Lagoudaki, E. (2006). Translation memories survey 2006: Users' perceptions around TM use. In *Proceedings of the ASLIB International Conference Translating & the Computer 28* (pp. 1–21). London: ASLIB Publications.
- LISA. (1995). *LISA QA MODEL Reference Manual* <<http://www.lisa.org/>>.
- LISA. (2007). *LISA QA Model 3.1 User's Manual* <<http://www.lisa.org/LISA-QA-Model-3-1.124.0.html>>.
- Oates, B. J. (2006). *Researching Information Systems and Computing*. London: Sage.
- Pym, A. (2011). Democratizing translation technologies: The role of humanistic research. *Luspio Translation Automation Conference*, Rome, Italy, 5 April, <http://usuaris.tinet.cat/apym/on-line/research_methods/2011_rome.pdf>.
- Raya, R. (2004). *XML in Localisation: A Practical Analysis*. <<http://www.ibm.com/developerworks/xml/library/x-localis>>.
- Raya, R. (2005). XML in localisation: Reuse translations with TM and TMX. <<http://www.ibm.com/developerworks/library/x-localis3>>.
- W3C. (2008). *Extensible Markup Language (XML) 1.0* (5th edn) <<http://www.w3.org/TR/REC-xml/#sec-pi>>.
- Warren, S. (2007). The Zachman enterprise framework', *Cambridge Technical Communicators*, <http://www.technical-communicators.com/articles/zachman_framework.pdf>.

- Wright, S. E. (2006). The creation and application of language industry standards. In K. J. Dunne (Ed.), *Perspectives on Localization* (pp. 241–278). Amsterdam: John Benjamins.
- XLIFF Technical Committee. (2003). *A White Paper on Version 1.1 of the XML Localization Interchange File Format (XLIFF)*. <http://www.oasis-open.org/committees/download.php/3110/XLIFF-core-whitepaper_1.1-cs.pdf/>.
- XLIFF Technical Committee. (2008a). *XLIFF Version 1.2 Specification*. <<http://docs.oasis-open.org/xliff/xliff-core/xliff-core.pdf>>.
- XLIFF Technical Committee. (2008b). *Frequently Asked Questions*. <<http://www.oasis-open.org/committees/xliff/faq.php#XLIFFvsTMX>>.