[*First International Conference on Language Resources & Evaluation*, Granada, Spain, 28-30 May 1998]

Localisation is good for you: The Localisation Resources Centre as an Example for the Successful Co-operation between Industrial Users and Academic Researchers

Reinhard Schäler
University College Dublin / Localisation Resources Centre
Belfield, Dublin 4, Ireland
[Reinhard.Schaler@ucd.ie]

Abstract

The Localisation Resources Centre (LRC) at University College Dublin (UCD) has been established as the focus point and the research and support centre for the localisation industry. It also co-ordinates the activities of the Software Localisation Interest Group (SLIG), an association of publishers, manufacturers, localisation vendors, and localisation tools developers. The LRC and SLIG operate a joint membership.

Through its unique position at the intersection of industrial user requirements and academic research interests it has initiated a number of research projects with a very clear focus on industrial user requirements.

This article provides some background information on the localisation industry and gives an overview of the aims and the activities of the Localisation Resources Centre. It concentrates on one particular aspect of its operation, the evaluation of translation technology for localisation, using the ETAT project as an example.

It will be become evident that 'localisation is good for you' not only if you are directly involved in the localisation business as a software publisher or localisation vendor, but also if you are a researcher in areas related to language engineering. Few industries have been so receptive to innovative approaches to natural language processing as the localisation industry. Few industries have shown such a degree of interest not only in the **use** of new language technologies, but also in the research of new approaches.

Background

Ireland is, after the U.S., the world's second biggest exporter of software. In 1995, Ireland produced 5 billion pounds worth of software, almost all of which was exported. Over 40% of all packaged software and 60% of business applications sold in Europe are produced in Ireland.

A key element of the industry in Ireland and a principal factor behind its success is localisation: the adaptation of a product to the requirements of a foreign market, its language and 'locale'[1]. The process of localising a software product is complex and extends far beyond translation. Cultural factors must be considered, such as the use of colour, style, forms of address, and the selection of images and graphical representations. Very practical considerations require the conversion of units of measure and standards such as weights and currencies. However, the single most expensive activity when localising a software product is still translation, closely followed by testing.

Since the mid 1980's, Ireland has become a world centre for software localisation. There are around 5,000 people directly employed in the industry, with approximately twice that number of people working in related industries. Companies involved in localisation in Ireland include software publishers (Microsoft, Oracle, Lotus, Corel etc.), manufacturers (Gateway 2000, Iomega, Creative Labs etc.), and service providers (Berlitz, Bowne Global Solutions, LioNBRIDGE, ITP etc.).

The Localisation Resources Centre (LRC)

The LRC was established under Forbairt's Technology Centres Programme in December 1995 at University College Dublin (UCD) with support from the European Regional Development Programme. (Forbairt is the Irish state agency for industrial development and scientific research.)

Thus the Irish government was recognising the substantial contribution of the field to the national economy and responding to the call by the Software Localisation Interest Group (SLIG) to establish a focus point, and a research and support centre for the Irish based localisation
industry.

The LRC has become a Centre of Excellence providing research and support for the localisation industry. It has the support of dozens of companies involved in localisation (software publishers, manufacturers, localisation service providers) and the development of localisation tools (mainly machine translation, computer assisted translation, and automatic testing system developers). It co-operates at national and international level with researchers and students, the media, consultancy firms, government agencies and the European Commission.

In its two years of existence, the LRC has been described as an example of how government money should be invested sensibly (Perkin, 1996) and has helped in attracting large investment into the country.

The activities of the centre currently cover:

1 A locale is "the features of the user's environment that are dependent on language, country, and cultural conventions." (Kano, 1995)

- Research, development and evaluation of electronic tools and resources relevant to the localisation industry.

- Consultancy and training for industry, academic and official/state institutions, including the European Commission.

- Infrastructural support for the localisation industry, e.g. regular publications in printed and electronic formats, Localisation Information Exchange Point (LIEP, the web-based database for information exchange), and the co-ordination of the Software Localisation Interest Group (SLIG)[2].

Co-operation between industry and academic partners: Evaluation of translation technology for localisation

One of the main reasons for the establishment of the LRC was the expressed requirement of the industry to automate at least some of the most costly parts of the localisation process. Despite its image as a high-tech industry, a great number of tasks in the localisation process are still carried out manually and are, as a result, very labour intensive and costly. Irish-based localisation companies have over the past years carried out expensive market research and evaluation exercises to identify suitable 3rd party tools. It is widely accepted in the industry that in order to avoid unnecessary duplication of efforts and to make optimal use of existing resources, this work should be carried out by one independent centre.

Three-step approach

The LRC addressed this urgent need for more automation - particularly in the single most expensive task in the localisation process, translation - in a three-step approach:

Localisation Tools Library For most companies, it is not possible to dedicate sufficient resources to keep up-to-date with the rapid developments in the area of localisation tools and have direct access to these tools. The LRC, therefore, decided to establish a Localisation Tools Library to cater for the needs of these companies and provide access to a wide range of localisation tools.

Workshops The LRC organized a series of workshops around the theme "Criteria for the evaluation of translation tools for localisation". These workshops were designed to gather ideas from expert users in the field on criteria for the selection and use of translation technology for their localisation projects. Over a period of little more than a year, a total of 100 representatives from different localisation companies attended two workshops. As a direct result of these workshops the LRC formed a Working Group on Evaluation and produced a list of evaluation criteria based on:

(i) the results of the workshops (LRC 1996/97);

(ii) the criteria used during an internal evaluation carried out by a large tools developer[3];

(iii) the report by the European EAGLES consortium (EAGLES 1996).

Tools Development The LRC also developed a pre-prototype of the Evaluation Tool for Automatic Translation (ETAT) designed to help translation managers decide whether a particular project should be translated using computer assisted translation (CAT) tools. ETAT's design is based on the results of **initial** discussions between a number of localisation companies. which were further developed during the LRC workshops.

Evaluation Tool for Automatic Translation (ETAT)

One example of the successful co-operation between industrial users and academic researchers is the development of the ETAT pre-prototype by the LRC. The project illustrates how researchers active in language engineering can offer directly applicable solutions to the problems encountered by the localisation industry.

Genesis and Purpose

Over the past years the industry has come under enormous pressure for mainly economic reasons to release a growing number of more complex products in less time without an increase in the cost of localisation. Ideally, there should be no 'delta', i.e. the time between the release of the original version and the release of the localised versions of a product should equal zero. SimShip, the simultaneous shipment of different local versions of products has become one of the most widely used buzz-words in the industry.

In order to achieve SimShip without an increase in cost or a decrease in quality, companies needed to automate the localisation process. In the area of translation, it was one newly developed technology in particular which seemed to be well-suited to achieve these aims: translation memories. Berlitz, then Softrans-Berlitz, was one of the first companies to evaluate a number of systems and use them for large multilingual localisation projects. Other companies have also since used translation memory applications, such as the TRADOS Translator's Workbench, STAR's TRANSIT, or IBM's Translation Manager.[4]

While many localisation projects have been successfully completed with these tools and their usefulness in the context of localisation has now been firmly established, it has also become apparent that:

- Some translation memory applications are better suited for specific localisation projects than others.

- Some localisation projects are not suited to automatic translation involving translation memory technology.

[2] Since 1995 the manager of the LRC has also been the chairperson of SLIG.

3 This unpublished report has been made available to the LRC by a leading tools developer.

[4] See also Schäler, 1996.

Substantial expertise is required to decide which projects should (or should not) be localised using translation memory technology and which translation memory offers the best solution for a given localisation scenario.

This expertise is not always readily available. Experts in this new technology are still scarce on the ground. In addition, it is usually the translation manager, i.e. a non-technical person, who takes the decision. There are hundreds of translation managers, particularly in large organisations, such as Microsoft, Oracle, Lotus, Corel, Berlitz and ITP, who on a daily basis and within a very short period of time have to decide whether or not to use translation technology for particular projects.

While many companies have chosen a translation memory system that suits their localisation environment particularly well, the expertise and information necessary to decide in which case to use it are generally not available to translation managers.

Against this background a number of expert users and managers from large software publishers, localisation service providers and the LRC came together to discuss ways to provide translation managers with an automated decision support tool.[5]

Approach

An initial meeting was followed in October 1996 by a SLIG/LRC workshop at which Gunnie Jacobsson, then Vendor Manager at Microsoft WPGI, presented a proposal for a spreadsheet based evaluation checklist. This idea was further developed by the LRC and has lead to the development of the pre-prototype of the Evaluation Tool for Automatic Translation (ETAT). ETAT was presented to a follow-up workshop organized by the LRC in April 1997. The aim of this prototype was to solicit feedback from industry experts on the viability of the general approach to the development of a decision support tool for translation managers.

The pre-prototype implementation of ETAT evaluates the suitability of a specific translation memory (TM) tool for a given project. This calculation is based on detailed project information and TM tool information. Currently, both have to be supplied manually. ETAT then calculates a percentage value as a 'usefulness indicator' which suggests whether a TM tool could be useful for a particular project. Moreover, ETAT generates a report supplying the user with the rationale behind the calculations.

In practical terms, ETAT

- Asks the user a series of questions about their project.
- Compares this information with information that has been supplied to it about TM technology in general and a specific system in particular.
- Calculates usefulness values for each section of the project information supplied to it, and an overall usefulness value.
- Generates reports for each section identifying the main areas that reduced the usefulness values and pointing out areas of incompatibility between a specific TM tool and a project.

Under the following headings each of the three most important modules of ETAT will be described in more detail: (i) project information (ii) TM tool information (iii) calculation of the recommendation by the decision kernel.

Project Information

Users, in most cases translation managers, have to supply detailed information on a specific project to ETAT. Currently this information has to be input manually and on a project by project basis. The information is not stored and is only used to calculate usefulness indicators.

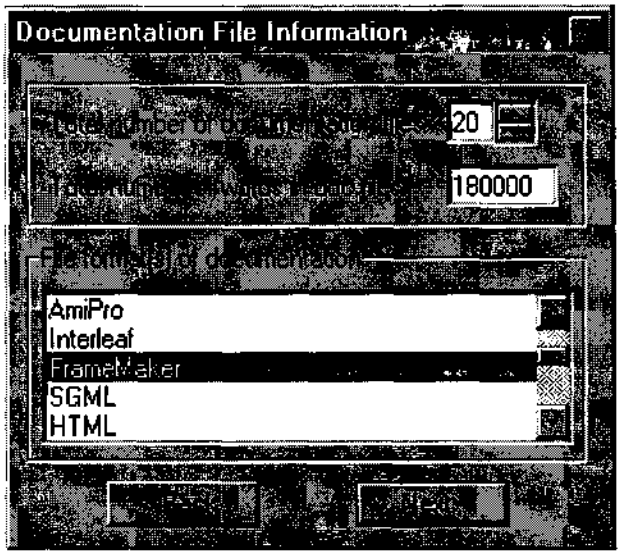

Figure l:One of several Project Information Entry Dialogs

Project information is collected under five headings[6]:

- General project information, i.e. expected number of versions over a given period (update cycle), number of languages, etc.
- Printed documentation, i.e. format of data files, number of words, etc.
- Online help system, i.e. format of data files, number of words, etc.
- User interface, i.e. type of UI, format and type of resources (e.g. Windows 3.1, '95/NT), etc.

5 See the acknowledgement section at the end of this article for details on the individuals and companies involved.

6 The current implementation of ETAT makes use only of the most important fields from the project information template, i.e. expected update cycle, file format, languages, etc.

- Glossary and terminology information, i.e. format of existing terminology, etc.

Once this information has been input, a usefulness value or indicator is calculated by comparing the project information with the detailed description of a particular TM system.

TM Tool Information

ETAT contains a generic template for the description of the most important features of TM applications in the context of software localisation[7]. This template currently addresses about 150-200 different features of TM applications, among them:

- General tool information, i.e. file formats supported, languages supported, import/export features, etc.
- Terminology management information, i.e. availability of integrated terminology management, terminology formats supported, support for glossary/terminology list creation, support for active terminology recognition and options for the automatic substitution or insertion of terms, etc.
- Alignment features i.e. support for automatic alignment of previously translated files, description of features and limitations of the alignment function, etc.

The current implementation of ETAT only supports one tool description at a time.

The description of TMs is stored in a template and can easily be edited or updated. However, this information will only change with updates to existing applications. Therefore TM descriptions can be updated and maintained at regular intervals by translation technology experts.

These descriptions can be re-used for the analysis of different projects. Indirectly, they also provide translation managers access to expert up-to-date information on all aspects of an individual translation tool that impact on its 'usefulness' in the context of a given project.

Decision kernel

During the development of ETAT much attention to detail was given to the compilation of the project and tool description templates. These templates form the basis for the calculation of the usefulness value by the system and the recommendation generated by it regarding the suitability of a TM application for a localisation project. It is expected that these templates will be re-used in future implementations of the decision support tool.

Given the practical constraints under which the development of ETAT had to be carried out and given its primary aim, i.e. the production of a pre-prototype to demonstrate a particular approach to the solution of a problem specified by industry experts, less resources were allocated to the development of the decision kernel which calculates the recommendation based on the project and tool information made available to it.

While the implementation of a more sophisticated decision engine, including a fully developed neural network, an expert system/ inference engine, or, indeed, a hybrid system consisting of elements of both, was discussed; it was not adopted. On the recommendation **of** the users, who preferred a relatively simple, but pragmatic, viable and easily implemented solution, ETAT currently uses relatively simple algorithms.

Example calculation When the usefulness value of a TM tool for the translation of *printed documentation* files is calculated, ETAT uses percentage values for its calculations. Different values are allocated for a number of different criteria according to their importance for the usefulness of the TM application in the context of this specific project[8]:

Criteria	**Percentage value allocated**
Support for document file format	10% (No support: 0%; Indirect support that involves some other process: 5 %; Full support: 10%)
Internal repetition (within the document)	20% (depending on % value of fuzzy matches, FM, and exact matches, EM)
Translation Memory (matches and quality)	70%
Quality of existing Translation memory	*21%* *(ranging between not approved, partially approved and fully approved)*
Matches with existing Translation Memory	*49%* *(using a function to adjust the linear match value for the allocation of percentage points)*

Table 1: Calculation of usefulness indicator for printed documentation files

When the usefulness value for translation memories is calculated, a function is used to adjust and increase the linear percentage value allocation in the case of an otherwise 'low' percentage match. For example, if there are 2,000 exact matches in 10,000 words, this would render a 20% match value. Even though there are a significant number of matches, this would not return a sufficiently high usefulness value.

Using a function as illustrated in figure 2, a more realistic value of about 44% is returned. The function returns a

[7] See also McDonagh/Schäler, 1997.

[8] These percentage values have been assigned by the developers based on their experience and based on the projects usually dak with. They might have to be readjusted given different circumstances.

relatively high value even for a relatively 'low' number of matches and, therefore returns more realistic values for the usefulness indicator.

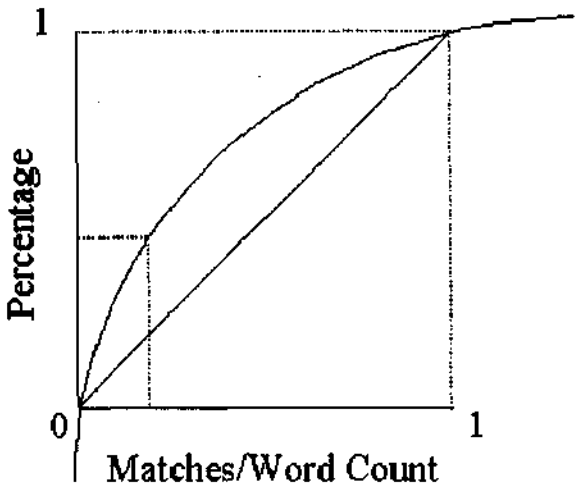

Figure 2: Adjustment of usefulness indicator

The current implementation of the decision kernel incompatibility only refers to what is considered to be the more important elements of the project information. It is planned to extend this coverage in future implementations.

Overall usefulness value After the calculation of values for the five individual sections of project information, an overall value is calculated based on weights allocated to each individual section. These weights are "customizable".

In the current implementation, the documentation, help and user interface sections have the same relative weight. This weight is adjusted according to word count, i.e. the usefulness value for the printed documentation section has a greater impact on the overall value than the user interface section if the printed documentation has a higher word count than the user interface.

Translation Router (TransRouter)

Following the highly successful demonstration of ETAT to localisation companies at the LRC workshop, the LRC looked for a suitable framework to develop this pre-prototype. The aim was the development of a prototype which would have more and better developed features and which would be stable enough to be tested by its potential users under industrial project conditions.

Together with other European partners - ISSCO, the University of Edinburgh, the University of Regensburg, the Centre for Language Technology (CST) Denmark, the Gesellschaft für Multilinguale Systeme (GMS) Germany, and Berlitz Ireland - the LRC proposed a project called " Translation Router" (TransRouter) under the 4th Call for Proposals of the Telematics Applications Programme, Language Engineering Sector. This project was to build on the ETAT pre-prototype and related work by other partners. It was to develop a decision support tool for translation managers who have to decide whether a specific project should be translated by a human translator, a human translator assisted by computer aided translation or by machine translation.

TransRouter was approved by the European Commission in late 1997 and started in early 1998.

For further information on TransRouter, please see the project's web site (TransRouter, 1998).

Conclusion

We consider that the LRC has shown how industrial and academic partners can work together successfully on many different levels. A prime example for this highly successful co-operation is:

- The establishment of the LRC's Tools Library,
- The subsequently held industry workshops focusing on criteria for the evaluation of translation technology systems (especially MT and translation memory systems), and following these
- The development of the ETAT pre-prototype and the subsequent TransRouter project proposal.

Already more than a dozen large developers and service providers have expressed their interest in joining the project through its user group. They would like to gain from the results of this highly practical and industry-oriented project which is based on the requirements specified by their peers and the work of some of Europe's most respected researchers.

We hope to have shown that 'localisation is good for you' not only if you are directly involved in the localisation business as a software publisher or localisation vendor, but also if you are a researcher in areas related to language engineering. Few industries have been so receptive to innovative approaches to natural language processing as the localisation industry. Few industries have shown such a degree of interest not only in the use of new language technologies, but also in the research of new approaches.

Bridges between industrial and academic interests have been built. We believe that organisations such as the LRC and SLIG are excellent examples for the successful co-operation between industrial users and academic researchers, which could be followed in other sectors such as the emerging call centre industry.

Acknowledgements

The Localisation Resources Centre (LRC) is part funded by the European Regional Development Fund through the Irish Office of Science and Technology.

The development of ETAT was based on ideas initially discussed during a meeting held at Microsoft WPGI which was attended by Catherine Gavin (Berlitz), Orlagh Neary (Corel), Sharon O'Brien (ITP), Reinhard Schäler (LRC), Gunnie Jacobsson (MS WPGI) and a number of other Microsoft representatives. The prototype was implemented by Darren Hogan (LRC) and Reinhard Schäler (LRC).

Translation Router (TransRouter) is a project funded by the European Commission's DGXIII under the Telematics Applications Programme, Language Engineering Sector (LE4-8345).

Bibliographical References

EAGLES (1996). Expert Advisory Group on Language Engineering Standards: *Evaluation of Natural Language Processing Systems*. 09 June 1996. http://www.ilc.pi.cnr.it/EAGLES96/home.html.

Kano, Nadine (1995). *Developing International Software for Windows 95 and NT*. Redmond, Washington (U.S.A.): Microsoft Corporation.

LRC (1996/97). Report on workshops: *Evaluation Criteria for Automatic Translation Tools for Localisation*, http://lrc.ucd.ie.

McDonagh, B. and Schäler, R. (1997). *Feature Check List for the Evaluation of Translation Memories with Test Suites*. Internal working paper available upon request from the LRC, Dublin (email: LRC@ucd.ie).

Perkins, Julian (1996). Found in translation. *Financial Times Survey: Ireland*, Financial Times, Monday 24 June 1996.

Schäler, Reinhard (1996). A Practical Evaluation of an Integrated Translation Tool during a Large Scale Localisation Project. In *Proceedings of the 4^{tn} Conference on Applied Natural Language Processing (ANLP-94)*, Stuttgart, Germany (October 13-15, 1994).

TransRouter (1998). The web site of the European Translation Router project (LE4-8345), Telematics Applications Programme, Language Engineering Sector, http://lrc.ucd.ie/TransRouter.