

The Language Translation Interface

Dominique Estival

Defence Science and Technology Organisation, Australia
Dominique.Estival@dsto.defence.gov.au

Abstract. The Language Translation Interface (LTI) is a prototype developed for the Australian Defence Organisation. The aim is provide a single, simple, interface to a variety of MT tools and utilities for personnel who need to produce translations when they have no easy access to human translators. Now that the LTI has been demonstrated and trialled at several military exercises, we are gathering user requirements to further develop it as the Language Translation Tools Suite. This paper describes the functionalities of the LTI and reports on our experience with users during development, leading to future improvements.

1. Introduction

I am very pleased to have been invited to give this opening talk at EAMT 2005, although I regret that Harry Somers cannot be with us in Budapest. Australia is a long way from Hungary and he is still enjoying his sabbatical there, but it was when I was talking to Harry a few months ago about some aspects of my work on MT at DSTO that he thought it would be interesting for the EAMT audience. I want to tell you about how we have been able to build a translation system without doing MT and about the way we dealt with the difficulties of getting access to users in our specific environment.

Before I tell you what I do at DSTO, I need to say a few words about what it is. DSTO stands for Defence Science and Technology Organisation and is it the R&D organisation for the Australian Defence Organisation. Our customers and end users are primarily the ADF (Australian Defence Force, the military side of Defence) and the ADO (Australian Defence Organisation, which also includes the civilian side of Defence), but also the Australian Government more generally. As an R&D organisation, DSTO may be more “small r and big D” than in earlier times, but it is committed to exploring and utilising technical innovations. In fact, our main role is to give advice on new technologies and to build prototypes to show what advantages these technologies can bring to the end users.

I am sure everyone knows where Australia is, but it is always interesting to see on a map

what the world looks like from our perspective. Australia is geographically isolated, certainly far away from Europe and North America, and further from Japan than people realise. We are part of the Pacific-Asia region, with very different linguistic neighbours than our traditional allies, the UK and the US. Australia is a “small” country in spite of its size, with a population of just over 20 million. Our resources are not huge, especially in terms of personnel. The environment, which is very harsh on most of the continent, also means that our technological requirements – and traditions – are quite different from those of European or North American countries. This is one of the reasons why Australia has a Defence R&D organisation, because technologies that may be appropriate for other countries need to be evaluated for our environment and sometimes new solutions need to be developed to meet Australian requirements. And we can argue that this is in fact the case when we look at MT and our linguistic environment. The languages spoken in our part of the world are not those that have traditionally been worked on for MT and, for many of them, NLP tools or resources are not even available.

In that context, at DSTO, I am now leading a research programme in language technologies for a variety of purposes. These include spoken dialogue systems (Estival et al., 2003), multi-modal interaction in a virtual environment (Estival et al., 2004), document classification (Carr & Estival, 2003), semantic clustering and language translation tools, which is the one I will

discuss today. That particular project started very small and was not considered particularly important in the beginning. It is still quite a small project in terms of its size, but I think it is interesting to see what we have been able to accomplish in that area, because that can be seen as an indication of the need for MT and of the value MT can bring to organisations which may not have been aware of their needs for it. This is where I hope my talk today will be of most interest to you, because I will talk about our experience in bringing some awareness of MT and of the need for MT to an organisation with no previous history of using language processing tools. I will describe the prototype tool we developed and how we went about to show it to potential users who do not have much time to play with software.

2. Project scope

Initially this project was only a small part of a larger project on Speech and Language Technologies¹ which was very much focussed on the speech aspects and which aimed at delivering speech interfaces in Headquarters environments. During demonstrations of the speech interface, people would ask the usual question: Can you do that in other languages? So, when I arrived at DSTO in March 2002, I was given the responsibility for looking at how that might be done and for assessing whether there was any potential for MT tools in the ADO. I was very lucky that at the same time, one of the students that we take every year for year-long projects, Jennifer Biggs, started at the same time and that she was very interested in that topic. At the time, Jenny had no background in machine translation, language processing, linguistics or computational linguistics and there was no one else working with us on this project. So I am fairly proud of the fact that, three years later, Jenny is still with me on a full-time contract, the LTI has been successfully demonstrated and it is being adopted by some sections of the ADO. In fact, Jenny is the person who actually built the LTI and, without her, the project would not have got off the ground. So what I want to talk about is what we did and how we did it.

¹ This project was initiated and led by Dr Ahmad Hashemi-Sakhtsari.

Having worked in MT before – in industry at Weidner in the US in the late 80s and in research at ISSCO in Geneva in the 90s –, my first assessment was that there was no point in us trying to build an MT system. We would have failed and not produced anything worthwhile. MT evaluation was another option, where the aim would have been to provide advice on what MT systems to purchase. That was not a very satisfactory proposition either: there was not enough funding to purchase systems to evaluate and more importantly, not enough trained personnel to perform the evaluation. I can look at French of course, Jenny knows Japanese and the task manager (Dr Ahmad Hashemi-Sakhtsari) could deal with Farsi, but again, it would have been very time-consuming and the results would probably not have even been worth reporting. So we settled for a survey of the tools available and for designing a way to make some of those tools accessible to our potential users. This resulted in the LTI (Language Translation Interface) and the LTDB (Language Translation DataBase).

The LTDB was a useful exercise for finding out what was available and we now use the information we collected to choose appropriate systems from within the LTI.

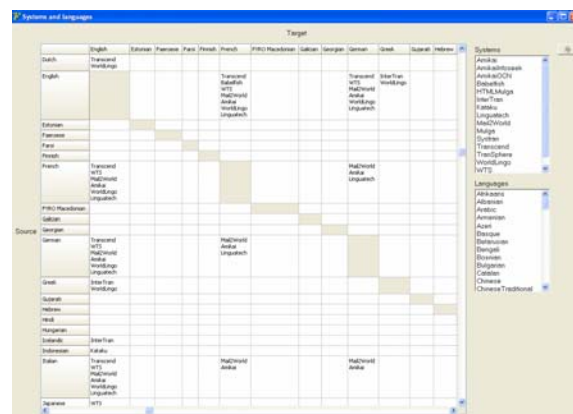


Figure 1. LTDB: Matrix of systems for language pairs

In the rest of this talk, I want to tell you about the technical design of the LTI and describe the functionalities of the two prototypes we developed; then I will talk about our experience in scoping out user requirements and setting up a trial system. I will conclude with what we have learned so far and where we are going with the continuation of this project, which we are now calling the Language Translation Tool Suite

(LTTS). But, first, I want to discuss why it would be worthwhile for the ADO to have translation tools in the first place.

3. Why would the ADO want Language Translation Tools?

The first point to make is that there is a growing recognition of the need for translation services in the ADO. This is a global issue which has only recently started to affect Australia but the ADO, like other Australian government agencies, is facing an increased demand for dealing with documents and information in languages other than English. This is especially true because the shift of focus for the ADO from “Defence of Australia” to “National Security” implies an increased awareness of the international environment around Australia. Other sources of demands for dealing with documents or communications in foreign languages include: intelligence gathering, coalition operations and foreign operations.

Intelligence gathering

I will not discuss intelligence gathering in great detail here, I imagine everyone in 2005 is aware of the intelligence failures which have been shown to precede the tragedy of 9/11 in the US and the ensuing discussions about the urgent need for better and more timely intelligence. The requests for more translators and for tools to help them have been widely publicised and Australia is in the same situation as all other countries in this respect. Of course, the Bali bombing in September 2002 and the bombing of the Australian embassy in Jakarta in October 2004 mean that there are also specific threats and concerns for Australia, with particular linguistic implications for us.

Coalition operations

Traditionally, our main allies are other English-speaking countries, such as the UK, the US, Canada and New-Zealand and, apart from the regular jokes about mutual unintelligibility of the various English dialects, there is not much need for translation between those countries. However, Australia also has strong ties with other nations in the Pacific region, and these countries do not all have English as their first language. It is also the case that military exercises have become increasingly multi-national and

that Australia is often involved in operations with a number of coalition partners whose first language is not English. Recent international exercises have included such countries as Japan, South Korea, Thailand or France, to name only a few.

The need for translation is not greatly felt in those exercises, because communications are assumed to be conducted in English. However, now that the technology allows e-mail communication not only in other languages but also, crucially, in other scripts, it is no longer the case that all communications during an operation will necessarily all be conducted in English, and it can be argued that Australians who are monolingual speakers of English will find themselves at a disadvantage when their coalition partners can choose to communicate in several other languages.

Foreign operations

Although Australia has participated in both Gulf Wars, over the past couple of decades, the ADF has been more involved in peace-keeping, humanitarian and relief operations in the Asia-Pacific region than in combat operations. For instance in the past few years, there have been operations in the Solomon Islands, in East Timor and in Aceh (Indonesia) after the tsunami. In this type of operations, there is a need not only to communicate with the population, but also to disseminate information, for instance by distributing leaflets or making radio broadcasts. From a technological point of view, the problem is that many, if not most, languages of the region are not covered by developments efforts for NLP and there are few, if any, computational linguistic resources for those languages. From the point of view of MT, it is not even possible to resort to building Translation Memories because there may not be enough texts available to build Translation Memories.

During foreign operations, there may also be situations on the ground where defence personnel might come into possession of documents or media (CDs, diskettes, computer hard-drives, etc.) which may contain crucial information. For example, when entering a building and seizing computers or filing cabinets. One issue here is the identification of the language or languages prior to translation, but there is also the issue of speed of access to translation services, whether

it be sending the document to a human translator in the field or back at home, or access to tools that could be used in the field or over a network.

So it is clear that there are great translation needs for an organisation like the ADO, and these needs have become apparent even to the more old-fashioned officers from a generation that used to consider English was all they needed. The question is: Can these needs be met by human translators?

First, we can make a comparison with the US. The US Department of Defence has a long tradition of training linguists and language specialists at the Monterey Defence Language Institute and, after 9/11, the FBI set up the National Virtual Translation Center to serve as a “clearinghouse for human translators” to “provide translation of foreign intelligence”. Nevertheless, the DoD also saw the need for developing the Phraselator (followed by the Speechlator), a PDA with limited speech translation capabilities which was first used in Afghanistan in 2002. The “Basic Language Translation Service” (BLTS) project, which is part of the larger “Horizontal Fusion” programme, now aims at developing automated language translation capabilities to meet the growing need for language translation in the battlefield (DoD, 2004). Looking at future research, on 18 March 2005, DARPA issued a Call for Proposals for a new research project, GALE (Global Autonomous Language Exploitation), whose goals are phrased as “eliminating the need for linguists and analysts” and “automatically ... interpret[ing] huge volumes of speech and text in multiple languages” (GALE, 2005).

In Australia, the ADO has also long recognised the need for personnel with linguistic skills and has its own training of linguists and translators, at the ADF School of Languages. Personnel receive training for spoken and written language skills in a number of languages that have been recognised to be of interest. However, these skills are mainly geared towards field operations and the training does not necessarily equip the linguists with specific translation skills.

As we all know, with the advent of e-mail and the internet, the number of documents which are of potential interest for intelligence gather-

ing has increased exponentially in the past decade. At the same time, the global growth of the internet and the development of electronic media for a large number of languages have eroded the dominance of English: although English is still the language of the majority of web pages, it is no longer the first language of the majority of web users. Many web sites and electronic communication channels (email, chat rooms, etc) now use other languages. These constitute sources of information which have to be taken into account by analysts. At the same time, these new media also constitute alternative channels for the dissemination of information to local populations during humanitarian and relief operations.

The problem is that it is not possible for the ADF School of Languages to train new linguists and translators for all the languages that might be of interest in the future.² It takes one to two years to train a linguist to attain a level of fluency in a language such that they can function using the spoken language. Training a translator/interpreter who can produce good translations may take another two to three years, depending on the language. However, it is very difficult to predict which languages are going to be of interest in a three year timeframe and even more difficult to predict the extent of the potential demand for translation for those languages. It would be impractical to train linguists in all the languages that might become of interest. Without expanding the size of the ADF, it is not possible to increase the number of recruits to be trained as linguists, because the existing personnel are already needed for other tasks and operations. However, the population of Australia is not of a size that can support a larger ADF at this time. In summary, given the size of the Australian population, there will never be enough personnel available to be trained and the range of languages of interest cannot be predicted

² In this respect, it is interesting to note the wide variety of languages spoken in Australia. This is not only due to the number of Aboriginal languages, which are of great interest linguistically but not so relevant for us at this time, but because of the large immigration from all over the world. As a result, there is in fact a sizeable pool of native speakers for many languages in Australia, but they would not all be available as translators for the ADO and their languages may not be those that are of interest.

in time to perform the training required to produce skilled translators in those languages.

So, given this situation, we have argued that automated translation tools can alleviate that problem by providing rough but usable translations which can either be used directly, for instance in the case of information gathering or of coalition operations, or which can be sent to a human translator for further editing if necessary, for instance in the case of foreign operations. Fortunately, this fits in quite well with recent ADO requirements for “increased efficiency through the use of automation in headquarters” and “the ability to work in multilingual environments”. This has been expressed as “computer aided comprehension of languages other than English”, and this is now part of the description of our project deliverables.

Since the start of this project, an overriding issue has been the constraint that neither the ADO nor the DSTO can realistically envisage to develop their own machine translation systems. Therefore we are limited to using existing systems, whether commercial off-the shelf (COTS) or freely available. Our focus is on developing easy access to existing translation engines and our main concern has been to make that access transparent to the users. The intention is to make available to ADO personnel existing tools which may increase the efficiency of current translation work and which would be appropriate in situations where there is a need for rapid translation and where no human translators are readily available.

4. The LTI

We have now produced and demonstrated several versions of the LTI. Two of them, the Translation Comparison Tool and the Web Translation Tool, deserve to be described separately because they illustrate quite different functionalities and because their interfaces look very different. We demonstrated them at several events within the last year and, after I describe the functionalities of the LTI, I will explain what those events were, who our audience was and what the outcomes were.

First, as I mentioned before, the LTI is not a translation system, but an interface to translation tools (Biggs and Estival, 2002; Estival and Biggs, 2003). The main idea was to provide a

single, simple, interface to as many translation systems as possible. We did not want to assume that our users would be trained translators, that they would know any other language besides English, or that they would be computer experts. We expect our users to be military or defence personnel, who are computer literate in that they know how to use a computer for basic e-mail, word processing and data entry, but not necessarily more. We first defined our users to be personnel who find themselves in positions where they have to get a translation for some form of document (for instance, participating in a coalition exercise or in a foreign operation) and in situations where they may not have access to a human translator (for instance, if there are no translators in the ADO for that language, or when there is not enough time to send the documents to a human translators). We also wanted the same tool to be useful to translators (military “linguists”) who could use it to get quick translation drafts and to build translation memories.

The first version of the LTI was the Translation Comparison Tool (TCT), shown in Figure 2.

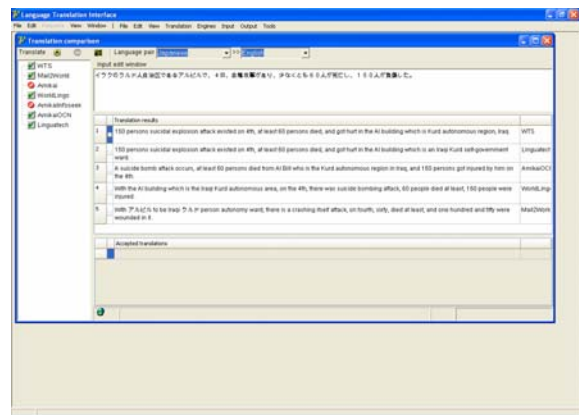


Figure 2. The LTI: Translation Comparison Tool

With this interface, the aim is to provide a set of translation results from as many translation systems as are available for the required language pair. The idea here is that if the users can view results from a number of systems, even if they have no knowledge of the other language, they may be able to make some useful comparison and select the most likely translation output. This is fraught with potential problems, which I

do not have to detail to you,³ but the main idea remains sound and it was met with great interest when we demonstrated it. The point here is not to dwell on the shortcomings of the individual systems, but to build upon the useable parts, if any, of the different outputs.

The emphasis for this tool was on the ease of use, at the three different stages of 1) input, 2) processing and 3) output. For ease of input, the user can choose to type text directly in the input window, or either load from a file or cut and paste from a file, or load a web page or an email message. Since most translation systems work best if the input is segmented into discrete sentences, when a file is loaded, it is first passed to a sentence segmenter. The sentence segmenter produces a list of sentences which are then used as input to the translation systems.

For ease of processing, all the systems are accessed in the same way. That is, from the user's point of view, by ticking the systems that are shown as available for that language pair. From the point of view of the LTI, the access to all the available systems is specified in an "ini" file, which gives all the information necessary so the users do not have to know how to access each separate system. For instance, access to Babelfish over the internet or access to the Indonesian-English Katakana system, which has to be installed on a Linux machine on a local network, look exactly the same to the users and the users do not have to know the difference. In the list of systems available for a language pair, we include the use of Translation Memories which may have been built for that language pair and which, from the point of view of the user, are just another translation tool.

Regarding the production of the translation output, the main issue has been the design of the output document. First, although translation is performed sentence by sentence, the user can choose to have the results presented either as continuous input and output texts or sentence by sentence. Second, the user can choose to accept all the translation results at once and then edit the output file. We found that this is what our users preferred to do when there is only one

translation system available. Alternatively, they can edit the translation results sentence by sentence within the LTI and then send the edited result to the output file. This is the mode which is probably the best when there are two or more translation systems available for a language pair.

In the LTI screen shown in Figure 2 above, there are several translation outputs, with one of the outputs highlighted for editing. The user can then choose one of the translation results for each input sentence and editing in place seems to be more convenient. When the user edits the translation result through the LTI, this is recorded in the output file. Figure 3 shows the default layout for the output document which is automatically generated when the translation results have been accepted. In this example, we have three translations for the Japanese input and the output document records information about which translation systems have been used and whether the output has been post-edited. Our users have already asked that the segments that have been post-edited be indicated in a different colour or highlighted, and this will be done for the next version.

Users can choose to have all the translation results included in the output file, or just the result they consider the best. The flexibility we wanted to offer the users reflects the range of possible situations in which they would need to produce a translation and the range of language skills they might have.

<p>Language Translation Interface session output: 5/05/2005 11:23:04 AM User: biggsj</p> <p>Source: Japanese 自爆攻撃は、警官募集にも使われているクルド民主党の事務所で起きた。</p> <p>Target: English; Translation engine: WTS; Post editor: BiggsJ</p> <p>The suicidal attack occurred in the Kurd Democratic Party office which is also currently used as a policeman enlistment post.</p> <p>Target: English; Translation engine: WTS; Post editor: None</p> <p>Suicidal explosion attack occurred in the office of Kurd Democratic Party which is used also as policeman collection.</p>
--

(continued on next page)

(Continued from previous page)

³ This process would be worth studying and we intend to include an evaluation of its merits when we gather user requirements in the next phase of the project.

Target: English; Translation engine: Linguattech; Post editor: None
Suicidal explosion attack occurred in the office of the ??? Democratic Party which is used also as policeman collection. .
Target: English; Translation engine: AmikaiOCN; Post editor: None
The suicide bomb attack broke out in the office of the Kurd Democratic Party currently used also for policeman collection.
Target: English; Translation engine: WorldLingo; Post editor: None
Suicide bombing attack occurred with the office of the Kurd Democratic party which is used even in officer collection.
Target: English; Translation engine: Mail2World; Post editor: None
A crashing itself attack occurred in an office of クルド Democratic Party used by police officer enlistment also.

Figure 3. Example output document

We presented the TCT version of the LTI at a multi-nation military exercise in June 2004. I will explain in more detail what was involved, but the main point is that these exercises serve as a trial for new technologies and the LTI was one of two systems presented by DSTO for Australia. The actual exercise takes place over a period of three weeks, but the preparation of this trial took several months and that in itself gave us a good exposure. The result from the exercise, that is the feedback we collected during and after it, was then the starting point for the next development of the tool. We had designed the TCT to be as widely useful as possible and, during that exercise, we showed that it could be used in a range of situations and for a range of purposes: coalition exercises with the translation of email from a South Korean ship (South Korea being one of the exercise coalition partners), information gathering for situation awareness for regional exercises with the translation of news sites from Arabic and Indonesian, and humanitarian operations with the production of a draft pamphlet in Tetun (one of the national languages of East Timor).⁴ Those language pairs were chosen both for experimental purposes, taking into account the availability of MT tools and the environment, and to fit in with the general exercise scenario.

⁴ While Portuguese is the official language of East Timor, Tetun serves as the lingua franca and Bahasa Indonesia is another language used in the area.

During the three weeks of the exercise, we collected feedback from both users and visitors to the exercise. Then, a new version of the tool was specifically developed for a particular environment, with automated web access being the main priority. This is the Web Translation Tool, shown in Figure 4.

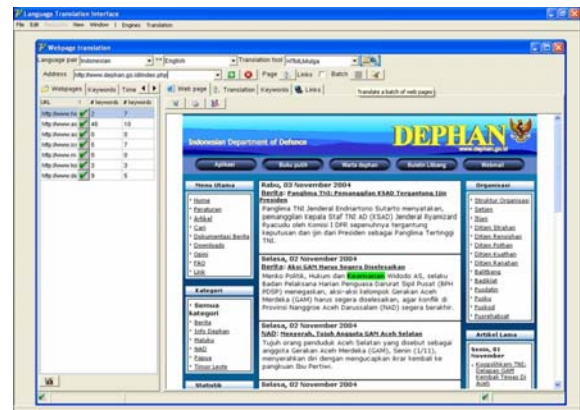


Figure 4. The LTI: Web Translation Tool

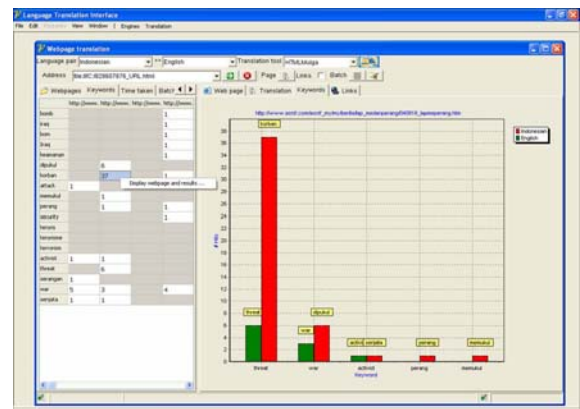


Figure 5. Keyword statistics

This second version of the LTI answered specific requests from users for new functionalities. With the Translation Comparison Tool, we had concentrated on the access to translation systems and on making it simple for users to deal with different types of input: typing input directly in the input window, loading a file or automatic access to e-mail. With the Web Translation Tool, the emphasis is on automating access to web pages and producing batch translation of those web pages. The users wanted to be able to have a list of web pages to be accessed and translated regularly. Also in answer to requests for new functionalities, we added other utilities so that users can create lists of keywords they want to monitor and they can get statistics on those

keywords when they are found in the source documents or in their translation (see Figure 5).

5. Access to users: Exercises and trials

This is always a problem for software designers: how do you get access to real users when you don't have a real system for them to try out? I have experienced that problem in a number of other projects, especially in research projects when you don't even necessarily know who the potential users might be, but also in industry when you already have a user base. There, the main issue is often that users are too busy to be interviewed or to be asked to participate in trials. For speech recognition, for instance, you may have to organise data collection projects, and you may try to entice users to give some of their time in return for a prize in a prize draw. In our case, not only are our users too busy in their daily job to be asked to participate in surveys or experiments, but they also change all the time. This is because of the posting cycle in the military. People may be assigned to positions for 12, 18 or 24 months and they may not be there when you come back to talk to them.

So, in our case, we first had to imagine who our users might be, then try to understand what they would want, and then take advantage of opportunities to get some of them to try out the system. We also have to organise how we can take advantage of those opportunities and make sure the few users we reach can give us feedback, so we can see whether we are on the right track. These opportunities to reach our users are demonstrations and trials during “open days” and military exercises where new technology is presented to various levels of the military.

Our first opportunity was a multi-nation coalition exercise (JWID 2004)⁵. In this exercise, our users were military personnel untrained in the use of language tools, who as part of their role-playing in the exercise had to produce translations for different types of input texts. This exercise runs on a scenario which is broken down into a number of “events” which recur at specific times throughout the five days of the

demonstration. Each event demonstrates a particular capability for the trials.

The LTI was an Australian trial and the LTI events only concerned Australian role players within the Australian exercise Headquarters. The four events for demonstrating translation capabilities were chosen to exemplify a range of situations where translation would be useful or even necessary:

- coalition exercises, with the translation of email from Korean;
- information gathering for situation awareness, with the translation of web news articles from Arabic and Indonesian;
- humanitarian operations, with the production of a draft pamphlet in Tetun, giving information on voting procedures.

This gave us four events with four language pairs. For some events, only one system was available for that language pair, e.g. only TM for English-Tetun, while for others we had two translation outputs, e.g. both TM and *Kataku* for Indonesian-English.

The whole exercise runs for three weeks. The first week is for training of the role players and rehearsal of all the events. In the second week, the role-players run through the complete scenario over five days. In the third week, they run through the complete scenario again, but with visitors attending and being given demonstrations throughout the events. Our users, the role-players who were running through the translation events, were monolingual English speakers who had never thought about translation. We had ample time to get to know them during the first week of training and rehearsal and to appreciate the job they were doing and what their background was. Although they were “role-playing”, they were representative of our intended users in real operations and their comments and feedback was extremely valuable.

They were very interested in the trial, found the LTI very easy to use, and said they could see that if they were asked to perform those duties, the LTI would be useful to them. The LTI also attracted a lot of interest from other role-players, who were not meant to have to use it but who asked to try it for themselves during down-time. Those other people also gave us useful feedback and suggestions. At the end of

⁵ <http://jitic.fhu.disa.mil/washops/jtca/jwid.html>.

the exercise, there was a formal assessment report, compiling comments collected from an on-line questionnaire. The assessment for the LTI was that it was considered to be of “significant value” and that the trial yielded “useful results”, with recommendation for further development and integration.

6. MT Tools available via the LTI

Turning to the MT tools we have made available with the LTI, the first point is that, as I mentioned at the beginning of this talk, we did not have the resources to buy many MT systems for demonstration, so we focussed on providing uniform access to as many free systems as possible. The drawback, of course, is that the translation quality is not as high as with commercial systems, but we managed to keep the emphasis on the flexibility and useability of the tools. We insisted on the fact that this was a demonstration prototype for an interface, not a testbed for people to evaluate the quality of the translation. So we provided access to a fairly large number of systems over the internet and concentrated on the issues of ensuring the smooth input and display of all writing systems, with all character encodings made available. Again, the emphasis was on making this invisible to the user, so they do not have to know how to switch between Arabic and Latin characters or between the different character encoding systems for Japanese. This effort has paid off because, predictably, it is always the first question we are asked: “Can you deal with other writing systems?”. So, our standard demo is to show Japanese and Arabic, as well as Indonesian – or French when we cannot access our Indonesian MT system. This leads me to my second point about the tools we have made available through the LTI and that is the issue of network constraints.

Our first prototype could only access free MT systems over the internet⁶ but it soon became obvious that this was never going to be the way it would be used in reality. In fact, it could not even be used that way when we got to the point of participating in trials and military exercises. One reason is that Defence uses secure net-

⁶ For example, the always popular Babelfish: <http://babelfish.altavista.com/>.

works and does not allow unrestricted access to the internet. Another is that it would not meet security requirements to send potentially sensitive data over the internet to be translated on a public site and then sent back to us.

For our first trial with real users, we had to run the whole exercise on a secure local network. Access to the internet was out of the question, so we had to find systems that we could either integrate on a local machine or access over that secure local network. In the end, we were able to have local access to the IBM *WebSphere Translation Server*;⁷ we were able to buy two language pairs from a commercial-off-the shelf system, AppTek's *Transphere*;⁸ and we were also able to use a research license for access to another commercial system, ToggleText's *Kataku*.⁹ In addition, with the purchase of *Wordfast*, we were able to build and demonstrate the use of Translation Memories.¹⁰ So, currently, we are able to demonstrate the LTI with the following translation systems:

- IBM WebSphere Translation Server (*WTS*), under a Defence-wide license. *WTS* provides translation for a number of languages, including Korean and Japanese.
- AppTek *TransSphere*, for Korean and Arabic. We purchased the English/Arabic and English/Korean language pairs and were granted a free temporary license for the API for the few months leading to JWID and for use during JWID.
- *Wordfast*, a Translation Memories (TM) application operating within Microsoft Word. We bought a license and we built TMs for Indonesian/English and for Tetun/English.
- ToggleText *Kataku* for Indonesian-English. We have an NDA with ToggleText, an Australian company, for research at DSTO and they have granted us permission to use *Kataku* during demonstrations and exercises.

The first three systems are all available on a single workstation, while *Kataku* is accessed over

⁷ IBM WebSphere: http://www-3.ibm.com/software/pervasive/ws_translation_server

⁸ AppTek: <http://www.apptek.com/>

⁹ ToggleText: <http://www.toggletext.com>

¹⁰ Wordfast: <http://www.wordfast.net>

a Local Area Network (LAN) via scripting commands within a telnet connection. For installation at a customer's site, of course, commercial licenses have to be purchased. In our role of providing advice on the choice and purchase of technology for the ADO, we are still looking at other systems which might be better suited to our customers' requirements for specific language pairs.

Besides the question of which MT tools we can make available through the LTI, another important issue for us is access to users, so we can assess their actual needs and requirements.

During the first two weeks of the exercise, we had shown the LTI to most of the participants in the Australian exercise Headquarters and, apart from some network connection issues, there had been no problem for any of the LTI demonstrations. That in itself and the acceptance by the military personnel were very positive results. Then, during the last week, when the role-players themselves had to do the demonstrations for the visitors (developers are not allowed to intervene), we received more feedback and very positive responses.

We then developed the Web Translation Tool I described earlier, to meet the specific user requirements which we received as a result of that first trial. A few months later, we brought the new Web Translation Tool in the particular Headquarters where people had said they wanted to try it for real. This allowed us to have new users try it in their own environment. They used it to translate websites that were of interest to them and they were positive about the results they were getting. What was interesting and very encouraging was that, although the system could not have been trained or tuned to the documents they wanted to translate, they found that the quality of the translation was enough for them to get the information they needed.

We had several opportunities to demonstrate the LTI again, first at a multi-national coalition exercise and then in a Headquarters exercise. These exercises did not involve users trying out the system, but we received very positive responses from the higher-level people who were attending. We made further improvements to the LTI, mostly to ensure the system was more secure and reliable, and we were then able to run a new trial in the Headquarters which had ex-

pressed strong interest in it. This time, the LTI was used over a period of several days by different analysts than those who had tried it earlier.

When I talked to Harry Somers about this project, I wanted to ask his advice on how best to utilise the opportunities we get to have users try the LTI for themselves. His first comment was that one must provide MT users with background reading on MT pitfalls and shortcomings and that one must give them training before letting them loose. I agreed this would be ideal but unfortunately this is not always feasible, for our users do not have much time to read background material before using new tools, they expect the tools to be useable right away. To help with that problem, we have produced very short user guides, in which we do warn users about what can go wrong, but it would be unrealistic to expect that the users will devote much time to those tools, at least until there is wider acceptance of the technology and the directive comes from the top that those tools must be used.

7. Conclusions from user trials and experiments

The main goals of the Language Translation Interface (LTI) project were the identification of requirements for automated translation within the ADF and the development of tools to meet these requirements. The development of a new translation engine requires enormous efforts and resources and is beyond the scope of a research project at DSTO. In any case it is not possible to predict which languages might become of interest and the results of such efforts would most likely not meet actual needs. It is interesting to note that these two issues are exactly parallel to the problems faced by the training of linguists and translators: it takes between one and three years to train a linguist in a new language, and languages of interest change according to world events and demands on the ADF. We had identified the development of a single interface to existing translation tools as filling the need for rapid and easy automated translation tools when human translators are not available and the LTI was first developed as a concept demonstrator providing users with a single interface for accessing a range of language translation systems and tools.

Following participation in military exercises and trials, two versions of the LTI, the Translation Comparison Tool and the Web Translation Tool are now fully integrated into one seamless system. However, the most important results from this interaction with potential users have been the exposure of the technology to those prospective users and the feedback we received from them. This exposure has raised awareness of the need for access to information in languages other than English, even in an English-speaking country such as Australia. To ADO members who were already aware of this need, but who had previously been reliant on human-only translations, this was an opportunity to see what can already be achieved with computer-assisted language translation and it raised an awareness of the need to develop tools to process documents in other languages. Finally, we were able to establish fruitful contacts with a user community for the LTI and we are now building upon them.

8. Future Work

Following these positive contacts, the main issue is to manage customers' and users' expectations. We have argued that relying on human translators to meet all the translation requirements of the ADO is not a viable option in the long run and that translation tools would help meet those needs. What we are now proposing is the development of a prototype with new capabilities, the Language Translation Tool Suite (LTTS), which will build on and extend the LTI. So, what we need to show is that the LTTS will in fact help meet those needs and not put an increased burden on the ADO translators or on defence personnel who would use the LTTS.

The prototype LTI tool is being tested in Headquarters, and this will give us more feedback from military personnel. We have also recently started the process of gathering user requirements for the new LTTS at the ADF School of Languages, with the goal of ensuring that these requirements coincide with those already established in Headquarters. The next phase of the project is to validate all these user requirements and to produce a report for transition to an operational system. We can already say that, from the point of view of the ADO, the LTTS is in line with the requirements for increased auto-

mation in Headquarters and for the ability to work in multi-lingual environments. It would also meet the stated ADF School of Languages goals of improving language training and efficiency, by improving the range of language skill training for students, with limited operational costs and limited additional training.

We are arguing that the LTTS would be a superior option to acquiring individual MT systems when the need for tools for a particular language pair arises, because the installation of the LTI (or the LTTS) is a one-off operation, which gives seamless access to all subsequent MT systems that might be added for new translation requirements. Furthermore, training of operators and of language students would be substantially lower than if separate systems were purchased on an ad hoc basis, because the same interface will be used for all systems, so the initial training for using the LTTS would cover all additional translation technologies accessed through the LTTS.

Another important advantage of the LTTS over having separate translation tools is the ability to combine the outputs of several MT systems for a language pair. We expect that this will increase the quality of translation output, especially when the systems also include Translation Memories. Although the MT systems we have so far made available through the LTI are primarily translation engines, we have been arguing that Translation Memory technology should be an important component of the LTTS: TMs give the ability to store translations and re-use them and this will both reduce translation time and contribute to building an organisation-wide database of translations. Use of this database will increase the translators' efficiency and improve the quality of translation. It is true that the creation of TMs requires resources but we hope that once the general tool is adopted, the translators will be able to build their own TMs and share them with other. More importantly, TMs will allow us to build new translation systems for languages with no existing MT engines, which is the case for many of the languages of interest in our region.

On the technical front, we plan to make improvements to the interface after feedback from users. We already know that this will include the addition of a number of utilities, in particu-

lar keyword synonym recognition, in addition to the keyword facility I mentioned we already developed for the WTC. We also plan to obtain significant improvements to translation quality by better text pre-processing. This will include spell-checking and limited named entity recognition, e.g. place names, dates, groups and individuals in the languages of interest. However, the first item on our list has to be the integration of military vocabulary, including acronyms, first for English then for the other languages of interest. This leads to a very challenging area of research, because what we want to develop is a general “Vocabulary Update” functionality which would provide users with the same simple interface to enter new vocabulary in the same way for all the systems. This can be seen as an extension of the idea of sharing linguistic data, either dictionaries or previous translations, between the tools for one language pair.

Users have already requested that the system be able to take input from OCR and we plan to include OCR and spell-checkers utilities. We demonstrated a couple of years ago the use of output from speech recognition, but I am not convinced that we are ready to offer this functionality yet. However, a simple utility to add is language detection and automatic selection of the appropriate translation tools and this will render the LTTS absolutely transparent to the users.

Further extensions include the integration of multilingual linguistic resources, e.g. dictionaries, Part-of-Speech taggers and extension of multi-lingual capabilities for named entity recognition. Ultimately, what we aim to do is no less than cross-language information retrieval and multi-lingual document classification, but we still have some way to go.

9. References

- CARR, Oliver and Dominique ESTIVAL. (2003). “Document Classification in Structured Military Messages”. *Australasian Language Technology Workshop*, Melbourne, Australia. pp-73-81.
- DoD/CIO (2004). Horizontal Fusion FY2004 After Action Report Version 1.0, Department of Defense, Assistant Secretary of Defence for Networks and Information Integration / DoD CIO: 83.
- ESTIVAL, Dominique and Jennifer BIGGS. (2003). “The Language Translation Interface and automated language translation tools for the ADO”. Eighth International Command and Control Research and Technology Symposium, Washington, DC, USA.
- ESTIVAL, Dominique, Michael BROUGHTON, Andrew ZSCHORN and Elizabeth PRONGER. (2003). “Spoken Dialogue for Virtual Advisers in a semi-immersive Command and Control environment”. *4th SIGdial Workshop on Discourse and Dialogue*. ACL 2003, Sapporo, Japan. pp-125-134.
- ESTIVAL, Dominique, Chris NOWAK and Andrew ZSCHORN. (2004). “Towards Ontology-based Natural Language Processing”. *RDF/RDFS and OWL in Language Technology: 4th Workshop on NLP and XML (NLPXML-2004)*. ACL 2004, Barcelona, Spain. pp.59-66.
- GALE (2005). <http://www2.eps.gov/spg/ODA/DARPA/CMOBAA05%2D28/SynopsisP.html>
- Hyperwave (2005). Hyperwave and the Horizontal Fusion Portfolio Initiative. Washington DC, Hyperwave Government Solutions: 7.
- JOHNSON, T. (2003). ARL technology to protect Soldiers, improve communication. RDECOM. December 2003: 12.
- OSTERHOLZ, J. L. (2003). Net-Centric Operations & Warfare. 6th Annual System Engineering Conference, San Diego, CA (DTIC – external site), National Defence Industrial Association (NDIA).