AI Magazine Volume 28 Number 3 (2007) (© AAAI)

Machine Translation for Manufacturing:

A Case Study at Ford Motor Company

Nestor Rychtyckyj

■ Machine translation (MT) was one of the first applications of artificial intelligence technology that was deployed to solve real-world problems. Since the early 1960s, researchers have been building and utilizing computer systems that can translate from one language to another without requiring extensive human intervention. In the late 1990s, Ford Vehicle Operations began working with Systran Software Inc. to adapt and customize its machine-translation technology in order to translate Ford's vehicle assembly build instructions from English to German, Spanish, Dutch, and Portuguese. The use of machine translation was made necessary by the vast amount of dynamic information that needed to be translated in a timely fashion. The assembly build instructions at Ford contain text written in a controlled language as well as unstructured remarks and comments. The MT system has already translated more than 7 million instructions into these languages and is an integral part of the overall manufacturing process-planning system used to support Ford's assembly plants in Europe, Mexico and South America. In this paper, we focus on how AI techniques, such as knowledge representation and natural language processing can improve the accuracy of machine translation in a dynamic environment such as auto manufacturing.

Ford Motor Company has been manufacturing and selling automobiles outside the United States since the early 1900s. As a global company, Ford currently has assembly plants located throughout the world, including many locations where English is not spoken by our assembly employees. These requirements have motivated us to explore new technologies, such as machine translation (MT), in order to translate and disseminate critical information in languages other than English. Since 1998, Ford Vehicle Operations has been utilizing a machine-translation system in order to translate our process assembly build instructions from English to German, Spanish, Portuguese, and Dutch. This system was developed in conjunction with Systran Software Inc. and is an integral part of our worldwide process-planning system for manufacturing assembly. The input to our system is a set of process build instructions that are written using a controlled language known as Standard Language. The process sheets are read by an artificial intelligence (AI) system that parses the instructions and creates detailed work tasks for each step of the assembly process (Rychtyckyj 1999). These work tasks are then released to the assembly plants where specific workers are allocated to perform each task. In order to support the assembly of vehicles at plants where the workers do not speak English, we utilize MT technology to translate these instructions into the native languages of these workers. Standard Language is primarily a restricted subset of English and contains a limited vocabulary of about 5,000 words that also include acronyms, abbreviations, proper nouns, and other Ford-specific terminology. In addition, Standard Language allows the process sheet writers to embed comments within Standard Language sentences and to attach explanatory remarks to the instructions. These comments and remarks are ignored by the AI system during processing, but have to be translated by the MT system. Standard Language also utilizes some structures that are grammatically incorrect, which creates problems for the MT translation process. Based on our experience, we have concentrated on two different approaches to improve the quality of machine translation: (1) develop a

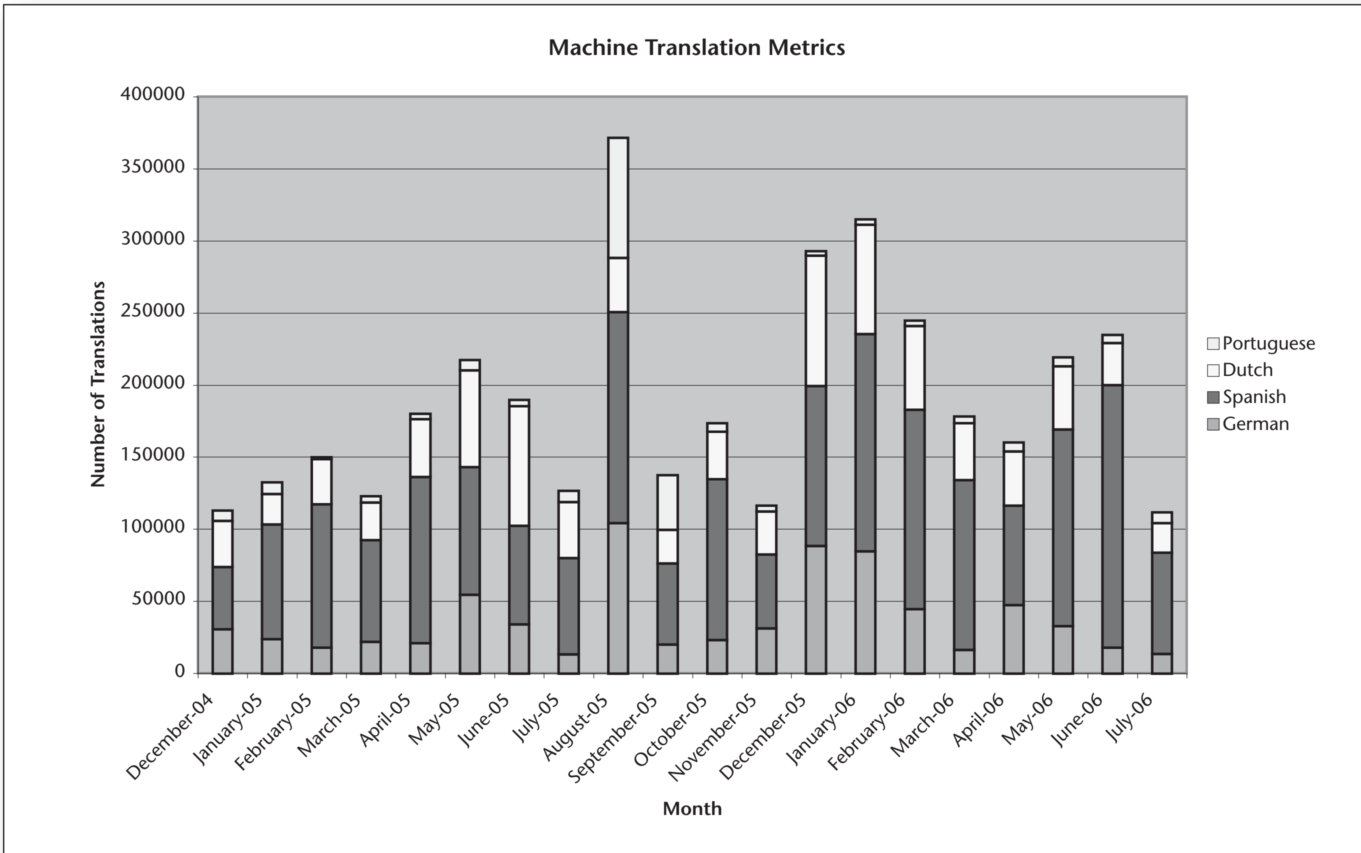

Figure 1. Monthly Translation Counts.

process and methodology to build, test, and maintain translation glossaries that contain the specific terminology that needs to be translated; and (2) develop a process to analyze and convert the source text into a format that is much more understandable to the MT system and produces more accurate results.

Problem Description

Ford Motor Company operates vehicle assembly plants all over the world, including locations in Germany, Spain, Belgium, Mexico, Brazil, and Venezuela. The assembly-line workers at these plants generally do not speak English, but the assembly build instructions are always written in English. Therefore, these instructions need to be translated into the home language of the workers who will actually be following these instructions. The standard process-planning document, the process sheet, is the primary means for conveying the assembly information from the initial process-planning activity to the assembly plant. A process sheet contains the detailed instructions needed to build a portion of a vehicle. A single vehicle may require several thousand process sheets to describe its assembly. An engineer writes a process sheet describing a portion of the assembly work utilizing a restricted subset of English known as Standard Language. Standard Language allows an engineer to write clear and concise assembly instructions that are machine readable. The process sheets also contain embedded comments and associated remarks that need to be translated. In addition, changes to the process build instructions are frequent and this necessitates the retranslation of those instructions. In a typical month, we may need to translate more than 150,000 records from English into our target languages of Spanish, Portuguese, Dutch, and German. Figure 1 displays our monthly translation metrics from 2004–2006. Since the initial deployment of this system, we have translated more than 7 million records. The sheer volume, quick turnaround, and cost required precluded the use of human translators on this project. The use of a controlled language, such as Standard Language, also gave us impetus to find an automated solution. The specific terminology required to describe the automotive assembly

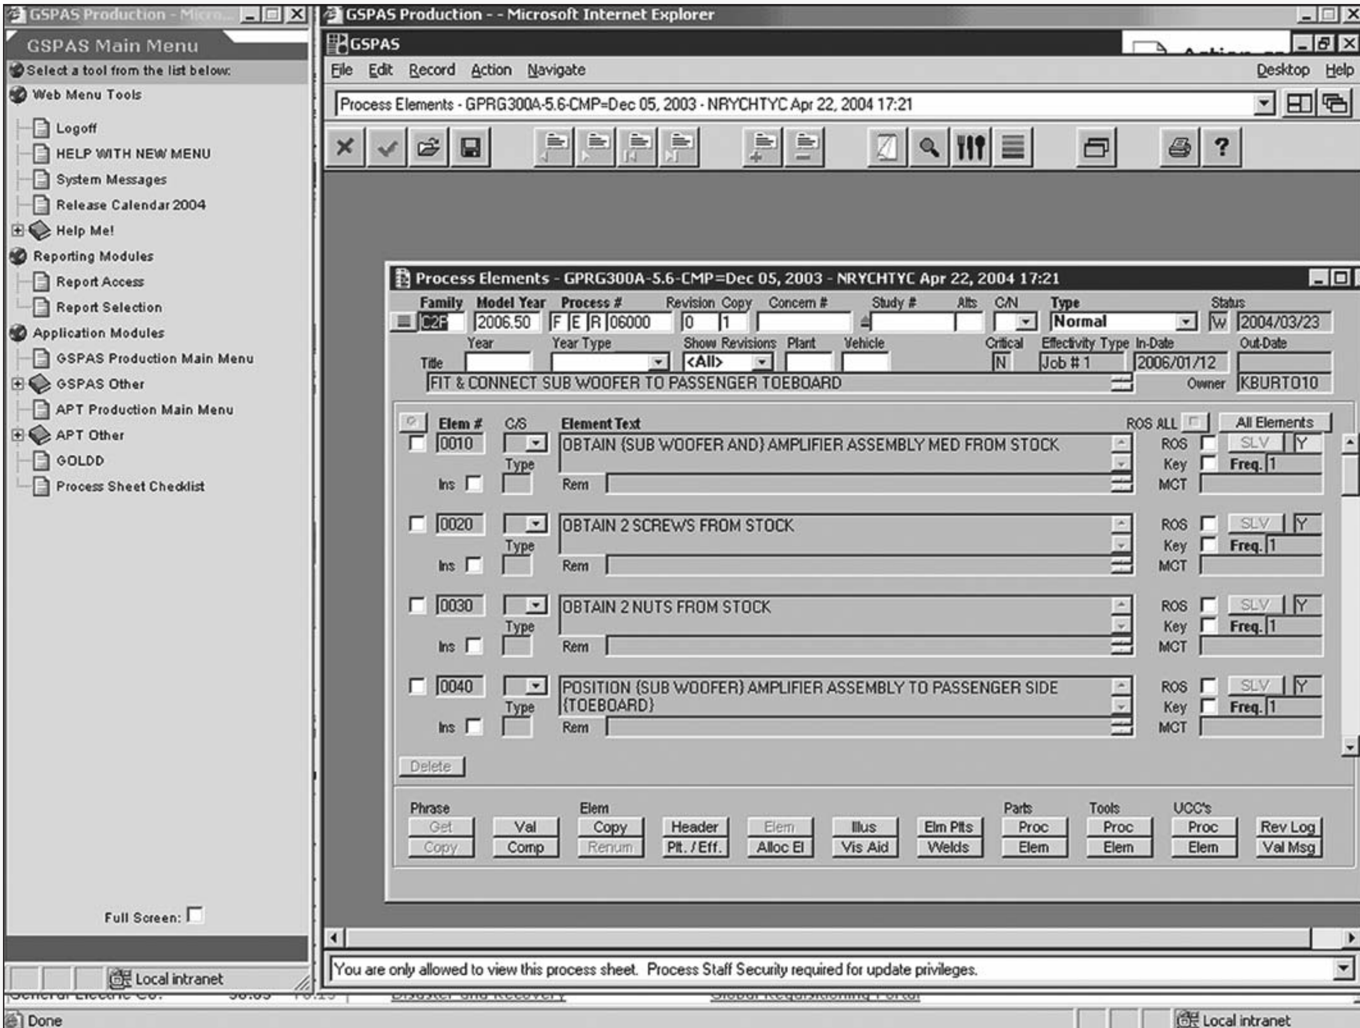

Figure 2. Machine Translation in GSPAS.

and engineering methodology at Ford required us to develop technical glossaries that could accurately translate text containing these terms. However, we also learned that machine-translation accuracy can be greatly improved by analyzing and modifying the source text to improve the quality of the translation output. In the next section, I will describe in more detail how we combined natural language processing and knowledge representation and reasoning to build and deploy a machine-translation system.

Application Description

The machine-translation system utilized at Ford is integrated into the Global Study Process Allocation System (GSPAS). The goal of GSPAS is to incorporate a standardized methodology and a set of common business practices for the design and assembly of vehicles to be used by all assembly plants throughout the world. GSPAS allows for the integration of parts, tools, process descriptions, and all other information required to build a motor vehicle into one system. It also provides the engineering and manufacturing communities with a common platform and toolset for manufacturing process planning. GSPAS utilizes Standard Language as a requirement for writing process build instructions, and we have deployed an MT solution for the translation of these process build instructions.

The translation process at Ford for our manufacturing build instructions is fully automated and does not require human manual intervention. All of the process build instructions are stored within an Oracle database; they are written in English and validated by the AI system. AI validation consists of parsing the Standard

Language sentence, analyzing it, and creating the appropriate work description based on the information in the knowledge base. The system then creates an output set of work instructions and assigns their associated MODAPTS (Modular Arrangement of Predetermined Time Standards) codes. MODAPTS codes are used to calculate the time required to perform these actions. MODAPTS is an industrial measurement system used around the world (Carey 2001). A more complete description of the GSPAS AI system can be found in Rychtyckyj (1999). A sample of a GSPAS process sheet is shown in figure 2.

After a process sheet is validated and the AI system generates the appropriate MODAPTS codes and times, a process engineer will release the process sheet to the appropriate assembly plants. A vehicle that is built at multiple plants needs to have these process sheets sent to each of these assembly plants. The information about each local plant is stored in the database, and those plants that require translation are identified by the system. The system then selects the process sheets that require translation and starts the daily translation process for each language. Currently we translate the process build instructions for 32 different vehicles into the appropriate language. English-Spanish is the most commonly used language pair, as it is utilized at our assembly plants in Spain, Mexico, and South America. However, we have recently developed and deployed a separate technical glossary for the English-Spanish translation system for our plants in Mexico due to the differences in the translated terminology between "Mexican Spanish" and regular Spanish.

The machine-translation system was integrated into GSPAS through the development of an interface to the Oracle database. Our translation programs extract the data from an Oracle database, modify the source text to improve translation accuracy, utilize the SYSTRAN system to perform some postprocessing, and then send the data back to the Oracle database.

Our user community is located globally. The translated text is displayed on the user's PC or workstation through the use of a graphical user interface to the GSPAS system. The Ford multitargeted customized dictionary that contains Ford technical terminology was developed in conjunction with Systran and Ford, based on input from engineers and linguists familiar with Ford's terminology.

One of the most difficult issues in deploying any translation is the need to obtain consistent and accurate evaluation with regard to the quality of translations (both human and machine). We are using the J2450 metric developed by the Society of Automotive Engineers (SAE) as a guide for our translation evaluators (SAE 2002). The J2450 metric was developed by an SAE committee consisting of representatives from the automobile industry and the translation community as a standard measurement that can be applied to grade the translation quality of automotive service information. This metric provides guidelines for evaluators to follow, describes a set of error categories, specifies the weight of the errors found, and calculates a score for a given document. The metric does not attempt to grade style, but focuses primarily on the understandability of the translated text. The utilization of the SAE J2450 metric has given us a consistent and tangible method to evaluate translation quality and identify which areas require the most improvement.

We have also spent substantial effort in analyzing the source text in order to identify which terms are used most often in Standard Language so that we can concentrate our resources to correctly translate those most common terms (Manning and Schulze 2000). This process was accomplished by using the parser from our AI system to store parsed sentences into the database. Periodically, we run an analysis of our parsed sentences and create a table where our terminology is listed in order of usage frequency. This table is then compared to the technical glossary to ensure that the most commonly used terms are being translated correctly. The frequency analysis also allows us to calculate the number of terms that need to be translated correctly to meet a given translation accuracy threshold. For example, we can calculate that 80 percent translation accuracy (based on terminology) requires that the most-frequently used 200 terms need to be inserted into the translation glossary. An example of this type of analysis is shown in figure 3. We perform this analysis on individual terms and on distinct noun phrases that are identified in the system.

A machine-translation system, such as the one we utilize from Systran, translates text sentence by sentence. In Standard Language, each sentence is self-contained, and users cannot use pronouns to refer back to objects that may have been described in a previous sentence. A single term by itself cannot be translated accurately because it may correspond to different parts of speech depending on the context. Therefore, it is necessary to build sample test cases for each word or phrase that we will need to test for translation accuracy. This test case utilizes that term in its correct usage within the sentence. A file containing these translated

Translation Frequency Usage

Noun Phrases Sorted by Usage	Count	Pct of Total	Running Pct
SPOT	10441	3.786071203	3.786071203
STOCK	9678	3.509395374	7.295466578
PART	7850	2.846533756	10.14200033
FIXTURE	6966	2.52598142	12.66798175
SCREW	4719	1.711183795	14.37916555
SPOT-WELD GUN	4701	1.704656712	16.08382226
HOLE	4663	1.690877313	17.77469957
BRACKET	3844	1.393895001	19.16859457
NUT	3504	1.270605641	20.43920021
SPOT-WELD-GUN	3293	1.194093714	21.63329393
BOLT	3112	1.128460261	22.76175419
PALM BUTTON	2782	1.008797058	23.77055125
VEHICLE	2557	0.927208511	24.69775976
CLAMP	2552	0.925395432	25.62315519
CLIP	2461	0.892397398	26.51555259
HAND-TOOL	2270	0.823137787	27.33869038
ASSEMBLY	2171	0.787238826	28.1259292
BODY	1610	0.583811382	28.70974058

Figure 3. Translation Frequency Usage.

sentences (known as a test corpus) is used as a baseline for regression testing of the translation dictionaries. After the dictionary is updated, the test corpus of sentences is retranslated and compared against the baseline. Any discrepancies are examined and a correction is made to either the baseline (if the new translation is correct) or to the dictionary (if the new translation is incorrect). We also designate a person for each language who has the final responsibility for the given language pair; any discrepancies or differences as to the correct translation will be decided by this language coordinator.

Our system allows the users to override any machine-generated translation with a manual translation. This manual translation will remain current until the underlying English text is modified. When the English text is changed, the system automatically deletes all the existing translations. We keep a copy of the manual translations and spend considerable time in analyzing these manual translations to determine if they could be used to improve the machine-translation quality. Unfortunately, there are several problems with trying to use unedited manual translations. Many of the users would be inconsistent in their usage terminology for the same English word. A more critical problem would result when users would add or delete content from an English sentence as part of the translation process. This would be done on an ad hoc basis and would make the manual translations extremely difficult to use. We found that the manual translation process would need to be strictly regulated to produce usable results and this is not feasible in our production environment. Therefore, in practice, our translation system automatically translates all of the assembly build instructions required for a given assembly plant without any manual human intervention.

Uses of AI Technology

It has been known that improving machine-translation quality can often be done most effectively by focusing on the source text

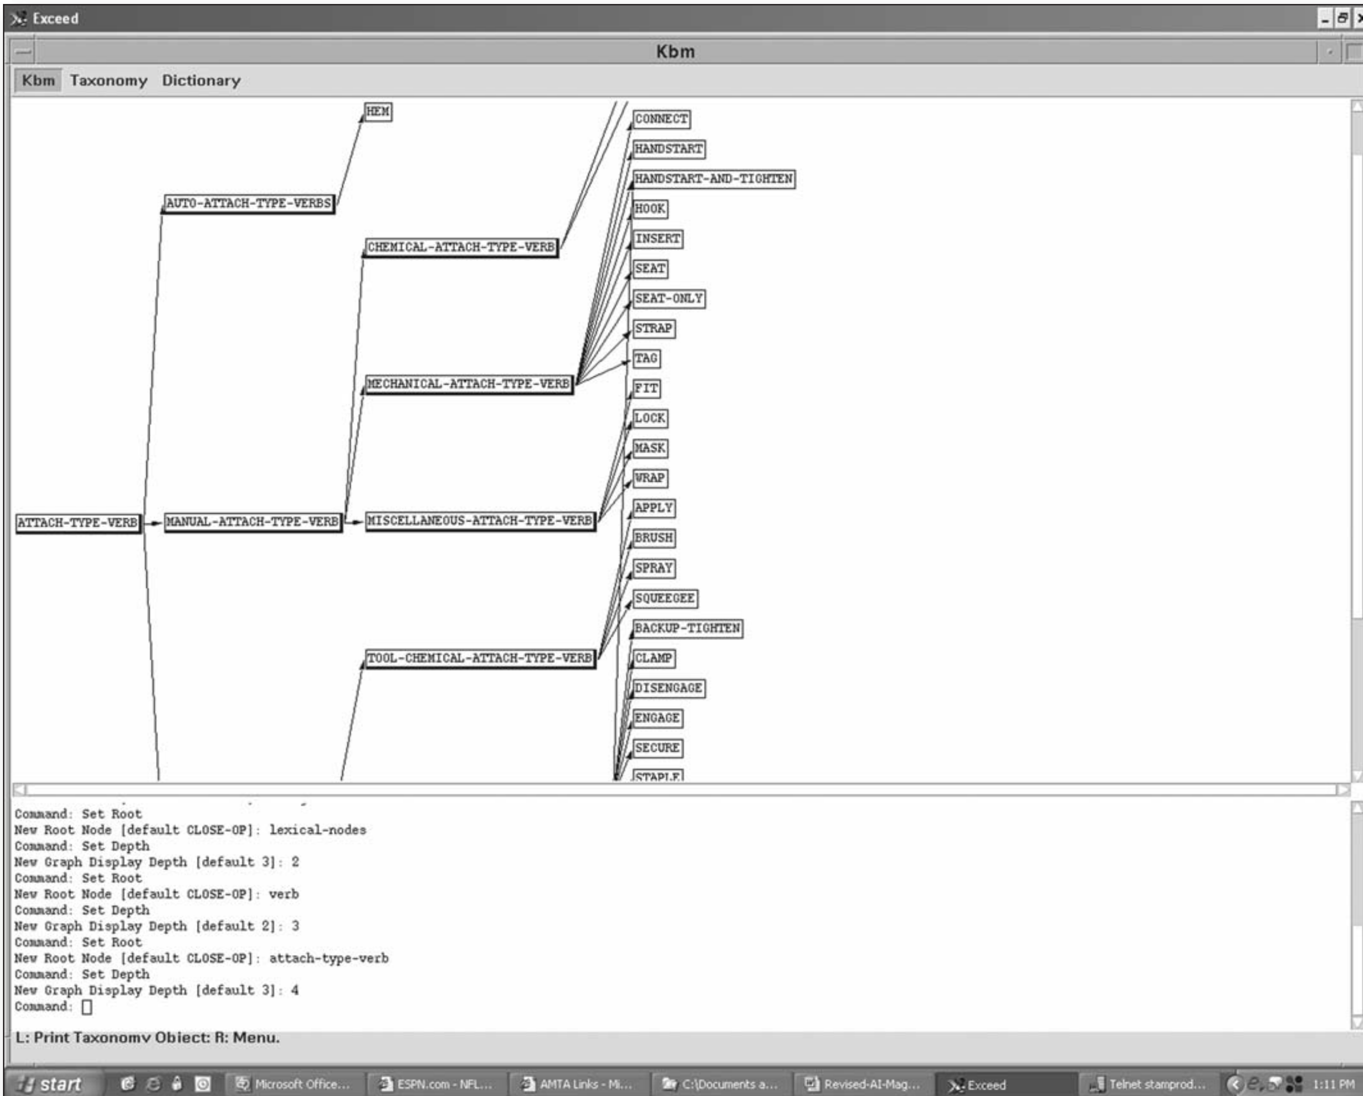

Figure 4. A Portion of the GSPAS Automotive Ontology.

(Hutchins and Somers 1992). In most cases, the preediting of text is performed by a human editor, who verifies and modifies the text before it is sent to the translation system. In our case, the source text is a combination of a controlled language and free-form text. Each of these must be treated in a somewhat different fashion in order to get the most accurate translation results. This can be done by applying natural language processing along with knowledge representation and reasoning to convert the source text to an equivalent form that can be processed more accurately by the machine-translation engine.

The first step in applying MT technology is to analyze the existing text in order to understand exactly what terminology needs to be translated and how the source text is structured. The terminology analysis is performed by running all of the source text through a program that retrieves each individual token and looks up the token in the automotive ontology that we have developed for Standard Language as part of the GSPAS project. The automotive ontology utilized is a semantic network that contains more than 10,000 concepts related to automotive assembly at Ford Motor Company. All of the associated knowledge about Standard Language, tools, parts, and everything else associated with the automobile assembly process, is contained in the DLMS knowledge base or ontology (Rychtyckyj 2006). This knowledge base structure is derived from the KL-ONE family of semantic network structures

(Brachman and Schmolze 1985) and is an integral component of the GSPAS system. Figure 4 shows a portion of the GSPAS automotive ontology.

A Standard Language sentence can be parsed and understood by the GSPAS AI system; therefore, each token in the sentence has the relevant information (part of speech, usage, size, and so on) available in the ontology. In addition, the ontology provides us with a method to identify phrases that need to be translated as an entity rather than as a collection of single words. The analysis of free-form text is substantially more difficult. We have discovered that a vast majority of the terms (87 percent) can be identified using the GSPAS ontology; however most of the free-form comments and remarks cannot be parsed successfully.

Along with the need for special technical glossaries for translation, we utilize a variety of approaches that take advantage of the natural language-processing and the knowledge representation technologies to convert the source text into a form that is much more likely to lead to a better translation. This is based on the fact that MT systems expect source text to conform to some specific rules including the following five: (1) Simple, unambiguous sentence structures (shorter sentences usually translate much better than long, complicated sentences). Many authoring systems put a strict limit on the length of a sentence. (2) It is preferable to put articles in front of nouns and noun phrases as it helps the MT system identify the proper part of speech and create a more understandable translation. (3) The regular grammar rules of capitalization and punctuation need to be observed. In general, a sentence that is written according to the structured rules of English grammar will be translated more accurately than one that is not. (4) Acronyms, abbreviations, and proper nouns need to be identified unambiguously; this is where the ontology is most useful. For example, the system needs to know that "ABS" is an abbreviation for ANTI-LOCK BRAKE SYSTEM and not for ABSOLUTE. (5) The MT system will utilize any additional information about the source text that can be gleaned from the system; in our case we utilize XML tags to identify certain properties of the source text, such as part of speech and its usage in this context.

Therefore, we have deployed a pretranslation component into our system that reads in the source text as it is written by the process engineers, converts the source text into a more MT-friendly form, and then submits the reformulated text to the translation engine. This reformulation process begins by using the ontology and AI parser to process the input text. At this point, the ontology is referenced to determine if any acronyms, abbreviations, or terms need to be replaced by a synonym, which will always translate correctly. Other changes to the Standard Language text are also performed to enhance the structure of the source text. For instance, articles are added into the text in front of noun phrases except in circumstances where the noun phrase would never expect an article. The sentence "SECURE BRACKET TO BUMPER" is converted to "SECURE THE BRACKET TO THE BUMPER," but "DRIVE VEHICLE 60 FEET" is not converted to "DRIVE THE VEHICLE THE 60 FEET."

In Standard Language, we allow the engineers to use ungrammatical structure in some cases, and this needs to be corrected before the sentence can be translated. A process writer may put a size adjective after a part to override the existing size of the part as in the following example: "OBTAIN BRACKET ASSEMBLY VERY-LARGE." In this case, the system uses the term "VERY-LARGE" to override the existing size of the "BRACKET ASSEMBLY." The sentence is then converted to "OBTAIN THE VERY-LARGE BRACKET ASSEMBLY" before it is sent to the translation engine. Similar types of text reformulation are performed when handling plurals, numeric constants, and special cases where the Standard Language text cannot be translated accurately.

As I mentioned previously, we also need to translate embedded remarks and comments that are not in Standard Language and contain free-form text. In this case, we rely on embedded XML tags to assist the MT program in the translation process (Senellart, Boitet, and Romary 2003). First, we identify the free-form remarks that are embedded in the Standard Language text. We then utilize the ontology to analyze the terminology that is contained within the remarks and replace any abbreviations or acronyms with the proper unambiguous Standard Language term. The system then looks at the length of the embedded remark and places the appropriate tag around the remark; we have found that very short remarks (one or two) words are generally modifiers, while longer remarks are self-contained phrases that should be translated as such. In effect, the XML tagging uses the benefits of the natural language processing and ontology from the AI system to assist the MT program in creating a more accurate translation. We are currently working on expanding the scope of the tagging process to incorporate additional information, such as part of speech tagging, to further enhance the translation accuracy.

Application Use and Payoff

The machine-translation system has been deployed at Ford for more than seven years. The impact of this system can be summarized as follows. First, we have translated more than 7 million records from English to Spanish, German, Portuguese, and Dutch. Second, the user community has access to translations of assembly instructions in their home language within 24 hours of the process sheet being written and completed. Third, we have created Ford-specific translation glossaries for each of the language pairs for which we need to translate our assembly instructions. The translation glossaries contain a significant number of part description phrases that need to be translated as a single entity and, consequently, contain up to 6,000 entries. Fourth, we have worked with Systran to deploy a web-based process that makes it possible for us to maintain and update the Ford-specific technical glossaries on a timely basis. Fifth, we have built a process that allows the assembly plant personnel to manually override the translations when necessary. These human translations will remain in the system as along as the underlying English source text is not modified. Finally, we have developed a process to retranslate the process sheets when an updated technical glossary is deployed; this ensures that the users will have the benefit of the latest version of the translations available.

The easiest way to calculate the benefits of using the machine translation is to compare the costs of human translation versus the cost of developing an MT solution that can generate translations with the same accuracy. A machine-translation system, even in a semicontrolled setting, will not generate translations that are as accurate as those completed by a trained human translator. We can develop translations that are highly accurate (our English-German is more than 90 percent correct), but this is directly dependent on the involvement of the bilingual technical people with the creation of technical glossaries. The English-German glossary is much more complete than English-Portuguese, so our translations are more accurate into German than into Portuguese. However, the huge amount of data that we need to translate precludes the use of human translators. Our goal in this project was to develop translations that are understandable to the operators at the assembly plants. These translations may not be as natural as those provided by human translators, but they will provide the correct information to the users. Since Standard Language is always evolving, the technical glossaries must always be modified to keep them current. The main payoff for this project is that we are able to provide "understandable" translations to our users around the world in a timely manner without utilizing any direct human intervention.

Application Development and Deployment

The artificial intelligence development for our applications here at Ford Manufacturing Engineering Systems is based on the Hewlett-Packard UNIX (HP-UX) platform utilizing the Lispworks and Knowledgeworks tools from Lispworks Inc.[1] We have found that this tool provides a flexible and powerful development environment while providing access to our Oracle database through an SQL interface. We have worked closely with Systran in seamlessly integrating their translation programs into our translation process. The largest amount of effort that we spent was to develop the customized translation glossaries for each of the four language pairs that we need to translate. This development work required the efforts of internal Ford bilingual subject matter experts, the use of retired and external people who understand Ford and automotive technical terminology, the use of linguistic experts from Systran, and our own expertise in bringing all of these knowledge sources together.

The actual translation process is shown in figure 5; the entire process is fully automated. Each evening, a batch run scans the database for those process sheets that need to be translated based on the assembly plant in which the vehicle is built. At this point, the element text has been reformulated into a more translation-friendly format by the AI system, and our translation programs selects the records from the database that need to be translated. The appropriate XML tags are added, and the record is then translated for each target language. The translated record is then written into the database. The translation process utilizes three different glossaries: a customized Ford-specific glossary, a generic automotive glossary, and a general-purpose glossary. The "translation parameters" file contains specific information about the translation processing for each language. For example, English-German is translated in imperative form while English-Spanish is translated in infinite form.

The initial application deployment and development took about six months to accomplish; this included writing the software that would interface with the translation engines and update the database as needed. These initial translations were of very poor quality and were not acceptable to the user community. At

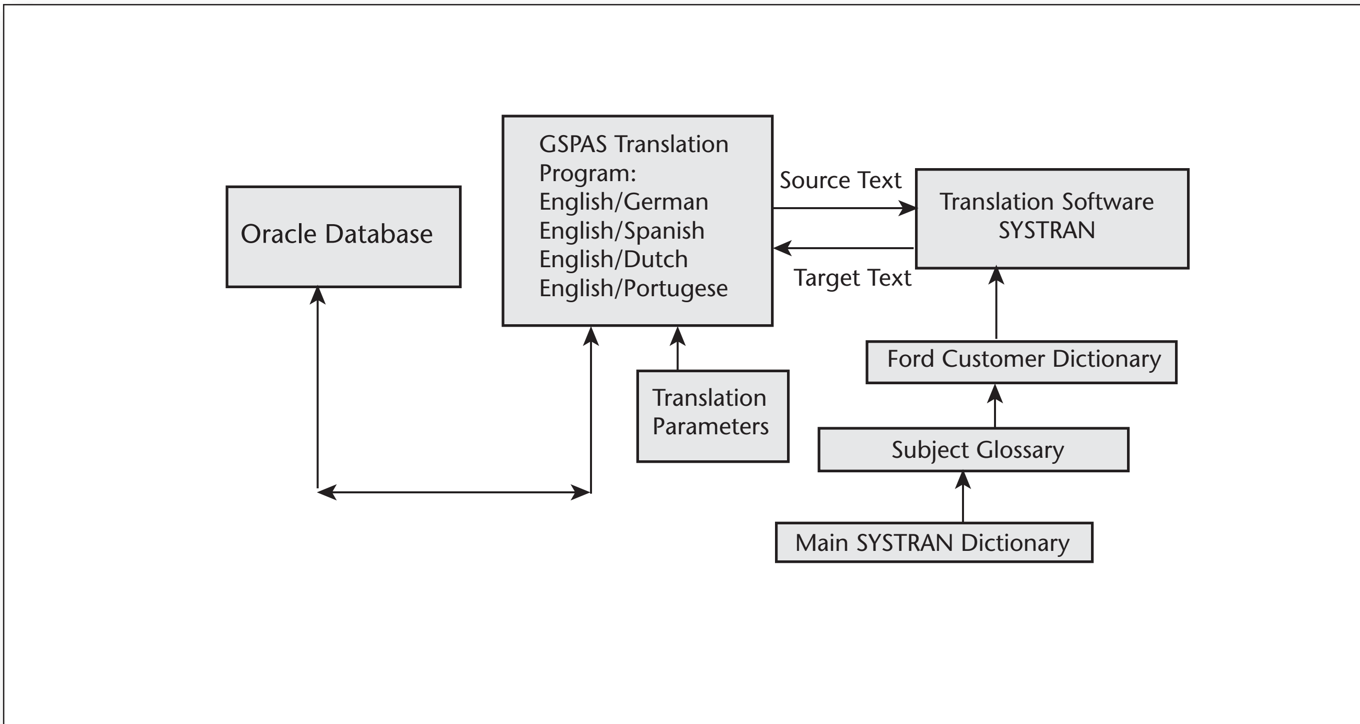

Figure 5. Actual Translation Process.

that point we started working to improve the translation quality by building up the technical glossaries and building a process to improve the source text before it is translated. This was accomplished by creating utilities to analyze the source text and identify the terminology that was causing translation problems. Changes were also implemented to the translation process to allow our users the ability to override the automated translations manually when necessary. Another important issue that had to be addressed was to ensure that the translated text could be properly displayed to the users because of the special characters that are required and the extra space that is often needed. The accuracy of the translations increased as we built up the technical glossaries and improved the text reformulation process. Over the next few years, Systran spent considerable time and effort to streamline and improve their translation programs, and as a result, we deployed more than 10 versions of the technical glossaries. Our translation accuracy improved noticeably with the English-German and English-Spanish, as it was much easier to find people who could work on these glossaries. The amount of maintenance required is also directly proportional to the size of the technical glossaries. This system has been in production since 1998, but we are still spending a considerable amount of effort maintaining and enhancing the system both through advances in technology and with the creation of more complete technical glossaries.

We have also studied the possibility of expanding the machine-translation approach beyond just manufacturing assembly instructions. There are other types of automotive information, such as technical service bulletins and warranty claims that need to be translated in a timely and accurate manner. This type of source information is much less controlled and contains more ambiguity than the assembly build instructions. In addition, the terminology glossaries will need to be refined and updated to improve the quality of these translations. However, we believe that further advancements in MT technology, including "part of speech tagging," statistical analysis, and learning techniques will increase the use of machine translation for other less-structured problem domains and applications.

Maintenance

As previously discussed, we have spent considerable time and effort to create a set of customized technical glossaries that are used during the translation process. These glossaries were developed in conjunction with Systran

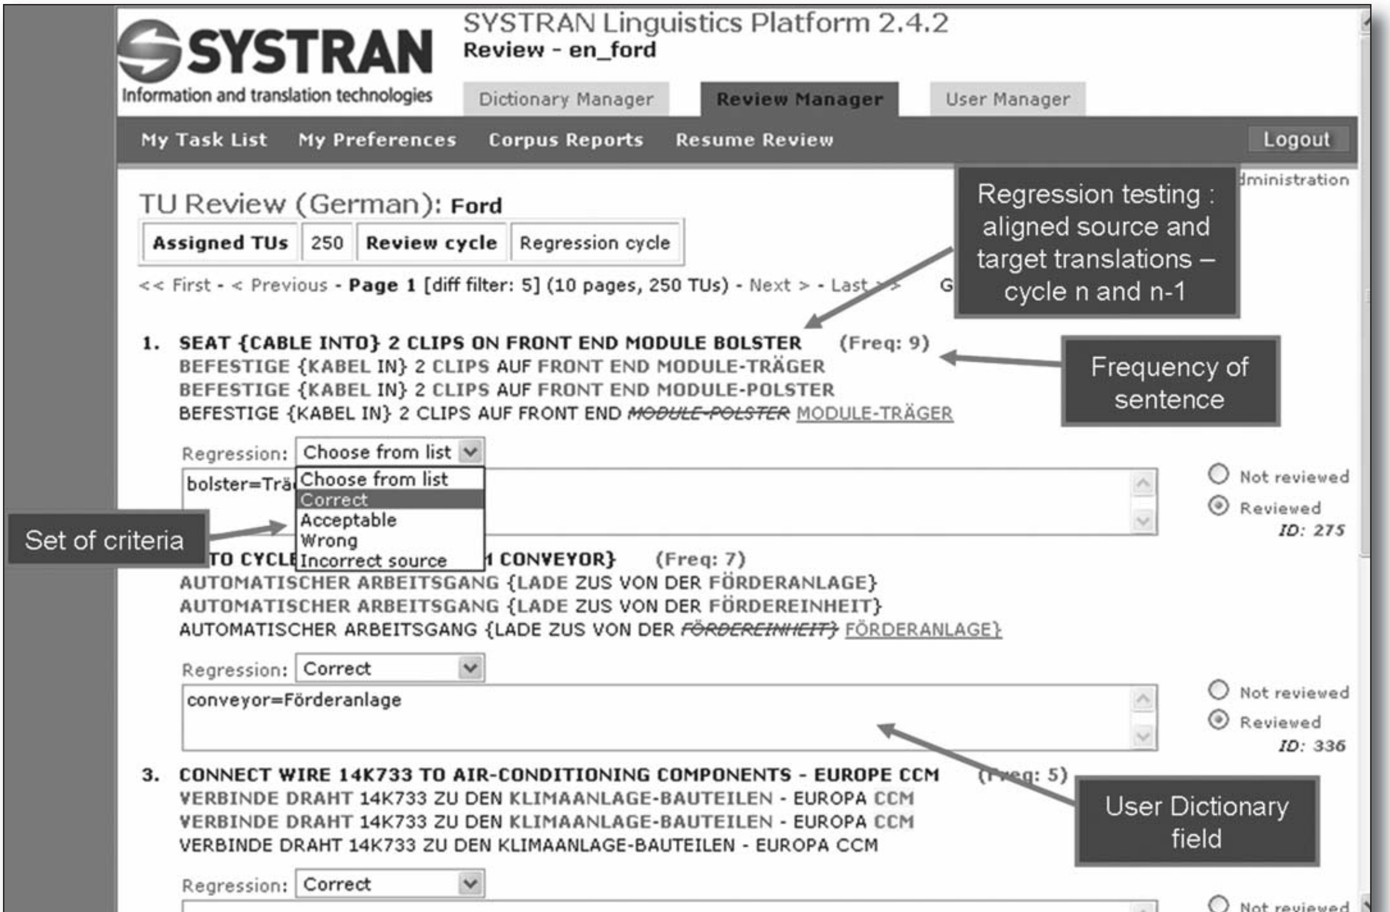

Figure 6. SYSTRAN Review Manager.

and with subject matter experts from Ford Motor Company. However, since Standard Language and Ford terminology are always evolving, it soon became obvious that we needed to develop a process to modify and add terminology to our technical glossaries in a timely manner.

The initial release of our MT system was designed so that all updates to the technical glossaries required Systran to create and compile a new set of dictionaries that would include the new changes. Systran would need to test these dictionaries through their internal quality control program and then deliver the updates to Ford. We would also need to test the updates against our internal benchmarks and deploy them into production if the result of the testing was acceptable. The entire process would be delayed if any problems were discovered during testing. This approach was too cumbersome and time-consuming and was not viable for the long term.

Systran developed a web-based system known as the Systran Review Manager (SRM) (Costa and Panissod 2003) that addressed all of these shortcomings. The SRM was deployed on a Ford internal server that allowed us to control and monitor the access to the application. Figure 6 shows a screen print of the SRM that displays how a user would deal with a term that was not found in the translation glossary. Figure 7 demonstrates how the user can review a sample corpus for translation accuracy. Our user community was trained to use this tool, and it gave them the a number of benefits, such as automation of the testing process (a user could make a change to the technical glossaries and immediately run a translation that would test to see how the change would impact the translation quality). The SRM allows users to create and modify different versions of user-defined dictionaries without impacting changes that are being made by a different user. Test corpora can be loaded and analyzed direct-

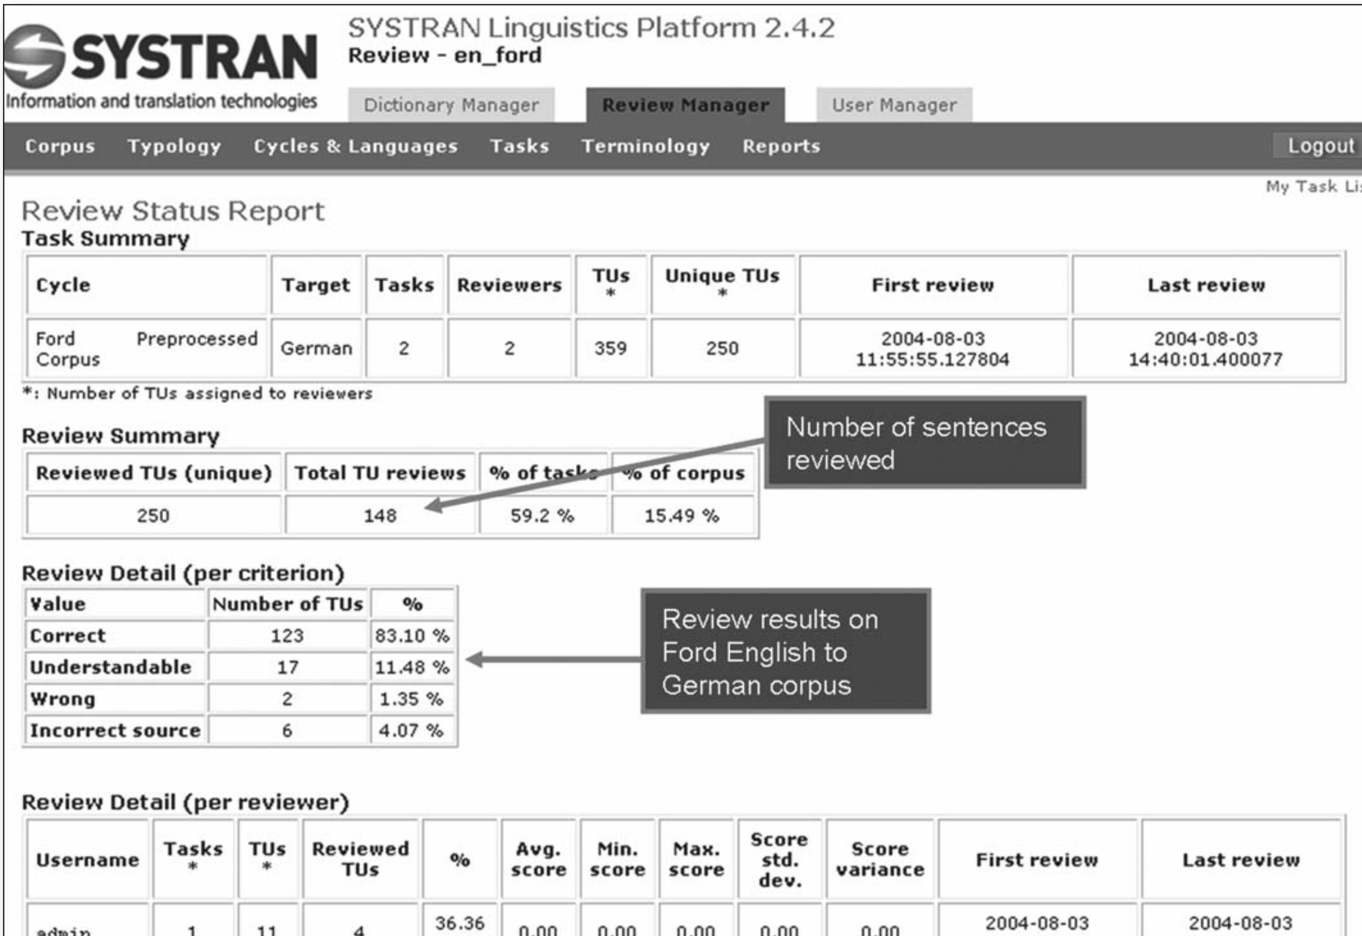

Figure 7. Review of a Sample Corpus for Translation Accuracy

ly within the SRM. The web-based architecture of the SRM allows our users to access the system without any additional software or hardware requirements. The SRM provides very quick turnaround time for the process of modifying and deploying an updated translation glossary.

Another important facet in dictionary maintenance involves the analysis and customization of the source text. We have previously described some of the techniques we have been using to clean up the source text to improve translation quality. In this section we will discuss additional capabilities we have added into the system to improve the translation of the free-form text. A Standard Language element may contain embedded free-form text that is ignored by the AI system; however this text must be translated and sent to the assembly plants. This free-form text usually consists of additional information that may be useful to the operator on the assembly line. These embedded remarks may contain non–Standard Language terminology or they may be separate phrases or sentences that describe specific circumstances for this process work. Our analysis has shown that the embedded remarks needed to be treated separately from the Standard Language text in order to create accurate translations. In many cases, a single embedded remark that looks innocuous inside a Standard Language element would lead to an incorrect translation. Therefore, we decided that the best solution would be to separate the embedded remarks from the Standard Language text and translate them separately. The following example shows how this process would take place.

PLACE TWO MOULDINGS INSIDE HEATER {TAPE SIDE UP}

The text inside the curly brackets {TAPE SIDE UP} is not really part of the sentence; it actually describes the position of the "mouldings."

Therefore, a translation system that processes this sentence as one entity would not generate an accurate translation. We need to be able to tell the system that the clause inside the curly brackets should be treated independently from the rest of the sentence. This problem is solved by embedding tags into the source text before it gets translated. These tags identify comments and provide the translation program with information about how these comments should be translated. Short comments are processed differently from long comments within Standard Language regarding translation parameters (dictionaries and segmentation). The above comment with embedded tags will look like the following:

PLACE TWO MOULDINGS INSIDE HEATER <comment> TAPE SIDE UP </comment>

Another facet of system maintenance addresses the underlying software architecture that supports our translation system. Translation in GSPAS involves a set of programs that communicate with a database as well as with the translation engines and technical glossaries. Most changes to the translation engine processing also require changes to the translation preprocessing programs. In addition, modifications to the database model or upgrades to the operating system require extensive testing and validation of the translation results. The testing needs to identify both the translation issues in both the Standard Language and the non–Standard Language components of the source text. The results of the translation tests focus on two types of potential problems: terminology and grammar. Terminology errors are almost always fixed just by adding the correct translation for the problem term into the appropriate translation glossary. The grammar errors are more complex; they may require changes to the translation engine itself.

Conclusions and Future Work

In this article I discussed some of the issues related to the maintenance of a machine-translation application at Ford Motor Company. This application has been in place since 1998, and we have translated more than 7 million records describing build instructions for vehicle assembly at our plants in Europe, Mexico, and South America. The source text for our translation consists of a controlled language, known as Standard Language, but we also need to translate free-form text comments that are embedded within the assembly instructions. The most difficult issue in the development of this system was the construction of technical glossaries that describe the manufacturing and engineering terminology in use at Ford. Our application uses a customized version of the Systran translation system coupled with a set of Ford-specific dictionaries that are used during the translation process. The automotive industry is very dynamic, and we need to be able to keep our technical glossaries current and to develop a process for updating our system in a timely fashion.

The solution to our maintenance issues was the development and deployment of the Systran Review Manager. This web-based tool allows our users the capability to test and update the technical glossaries as needed. This has reduced our turnaround time for deploying changes to the dictionaries from two months to less than 48 hours. The Systran Review Manager runs on an internal Ford server and is available for use by our internal customers.

System maintenance is an ongoing issue. We still require additional capabilities to improve our translation accuracy and to expand our system to other types of source data, including part and tool descriptions. We have already introduced XML tagging into our free-form comment translation and are working with SYSTRAN to enhance that capability and improve translation accuracy. Our current AI system in GSPAS already parses Standard Language into its components, and we plan to pass the information obtained during parsing to the translation system to improve the sentence understanding that should lead to higher accuracy. One of the unique advantages that we have on this project is the automotive ontology that we have developed for our manufacturing processes at Ford. This ontology allows us to retrieve knowledge and infer context information about the source text that needs to be translated. Our challenge is to leverage this background knowledge and integrate the context information into the translation process.

This project has given us a unique perspective into the culture and business processes of our fellow Ford employees around the world. We allow the users to override the translations manually when they are unacceptable and also provide a feedback mechanism to measure the accuracy of these translations. We have been surprised to see that, in many cases, our users prefer that we utilize an English acronym or term rather than the correct translated word. We have also discovered that even in a technical domain such as automobile assembly, there still exists some variation between Spanish in Spain, Mexico, Argentina, and Venezuela. The proliferation of free web-based translation engines has proven to be both a blessing and a curse for our project. In some cases, users would not even consider using MT after trying out these web services; in other cases users were perfectly satisfied with the quality of these translations and did not see the need for any customization work. Perhaps one of our biggest challenges is to properly educate and manage the expectations of the user community when exposing them to this technology.

Our experience with machine-translation technology at Ford has been positive; we have shown that customization of a commercial translation system can lead to very positive results. It is also essential to put a process in place that allows for the timely testing and upgrades to the technical glossaries. We are confident that further enhancements to the technology, such as tagging of terminology, will lead to better results in the future and improve the use and acceptance of machine translation in the corporate world.

Acknowledgements

I thank the *AI Magazine* reviewers for their insightful comments; in addition, I would like to thank Mike Rosen and Rosemarie Janisse from Ford and Christiane Pannisod, John Paul Barazza, and Jean Senellart from Systran Software Inc. for their work on this project. I would also like to thank Erica Klampfl and Reba Sitzer for their assistance in the preparation of this article.

Note

1. www.lispworks.com.

References

Brachman, R., and Schmolze, J. 1985. An Overview of the KL-ONE Knowledge Representation System. *Cognitive Science* 9(2): 171–216.

Carey, P.; Farrell, J.; Hui, M.; and Sullivan, B. . *Heyde's Modapts: A Language of Work*. Brighton, Victoria, Australia: Heyde Dynamics Pty. Ltd.

Costa, J.-C., and Panissod, C. 2003. SYSTRAN Review Manager. In *Proceedings of the Ninth Machine Translation Summit,* 451–454. Stroudsburg, PA: Association for Machine Translation in the Americas.

Gazdar, G., and Mellish, C. 1989. *Natural Language Processing in LISP.* Reading, MA: Addison-Wesley.

Hutchins, W., and Somers, H. 1992. *Introduction to Machine Translation.* London: Academic Press.

Iwanska, L., and Shapiro, S., eds. 2000. *Natural Language Processing and Knowledge Representation: Language for Knowledge and Knowledge for Language*. Menlo Park, CA: AAAI Press.

Manning, D., and Schutze, H. 2000. *Foundations of Statistical Natural Language Processing*. Cambridge, MA: The MIT Press.

Rychtyckyj, N. 1999. DLMS: Ten Years of AI for Vehicle Assembly Process Planning. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence and the Eleventh Innovative Applications of Artificial Intelligence Conference,* 821–828. Menlo Park, CA: AAAI Press.

Rychtyckyj, N. 2002. An Assessment of Machine Translation for Vehicle Assembly Process Planning at Ford Motor Company. In *Machine Translation: From Research to Real Users, Proceedings of the Fifth Conference of the Association for Machine Translation in the Americas,* 207–215,. Berlin: Springer-Verlag.

Rychtyckyj, N. 2004. Maintenance Issues for Machine Translation Systems. In *Machine Translation: From Real Users to Research: Proceedings of the Sixth Conference for the Association of Machine Translation in the Americas,* Lecture Notes in Computer Science Vol. 3265, 252–261. Berlin: Springer-Verlag,.

Rychtyckyj, N. 2006. Measuring Long-Term Ontology Quality: A Case Study from the Automotive Industry. In *Proceedings of the Nineteenth International FLAIRS Conference (FLAIRS-2006),* 147–152. Menlo Park, CA: AAAI Press.

Senellart, J., Boitet, C., Romary, L. 2003. SYSTRAN New Generation: The XML Translation Workflow. In *Proceedings of the Ninth Machine Translation Summit,* 338–345. Stroudsburg, PA: Association for Machine Translation in the Americas.

Society of Automotive Engineers. 2002. *J2450 Quality Metric for Language Translation*. Warrendale, PA: SAE International.

Nestor Rychtyckyj is a technical expert in artificial intelligence at Ford Motor Company in Dearborn, Michigan, in advanced and manufacturing engineering systems. He received his Ph.D. in computer science from Wayne State University in Detroit, Michigan. His research focuses on the application of knowledge-based systems for vehicle assembly process planning, ergonomics, and adaptive in-vehicle systems. Currently his responsibilities include the development of automotive ontologies, intelligent manufacturing systems, controlled languages, machine translation, and corporate terminology management. He is a member of AAAI, ACM, and the IEEE Computer Society. His email address is nrychtyc@ford.com.