

Translation Methodology in the Spoken Language Translator: An Evaluation

David Carter
Ralph Becket
Manny Rayner

Robert Eklund
Catriona MacDermid
Mats Wirén

Sabine Kirchmeier-Andersen
Christina Philp

SRI International
Suite 23, Millers Yard
Cambridge CB2 1RQ
United Kingdom

Telia Research AB
Spoken Language Processing
S-136 80 Haninge
Sweden

Handelshøjskolen i København
Institut for Datalingvistik
Dalgas Have 15
DK-2000 Frederiksberg
Denmark

dmc@cam.sri.com
rwab1@cam.sri.com
manny@cam.sri.com

Robert.H.Eklund@telia.se
Catriona.I.Macdermid@telia.se
Mats.G.Wiren@telia.se

sabine.id@cbs.dk
cp.id@cbs.dk

Abstract

In this paper we describe how the translation methodology adopted for the Spoken Language Translator (SLT) addresses the characteristics of the speech translation task in a context where it is essential to achieve easy customization to new languages and new domains. We then discuss the issues that arise in any attempt to evaluate a speech translator, and present the results of such an evaluation carried out on SLT for several language pairs.

1 The nature of the speech translation task

Speech translation is in many respects a particularly difficult version of the translation task. High quality output is essential: the speech produced must sound natural if it is to be easily comprehensible. The quality of the translation itself must also be high, in spite of the fact that, by the nature of the problem, no post-editing is possible. Things are equally difficult on the input side: pre-editing, too, is difficult or impossible, yet ill-formed input and recognition errors are both likely to be quite common. Thus robust analysis and translation are also required. Furthermore, any attempted solutions to these problems must be capable of operating at a speed close enough to real time that users are not faced with unacceptable delays.

Together, these factors mean that speech translation is currently only practical for limited domains, typically involving a vocabulary of a few thousand words. Because of this, it is desirable that a speech translator should be easily portable to new domains. Portability to new lan-

guages, involving the acquisition of both monolingual and cross-linguistic information, should also be as straightforward as possible. These ends can be achieved by using general-purpose components for both speech and language processing and training them on domain-specific speech and text corpora. The training should be automated whenever possible, and where human intervention is required, the process should be deskilled to the level where, ideally, it can be carried out by people who are familiar with the domain but are not experts in the systems themselves.

These points will be discussed in the context of the Spoken Language Translator (SLT) (Rayner, Alshawi *et al.*, 1993; Agnäs *et al.*, 1994; Rayner and Carter, 1997), a customizable speech translator built as a pipelined sequence of general-purpose components. These components are: a version of the Decipher (TM) speech recognizer (Murveit *et al.*, 1993) for the source language; a copy of the Core Language Engine (CLE) (Alshawi (ed), 1992) for the source language; another copy of the CLE for the target language; and a target language text-to-speech synthesizer.

The current SLT system carries out multilingual speech translation in near real time in the ATIS domain (Hemphill *et al.*, 1990) for several language pairs. Good demonstration versions exist for the four pairs English → Swedish, English → French, Swedish → English and Swedish → Danish. Preliminary versions exist for five more pairs: Swedish → French, French → English, English → Danish, French → Spanish and English → Spanish.

We describe the methodology used to build the SLT system itself, particularly in the areas of customization (Section 2), robustness (Section 3), and multilinguality (Section 4). For further details on the topics of customization and multilin-

quality, see (Rayner, Bretan *et al*, 1996; Rayner, Carter *et al*, 1997); and on robustness, see (Rayner and Carter, 1997). We then discuss the evaluation of speech translation systems. This is an area that deserves more attention than it has received to date; indeed, it is not obvious how best to perform such an evaluation so as to measure meaningfully the performance both of the overall system and of each of its components. In Sections 5 and 6 of this paper, we therefore consider the characteristics an evaluation should have, and describe one we have carried out, discussing the extent to which it meets the desired criteria.

2 Customization to languages and domains

In the Core Language Engine, the language processing component of the Spoken Language Translator system, we address the requirement of portability by maintaining a clear separation between (1) the system code; (2) linguistic rules, including lexicon entries, to generate possible analyses and translations non-deterministically; and (3) statistical information, to choose between these possibilities. The practical advantage of this architecture is that most of the work involved in porting the system to a new domain is concerned with the parts of the system that can be modified by non-experts: the central activities are addition of new lexicon entries, and supervised training to derive the statistical preference information. Porting to new languages is a more complex task, but still only involves modifications to a relatively small subset of the whole system. In more detail:

(1) The *system code* is completely general-purpose and does not need any changes for new domains or, other than in exceptional cases,¹ for new languages.

(2) The more complex of the *linguistic rules* for a given language are the grammar, the function word lexicon, and the macros defining common content word behaviours (count noun, transitive verb, etc). These are defined using explicit feature-value equations which must be written by a skilled grammarian. For a given language *pair*, the more complex transfer rules, which tend to be for function words and other commonly-occurring, idiosyncratic words, can also involve arbitrarily large, recursive structures. However, nearly all of these monolingual and bilingual rules are domain-independent.

On the other side of the coin, the main domain-dependent aspects of a linguistic description are

¹E.g. in our initial extension from English to languages with more complicated morphology, which necessitated the development of a morphological processor based on the two-level formalism (see (Carter, 1995)).

lexicon entries defining content words in terms of existing behaviours, and simple (atomic-to-atomic) transfer rules. These do need to be created manually for each new domain, but they are simple enough to be defined by non-experts with the help of relatively simple graphical tools. See Figures 1 and 2 for some examples of these two kinds of rule (the details of the formalism are unimportant here, we intend simply to illustrate the differences in complexity).

When moving to a new language, more expert intervention is typically required than for a new domain, because many of the complex rules do need some modifications. However, we have found that the amount of work involved in developing new grammars for Swedish, French, Spanish and most recently Danish has always been at least an order of magnitude less than the effort required for the original grammar (Gambäck and Rayner, 1992; Rayner, Carter and Bouillon, 1996; Rayner, Carter *et al*, 1997).

(3) The *statistical information* used in analysis is entirely derived from the results of supervised training on corpora carried out using the TreeBanker (Carter, 1997), a graphical tool that presents a non-expert user with a display of the salient differences between alternative analyses in order that the correct one may be identified. Once a user has become accustomed to the system, around two hundred sentences per hour may be processed in this way. This, together with the use of representative subcorpora (Rayner, Bouillon and Carter, 1995) to allow structurally equivalent sentences to be represented by a single example, means that a corpus of many thousands of sentences can be judged in just a few person weeks. The principal information extracted automatically from a judged corpus is:

- Constituent pruning rules, which allow the detection and removal, at intermediate stages of parsing, of syntactic constituents occurring in contexts where they are unlikely to contribute to the correct parse. Removing these constituents significantly constrains the search space and speeds up parsing (Rayner and Carter, 1997).
- An automatic tuning of the grammar to the domain using the technique of Explanation-Based Learning (van Harmelen and Bundy, 1988; Rayner, 1988; Samuelsson and Rayner, 1991; Rayner and Carter, 1996). This rewrites it to a form where only commonly-occurring rule *combinations* are represented, thus reducing the search space still further and giving an additional significant speedup.
- Preference information attached to certain characteristics of full analyses of sentences – the most important being semantic triples of head, relationship and modifier – which allow

Syntax rule for $S \rightarrow NP VP$:

```
syn(s_np_vp_Normal, core,
  [s:[@s_np_feats(MMM), @vp_feats(MM),
    sententialsubj=SS,sai=Aux, hascomp=n,conjoined=n],
  np:[@s_np_feats(MMM),vform=(fin\to), relational=_,temporal=_,agr=Ag,
    sentential=SS, wh=_, whmoved=_,pron=_,nform=Sfm],
  vp:[@vp_feats(MM),vform=(\en),agr=Ag,sai=Aux, modifiable=_,
    mainv=_,headfinal=_,subjform=Sfm]]).
```

Macro definition for syntax of transitive verb:

```
macro(v_subj_obj,
  [v:[vform=base,mhdf1=A,passive=A,gaps=B,conjoined=n,
    subcat=[np:[relational=_,passive=A,wh=_,gap=_,gaps=B,
      temporal=_,pron=_,case=nonsubj]]]]).
```

Transfer rule relating English adjective “early” and French PP “de bonne heure”:

```
trule([eng,fre],semi_lex(early-de_bonne_heure),
  [early_NotLate,tr(arg)]
  ==
  @form(pre('de bonne heure_Early'),_,
    P^[P,tr(arg),
      @term(ref(pro,de_bonne_heure,sing,_),
        V,W^[time,W]+_)]).
```

Figure 1: Complex, domain-independent linguistic rules

a selection to be made between competing full analyses. See (Alshawi and Carter, 1994) and (Carter, 1997) for details.

A similar mechanism has been developed to allow users to specify appropriate translations, giving rise to preferences on outcomes of the transfer process. Work on this continues.

3 Robustness

Robustness in the face of ill-formed input and recognition errors is tackled by means of a “multi-engine” strategy (Frederking and Nirenburg, 1994; Rayner and Carter, 1997), combining two different translation methods. The main translation method uses transfer at the level of QLF (Alshawi et al., 1991; Rayner and Bouillon, 1995); this is supplemented by a simpler, glossary-based translation method. Processing is carried out bottom-up. Roughly speaking, the QLF transfer method is used to translate as much as possible of the input utterance, any remaining gaps being filled by application of the glossary-based method.

In more detail, source-language parsing goes through successive stages of lexical (morphological) analysis, low-level phrasal parsing to identify constituents such as simple noun phrases, and finally full sentential parsing using a version of the original grammar tuned to the domain using explanation-based learning (see Section 2 above). Parsing is carried out in a bottom-up mode. After each parsing stage, a corresponding translation operation takes place on the resulting constituent lattice. Translation is performed by using the

glossary-based method at the early stages of processing, before parsing is initiated, and by using the QLF-transfer method during and after parsing. Each successful transfer attempt results in a target language string being added to a target-side lattice. Metrics are then applied to choose a path through this lattice. The criteria used to select the path involve preferences for sequences that have been encountered in a target-language corpus; for the use of more sophisticated transfer methods over less sophisticated; and for larger over smaller chunks.

The bottom-up approach contributes to robustness in the obvious way: if a single analysis cannot be found for the whole utterance, then translations can be produced for partial analyses that have already been found. It also contributes to system response in that the earlier, more local, shallower methods of analysis and transfer usually operate very quickly to produce an attempt at translation. The target-language user may interrupt processing before the more global methods have finished if the translation (assuming it can be viewed on a screen) is adequate, or the system itself may abandon a sentence, and present its current best translation, if a specified time has elapsed.

Figure 3 exemplifies the operation of the multi-engine strategy as well as of the preferences applied to analysis and transfer.² The N-best list

²The example chosen was the most interesting of the dozen or so in our most recent demonstration session, and the intermediate results have been repro-

Lexicon entry, using transitive verb macro, for “serve” as in “Does Continental serve Atlanta?”:

```
lr(serve,v_subj_obj,serve_FlyTo).
```

Transfer rule relating that sense of “serve” to one sense of French “desservir”:

```
trule([eng,fre],lex(simple),serve_FlyTo==desservir_ServeCity).
```

Figure 2: Simple, domain-dependent linguistic rules

delivered by the speech recognizer contains the sentence actually uttered, “Could you show me an early flight please?”, but only in fourth position.

- Before any linguistic processing is carried out, the word sequence at the top of the N-best list is the most preferred one, as only recognition preferences (shown by position in the list) are available. This sequence is translated word-for-word using the glossary method, giving result (a) in the figure.
- After lexical analysis, which effectively includes part-of-speech tagging, it is determined that the word “a” is unlikely to precede “are”, and so “a” is dropped from the translated sequence (b) – thus translating recognizer hypothesis 2, using the glossary-based method.
- Phrasal parsing identifies “an early flight” as a likely noun phrase, so that this is for the first time selected for translation, in (c). Note that the system has now settled on the correct English word sequence. QLF-based transfer is used for the first time, and the transfer rule in Figure 1 is used to translate “early” as “de bonne heure” which, because it is a PP, is placed after “vol” (flight) by the French grammar.
- Finally, as shown in (d), an analysis and a QLF-based translation are found for the whole sentence, allowing the inadequate word-for-word translation of “could you show me” as “*pourriez vous montrez moi” to be improved to a more grammatical “pourriez-vous m’indiquer”.

We thus see the results of translation becoming steadily more accurate and comprehensible as processing proceeds.

4 Multilinguality, interlinguas and the “N-squared problem”

While using an interlingual representation would seem to be the obvious way to avoid the “N-squared problem” (translating between N languages involves order N^2 transfer pairs), we are sceptical about interlinguas for the following reasons.

duced from the system log file without any changes other than reformatting.

Firstly, doing good translation is a mixture of two tasks: semantics (getting the meaning right) and collocation (getting the appearance of the translation right). Defining an interlingua, even if it is possible to do so for an increasing number N of languages, really only addresses the first task. Interlingual representations also tend to be less portable to new domains, since they if they are to be truly interlingual they normally need to be based on domain concepts, which have to be redefined for each new domain – a task that involves considerable human intervention, much of it at an expert level. In contrast, a transfer-based representation can be shallower (at the level of linguistic predicates) while still abstracting far enough away from surface form to make most of the transfer rules simple atomic substitutions.

Secondly, systems based on formal representations are brittle: a fully interlingual system first needs to translate its input into a formal representation, and then realise the representation as a target-language string. An interlingual system is thus inherently more brittle than a transfer system, which can produce an output without ever identifying a “deep” formal representation of the input. For these reasons, we prefer to stay with a fundamentally transfer-based methodology; none the less, we include some aspects of the interlingual approach, by regularizing the intermediate QLF representation to make it as language-independent as possible consonant with the requirement that it also be independent of domain. Regularizing the representation has the positive effect of making the transfer rules simpler (in the limiting case, a fully interlingual system, they become trivial).

We tackle the N-squared problem by means of *transfer composition* (Rayner, Carter and Bouillon, 1996; Rayner, Carter *et al*, 1997). If we already have transfer rules for mapping from language A to language B and from language B to language C , we can compose them to generate a set to translate directly from A to C . The first stage of this composition can be done automatically, and then the results can be manually adjusted by adding new rules and by introducing declarations to disallow the creation of implausible rules: these typically arise because the contexts in which $\alpha \in A$ can correctly be translated to $\beta \in B$ are disjoint from those in which β can be translated into $\gamma \in C$. As with the other cus-

N-best list (N=5) delivered by speech recognizer:

- 1 could you show me a are the flight please
- 2 could you show me are the flight please
- 3 could you show me in order a flight please
- 4 could you show me an early flight please
- 5 could you show meals are the flight please

(a) Selected input sequence and translation after surface phase:

could	you	show	me	a	are	the	flight	please
pourriez	vous	montrez	moi	un	sont	les	vol	s'il vous plaît

(b) Selected input sequence and translation after lexical phase:

could	you	show	me	are	the	flight	please
pourriez	vous	montrez	moi	sont	les	vol	s'il vous plaît

(c) Selected input sequence and translation after phrasal phase:

could	you	show	me	an early flight	please
pourriez	vous	montrez	moi	un vol de bonne heure	s'il vous plaît

(d) Selected input sequence and translation after full parsing phase:

could you show me an early flight please
pourriez-vous m'indiquer un vol de bonne heure s'il vous plaît

Figure 3: N-best list and translation results for "Could you show me an early flight please?"

tomization tasks described here, the amount of human intervention required to adjust a composed set of transfer rules is vastly less, and less specialized, than what would be required to write them from scratch.

In the current version of SLT, transfer rules were written directly for neighbouring languages in the sequence Spanish - French - English - Swedish - Danish (most of these neighbours being relatively closely related), with other pairs being derived by transfer composition. Further details can be found in (Rayner, Carter *et al*, 1997).

5 Evaluation of speech translation systems: methodological issues

There is still no real consensus on how to evaluate speech translation systems. The most common approach is some version of the following. The system is run on a set of previously unseen speech data; the results are stored in text form; someone judges them as acceptable or unacceptable translations; and finally the system's performance is quoted as the proportion that are acceptable. This is clearly much better than nothing, but still contains some serious methodological problems. In particular:

1. There is poor agreement on what constitutes an "acceptable translation". Some judges regard a translation as unacceptable if a single word-choice is suboptimal. At the other end of the scale, there are judges who will accept

any translation which conveys the approximate meaning of the sentence, irrespective of how many grammatical or stylistic mistakes it contains. Without specifying more closely what is meant by "acceptable", it is difficult to compare evaluations.

2. Speech translation is normally an interactive process, and it is natural that it should be less than completely automatic. At a minimum, it is clearly reasonable in many contexts to feed back to the source-language user the words the recognizer believed it heard, and permit them to abort translation if recognition was unacceptably bad. Evaluation should take account of this possibility.
3. Evaluating a speech-to-speech system as though it were a speech-to-text system introduces a certain measure of distortion. Speech and text are in some ways very different media: a poorly translated sentence in written form can normally be re-examined several times if necessary, but a spoken utterance may only be heard once. In this respect, speech output places heavier demands on translation quality. On the other hand, it can also be the case that constructions which would be regarded as unacceptably sloppy in written text pass unnoticed in speech.

We are in the process of redesigning our translation evaluation methodology to take account of all of the above points. Currently, most of our

empirical work still treats the system as though it produced text output; we describe this mode of evaluation in Section 5.1. A novel method which evaluates the system's actual spoken output is currently undergoing initial testing, and is described in Section 5.2. Section 6 presents results of experiments using both evaluation methods.

5.1 Evaluation of speech to text translation

In speech-to-text mode, evaluation of the system's performance on a given utterance proceeds as follows. The judge is first shown a text version of the correct source utterance (what the user actually said), followed by the selected recognition hypothesis (what the system thought the user said). The judge is then asked to decide whether the recognition hypothesis is acceptable. Judges are told to assume that they have the option of aborting translation if recognition is of insufficient quality; judging a recognition hypothesis as unacceptable corresponds to pushing the 'abort' button.

When the judge has determined the acceptability of the recognition hypothesis, the text version of the translation is presented. (Note that it is not presented earlier, as this might bias the decision about recognition acceptability.) The judge is now asked to classify the quality of the translation along a seven-point scale; the points on the scale have been chosen to reflect the distinctions judges most frequently have been observed to make in practice. When selecting the appropriate category, judges are instructed only to take into account the actual spoken source utterance and the translation produced, and ignore the recognition hypothesis. The possible judgement categories are the following; the headings are those used in Tables 1 and 2 below.

Fully acceptable. Fully acceptable translation.

Unnatural style. Fully acceptable, except that style is not completely natural. This is most commonly due to over-literal translation.

Minor syntactic errors. One or two minor syntactic or word-choice errors, otherwise acceptable. Typical examples are bad choices of determiners or prepositions.

Major syntactic errors. At least one major or several minor syntactic or word-choice errors, but the sense of the utterance is preserved. The most common example is an error in word-order produced when the system is forced to back up to the robust translation method.

Partial translation. At least half of the utterance has been acceptably translated, and the rest is nonsense. A typical example is when most of the utterance has been correctly recognized and translated, but there is a short

'false start' at the beginning which has resulted in a word or two of junk at the start of the translation.

Nonsense. The translation makes no sense. The most common reason is gross misrecognition, but translation problems can sometimes be the cause as well.

Bad translation. The translation makes some sense, but fails to convey the sense of the source utterance. The most common reason is again a serious recognition error.

Results are presented by simply counting the number of translations in a run which fall into each category. By taking account of the "unacceptable hypothesis" judgements, it is possible to evaluate the performance of the system either in a fully automatic mode, or in a mode where the source-language user has the option of aborting misrecognized utterances.

5.2 Evaluation of speech to speech translation

Our intuitive impression, based on many evaluation runs in several different language-pairs, is that the "fine-grained" style of speech-to-text evaluation described in the preceding section gives a much more informative picture of the system's performance than the simple acceptable/unacceptable dichotomy. However, it raises an obvious question: how important, in objective terms, are the distinctions drawn by the fine-grained scale? The preliminary work we now go on to describe attempts to provide an empirically justifiable answer, in terms of the relationship between translation quality and comprehensibility of output speech. Our goal, in other words, is to measure objectively the ability of subjects to understand the content of speech output. This must be the key criterion for evaluating a candidate translation: if apparent deficiencies in syntax or word-choice fail to affect subject's ability to understand content, then it is hard to say that they represent real loss of quality.

The programme sketched above is difficult or, arguably, impossible to implement in a general setting. In a limited domain, however, it appears quite feasible to construct a domain-specific form-based questionnaire designed to test a subject's understanding of a given utterance. In the SLT system's current domain of air travel planning (ATIS), a simple form containing about 20 questions extracts enough content from most utterances that it can be used as a reliable measure of a subject's understanding. The assumption is that a normal domain utterance can be regarded as a database query involving a limited number of possible categories: in the ATIS domain, these are concepts like flight origin and destination, departure and arrival times, choice of airline, and

so on. A detailed description of the evaluation method follows.

The judging interface is structured as a hypertext document that can be accessed through a web-browser. Each utterance is represented by one web page. On entering the page for a given utterance, the judge first clicks a button that plays an audio file, and then fills in an HTML form describing what they heard. Judges are allowed to start by writing down as much as they can of the utterance, so as to keep it clear in their memory as they fill in the form.

The form is divided into four major sections. The first deals with the linguistic form of the enquiry, for example, whether it is a command (imperative), a yes/no-question or a wh-question. In the second section the judge is asked to write down the principal "object" of the utterance. For example, in the utterance "Show flights from Boston to Atlanta", the principal object would be "flights". The third section lists some 15 constraints on the object explicitly mentioned in the enquiry, like "...one-way from New York to Boston on Sunday". Initial testing proved that these three sections covered the form and content of most enquiries within the domain, but to account for unforeseen material the judge is also presented with a "miscellaneous" category. Depending on the character of the options, form entries are either multiple-choice or free-text. All form entries may be negated ("No stopovers") and disjunctive enquiries are indicated by dint of indexing ("Delta on Thursday or American on Friday"). When the page is exited, the contents of the completed form are stored for further use.

Each translated utterance is judged in three versions, by different judges. The first two versions are the source and target speech files; the third time, the form is filled in from the *text* version of the source utterance. (The judging tool allows a mode in which the text version is displayed instead of an audio file being played.) The intention is that the source text version of the utterance should act as a baseline with which the source and target speech versions can respectively be compared. Comparison is carried out by a fourth judge. Here, the contents of the form entries for two versions of the utterance are compared. The judge has to decide whether the contents of each field in the form are compatible between the two versions.

When the forms for two versions of an utterance have been filled in and compared, the results can be examined for comprehensibility in terms of the standard notions of precision and recall. We say that the recall of version 2 of the utterance with respect to version 1 is the proportion of the fields filled in version 1 that are filled in compatibly in version 2. Conversely, the precision is the proportion of the fields filled in in version 2 that are filled

in compatibly in version 1.

The recall and precision scores together define a two-element vector which we will call the *comprehensibility* of version 2 with respect to version 1. We can now define C_{source} to be the comprehensibility of the source speech with respect to the source text, and C_{target} to be the comprehensibility of the target speech with respect to the source text. Finally, we define the quality of the translation to be $1 - (C_{source} - C_{target})$, where $C_{source} - C_{target}$ in a natural way can be interpreted as the extent to which comprehensibility has degraded as a result of the translation process. At the end of the following section, we describe an experiment in which we use this measure to evaluate the quality of translation in the English \rightarrow French version of SLT.

6 An evaluation of the Spoken Language Translator

We begin by presenting the results of tests run in speech-to-text mode on versions of the SLT system developed for six different language-pairs: English \rightarrow Swedish, English \rightarrow French, Swedish \rightarrow English, Swedish \rightarrow French, Swedish \rightarrow Danish, and English \rightarrow Danish. Before going any further, it must be stressed that the various versions of the system differ in important ways; some language-pairs are intrinsically much easier than others, and some versions of the system have received far more effort than others.

In terms of difficulty, Swedish \rightarrow Danish is clearly the easiest language-pair, and Swedish \rightarrow French is clearly the hardest. English \rightarrow French is easier than Swedish \rightarrow French, but substantially more difficult than any of the others. English \rightarrow Swedish, Swedish \rightarrow English and English \rightarrow Danish are all of comparable difficulty. We present approximate figures for the amounts of effort devoted to each language pair in conjunction with the other results.

We evaluated performance on each language-pair in the manner described in Section 5.1 above, taking as input two sets of 200 recorded speech utterances each (one for English and one for Swedish) which had not previously been used for system development. Judging was done by subjects who had not participated in system development, were native speakers of the target language, and were fluent in the source language. Results are presented both for a fully automatic version of the system (Table 1), and for a version with a simulated 'abort' button (Table 2).

Finally, we turn to a preliminary experiment which used the speech-to-speech evaluation methodology from Section 5.2 above. A set of 200 previously unseen English utterances were translated by the system into French speech, using the same kind of subjects as in the previous experiments. Source-language and target-language

speech was synthesized using commercially available, state-of-the-art synthesizers (TrueTalk from Entropics and CNETVOX from ELAN Informatique, respectively). The subjects were only allowed to hear each utterance once. The results were evaluated in the manner described, to produce figures for comprehensibility of source and target speech respectively. The figures are presented in Table 3; we expect to be able to present a more detailed discussion of their significance by the time of the workshop.

In summary, we have improved the standard evaluation method for speech translation by developing a feasible alternative with a more fine-grained taxonomy of acceptability. In order to make the task of evaluation more realistic, we have also created a method in which instead of textual translations it is the spoken form that is judged. This method is currently in embryonic form, but the pilot experiment described here leads us to think that the method shows promise for further development.

An interesting future task would be to investigate the significance of various kinds of written-language translation errors in terms of reducing comprehensibility of the spoken output. This would amount to systematically comparing C_{target} with results obtained in speech-to-text evaluations, divided up according to error categories such as those in our taxonomy.

Acknowledgements

The Danish-related work reported here was funded by SRI International and Handelshøjskolen i København. Other work was funded by Telia Research AB under the SLT-2 project. We would like to thank Beata Forsmark, Nathalie Kirchmeyer, Carin Lindberg, Thierry Reynier and Jennifer Spenader for carrying out judging tasks.

References

- Agnäs, M-S., Alshawi, H., Bretan, I., Carter, D., Ceder, K., Collins, M., Crouch, R., Digalakis, V., Ekholm, B., Gambäck, B., Kaja, J., Karlgren, J., Lyberg, B., Price, P., Pulman, S., Rayner, M., Samuelsson, C. and Svensson, T. 1994. Spoken Language Translator: First Year Report. SRI Cambridge Technical report CRC-043.³
- Alshawi, H. (ed.) 1992. The Core Language Engine. MIT Press.
- Alshawi, H., D. Carter, B. Gambäck and M. Rayner. 1991. Transfer through Quasi Logical Form. Proceedings of ACL-91. Also SRI Cambridge Technical report CRC-021.
- Alshawi, H., and D. Carter. 1994. Training and Scaling Preference Functions for Disambiguation Computational Linguistics 20:4, 635-648. Also SRI Cambridge Technical report CRC-041.
- Carter, D. 1995. Rapid Development of Morphological Descriptions for Full Language Processing Systems. Proceedings of 7th European ACL. Also SRI Cambridge Technical Report CRC-047.
- Carter, D. 1997. The TreeBanker: a Tool for Supervised Training of Parsed Corpora. Proc. ACL/EACL workshop "Computational Environments for Grammar Development and Linguistic Engineering", Madrid. Also SRI Cambridge Technical report CRC-068.
- Frederking, R. and S. Nirenburg. 1994. Three heads are better than one. Proc. 4th ANLP, Stuttgart, Germany, pp 95-100.
- Gambäck, B., and M. Rayner. The Swedish Core Language Engine. Proc. 3rd Nordic Conference on Text Comprehension in Man and Machine, Linköping, Sweden. Also SRI Cambridge Technical Report CRC-025.
- van Harmelen, F., and A. Bundy. 1988. Explanation-Based Generalization = Partial Evaluation (Research Note). Artificial Intelligence 36, pp. 401-412.
- Hemphill, C.T., J.J. Godfrey and G.R. Doddington. 1990. The ATIS Spoken Language Systems pilot corpus. Proc. DARPA Speech and Natural Language Workshop, Hidden Valley, Pa., pp. 96-101.
- Murveit, H., Butzberger, J., Digalakis, V. and Weintraub, M. 1993. Large Vocabulary Dictation using SRI's DECIPHER(TM) Speech Recognition System: Progressive Search Techniques. Proc. Inter. Conf. on Acoust., Speech and Signal, Minneapolis, Mn.
- Pulman, S. 1992. Unification-Based Syntactic Analysis. In (Alshawi (ed), 1992).
- Rayner, M. 1988. Applying Explanation-Based Generalization to Natural-Language Processing. Proc. the International Conference on Fifth Generation Computer Systems, Kyoto, pp. 1267-1274.
- Rayner, M., Alshawi, H., Bretan, I., Carter, D.M., Digalakis, V., Gambäck, B., Kaja, J., Karlgren, J., Lyberg, B., Price, P., Pulman, S. and Samuelsson, C. 1993. A Speech to Speech Translation System Built From Standard Components. Proc. 1st ARPA workshop on Human Language Technology, Princeton, NJ. Morgan Kaufmann. Also SRI Technical Report CRC-031.
- Rayner, M. and Bouillon, P. 1995. Hybrid Transfer in an English-French Spoken Language Translator. Proceedings of IA '95, Montpellier,

³All SRI Cambridge technical reports are available through WWW from <http://www.cam.sri.com>

- France. Also SRI Cambridge Technical Report CRC-056.
- Rayner, M., P. Bouillon and D. Carter. 1995. Using Corpora to Develop Limited-Domain Speech Translation Systems. Proc. Translating and the Computer 17. Also SRI Cambridge Technical Report CRC-059.
- Rayner, M., I. Bretan, M. Wirén, S. Rydin and E. Beshai. 1996. Composition of Transfer Rules in a Multi-Lingual MT System. Proc. Workshop on Future Issues for Multilingual Text Processing. Also SRI Cambridge Technical Report CRC-063.
- Rayner, M., D. Carter and P. Bouillon. 1996. Adapting the Core Language Engine to French and Spanish. Proc. NLP-IA, Moncton, New Brunswick. Also SRI Cambridge Technical Report CRC-061.
- Rayner, M. and D. Carter. 1996. Fast Parsing using Pruning and Grammar Specialization. Proc. ACL-96, Santa Cruz, Ca. Also SRI Cambridge Technical Report CRC-060.
- Rayner, M., and D. Carter. 1997 Hybrid language processing in the Spoken Language Translator. Proc. ICASSP-97, Munich, Germany. Also SRI Cambridge Technical Report CRC-064.
- Rayner, M., D. Carter, I. Bretan, R. Eklund, M. Wirén, S. Hansen, S. Kirchmeier-Andersen, C. Philp, F. Sørensen and H. Erdman Thomsen. 1997. Recycling Lingware in a Multilingual MT System Proc. ACL/EACL workshop "From Research to Commercial Applications", Madrid. Also SRI Cambridge Technical Report CRC-067.
- Samuelsson, C., and M. Rayner. 1991. Quantitative Evaluation of Explanation-Based Learning as an Optimization Tool for a Large-Scale Natural Language System. Proc. 12th IJCAI, Sydney, pp. 609-615.

Table 1: Translation results for six language pairs on 200 unseen utterances, all utterances in test set counted. Note that in both tables on this page, the “effort” figures refer specifically to translation work for the language pair in question, and exclude work on grammar and lexicon development for the individual languages.

Source language	English	English	Swedish	Swedish	Swedish	English
Target language	Swedish	French	English	French	Danish	Danish
Effort (person-months)	8-10	3-5	3-5	1-2	0.5	<0.5
Fully acceptable	46.0%	52.0%	45.0%	19.0%	36.5%	27.0%
Unnatural style	14.0%	10.5%	4.5%	15.0%	0.0%	0.0%
Minor syntactic errors	12.0%	3.5%	12.0%	13.0%	37.5%	28.0%
Clearly useful	72.0%	66.0%	61.5%	47.0%	74.0%	55.0%
Major syntactic errors	7.0%	2.5%	7.5%	13.0%	0.0%	0.0%
Partial translation	6.5%	11.5%	14.5%	17.5%	1.5%	1.5%
Borderline	13.5%	14.0%	22.0%	30.5%	1.5%	1.5%
Nonsense	7.5%	13.0%	10.5%	18.0%	13.0%	30.5%
Bad translation	5.0%	5.5%	4.5%	3.5%	9.0%	10.5%
No translation	2.0%	1.5%	1.5%	1.0%	2.5%	2.5%
Clearly useless	14.5%	20.0%	16.5%	22.5%	24.5%	43.5%

Table 2: Translation results for six language pairs on 200 unseen utterances, ignoring utterances judged as recognition failures.

Source language	English	English	Swedish	Swedish	Swedish	English
Target language	Swedish	French	English	French	Danish	Danish
Effort (person-months)	8-10	3-5	3-5	1-2	0.5	<0.5
Fully acceptable	55.8%	65.8%	60.7%	23.1%	49.0%	35.9%
Unnatural style	15.8%	12.9%	6.4%	19.2%	0.0%	0.0%
Minor syntactic errors	12.1%	3.2%	11.4%	15.4%	38.1%	35.9%
Clearly useful	83.7%	81.9%	78.5%	57.7%	87.1%	71.8%
Major syntactic errors	7.9%	2.6%	10.0%	12.8%	0.0%	0.0%
Partial translation	2.4%	5.8%	5.0%	14.1%	0.7%	2.1%
Borderline	10.3%	8.4%	15.0%	26.9%	0.7%	2.1%
Nonsense	3.0%	4.5%	2.9%	11.5%	4.8%	12.4%
Bad translation	1.2%	3.2%	2.1%	2.6%	5.4%	11.7%
No translation	1.8%	1.9%	1.4%	1.3%	2.0%	2.1%
Clearly useless	6.0%	9.6%	6.4%	15.4%	12.2%	26.2%
(Utterances ignored)	35	45	60	44	53	55

Table 3: Relative comprehensibility of source and target speech for English → French test on 200 unseen utterances.

	Source	Target	Difference	Quality
Precision	97.6%	86.0%	11.6%	88.4%
Recall	97.5%	84.0%	13.5%	86.5%