

# Rapid Portability among Domains in an Interactive Spoken Language Translation System

**Mark Seligman**  
Spoken Translation, Inc.  
Berkeley, CA, USA 94705  
mark.seligman  
@spokentranslation.com

**Mike Dillinger**  
Spoken Translation, Inc.  
Berkeley, CA, USA 94705  
mike.dillinger  
@spokentranslation.com

## Abstract

Spoken Language Translation systems have usually been produced for such specific domains as health care or military use. Ideally, such systems would be easily portable to other domains in which translation is mission critical, such as emergency response or law enforcement. However, porting has in practice proven difficult. This paper will comment on the sources of this difficulty and briefly present an approach to rapid inter-domain portability. Three aspects will be discussed: (1) large general-purpose lexicons for automatic speech recognition and machine translation, made reliable and usable through interactive facilities for monitoring and correcting errors; (2) easily modifiable facilities for instant translation of frequent phrases; and (3) quickly modifiable custom glossaries. As support for our approach, we apply our current SLT system, now optimized for the health care domain, to sample utterances from the military, emergency service, and law enforcement domains, with discussion of numerous specific sentences.

## 1 Introduction

Recent years have seen increasing research and commercial activity in the area of Spoken Language Translation (SLT) for mission-critical applications. In the health care area, for instance, such products as *Converser* (Dillinger & Seligman, 2006), *S-MINDS* ([www.fluentia.com](http://www.fluentia.com)), and *Med-SLT* (Bouillon et al, 2005) are coming into use. For military applications, products like *Phraselator* ([www.phraselator.com](http://www.phraselator.com)) and *S-MINDS* ([www.fluentia.com](http://www.fluentia.com)) have been deployed. However, the demand for real-time translation is by no means restricted to these areas: it is clear in numerous other areas not yet extensively addressed – emergency services, law enforcement, and others.

Ideally, a system produced for one such domain (e.g., health care) could be easily ported to other domains. However, porting has in practice proven difficult. This paper will comment on the sources of

this difficulty and briefly present an approach to rapid inter-domain portability that we believe is promising. Three aspects of our approach will be discussed: (1) large general-purpose lexicons for automatic speech recognition (ASR) and machine translation (MT), made reliable and usable through interactive facilities for monitoring and correcting errors; (2) easily modifiable facilities for instant translation of frequent phrases; and (3) quickly modifiable custom glossaries.

As preliminary support for our approach, we apply our current SLT system, now optimized for the health care domain, to sample utterances from the military, emergency service, and law enforcement domains.

With respect to the principal source of the porting problems affecting most SLT systems to date: most systems have relied upon statistical approaches for both ASR and MT (Karat and Nahamoo, 2007; Koehn, 2008); so each new domain has required extensive and high-quality in-domain corpora for best results, and the difficulty of obtaining them has limited these systems' portability. The need for in-domain corpora can be eliminated through the use of a quite general corpus (or collection of corpora) for statistical training; but because large corpora give rise to quickly increasing perplexity and error rates, most SLT systems have been designed for specialized domains.

By contrast, breadth of coverage has been a central design goal of our SLT systems. Before any optimization for a specific domain, we “give our systems a liberal arts education” by incorporating very broad-coverage ASR and MT technology. (We presently employ rule-based rather than statistical MT components, but this choice is not essential.) For example, our MT lexicons for English<->Spanish translation in the health care area contain roughly 350,000 words in each direction, of which only a small percentage are specifically health care terms. Our translation grammars (presently licensed from a commercial source, and further developed with our collaboration) are similarly designed to cover the structures of wide-ranging general texts and spoken discourse.

To deal with the errors that inevitably follow as coverage grows, we provide a set of facilities that enable users from both sides of the language barrier to

---

© 2008. Licensed under the *Creative Commons Attribution-Noncommercial-Share Alike 3.0 Unported* license (<http://creativecommons.org/licenses/by-nc-sa/3.0/>). Some rights reserved.

interactively monitor and correct such errors. We have described these interactive techniques in (Dillinger and Seligman, 2004; Zong and Seligman, 2005; Dillinger and Seligman, 2006; and Seligman and Dillinger, 2006). With users thus integrated into the speech translation loop, automatically translated spoken conversations can range widely with acceptable accuracy (Seligman, 2000). Users can move among domains with relative freedom, even in advance of lexical or other domain specialization, because most domains are already covered to some degree. After a quick summary of our approach (in Section 2), we will demonstrate this flexibility (in Section 3).

While our system's facilities for monitoring and correction of ASR and MT are vital for accuracy and confidence in wide-ranging conversations, they can be time consuming. Further, interactivity demands a minimum degree of computer and print literacy, which some patients may lack. To address these issues, we have developed a facility called *Translation Shortcuts*<sup>TM</sup>, through which prepared translations of frequent or especially useful phrases *in the current domain* can be instantly executed by searching or browsing. The facility is described in (Seligman and Dillinger, 2006). After a quick description of the Translation Shortcuts facility (Section 4), this paper will emphasize the contribution of the Translation Shortcuts facility to domain portability, showing how a domain-specific set of Shortcuts can be composed and integrated into the system very quickly (Section 5).

Finally, while the extensive lexical resources already built into the system provide the most significant boost to domain portability in our system, it will always be desirable to add specialized lexical items or specialized meanings of existing ones. Section 6 will briefly present our system's *glossary import* facility, through which lexical items can be added or updated very quickly. Our concluding remarks appear in Section 7.

## 2 Highly Interactive, Broad-coverage SLT

We now briefly summarize our group's approach to highly interactive, broad-coverage SLT. Our systems stress interactive monitoring and correction of both ASR and MT.

First, users can monitor and correct the speaker-dependent speech recognition system to ensure that the text which will be passed to the machine translation component is as correct as necessary. Voice commands (*e.g.*, **Scratch That** or **Correct <incorrect text>**) can be used to repair speech recognition errors. Thus, users of our SLT systems in effect serve to enhance the interface between ASR and MT.

Next, during the MT stage, users can monitor, and if necessary correct, translation errors.

As an initial safeguard against translation errors, we supply a *back-translation*, or re-translation of the translation. Using this paraphrase of the initial input, even a monolingual user can make an initial judgment concerning the quality of the preliminary machine translation output. If errors are seen, the user can modify specific parts of the input and retranslate. (Other systems, *e.g.* IBM's MASTOR (Gao et al, 2006), have also employed re-translation. Our implementations, however, exploit proprietary technologies to ensure that the lexical senses used during back-translation accurately reflect those used in forward translation. We also allow users to modify part or all of the input before regenerating the translation and back-translation.)

In addition, if uncertainty remains about the correctness of a given word sense, we supply a proprietary set of Meaning Cues<sup>TM</sup> – synonyms, definitions, examples, pictures, etc. – which have been drawn from various resources, collated in a database (called SELECT<sup>TM</sup>), and aligned with the respective lexica of the relevant MT systems. (In the present English<>Spanish version of the system, this database contains some 140,000 entries, corresponding to more than 350,000 lexical entries. The cues are automatically grouped by meaning, and cue groups are automatically mapped to MT lexica using proprietary techniques – thus in effect retrofitting an MT system with the ability to explain to users the meanings of its pre-existing lexical items.) With these cues as guides, the user can monitor the current, proposed meaning and if necessary select a different, preferred meaning from among those available. Automatic updates of translation and back-translation then follow. (Our current MT vendor has modified its rule-based translation engine to allow specification of a desired sense when translating a word or expression; we provide guidelines for other vendors to do likewise. Comparable modifications for statistical MT engines will entail the setting of temporary weightings that will bias the selection of word or phrase translations for the current sentence only.) Future versions of the system will allow personal word-sense preferences thus specified in the current session to be optionally stored for reuse in future sessions, thus enabling a gradual tuning of word-sense preferences to individual needs. (However, such persistent personal preferences will still be applied sentence by sentence, rather than by permanently modifying lexica or phrase tables. Further, users will always be able to temporarily override, or permanently reset, their personal preferences.) Facilities will also be provided for sharing such preferences across a working group.

Given such interactive correction of both ASR and MT, wide-ranging, and even playful, exchanges become possible (Seligman, 2000). Such interactivity within a speech translation system enables increased accuracy and confidence, even for wide-ranging conversations.

### 3 Advantages of Very Broad Coverage for Domain Switching

This section discusses the advantages of very broad lexical coverage for rapid domain porting. Using our interactive SLT system in its present configuration, optimized for the health care domain but with a general-purpose foundation of over 60,000 lexical items for ASR and 350,000 lexical items for rule-based MT, we will test several input sentences from each of three distinct domains in which translation is mission-critical – military, emergency response, and law enforcement. The test sentences were invented by the authors; readers can judge their plausibility. They were pronounced by Seligman using the built-in microphone of a Motion Computing LE1600 tablet PC equipped with a push-to-talk button.

For each input, we will show (1) the English input, (2) the original Spanish translation, and (3) the English back-translation. We also comment on several factors:

- If ASR errors occurred, we describe their interactive resolution. (All inputs were corrected before proceeding with translation. All corrections were made by voice.)
- If our Meaning Cues facility indicated questionable meanings for any of the expressions in the input, we note the problems and describe the resolutions.
- Some problems in translation result from bugs or gaps in the translation component. These are

marked for repair. (Because our MT components are presently rule-based, we can address such problems individually and manually. If a statistical MT component were used instead, the recorded errors could guide the selection of texts for further training.)

As mentioned, in our system, the back-translation is designed to function as the first line of defense against inadequate translation. If an unsatisfactory back-translation is obtained, we advise users to rephrase the input and translate again until satisfied. (False negatives sometimes occur, though we work to eliminate them; however, it is best to err on the side of caution.) If the back-translation is satisfactory, we advise checking the Meaning Cues as a defense against false positives. These may result if an ambiguous English input word is translated into Spanish in the wrong sense (for instance, *bank* may be translated as *banco* ("financial institution") when *orilla del río* ("riverbank") is wanted), but is nevertheless retranslated as the same ambiguous English word (*bank*). We are experimenting with mechanisms to eliminate such cases by substituting non-ambiguous synonyms in the back-translation for ambiguous input words. In the current tests, if back-translations are judged insufficient to convey the intended meaning, paraphrases are substituted and any lexical translation errors are corrected until acceptable results are achieved. All such paraphrases are displayed below, whether they involve simple word substitutions or more extensive changes.

#### Military Domain

input	initial translation	back-translation	comments
Watch out for mines around here	tenga cuidado con minas por aquí	Be careful with mines around here	ASR: <i>mines</i> misrecognized as <i>minds</i> MT: <i>mines</i> mistranslated as <i>minas</i> ("reserve, mine"); meaning interactively changed to <i>minas</i> ("booby trap, ground-emplaced mine, land mine")
[retranslation]	tenga cuidado con minas por aquí	Be careful with mines around here	Spanish is unchanged, but correct meaning has been verified.
We confiscated their arms last week	confiscamos sus brazos la semana pasada	We confiscated their arms last week	MT: <i>arms</i> mistranslated as <i>brazos</i> ("brachia"); retranslated as <i>armas</i> ("arm, gun, weapon")
[retranslation]	confiscamos sus armas la semana pasada	We confiscated their weapons last week	
The operation is scheduled for oh 600	la operación es programada para oh 600	The operation is programmed for oh 600.	ASR: <i>The</i> misrecognized as <i>knee</i> . MT: Translation of <i>oh 600</i> is uncertain
The operation is scheduled for 6 a.m.	la operación es programada para 6 a.m.	The operation is programmed for 6 a.m.	MT: Translation of <i>6 a.m.</i> is still unclear.
The operation is scheduled for six o'clock in the morning	la operación es programada para las seis de la mañana	The operation is programmed for six in the morning	MT: Translation is now verified, given slight rewording ( <i>six</i> instead of <i>six o'clock</i> ).

We're training them as guerrillas	Los entrenamos como guerrillas	We train them like guerrillas	ASR: Correct spelling ( <i>c.f. gorillas</i> ) was produced. MT: Bug: tolerable back-translation error: <i>like</i> should be <i>as</i> .
-----------------------------------	--------------------------------	-------------------------------	--

NOTE: For the military domain and more generally, improved translation of *day times*, especially when expressed as *military time*, is clearly needed.

### Emergency Response Domain

input	initial translation	back-translation	comments
Tell them to drop the food at headquarters	Dícales a ellos que dejen caer la comida en cuartel general	Tell them to them that they drop the food in headquarters	MT: Bug: tolerable Spanish>English mistranslation of pattern " <i>dícales a ellos que &lt;action&gt;</i> " (" <i>tell them to &lt;action&gt;</i> "); <i>drop</i> mistranslated as "drop down, drop away, let fall, ...", but no suitable alternate meaning found; substituting <i>drop off</i>
... drop off ...	dícales a ellos que dejen caer la comida en cuartel general	Tell them to them that they drop the food in headquarters	MT: translation and back-translation unchanged; still no suitable meaning; substituting <i>leave</i>
... leave ...	Dícales a ellos que dejen la comida en cuartel general	Tell them to them that they leave the food at headquarters	MT: back-translation and Meaning Cues now okay
We need more shovels and crowbars right now	Necesitamos más palas y más palancas ahora mismo	we need more shovels and more levers right now	MT: back-translation <i>levers</i> is considered okay for <i>crowbars</i>
It's a matter of life and death	es cuestión de la vida y la muerte	it is issue of life and Death	MT: capitalization of <i>death</i> prompts uncertainty; rephrasing
It's absolutely critical.	Es absolutamente crítico.	it's absolutely critical	MT: meaning cues for <i>critical</i> are okay: "final, significant, key, crucial ..."
These people are desperately short of water	Estas personas andan desesperadamente escasas de agua.	These people are desperately scarce of water	MT: Spanish is okay, but poor back-translation of <i>escasas de</i> (should be "short of/low on") gives false negative, low confidence. Substituting <i>low on</i> .
.. low on ...	Estas personas andan desesperadamente de capa caída en agua.	These people incur in desperately on water.	MT: worse; rephrasing
These people are desperate for water	estas personas están desesperadas para agua.	These people are desperate for water.	MT: Preposition error in Spanish ( <i>para</i> should be <i>por</i> ) gives false positive, but meaning is clear

### Law Enforcement Domain

input	initial translation	back-translation	comments
Step away from the car	Aléjese del coche	Get away from the car	MT: <i>get away</i> is acceptable for <i>step away</i>
May I see your license, please	Que pueda ver su licencia, por favor.	That I can see your license, please.	MT: Unacceptable mistranslation of pattern " <i>que pueda &lt;action&gt;, por favor</i> " (" <i>may I &lt;action&gt;, please</i> "); rephrasing
Show me your license, please	Muéstreme su licencia, por favor.	Show me your license, please	
Keep your hands where I can see them	Conserve sus manos donde las puedo ver.	Preserve your hands where I can see them.	MT: <i>keep</i> mistranslated as <i>conserve</i> ("take, hold, maintain, save, retain, preserve, ..."); retranslated as <i>mantenga</i> ("keep")

[retranslation]	Mantenga sus manos donde las puedo ver	Keep your hands where I can see them	
How long have you been living at this address?	Cuánto tiempo usted ha vivido en esta dirección?	How long have you been living in this address?	MT: minor but tolerable error with prepositions
Who's your insurer	Quién es su asegurador	Who is your insurer	

NOTE: General-purpose Spanish>English pattern “*que pueda <action>, por favor*” (“*may I <action>, please*”) requires fix for all domains.

#### 4 Translation Shortcuts

Having summarized our approach to highly interactive speech translation and discussed the advantages of very broad lexical and grammatical coverage for domain switching, we now turn to the use of Translation Shortcuts™ in domain ports. This section briefly describes the facility; and Section 5 explains the methods for quickly updating Shortcuts as an element of a rapid port.

A Translation Shortcut contains a short translation, typically of a sentence or two, which has been pre-verified, whether by a human translator or through the use of the system’s interactive tools. Thus re-verification of the translation is unnecessary. In this respect, Translation Shortcuts provide a kind of translation memory. However, it is a handmade sort of memory (since Shortcuts are composed by linguists or explicitly saved by users) and a highly interactive sort as well (since users can browse or search for Shortcuts, can make and categorize their own Shortcuts, and are advised when the input matches a Shortcut). It is in the ease of composition or customization, as well as in the quality of the interaction, that innovation can be claimed.

We can consider the quality of interaction first. Access to stored Shortcuts is very quick, with little or

no need for text entry. Several facilities contribute to meeting this design criterion:

- A *Shortcut Search* facility can retrieve a set of relevant Shortcuts given only keywords or the first few characters or words of a string. The desired Shortcut can then be executed with a single gesture (mouse click or stylus tap) or voice command.

NOTE: If no Shortcut is found, the system automatically allows users access to the full power of broad-coverage, interactive speech translation. Thus, a seamless transition is provided between the Shortcuts facility and full, broad-coverage translation.

- A *Translation Shortcuts Browser* is provided, so that users can find needed Shortcuts by traversing a tree of Shortcut categories. Using this interface, users can execute Shortcuts by tapping or clicking alone.

Figure 1 below shows the Shortcut Search and Shortcuts Browser facilities in use.

- On the left, the Translation Shortcuts Panel contains the Translation Shortcuts Browser, split into two main areas, Shortcuts Categories (above) and Shortcuts List (below).




Figure 1: The Input Screen, showing the Translation Shortcuts Browser and Shortcut Search facilities. Note the new **Nutrition** category and the results of automatic Shortcut Search.

- The Categories section of the Panel shows current selection of the **Nutrition** category, containing frequently used questions and answers for a nutrition interview. This new category was created overnight, as described in Section 5, below. Currently hidden is its **Staff** subcategory, containing expressions most likely to be used by health care staff members. There is also a **Patients** subcategory, used for patient responses. Categories for **Background information, Directions**, etc. are also visible.
- Below the Categories section is the Shortcuts List section, containing a scrollable list of alphabetized Shortcuts. Double clicking on any visible Shortcut in the List will execute it. Clicking once will select and highlight a Shortcut. Typing **Enter** will execute any currently highlighted Shortcut.

We turn our attention now to the Input Window, which does double duty for Shortcut Search and arbitrary text entry for full translation. The search facility is also shown in Figure 1.

- Shortcuts Search begins automatically as soon as text is entered by any means – voice, handwriting, touch screen, or standard keyboard – into the Input Window.
- The **Shortcuts Drop-down Menu** appears just below the Input Window, as soon as there are results to be shown. The user has entered “Do you have”. The drop-down menu shows the results of a search within the new **Nutrition** category based upon these initial characters.

If the user goes on to enter the exact text of any Shortcut in this category, *e.g.* “Do you have any food allergies?,” the interface will show that this is in fact a Shortcut, so that verification of translation accuracy will not be necessary.

However, final text not matching a Shortcut, *e.g.* “Do you have any siblings?” will be passed to the routines for full translation with verification.

A *Personal Translation Shortcuts™* facility is in progress for future versions of the system: once a user has verified a translation via the interactive facilities described above, he or she can save it for future reuse by pressing a **Save as Shortcut** button. The new custom Shortcut will then be stored in a personal profile. Facilities for sharing Shortcuts will also be provided.

## 5 Rapid Customization of Translation Shortcuts for New Domains

Translation Shortcuts are stored and distributed as text-format XML files. Each file contains information about which categories (*e.g.* **Nutrition**) and subcategories (**Staff, Patient**, etc.) to which each phrase belongs. Since Shortcuts are stored as external data files, integration of new Shortcuts into the system is straightforward and highly scalable. Once we have built a database of frequently used expressions and their translations for a given domain (in which there may be thousands of expressions or just a few), we can automatically generate the associated files in XML format in minutes. Once this new file is added to the appropriate directory, the Shortcuts become usable in the next session for text- or voice-driven searching and browsing. The entire sequence can be completed overnight. In one case, the Nutrition Department of a major hospital submitted several pages of frequently asked questions, which were entered, translated, re-generated as an XML file, and integrated into the system for demonstration the next day.

```
<Category categoryName1="Nutrition" categoryName2="Alimentación">
  <Categories>
    <Category categoryName1="Staff" categoryName2="Personal">
      <Shortcuts>
        <Shortcut categoryPath="Nutrition\\Staff">
          <Language1Text>Do you have any food allergies?</Language1Text>
          <Language2Text>¿Tiene alguna alergia a alguna comida?</Language2Text>
        </Shortcut>
        <Shortcut categoryPath="Nutrition\\Staff">
          <Language1Text>Can you tolerate milk?</Language1Text>
          <Language2Text>¿Tolera la leche?</Language2Text>
        </Shortcut>
        <Shortcut categoryPath="Nutrition\\Staff">
          <Language1Text>Do you follow a special diet at home?</Language1Text>
          <Language2Text>¿Sigue alguna dieta especial en casa?</Language2Text>
        </Shortcut>
      </Shortcuts>
    </Category>
  </Categories>
</Category>
```

Figure 2: Sample fragment of an automatically formatted Translation Shortcuts file for the **Nutrition>Staff** category and subcategory.

## 6 Use of the Glossary Import for Quick Domain Switching

Similarly, our system includes a glossary import function which supports quick addition of domain-specific or other custom lexical information (*e.g.*, site-specific or client-specific vocabulary), once again in text format. This glossary file may provide additional terms or may stipulate preferred (and thus overriding) translations for existing terms. The glossary file is automatically generated from a simple, two-column text-format file in which each line contains the source-language and target-language terms. A system utility will then generate the necessary linguistic markup (in curly brackets in Figure 3) for each of the terms. (Markup can be elaborated as appropriate for the machine translation engine in use, *e.g.* to specify verb sub-categorization, semantic class, etc.) Like the XML file used for Translation Shortcuts, the resulting custom glossary file can simply be placed in the appropriate directory.

```
hemolítico { A, 11, 6, 0, } = hemolytic
hemolitopoyético { A, 11, 6, 0, } = hemolytopoietic
hemolizable { A, 11, 6, 0, } = hemolyzable
hemolización { N, 2, 2, 1, } = hemolyzation
hemolizar { V, 7, 0, 1, } = hemolyze
derecho { A, 11, 6, 0, } = right
```

Figure 3. Sample glossary-import entries for the health care domain.

Here, the entry for *right* establishes the "right-hand" sense as the system-wide default, overriding the current global default sense ("correct"). (The new global default can, however, be overridden in turn by a personally preferred sense as specified by a user's personal profile; and both kinds of preferences can be overridden interactively for any particular input sentence.) The other entries are domain-specific lexical additions for health care not in the general dictionary.

We make no claims for technical innovation in our Glossary Import facility, but simply point out its usefulness for rapid porting, in that new lexical items, or new preferred senses for old items, can be altered per user and from session to session.

## 7 Conclusion

The principal source of the porting problems affecting most SLT systems to date, we have observed, is that, given the general current reliance upon statistical approaches for both ASR and MT, each new domain has required an extensive and difficult-to-obtain new corpus for best results. One might consider the use of a single very large and quite general corpus (or collection of corpora) for statistical training; but large corpora engender quickly

increasing perplexity and error rates, so this very-broad-coverage approach has generally been avoided.

Our approach, however, has been to adopt a broad-coverage design nevertheless, and to compensate for the inevitable increase in ASR and MT errors by furnishing users with interactive tools for monitoring and correcting these mistakes. (We have to date used rule-based rather than statistical MT components, but comparable interactive facilities could be supplied for the latter as well. Operational prototypes for English<>Japanese and English<>German suggest that the techniques can also be adapted for languages other than English<>Spanish.) Because such interactive tools demand some time and attention, we have also put into place easily modifiable facilities for instant translation of frequent phrases (Translation Shortcuts). And finally, since even systems with very large lexicons will require specialized lexical items or specialized meanings of existing ones, we have implemented a quick glossary import facility, so that lexical items can be added or updated very easily.

Our current SLT system, optimized for health care, is now in use at a medium-sized hospital in New Jersey, with more than twenty machines installed. For this paper, we have applied the same system, without modifications, to sample utterances from the military, emergency service, and law enforcement domains. While this exercise has yielded no quantitative results, readers can judge whether it demonstrates that users can convey mission-critical information with acceptable reliability in multiple domains, even in advance of any porting efforts. Users do pay a price for this flexibility, since time and attention are required for monitoring and correcting to achieve reliable results. However, when users judge that accuracy is not crucial, or when they are unable to monitor and correct, they can simply accept the first translation attempt as is. (A bilingual transcript of each conversation, soon to optionally include the back-translation, is always available for later inspection.) They can also gain considerable time through the use of Translation Shortcuts.

## References

- Bouillon, P., Rayner, M., et al. 2005. A Generic Multi-Lingual Open Source Platform for Limited-Domain Medical Speech Translation. Presented at *EAMT 2005*, Budapest, Hungary.
- Dillinger, M. and Seligman, M. 2006. *Converser™* : highly interactive speech-to-speech translation for health care. *HLT-NAACL 2006: Proceedings of the Workshop on Medical Speech Translation* (pp.40-43). New York, NY, USA.
- Dillinger, M. and Seligman, M. 2004. System description: A highly interactive speech-to-speech translation system. In: Robert E. Frederking and Kathryn B. Taylor (Eds.),

*Machine translation: from real users to research: 6th conference of the Association for Machine Translation in the Americas -- AMTA 2004* (pp. 58-63). Berlin: Springer Verlag.

Gao, Y., Liang, G., Zhou, B., Sarikaya, R., et al. (2006). IBM MASTOR system: multilingual automatic speech-to-speech translator. In: *HLT-NAACL 2006: Proceedings of the Workshop on Medical Speech Translation* (pp.57-60). New York, NY, USA.

Karat, C-M. and Nahamoo, D. 2007. Conversational interface technologies. In A. Sears & J. Jacko (Eds.), *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies, and Emerging Applications*. Mahwah, NJ: L. Erlbaum.

Koehn, P. 2008. *Statistical Machine Translation*. New York: Cambridge University Press.

Seligman, M.. 2000. Nine Issues in Speech Translation. *Machine Translation*, 15, 149-185.

Seligman, M. and Dillinger, M. 2006. Usability issues in an interactive speech-to-speech translation system for health care. *HLT-NAACL 2006: Proceedings of the Workshop on Medical Speech Translation* (pp. 1-8). New York, NY, USA.

Zong, C. and Seligman, M. 2005. Toward Practical Spoken Language Translation. *Machine Translation*, 19, 113-137.