The South African Sign Language Machine Translation Project: Issues on Non-manual Sign Generation

LYNETTE VAN ZIJL and ANDRIES COMBRINK
Stellenbosch University

We report on the South African Sign Language Machine Translation project, and more specifically on our experience in extending a synchronous tree-adjoining grammar parser approach in order to generate non-manual signs and construct a suitable signing space. We show that post-processing of the target language tree, after transfer rules have been applied, results in a simple and efficient mechanism to generate information on non-manual signs for use in a signing avatar.

Categories and Subject Descriptors: I.2.7 [**Artificial Intelligence**]: Natural Language Processing—*Machine Translation*

General Terms: Machine Translation

Additional Key Words and Phrases: Generation of Non-manual Signs, Machine Translation, Sign Language

1. INTRODUCTION

The South African Deaf community is marginalized in its access to information, as interpreters are scarce and prohibitively expensive. The South African Sign Language Machine Translation (SASL-MT) project at Stellenbosch University is part of an effort to develop assistive technologies to bridge the communication gap between the hearing and the Deaf. Such technologies include tools to assist the hearing in learning South African Sign Language (SASL), as well as tools to assist in the machine translation of English text to SASL.

The machine translation of spoken languages to sign languages has received much attention over the past decade [Fotinea et al. 2005; Huenerfauth 2004; Suszczanska et al. 2002; Speers 2001]. The typical components of such a system comprise:

—parsing and analysing the input (source) language text,

—translating the source language text into the target language, and

—generating the resulting target language sentence by means of a graphical signing avatar.

The parsing method influences the translation phase to a large degree, and current systems show widely differing approaches in this regard, each with its own advantages and disadvantages (see [Huenerfauth 2003] for an overview). In addition, because sign languages are visual-spatial rather than written languages, some unique problems arise in the translation and generation phases of these systems. Amongst others, such problems include the construction of the signing space (the area in front of the signer where signs are executed), and the generation of so-called non-manual signs (facial expressions and body movement). Typically, successful translation in this regard would require semantic and discourse analysis, with varying degrees of success.

This article reports on the current status of the SASL-MT project, and specifically address our progress in the construction of the signing space, and in the generation of non-manual signs. We base our generation of non-manual signs on known techniques from text-to-speech systems, which is, to our knowledge, a novel approach to non-manual sign generation. In particular, we analyze the English input text to extract emotional content, stress and other semantic information in order to generate better non-manual signs.

The rest of this article is organized as follows: we briefly summarize related work in Section 2, followed by the current status of the SASL-MT project in Section 3. We then address our proposed approach to signing space construction and non-manual sign generation in Section 4, and discuss our results. We conclude in Section 5.

Author Addresses: Lynette van Zijl and Andries Combrink, Department of Computer Science, Stellenbosh University, P/Bag X1, Matieland, 7602, South Africa; lynette@cs.sun.ac.za. This research was supported by NRF grant #2053436.

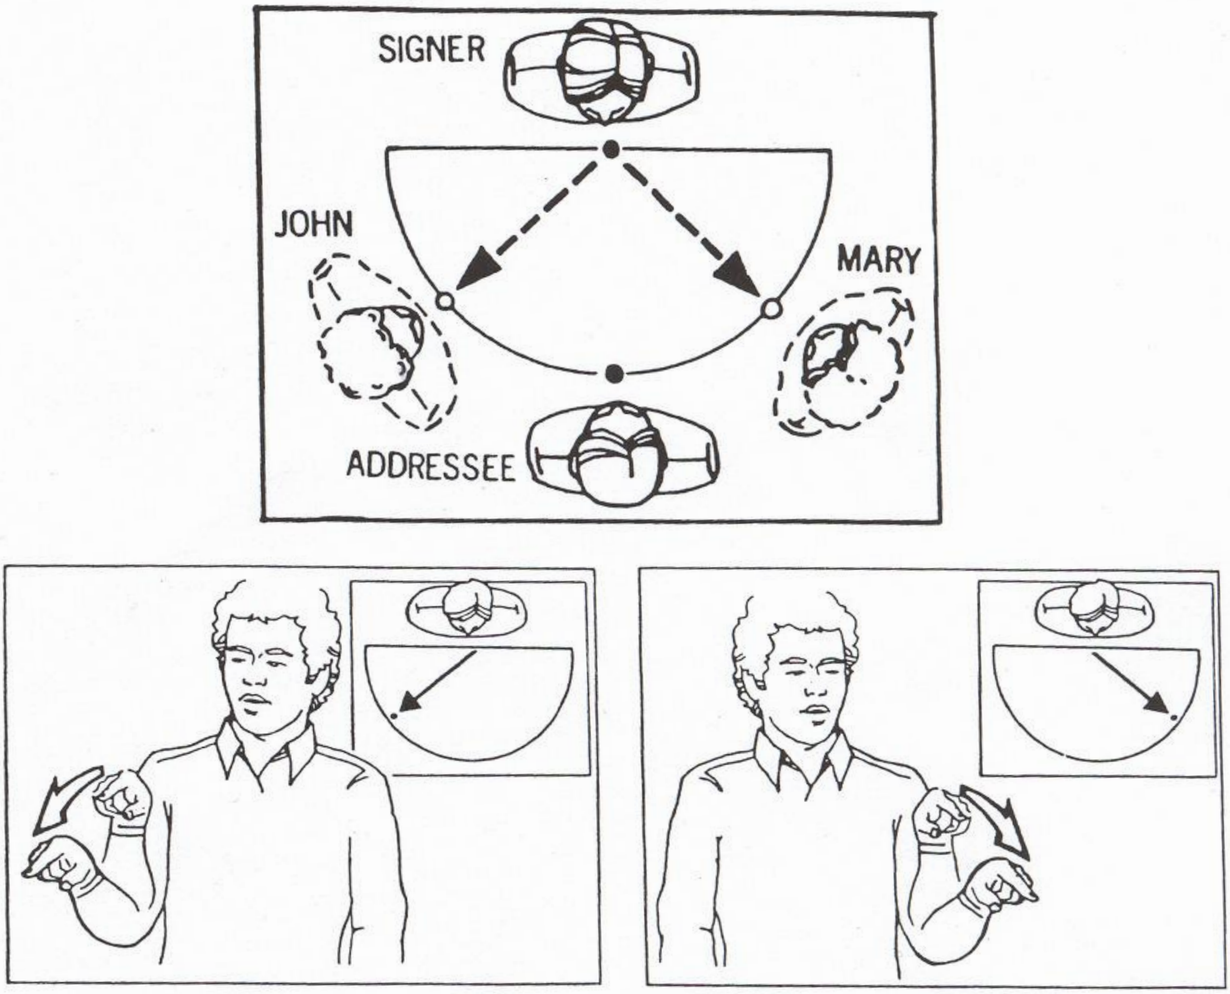

Figure 1. Pronoun indication in the signing space (reproduced from [Lillo-Martin et al. 1990]). The top diagram indicates the relative positions of persons *John* and *Mary*, while the bottom left diagram shows pronoun *he* for *John* and the bottom right diagram shows pronoun *she* for *Mary*.

2. RELATED WORK

Sign languages are natural languages, with their own syntax and own grammar. Two of the many interesting aspects of the linguistics of sign languages that are directly relevant to this paper, are their use of pronouns and of non-manual signs.

In a spoken language, proper nouns can be refered to by means of pronouns. For example, in the text

Harry ate a strawberry. He likes strawberries very much.

the pronoun *He* refers to the proper noun *Harry*. In sign languages, the analogous phenomenon is signed by indicating a position in the signing space for the first occurrence of the proper noun, and subsequent pronouns are then signed by pointing to the position of the proper noun in the signing space [Lillo-Martin and Klima 1990] – see Figure 1. In our translation system, we therefore have to translate pronouns into positions in the signing space. This implies that all the problems normally associated with pronoun resolution [Mitkov 2003] in a machine translation system, also occur in sign language machine translation systems.

Another issue in sign languages is the role played by the non-manual signs. In sign languages, a sign is characterized by the form of the hands accompanied by the movement of the arms. In addition, facial expression and body movement form an integral part of the meaning of a word or sentence. The facial expression and body movement are known as the non-manual part of the sign. It is important to note that non-manual signs form an integral part of the grammar of sign languages, and executing a sign without its non-manual part renders it nonsensical. In Figure 2, we illustrate two signs in SASL: here, the hand form and arm movements are identical for the signs for *maybe* and *weight*, but the non-manual questioning expression with *maybe* changes the meaning.

Figure 2. The SASL sign for *maybe* with questioning facial expression on the left, and for *weight* with neutral facial expression on the right.

In the machine translation of sign languages, the construction of the signing space and the question of non-manual signs remain active implementation and research issues.

Some early work with regards to the machine translation of sign languages was done by Grieve-Smith [Grieve-Smith 2002], at the University of New Mexico, with his SignSynth system. One of the first detailed linguistic analysis with regard to the issues that would appear in machine translation systems of sign languages, is that of Speers [Speers 2001]. A small prototype system based on synchronous tree adjoining grammars was developed at the University of Pennsylvania [Zhao et al. 2000] by Zhao, and this system is currently being redesigned and improved by Huenerfauth [Huenerfauth 2004]. Commercially, the best known systems are the Visicast system [Bangham et al. 2000] and the VCom3D system [VCom3D 2004]. Currently, the largest systems in development are that of Stavroula *et al* for the translation of Greek to Greek Sign Language [Fotinea et al. 2005], and a system for the translation of Polish to Polish Sign Language [Suszczanska et al. 2002]. Both of these two systems involve combined efforts of large teams of linguists and developers, and both have detailed linguistic information available about the respective sign languages. There is also an emerging effort towards the statistical machine translation of sign languages, in particular for German Sign Language [Bungeroth and Ney 2004]. An investigation into example-based translation of Irish Sign Language is also taking place [Morrissey and Way 2005].

The major issues to be considered in a machine translation system from a spoken language to a sign language, are the following:

—the linguistic information that is available about the sign language – the linguistic analysis of sign languages is not nearly as complete as for their spoken counterparts, and many sign languages around the world have not been linguistically investigated;

—the choice of parser for the input spoken language influences the semantic information that is available for the generation of the sign language – in particular, discourse analysis is typically needed to build a signing space and incorporate non-manual signs;

—the generation phase of the sign language requires a mapping from the translated output to a format suitable to drive a graphical avatar; and lastly,

—the implementation issues surrounding the graphical avatar.

In the next section, we shall refer back to each of these issues as it pertains to the development decisions taken in the SASL-MT project.

3. THE SASL-MT PROJECT

In this section, we briefly summarize the progress of the SASL-MT project, and highlight design decisions where appropriate.

At the inception of the SASL-MT project, no linguistic data was available for SASL. In fact, SASL is linguistically under-investigated, with an almost complete lack of published linguistic information. The project therefore had no established base of word lists, lexicons, or corpora to start off with, and we hence initiated our own data gathering. A large amount of effort was put into data capturing, data analysis, and the buildup of rudimentary word lists and phrase lists. Note that we based our data gathering as far as possible on the techniques suggested by Neidle [Neidle et al. 2000]. Trained linguists (but not SASL experts) analysed the captured data in cooperation with different native speakers of SASL. Through this process rudimentary word lists and phrase lists were constructed. As an aside, the reader may note that this lack of data about SASL directly implies that neither statistical machine translation nor example-based machine translation was an option for this project.

The manpower and other limitations of the SASL-MT project made it imperative that the development of the parser should be as quick and easy as possible. In addition, we wanted to re-use existing data and software for the analysis of the English text. For these reasons, we decided to implement a standard Synchronous Tree Adjoining Grammar (STAG) parser [Joshi and Schabes 1997]. The main advantages of a STAG parser for our project specifically, are

—the fact that an initial set of transfer rules could be developed, and later refined as more linguistic information becomes available;

—implementation time of the parser is quite short; and

—the approach does not restrict the choice of notation to drive the graphical signing avatar.

The STAG parser implementation has been completed and fully tested. Using the English trees developed by Zhao [Zhao et al. 2000], we are now able to analyse English text and build syntactic parse trees. Based on examples from our SASL data, we constructed a small set of rule trees for SASL, together with the transfer rules from the English trees to the SASL trees. As a simple example, there is a rule that removes all articles (that is, *a*, *an*, *the*) in the translation of the English to the SASL tree. Therefore, an English tree representing the sentence

The man eats.

is translated to the SASL tree representing the sentence

Man eats.

These rules were constructed with the full knowledge that they may possibly need improvement after extensive user testing. Therefore, given an English sentence, our system can now produce an output tree with leaves containing SASL glosses. Note that no discourse information has been extracted from the source text at this stage, but that the discourse information is necessary for constructing a signing space and non-manual signs[1].

As Huenerfauth [Huenerfauth 2003] points out, the major drawback of the original STAG-based system of the University of Pennsylvania is the fact that it lacks any direct output to indicate non-manual signs in the final translation. We illustrate in the next section that this disadvantage can be counter-acted, in that standard analysis techniques on the English input text can be used to modify/flag the original output tree.

In the second last phase of the SASL-MT system, a notational mapping from the SASL output tree to a graphical description suitable for input to a signing avatar, is performed. This part of the system is still under construction. For the last phase, a generic signing avatar has been developed and is fully functional [Fourie 2006].

4. NON-MANUAL SIGNS IN SASL-MT

In this section, we discuss our implementation of signing space construction and generation of non-manual signs. The signing space construction in many text to sign language machine translation systems operate on a sentence-by-sentence level, where individual persons or objects in the current sentence is placed anew in the signing space. Our approach is based on a history of sentences spanning a whole paragraph, and hence gives a more realistic representation of the signing space. For the non-manual sign generation, our approach is based on text-to-speech techniques. To our knowledge, no such approach had previously been investigated in sign language machine translation. Here, we analyze the emotive[2] content of the source text, and the result of this analysis is flagged onto the output parse tree.

4.1 Signing space construction

The construction of a signing space is simple, and only requires some geometric representation of the space in front of the signer. However, filling in the signing space (that is, placing objects in the space) in a meaningful manner is not trivial when all the subtle nuances of sign languages are taken into account.

The geometry of the signing space in our case is just a semi-circle in front of the signer, divided into six equally-sized slices. To place objects in the signing space, the input STAG syntax parse tree is traversed to find

[1]Zhao [Zhao et al. 2000] had limited success with constructing non-manual signs by using adverbial information in the source text.
[2]The emotive content of the text would influence aspects such as the speed of the signing movement, and the size of the movement, as part of the non-manual sign content.

the objects. In the current implementation, we only place people in the signing space[3]. For each new object found, a location is assigned, with either a *place* or *refer* flag. The *refer* flag indicates that the object already has a known position in the signing space.

The algorithm searches for noun phrases, nouns, or determiners that satisfy person agreement. A history list is kept to help assign areas to objects, and keep other information such as the so-called group identifier. The group identifier is used when conjunctive phrases occur. For example, the phrase *Harry and Andre* identifies two different positions in the signing space, but both positions need to be identified on the resolution of the pronoun *they*, referring to both *Harry* and *Andre*.

Next, pronouns are resolved, in order to determine their (pre-assigned) position in the signing space. First and second person singular pronouns are ignored, as these are special cases in SASL, which mostly influence the direction of the relevant verb (sign language verbs are directional, in the sense that the movement is from the subject to the object). Note that this is not true in the case of role playing, and that the SASL-MT system currently does not cater for role playing situations[4]

We currently clear the signing space data structure based on paragraph boundaries. This prevents the signing space from quickly becoming overfull, and although it can lead to somewhat stilted SASL, it does not seem to cause grammatical errors.

The pronoun resolution algorithm implemented in the SASL-MT system is a straightforward adaptation of known methods [Mitkov 2003]. The algorithm receives a STAG tree as input, and its first step is to tag the tree with morphosyntactic information. This includes person, number and gender classification for proper determiners, nouns and noun phrases (the morphosyntactic information is stored in a semantic relation database, which is an extension of WordNet [Fellbaum 1998]). The special cases for first and second person singular are treated separately, as well as the case where pronouns are used as determiners. During the traversal, various lists are updated. The first is a *history* list, which contains all the objects found. The second is a *recency* list, which contains only the objects from the current and previous two sentences. The recency list enforces focus, and speeds up the search for a possible antecedent when a pronoun is found. When a pronoun is encountered, a *candidate* list of possible antecedents is constructed based on the recency and/or history lists. For the candidate list, preferences are used to select an antecedent, by keeping a score for each candidate. Scores are calculated based on factors such as gender agreement, number agreement, and definiteness. Reflexive pronouns are resolved locally (within the current sentence).

Experimental analysis of the pronoun resolution algorithm shows performance comparable to existing algorithms (see Mitkov [Mitkov 2003] for an overview). However, the experiments also showed clearly that the signing space is highly dependent on noun phrase co-resolution [Ng and Cardie 2002] (that is, identifying synonomous noun phrases). Indeed, in our test cases, it was clear that pronoun resolution without noun phrase co-resolution results in object placement in the signing space that would be incomprehensible to a native signer. As an example of the extent to which noun phrase co-resolution influences the signing space, consider the following text:

Harry sat down on the couch. The cat watched the boy, timing his approach, and jumped onto his lap.

In our current implementation, *Harry* and *the boy* would each be (incorrectly) assigned different places in the signing space. The occurrence of co-referential noun phrases is so common that we consider it essential in the construction of a meaningful signing space. This topic requires computational linguistics discourse analysis techniques, and is not trivial to implement or solve successfully in all its generality. Our next step is to investigate the feasibility of incorporating this capability into the SASL-MT system.

4.2 Non-manual Sign Generation

In machine translation systems for sign languages, it is necessary to establish non-lexical information in order to generate non-manual signs. A similar problem exists in text-to-speech systems, where written text has to be rendered as speech with accent, pitch and intonation. These prosodic stress patterns are not well understood for all sign languages, and specifically not for SASL. However, it is clear that analogies between spoken language prosody and sign language stress patterns exist [Coulter 1997; Kingston 1999]. We therefore based our implementation of the emotive content analysis of SASL on known results for American Sign Language (ASL), and the next step in this process is a linguistic analysis to determine the SASL correctness of our emotive tagging.

Various methods for prosody analysis are known [Dutoit and Stylianou 2003; Hiyakumoto et al. 1997; Oehrle 1991; Steedman 2001], some based on syntactic information and others on semantic information. The latter gave

[3]The specific circumstances in SASL where other animate or inanimate objects need placement, require more careful linguistic analysis.

[4]Role playing occurs where the signer directly relays information contributable to another person. For example, when signing *John said:"I am going to town today."*, the signer physically moves to John's pre-assigned position in the signing space, and signs *I am going to town today.*

rise to the so-called concept-to-speech (CTS) systems, which is based on the premise that certain stress patterns in sentences give an indication of the meaning of the sentence [Hiyakumoto et al. 1997]. We therefore based our prosodic analysis on a CTS approach.

In the SASL system, we consider two levels of detail for sentences. The first is specific words in the sentence, and the second is the combination of words to form phrases. During prosodic generation, the output STAG trees are tagged with metadata, including so-called *expressiveness* information, *accent* for specific words, and lastly an indication of the phrase boundaries in sentences to indicate rhythm.

Expressiveness concerns firstly form – that is, the speed and range of the movement. Secondly, we define emotional classes, encompassing love, happiness, determination, seriousness, contentment, surprise, anxiety, fear, anger and hatred. These emotions are based on a psychological analysis of basic human emotions [van Zijl and Barker 2003]. We use WordNet [Fellbaum 1998] to tag the STAG trees with form and emotional class information.

In testing our system, we considered different types of documents ranging in emotive content. For example, academic texts showed neutral emotion (as expected), while newspaper articles showed higher and more fluctuating emotions. At this stage, our emotive tags are glossed-based on a word level. We recognize this as a shortcoming, and are working towards an improvement. In particular, we are currently merging the emotional class information with the output from the phrase boundaries class (see below), in order to apply the emotive tags to phrases instead of single words. However, we currently have no better results in this specific aspect than that of the original system of Zhao [Zhao et al. 2000]. Lastly, we identify the type of sentence – that is, assertion, question or exclamation, which generates differences in non-manual signs in sign languages. The sentence type is derived from punctuation.

Our analysis for accent of specific words entails an adaptation of the CTS model of Hiyakumoto, Prevost and Cassel [Hiyakumoto et al. 1997]. In essence, this model identifies new information, by keeping history lists and equivalence lists of nouns, verbs, adjectives and adverbs. Words are then marked as new words if they are not in the history list, and are not equivalent to any of the words in the history list. Otherwise, the word is considered to be inferable. The assumption is that new information is stressed more than other information. In addition, the text is subsequently scanned for contrasts, as contrasts also typically elicit more stress. Note that these stress points correspond to Coulter's analysis of emphatic stress in ASL [Coulter 1997]. Our experiments showed that the implementation of the stress component removes the need to store detailed variations of signs in the lexicon, and eases the requirements on the notation that generates the avatar movement instructions from the SASL gloss output.

We summarize the accent algorithms in Figure 3.

For each word w
  (1) If w is a noun, verb, adjective or adverb, and $w \notin \mathit{history}(w)$ and $w \notin \mathit{equiv}(x)$, for any $x \in \mathit{history}()$:
    (a) tag w as a focused item
    (b) add w to *history()*
    (c) create *equiv(w)*
  (2) If w is a noun, verb, adjective or adverb, and $w \in \mathrm{equiv}(x)$ tag w as inferable.

For each word u:
  (1) For each word v in the history list, from most to least recent:
    (a) For each word in $\{w : w \in \mathit{contrast}(v)\}$:
      i. if $u = w$ the mark u as contrastive focus; end the search
      ii. else add u to the history list;
        generate and store $\{w : w \in \mathit{contrast}(u)\}$

Figure 3. Accent algorithms from Hiyakumoto, Prevost and Cassell

The last part of our emotive content analyser finds phrase boundaries in sentences, mostly based on punctuation and conjunctions. We differentiate amongst simple co-ordinate conjunction, correlative co-ordinate conjunction and subordinate conjunction. A simple co-ordinate conjuncton connects two parts with the same importance (for example, *Harry and Andre went to town*). A correlative co-ordinate conjunction is similar to the case of a simple co-ordinate conjunction, but here a pair of words (*either-or*, *both-and*) connects two parts with the same importance. Lastly, if one phrase is subordinate to another, the conjunction is a subordinate conjunction (for example, *Andre went to bed after he got home*). Boundaries are then chosen after punctuation occurs, after the first part of a correlative co-ordinate conjunction, and before simple co-ordinate conjunctions. As was the case for the stress analysis, we found that knowledge about phrases and sub-phrases in a sentence eased the intricacies involved in the generation phase of the system.

In summary, our experiments showed that a simple STAG-based system can be implemented for a sign language machine translation system, and the main disadvantage of the glossed tree output can be counter-acted by more detailed analysis of the source and accompanying tagging of the glossed tree with metadata. From a software development viewpoint, this allows for a fast implementation, and for modularized implementation of different techniques for comparative purposes. From a linguistic point of view, this seems one of the most sensible approaches for linguistically under-investigated sign languages. We do acknowledge that the glossed tree with metadata may cause complications in the final generation of the output for the signing avatar, and are currently considering different solutions to this problem.

5. CONCLUSION

We reported on the South African Sign Language Machine Translation project. We showed that a STAG parser can be extended to include post-processing of the STAG output tree, in order to allow for meaningful signing space construction, and also additional emotive capabilities to include stress patterns which improve non-manual sign generation.

We are currently working on noun phrase co-reference resolution to further improve the signing space construction, and plan to extend our investigations into better production of non-manual signs.

ACKNOWLEDGMENTS

This research was supported by NRF grant #2053436. The financial assistance of the National Research Foundation towards this research is hereby acknowledged. Opinions expressed in this paper and conclusions arrived at are those of the authors and not necessarily to be attributed to the National Research Foundation.

REFERENCES

Bangham, J., Cox, S., Elliott, R., Glauert, J., and Marshall, I. 2000. Virtual signing: Capture, animation, storage and transmission – an overview of the VisiCAST project. In *IEE Seminar on Speech and Language Processing for Elderly and Disabled People*. http://www.visicast.co.uk.

Bungeroth, J. and Ney, H. 2004. Statistical sign language translation. In *Proceedings of the Workshop on Representation and Processing of Sign Languages, 4th International Conference on Language Resources and Evaluation (LREC2004)*. Lisbon, Portugal, 105–108.

Coulter, G. 1997. Emphatic stress in ASL. In *Theoretical Issues in Sign Language*. Springer-Verlag, 109–125.

Dutoit, T. and Stylianou, Y. 2003. Text-to-speech synthesis. In *Handbook of Computational Linguistics*, R. Mitkov, Ed. Oxford University Press.

Fellbaum, C. 1998. *WordNet: an electronic lexical database*. MIT Press, http://wordnet.princeton.edu.

Fotinea, S.-E., Efthimiou, E., Karpouzis, K., and Caridakis, C. 2005. Dynamic GSL synthesis to support access to e-content. In *Proceedings of the 3rd International Conference on Universal Access in Human-Computer Interaction*. Las Vegas, USA.

Fourie, J. 2006. The design of a generic signing avatar. Tech. rep., University of Stellenbosch, http://www.cs.sun.ac.za/˜ jfourie/.

Grieve-Smith, A. 2002. SignSynth: A sign language synthesis application using Web3D and Perl. In *Gesture and Sign Language in Human-Computer Interaction*, I. Wachsmuth and T. Sowa, Eds. Vol. 2298. Lecture Notes in Computer Science, Springer-Verlag, Berlin, 134–145.

Hiyakumoto, L., Prevost, S., and Cassell, J. 1997. Semantic and discourse information for text-to-speech intonation. In *Association for Computational Linguistics Workshop on Concept-to-Speech Generation*. Madrid, Spain.

Huenerfauth, M. 2003. A survey and critique of American Sign Language natural language generation and machine translation systems. Tech. Rep. MS-CIS-03-32, University of Pennsylvania, http://www.seas.upenn.edu/˜ matthewh/publications/huehnerfauth-2003-ms-cis-03-32-asl-nlg-mt-survey.pdf.

Huenerfauth, M. 2004. A multi-path architecture for machine translation of English text into American Sign Language animation. In *Proceedings of the Student Workshop at the Human Language Technology Conference/North American Chapter of the Association for Computational Linguistics (HLT-NAACL)*.

Joshi, A. and Schabes, Y. 1997. *Handbook of Formal Languages*. Vol. 3. Springer-Verlag.

Kingston, J. 1999. Extrapolating from spoken to signed prosody via laboratory phonology. *Language and Speech 42,* 2 (Apr.), 251–281.

Lillo-Martin, D. and Klima, E. 1990. Pointing out differences: Asl pronouns in syntactic theory. In *Theoretical Issues in Sign Language Research*, S. Fischer and P. Siple, Eds. The University of Chicago Press, 191–210.

Mitkov, R. 2003. *The Oxford Handbook of Computational Linguistics*. Oxford University Press, USA.

Morrissey, S. and Way, A. 2005. An example-based approach to translating sign language. In *Proceedings of the Workshop on Example-based Machine Translation (MT X-05)*. Phuket, Thailand, 109–116.

Neidle, C., Kegl, J., MacLaughlin, D., Bahan, B., and Lee, R. 2000. *The Syntax of American Sign Language: Functional Categories and Hierarchical Structure*. MIT Press, Cambridge, Massachusetts, USA.

Ng, V. and Cardie, C. 2002. Identifying anaphoric and nonanaphoric noun phrases to improve coreference resolution. In *Proceedings of the 19th International Conference on Computational Linguistics (COLING-02)*.

Oehrle, R. 1991. Prosodic constraints on dynamic grammatical analysis. In *Edinburgh Working Papers in Cognitive Science: Declarative Perspective on Phonology*, S. Bird, Ed. Vol. 7. 167–196.

Speers, D. L. 2001. Representation of American Sign Language for machine translation. Ph.D. thesis, Georgetown University, Washington, USA.

Steedman, M. 2001. Information-structural semantics for English intonation. In *LSA Summer Institute Workshop on Topic and Focus*. Santa Barbara, California, USA.

Suszczanska, N., Szmal, P., and Jaroslaw, F. 2002. Translating Polish text into sign language in the TGT system. In *Proceedings of the 20th IASTED International Multi-Conference on Applied Informatics*. Innsbruck, Austria, 282–287.

van Zijl, L. and Barker, D. 2003. A machine translation system for South African Sign Language. In *Proceedings of Afrigraph 2003, ACM Siggraph Conference Proceedings Series*. Cape Town, South Africa, 49–52.

VCom3D. 2004. *SignSmith Studio*. http://www.vcom3d.com.

Zhao, L., Kipper, K., Shuler, W., Vogler, C., Badler, N., and Palmer, M. 2000. A machine translation system from English to American Sign Language. Association for Machine Translation in the Americas, http://www.citeseer.ist.psu.edu/539314.html.