

Morphosyntactic Target Language Matching in Statistical Machine Translation

Simon Zwarts

Centre for Language Technology
Macquarie University
Sydney, Australia
szwarts@ics.mq.edu.au

Mark Dras

Centre for Language Technology
Macquarie University
Sydney, Australia
madras@ics.mq.edu.au

Abstract

While the intuition that morphological preprocessing of languages in various applications can be beneficial appears to be often true, especially in the case of morphologically richer languages, it is not always the case. Previous work on translation between Nordic languages, including the morphologically rich Finnish, found that morphological analysis and preprocessing actually led to a decrease in translation quality below that of the unprocessed baseline.

In this paper we investigate the proposition that the effect on translation quality depends on the *kind* of morphological preprocessing; and in particular that a specific kind of morphological preprocessing before translation could improve translation quality, a preprocessing that first transforms the source language to look more like the target, adapted from work on preprocessing via syntactically motivated reordering. We show that this is indeed the case in translating from Finnish, and that the results hold for different target languages and different morphological analysers.

1 Introduction

In many NLP applications, morphological preprocessing such as stemming is intuitively felt to be important, especially when dealing with morphologically rich languages. In fairly early work on morphological disambiguation of the agglutinative language Turkish, Hakkani-Tur et al. (2000) cite the note of Hankamer (1989) that a single Turkish root

can potentially have over a million inflected variants; such a proliferation of forms could exacerbate any data sparsity problems. Zwarts and Dras (2007a) showed, for languages that differ in morphological richness but are otherwise structurally similar, that, as expected, morphological preprocessing in an Statistical Machine Translation (SMT) environment confers greater benefits on the morphologically richer language than on the morphologically poorer one.

However, it is not always the case that morphological preprocessing of a morphologically rich language does provide a benefit. Again in an SMT context, Virpioja et al. (2007) preprocessed Finnish (a Uralic language, typologically between inflected and agglutinative, and consequently morphologically rich) for translation into Danish and Swedish (both Indo-European languages, quite different in many morphosyntactic respects from Finnish, including being morphologically less rich). In doing this, they found that translation quality was generally worse than baseline, and never better.

What this paper looks at is whether it is the way in which the morphology is preprocessed that is important. We draw on an idea from other work in machine translation first presented by Collins et al. (2005), where the source language is reordered so that its syntax is more like that of the target language, leading to an improvement in translation quality; here, we see whether that idea can apply at the level of morphology. In particular, where there are phenomena that are handled via morphology in one language and by syntax in another, we investi-

gate whether an adaptation of this reordering idea can improve translation quality.

In Section 2 we briefly review relevant literature. In Section 3 we describe the important characteristics of Finnish, followed by various models for morphological preprocessing, including our method for transforming the morphology of the Finnish source to more closely match the target language. In Section 4 we discuss the results, and in Section 5 we conclude.

2 Literature Review

2.1 Morphological Analysis

So many systems make some use of morphological preprocessing, particularly stemming, that we only point to a few specific instances here. In the context of parsing, morphological preprocessing has been shown to be necessary for the agglutinative languages Korean (Han and Sarkar, 2002) and Turkish (Eryiğit and Oflazer, 2006). In SMT, use of morphological preprocessing has been fairly ad hoc. One quite systematic comparison of morphological preprocessing parameters was carried out by Sadat and Habash (2006) for the language pair Arabic-English; their approach was just to search the whole space of parameter combinations, rather than looking at any characteristics of the pairing of the specific languages. One earlier work looking at Czech-English, Al-Onaizan et al. (1999), did carry out morphological analysis that looked at characteristics of the pair of languages, transforming some Czech morphemes into pseudo-prepositions.

The specific work we cited in the Introduction as having found that morphological preprocessing did not help, that of Virpioja et al. (2007), used Morfessor, a morphological analyser for Nordic Languages, as a preprocessor. They built translation systems between the Nordic languages Danish, Finnish and Swedish. They found that using Morfessor's morphemes as translation units instead of words degraded SMT quality¹ for all six language pairs and language directions possible (four statistically significantly). They tried unsupervised learning techniques to decide when to use mor-

¹Tables 6 and 7 in their paper

phological information and when to use normal phrases; this helped the system, but did not manage to beat a normal Phrase-Based Statistical Machine Translation (PSMT) system using words as the most fine-grained translation units.

2.2 Source-Side Reordering as Preprocessing

There are a number of different approaches to word reordering. It can be done based on rules over word alignment learnt statistically, for example Costa-Jussà and Fonollosa (2006). In this work an improvement in overall translation quality in a Spanish-English MT system was achieved by using statistical word classes and a word-based distortion model to reorder words in the source language. Reordering here is purely a statistical process and no syntactic knowledge of the language is used. Xia and McCord (2004) on the other hand use syntactic knowledge; they use pattern learning in their reordering system. In their work they parse and align sentences in the training phase and derive reordering patterns. From the English-French Canadian Hansard they extract 56,000 different transformations for translation. In the decoding phase they use these transformations on the source language. The main focus then is monotonic decoding. Both of these two cited works assume that explicitly matching the word order of the target language is the key.

The work that we draw on in this paper is that of Collins et al. (2005), which uses syntactically motivated rules based on clause restructuring. They define six hand-written rules for reordering source sentences in German for translation to English, which operate on the output of an automatic parser. The rules cover German phenomena such as the location of verbs, separable verb prefixes, negations and subjects, several of which represent long-distance relationships in German (e.g. where the inflected verb is in second position in the clause and its uninflected dependent verbs are at the end of the clause). The approach has also been applied successfully to Dutch-English (Zwarts and Dras, 2007b) and Chinese-English (Wang et al., 2007), among others.

Zwarts and Dras (2007b) found that there are at least two sources of the translation improvement: one is the explicit matching of target language syntax, while the other is the moving of heads and depen-

dants closer together to take advantage of the phrasal window of PSMT.

3 Morphological Target Language Matching

3.1 Languages

Taking the work of Virpioja et al. (2007) as a broad starting point, we use Finnish as our source language. As noted above, Finnish is part of the Finno-Ugric branch of the Uralic language family rather than part of the majority Indo-European family, and compared to languages like English is very rich in morphology. Finnish has fifteen noun cases: four grammatical cases, six locative cases, two essive cases and three marginal cases. It has three verb moods and on the imperative mood Finnish marks: 1st, 2nd or 3rd person, singular or plural, definite or indefinite and positive or negative. Verbs can have a morphological perfect, present or future tense. Finnish needs morphological agreement between words, for example noun and adjective agreement. All of this morphology is purely postfixing. The case system, which will be the focus in this paper, is described in more detail in Table 1.

For our main target language we use English, which is morphologically not at all a rich language. As a supplementary target language, to check results for Finnish-English, we use Dutch, which is morphologically quite similar to English, in terms of quantity and type of inflection.

3.2 Models

We hypothesise that an important reason for a gain in translation quality when using morphological preprocessing is the matching of the morphology and syntax of the target language. As in the work described in Section 2.2, where the source language was made to look more like the target language by applying grammatical rules, now we want to do this on a morphological level. As comparisons, we design several models which all differ in the fact that they have a different way of preprocessing the text.

To obtain morphological analyses we use Connexor²

²<http://www.connexor.com>

(Tapanainen and Järvinen, 1997) and Morfessor³ (Creutz et al., 2005) for Finnish. Connexor provides a per-word morphological analysis and it provides the stem of the tokens. Connexor indicates the Part-Of-Speech (POS) and depending on that other information. For example for a noun it gives the case, and whether it is singular or plural. It does not indicate morpheme boundaries; however, since it provides stems, boundaries are recoverable if the token in question has multiple stems. Connexor also provides other parse information, but since we are interested in morphology in this paper, we only use the morphological information seen in the previous example. Morfessor, on the other hand, only indicates morpheme boundaries and indicates per morpheme whether this morpheme is a suffix or a stem. Figure 1 shows an example of the first line in the Europarl Corpus.

We build models with different amounts of preprocessing, and different types of preprocessing, and investigate their effects: four models based on Connexor output, one straight PSMT baseline model, and two models based on Morfessor output. We first define the main model of interest, where the morphology of the source language is mapped to the morphology and syntax of the target language. Then as comparisons we look at baselines with no morphology, preprocessing with full morphology, and preprocessing with word stems only, both with compounds and without.

Model C1 - Noun Case matching This is the model with specific morphological matching to the target language. Finnish has many different noun cases, while the case system has almost completely disappeared in English and Dutch, as only the pronouns still exhibit some vestiges of a case system. Finnish however has many cases, even compared to other European languages which still have case systems (for example German, Greek etc.). A lot of cases in Finnish fall in the group of locative case, which is to indicate how the noun is located. In English and Dutch this is usually expressed with prepositions.

The preprocessing steps are then as follows:

³<http://www.cis.hut.fi/projects/morpho/>

Finnish cases				
Case	Suffix	English prep.	Sample	Translation
Grammatical				
nominatiivi (nominative)		-	talo	house
genetiivi (genitive)	-n	of	talon	of (a) house
akkusatiivi (accusative)	- or -n	-	talo or talon	house
partitiivi (partitive)	-(t)a	-	taloa	house (as an object)
Locative (internal)				
inessiivi (inessive)	-ssa	in	talossa	in (a) house
elatiivi (elative)	-sta	from (inside)	talosta	from (a) house
illatiivi (illative)	-an, -en, etc.	into	taloon	into (a) house
Locative (external)				
adessiivi (adessive)	-lla	at, on	talolla	at (a) house
ablatiivi (ablative)	-lta	from	talolta	from (a) house
allatiivi (allative)	-lle	to	talolle	to (a) house
Essive				
essiivi (essive)	-na	as	talona	as a house
(eksessiivi; dialectal) (exessive)	-nta	from being	talonta	from being a house
translatiivi (translative)	-ksi	to (role of)	taloksi	to a house
Marginal				
instruktiivi (instructive)	-n	with (the aid of)	taloin	with the houses
abessiivi (abessive)	-tta	without	talotta	without (a) house
komitatiivi (comitative)	-ne-	together (with)	taloineni	with my house(s)

Table 1: Finnish Case system, taken from Wikipedia. (http://en.wikipedia.org/wiki/Finnish_grammar)

Finnish: istuntokauden uudelleenavaaminen
Connexor: Lemma='istunto kausi' Morpho='N SG GEN'
Lemma='uudelleen avata' Morpho='V ACT INF4 NOM'
Morfessor: istu/STM n/SUF tokauden/STM
uude/STM lle/SUF en/SUF avaam/STM in/SUF en/SUF
English: resumption of the session

Figure 1: Morphological Output from Connexor and Morfessor

1. For every token in the sentence we retrieve the POS and case information.
2. For every token marked as N (noun) we replace this token by its stem.
3. For every token marked as N we replace all tokens directly preceding the noun by its stem if this token shares the same case marker as the noun.
4. We insert before each initial token in Step 3 a token marking the case.

To explain: To make the Finnish more like the target language, we need to remove case morphology from the noun itself as the target language does not have cases morphologically marked. Therefore in Step 2 we replace the token by its stem. Because Finnish is a language with agreement, if a token has a noun case, other tokens in agreement with this noun (for example adjectives) need to undergo the same case-morphology removal step; this is Step 3. Finally in the last step we need to insert a token, providing the information of case as a separate token. Usually, this will result in being translated as a preposition. This information is represented in a token before the actual noun, matching the English word ordering where the preposition is positioned before the noun. We chose to introduce this token as -T-<case> so that for a genitive case, for example, we introduce the token -T-GEN. Investigation shows that in this way there is no clash between these tokens and Finnish vocabulary. An example is provided in Table 2, showing the original Finnish, the preprocessed variant and an English reference.

We note that in general we would not expect this to move related morphemes into or out of the phrasal window used in PSMT, thus we would not expect to gain anything as a consequence of this source of translation improvement described in Zwarts and Dras (2007b). However, the other source of improvement, the explicit matching of the target language, is still possible.

Model C2 - Full morphological preprocessing

Part of our hypothesis is that it is not full morphological preprocessing which will bring the most improvement of translation quality, but rather the

morphological matching of the target language; so we construct a full preprocessing model for comparison. For our full morphological preprocessing model, we replace every token by its stem and insert after every token a token indicating morphology. For example the Finnish word *istuntokauden* ('resumption', from Figure 1) with Lemma='istunto kausi' and Morpho='N SG GEN' will result in the three tokens: *istunto kausi N-SG-GEN*. This is the same basic approach as that of Virpioja et al. (2007).

Model C3 - Stem Only - Compounds

In this model we investigate the effects of having no morphology in Finnish at all. Because Finnish is morphologically much richer than English, there are many more individual tokens in Finnish with the same base stem, while in English translation often only one token is there to express this concept. To make the Finnish language less sparse, and make it possible to do a more reliable alignment estimation from which we calculate probabilities for Finnish to English, we only use stems. For example, the Finnish word *vaunu* ('car') can take many different surface realisations depending on morphology (*vaunut, vaunuilla, vaunuja, vaunujen, vaunujensa* etc.). Without morphological preprocessing these are entirely different tokens. We map all these tokens onto the same base stem *vaunu*, so the probability estimation for *car* and *vaunu* should be more accurate. In this model we leave compounds (for example compound nouns) as one token. The word *istuntokauden* results in the token *istuntokausi*.

Model C4 - Stem Only - Compounds separated

This model aims to investigate the effect of compound splitting with stemming. This model is similar to the previous model, Model C3, except now we do split compounds. The word *istuntokauden* results in the two tokens *istunto* and *kausii*.

Koehn (2003) shows that compound splitting in Machine Translation (MT) is superior to leaving compounds intact. However, for completeness we include both models.

Model B - Baseline, no morphological preprocessing

Here we just use the straight PSMT output, taking the original source as the input.

Original Finnish	istuntokauden uudelleenavaaminen
Case Isolated Finnish	-T-GEN istuntokausi uudelleenavaaminen
English Reference	resumption of the session
Original Finnish	äänestimme -T-ELA asia -T-GEN keskustelu jälkeen
Case Isolated Finnish	äänestimme asiasta keskustelun jälkeen
English Reference	we then put it to a vote
Original Finnish	-T-INE ensimmäinen käsittely tekemämme -T-NOM tarkistusehdotus
Case Isolated Finnish	on otettu -T-POSS:SG1 mieli kiitettävästi huomioon
English Reference	our amendments from the first reading have i believe been taken into account very satisfactorily

Table 2: Noun Case Matching

Model M1 - Noun Case matching We also experimented with a different type of morphological information, to see if our findings still held up. Morfessor, as discussed, provides different information from Connexor. The baseline model (Model B) is the same for the Connexor models and the Morfessor models since in both cases it does not have any preprocessing. Other models are designed to target different morphological analysis.

This model is targeted to be the Morfessor equivalent of Model C1 for Connexor. The morphological preprocessing is targeted to match the morphology of the target language. As in the Connexor model we target the elaborate locative case system in the Finnish language. We identify which morphological suffixes are responsible for indicating a noun case. These morphemes⁴ are the only morphemes treated. By default we leave a token unchanged by the preprocessing algorithm. However if one of these morphemes appears in the token, we delete this morpheme from the token (with all its doubles, because we assume these are there for agreement reasons) and put that morpheme in front of the token we are preprocessing. On the suffix we leave the suffix identifier attached so these tokens do not clash with other words in the Finnish vocabulary.

⁴The complete list is: n/SUF, ssa/SUF, sta/SUF, en/SUF, an/SUF, lla/SUF, lta/SUF, lle/SUF, na/SUF, nta/SUF, ksi/SUF, tta/SUF, ne/SUF. See Table 1 for how these suffixes are acquired.

For example:

istuntokauden uudelleenavaaminen

turns into:

n/SUF istutokauden lle/SUF en/SUF uudeavaamin

Model M3 - Stemmed This model is the Morfessor equivalent of Connexor model C3. We directly input the output from Morfessor into the PSMT system. Morfessor identifies the different morphemes in the different tokens. For each morpheme Morfessor identifies whether it is a suffix or a stem morpheme. We leave this identifier attached to the morpheme.

For example:

istuntokauden uudelleenavaaminen

turns into:

istu/STM n/SUF tokauden/STM uude/STM lle/SUF en/SUF avaam/STM in/SUF en/SUF

3.3 Data

We use as our data the Europarl corpus, language pairs Finnish-English and Finnish-Dutch. In both

cases the baseline is a normal sentence-aligned corpus, which is trained with GIZA++. After this training, the phrases are derived via the standard method described in the Pharaoh manual (Koehn, 2004). The language model is trained on Europarl text only and is an n -gram language model. We use Pharaoh to translate a test set of 10k sentences while we have 774k sentences in the training corpus.

The same approach is used for the various other models, with the text passed to GIZA++ and Pharaoh being the preprocessed version of the source. As Virpioja et al. (2007), we do not perform Minimum Error Rate Training (MERT) (Och, 2003) or optimise variables.

4 Results and Discussion

The results, expressed in BLEU points, are shown in Table 3 for Finnish to English. For the translations to Dutch with the Connexor models, results are shown in Table 4, which follow the trends of the English models.

First of all, we note that the the poor performance of the full morphological analysis using Morfessor, Model M3, relative to the baseline, confirms the findings of Virpioja et al. (2007).

However, as can be seen from the tables, it is possible to outperform the baseline, but only where the specially targetted morphological analysis has been performed. This is true both where the morphological analysis has been done using Connexor (Model C1) and where it has been done using Morfessor (Model M1). All of the other preprocessings perform worse. This result also carries over to Dutch.

We had a closer look at Model B and Model C1, the baseline and the model with special targetted morphological analysis. Table 5 shows how the BLEU score is calculated for the baseline and the model where morphology has improved translation quality. We can see that in fact the token-for-token translation quality is not that different. The brevity penalty is what is responsible for most of the gain, which suggests that the original baseline is undergenerating words.

It is possible that the large brevity penalty is

Finnish	to English
B. Original (baseline)	0.1273
Connexor	
C1. Noun Case Isolated	0.1443
C2. Stemmed, separate morph.	0.1082
C3. Stemmed	0.1060
C4. Stemmed and de-compounded	0.1055
Morfessor	
M1. Noun Case Isolated	0.1312
M3. Stemmed	0.1073

Table 3: BLEU scores for different morphological analysis

Finnish	to Dutch
B. Original (baseline)	0.132
Connexor	
C1. Noun Case Isolated	0.138
C2. Stemmed, separate morph.	0.089
C3. Stemmed	0.101
C4. Stemmed and de-compounded	0.102

Table 4: BLEU scores for different morphological analysis

	Model B	Model C1
1-gram	0.5167	0.4945
2-gram	0.2276	0.2135
3-gram	0.1114	0.1029
4-gram	0.0564	0.0500
Brevity Penalty	0.7717	0.9453
BLEU	0.1273	0.1443

Table 5: Decomposition of BLEU scores for Model B and Model C1

a consequence of not carrying out MERT.⁵ Och (2003) notes that using standard optimisation criteria (rather than optimising for BLEU score) can “prefer shorter translations which are heavily penalized by the BLEU and NIST brevity penalty”. However, all of the models are comparable, being implemented without MERT, so they are all affected in the same way.

In attempting to understand it, we note that what the morphological preprocessing has achieved is to generate more correct words. We can conclude that in the normal (baseline) translation the problem is not so much to find the right translation for the words: even when morphology is attached the decoder still manages to find right translations, but fails to decode that part of information with morphology holds by itself.

Even PSMT is based on underlying token-to-token translation tools. Although the phrases can handle translation of multiple tokens to multiple tokens, the most common ways to derive these phrases are still based on single token to token probability estimation. GIZA++, which is used for phrase extraction, needs fertility to translate one token to multiple tokens in the target language. In our Finnish corpus, every sentence has on average 14.0 tokens per sentences against 20.3 for English. With the model C1’s preprocessed text (see Table 2) we have on average 18.4 tokens per sentence, which is much closer to the target language. So an explanation is that this closer matching of language sizes has led to this difference.

5 Conclusion

We have investigated whether the idea of reordering as preprocessing carries over to the interface between morphology and syntax. We have compared this to a series of models covering a range of morphological preprocessing, and shown that among these the reordering model, restructuring the morphology of the source language to look more like the target language, works the best, and is in fact the only model to outperform the baseline. In particular this model outperforms the model based on previous work where it is full morphological information

⁵Thanks to an anonymous reviewer for this suggestion.

that is added in the preprocessing step. These results hold across the target languages English and Dutch, and for two different morphological analysers.

In terms of future work, Section 4 raised the issue of MERT; the next step will involve looking at the effect, if any, of this. With respect to directions subsequent to this, the work in this paper targeted only one particular phenomenon on the interface of morphology and syntax, the correspondence of locative cases in Finnish to prepositions in English and Dutch. There are many potential other phenomena that could also be captured, and possibly combined with a broader class of reorderings, an area for future investigation. In addition, our work did not use any of the unsupervised learning methods that Virpioja et al. (2007) used to decide when to use morphological analysis, which for them improved results (although only up to the baseline); there would be scope to apply a similar idea here as well, perhaps adapting the work of Zwarts and Dras (2008) in identifying situations where syntactic reordering has led to a worse translation in the case of an individual sentence.

6 Acknowledgements

We would like to thank Tuomo Kakkonen, for helping us with Finnish advice and for supplying us with an automated morphological markup of our training text.

References

- Yaser Al-Onaizan, Jan Curin, Michael Jahr, Kevin Knight, John Lafferty, Dan Melamed, Franz-Josef Och, David Purdy, Noah A. Smith, and David Yarowsky. 1999. Statistical machine translation, final report, JHU Workshop, Dec.
- Michael Collins, Philipp Koehn, and Ivona Kucerova. 2005. Clause restructuring for statistical machine translation. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 531–540, Ann Arbor, Michigan, June.
- Marta R. Costa-Jussà and José A. R. Fonollosa. 2006. Statistical Machine Reordering. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, pages 70–76.
- Mathias Creutz, Krista Lagus, Krister Lindén, and Sami Virpioja. 2005. Morfessor and Hutmegs: Unsuper-

- vised Morpheme Segmentation for Highly-Inflecting and Compounding Languages. In *Proceedings of the Second Baltic Conference on Human Language Technologies*, pages 107–112.
- Gülsen Eryiğit and Kemal Oflazer. 2006. Statistical Dependency Parsing for Turkish. In *11th Conference of the European Chapter of the Association for Computational Linguistics: EACL*.
- Diiek Z. Hakkani-Tur, Kemal Oflazer, and Gokhan Tur. 2000. Statistical Morphological Disambiguation for Agglutinative Languages. In *Proceedings of the 18th International Conference on Computational Linguistics (COLING2000)*, pages 285–291.
- Chung-hye Han and Anoop Sarkar. 2002. Statistical Morphological Tagging and Parsing of Korean with an LTAG Grammar. In *Proceedings of the 6th International Workshop on Tree Adjoining Grammars and Related Formalisms*, pages 48–56.
- Jorge Hankamer. 1989. Morphological Parsing and the Lexicon. In *Lexical Representation and Process*. The MIT Press.
- Philipp Koehn. 2003. *Noun Phrase Translation*. Ph.D. thesis, University of Southern California.
- Philipp Koehn. 2004. Pharaoh: a Beam Search Decoder for Phrase-Based Statistical Machine Translation Models. In *Proceedings of Association for Machine Translation in the Americas (AMTA)*, pages 115–124.
- Franz Josef Och. 2003. Minimum error rate training in statistical machine translation. In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*, pages 160–167, Sapporo, Japan, July. Association for Computational Linguistics.
- Fatiha Sadat and Nizar Habash. 2006. Combination of Arabic preprocessing schemes for statistical machine translation. In *Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the ACL*, pages 1–8. Association for Computational Linguistics.
- Pasi Tapanainen and Timo Järvinen. 1997. A non-projective dependency parser. In *Proceedings of the 5th Conference on Applied Natural Language Processing*, pages 64–71.
- Sami Virpioja, Jaako J. Väyrynen, Mathias Creutz, and Markus Sadeniemi. 2007. Morphology-Aware Statistical Machine Translation Based on Morphs Induced in an Unsupervised Manner. In *Proceedings of MT Summit XIII*, pages 491–498.
- Chao Wang, Michael Collins, and Philipp Koehn. 2007. Chinese Syntactic Reordering for Statistical Machine Translation. In *Proceedings of Empirical Methods on Natural Language Processing (EMNLP)*, pages 737–745.
- Fei Xia and Michael McCord. 2004. Improving a statistical MT system with automatically learned rewrite patterns. In *Proceedings of the 20th International Conference on Computational Linguistics (Coling)*, pages 508–514.
- Simon Zwarts and Mark Dras. 2007a. Statistical Machine Translation of Australian Aboriginal Languages: Morphological Analysis with Languages of Differing Morphological Richness. In *Proceedings of the Australasian Language Technology Workshop*.
- Simon Zwarts and Mark Dras. 2007b. Syntax-Based Word Reordering in Phrase-Based Statistical Machine Translation: Why Does it Work? In *Proceedings of MT Summit XI*, pages 559–566.
- Simon Zwarts and Mark Dras. 2008. Choosing the Right Translation: A Syntactically Informed Approach. In *Proceedings of the 22nd International Conference on Computational Linguistics (Coling)*.