

Towards a humour switch for machine translation

Manfred Stede; Dan Tidhar

Technische Universität Berlin; Projektgruppe KIT; Franklinstr. 28/29; 10587 Berlin; Germany
stede|dan@cs.tu-berlin.de

Abstract

We argue that machine translation provides a useful testbed for developing computational theories of humour. Thus focusing on the notion of ‘translation humour’, where the humorous note arises from two sentences in different languages, we suggest a system architecture that extends the standard transfer-based machine translation paradigm with a humour processing component, which analyzes the full syntactic chart of the source sentence. An implementation within the existing MT system ‘KIT-FAST’ is under way. We illustrate our approach with an example.

1 Introduction

Countless students of computational linguistics have been told the legendary output of a machine translation (MT) system from the 1950s, which allegedly was confronted with the proverb *The spirit is willing, but the flesh is weak*, translated it into Russian and then back into English, which resulted in *The vodka is good, but the steak is lousy*. Regardless whether the legend is true (insiders seem to agree that it is not), the very fact that the example keeps coming up in folk discussions of machine translation indicates that it is an instance of successful humour. On the other hand, as far as we can tell today, the system (if any) that produced the translation probably did not realize that its output was considered funny. It presumably did not have the choice to produce either a serious or a funny translation — instead, it simply did its best to do a good job, and failed. The idea of this particular *choice*, however, is of interest to us: Can we equip an MT system with a parameter such that it can (try to) translate its input in a funny way rather than seriously, on demand?

For computational humour research, we think that MT provides a useful setting, since both input and output are in natural language (as opposed to language generation from conceptual representations) and therefore the performance of a system — in general as well as with respect to humour — can be judged easily, and shortcomings be identified. Thus, a working MT system can provide a testbed for experimenting with theories of humorous translation. While the cases of translation humour investigated here are certainly not instances of humour of the highest quality, they operate on a level where developing a computational account seems a feasible goal.

In this paper, we first present the results of a small case study on (non-)funny machine translation results. We look into the reasons why a translation is funny or not, and characterize what we mean by translation humour. Then, we describe a system architecture that integrates a ‘humour switch’ into a working MT system, and we dis-

cuss an example of automatically producing a humorous translation.

2 A small case study

In our experiment, we took actual MT output from the early stages of the German ‘Verbmobil’ project (Wahlster, 1993), which deals with translating spontaneous speech between English, Japanese, and German. Focusing on translations from German into English, we collected a corpus of 33 translations, all of which were obviously wrong. Ten subjects were asked to judge whether the cases were funny. Apart from the well-known fact that people occasionally disagree on what they find funny, several other factors, such as diverse levels of English fluency among the subjects, caused some disagreement among the obtained replies. Nevertheless, out of the 33 cases under discussion, a list of 6 best and 6 worst cases could be prepared with an almost unanimous agreement. The two lists can be used as first approximations to what can be considered funny machine translation, in contrast to the results of a merely false translation. Appreciating the examples presupposes some knowledge of German (and there is no point in giving ‘literal’ English back translations), but we will provide some explanations in the next section.

The six examples rated most funny:

- (1) Das paßt mir überhaupt nicht in den Kram /
That doesn’t generally suit me in the stuff
- (2) Da ist was dran /
Something is then at that
- (3) auf Wiedersehen /
On reunions
- (4) Wir packen es /
We pack it
- (5) Wie wäre es denn um drei Viertel elf /
How about at 11 o’clock in 3 quarters?

- (6) Ich weiß nicht so recht /
I so conveniently don't know

The six examples rated least funny:

- (1') Da bin ich nun völlig unabhkömmlich /
I can't absolutely then get away
- (2') Ich mache mal einen anderen Vorschlag /
I do another suggestion
- (3') Das finde ich ja wunderbar /
I think that perfectly
- (4') So kommen wir dann zu der Mitarbeiterbesprechung /
We come like that to the team meeting
- (5') noch /
Neither nor
- (6') Dann würde eigentlich nur noch der neunundzwanzigste
in Frage kommen /
The twenty-ninth would just only actually be possible

3 What makes a translation funny?

In the course of developing computational translation humour, the first task is to sort out the simply annoying examples of bad MT, and the second task is to distinguish the instances of simply laughing at the stupidity of the machine from the instances of “good” humour. In other words, one should reckon with a difference between ‘translation humour’ on the one hand and ‘machine translation humour’ on the other. The latter can be characterized as the man–machine variant of the hostility/superiority type of humour, which (Raskin, 1985, p. 36) describes as arising “by comparison with the infirmity of others.”

In this section, we first analyze the MT examples given above, then turn to instances of translation humour with no machines involved, and then sum up with a characterization of what we regard as ‘translation humour’.

3.1 Analysis of MT examples

When we asked our subjects about their reasons for making particular judgements, we were told, for example, that sentence (3) is funny because the machine is “not even able to process such a simple phrasal expression.” (*auf Wiedersehen* is the German *good bye*.)

This observation points to the general scheme underlying the majority of “good” examples in our test set: a phrasal expression in the source utterance (SU), with idiomatic meaning, is translated literally and thus incorrectly in the target language. This is not a sufficient condition for fun, though; the wrong target utterance (TU) should at least make some sense in the target language. Consider (5), where the German *dreiviertel elf* is a (dialect) way of saying ‘10.45am’, while the translation seems to suggest that clocktimes can be chopped in pieces.

In (2), on the other hand, the humour seems to stem from the fact that the translation is syntactically well-formed

yet clearly lacks meaning, due to the proliferation of pronouns. The German phrase *da ist (et)was dran* in actuality means *there is some truth in what you say*.

One other indication of humour, according to our subjects, is an inflation of words, or over-complicated constructions in the target language, as in:

können wir ja erst mal so festhalten

We can only certainly make a note of it like that at a time.

Here, the various German particles, most of which act as mere fillers, have been given an English “translation” and somehow added to the SL sentence, which resulted in an utterance not entirely nonsensical but highly awkward. It seems appropriate to assign such cases to the class of ‘machine translation humour.’

To further evaluate the results, we devised the following set of parameters:

- (1) Is the source sentence idiomatic?
- (2) Is the target sentence grammatical?
- (3) Does the target sentence make sense?
- (4) Is the target sentence idiomatic?
- (5) Does the target sentence successfully convey a different meaning?
- (6) Does the target sentence clearly belong to a different domain?
- (7) Is the false translation based on semantic/lexical ambiguity?
- (8) Is the false translation based on syntactic ambiguity?

When answering these questions for the most humorous and the least humorous cases, respectively, and then computing the difference between the resulting vectors, features 2, 5, 6, and 7 turn out to be the most discriminating ones. Thus: in humorous translations the TU tends to be grammatical, convey a different meaning than the SU, and belong to a different domain; the wrong translation tends to be based on semantic/lexical ambiguity. It should be kept in mind, though, that these results are based solely on the sample corpus of bad ‘Verbmobil’ translations — with other data, one might find different factors. Notice, for instance, the difference to the legendary example given at the very beginning, which is probably successful because it plays not only on idiomaticity but also on cultural stereotypes that many people associate with speakers of the target language.

3.2 Non-machine translation humour

We now discuss a few examples of translation-based jokes produced by human beings, i.e. with no connection to MT. There are quite a few popular examples that keep coming up in joke telling, and they can be found in several collections on the internet; the following instances were all taken from http://www.geocities.com/WestHollywood/Village/2744/humour_trans.html.

In a Paris hotel elevator:

Please leave your values at the front desk.

Playing on the phonological similarity between ‘values’ and ‘valuables’, this example seems to work almost independent of its translation context; i.e., the knowledge of the French origin does not seem really crucial. In French, ‘valuables’ are *objets de valeur*, ‘values’ are *les valeurs*¹, hence the ambiguity is the same as in English. If the example is authentic, then maybe the translator knew that English has a single word for the French phrase *objets de valeur*, but then picked the wrong word.

In a Bucharest hotel lobby:
The lift is being fixed for the next day. During that time we regret that you will be unbearable.

If this case were a result of mis-disambiguation, there should have been a translation fallacy between the Romanian words for ‘to carry’ and ‘to bear’. There is, however, no close similarity between ‘bearable’ (*supportabil*) and ‘carry-able’ (something that can be *carat* or *mutat* or *transportat*)². Hence, no specific Romanian–English problem can be discerned here. (Interestingly, the example would work slightly better in German, where *tragen* corresponds to ‘to carry’, and *er-tragen*, derived from the same root, is the translation of ‘to bear’. Only, *tragen* is not very common in the ‘escalator’ context, where one would prefer *befördern*.)

In the lobby of a Moscow Hotel across from a Russian Orthodox monastery:
You are welcome to visit the cemetery where famous Russian and Soviet composers, artists, and writers are buried daily except Thursday.

In English, relative clauses need not be marked by a comma at beginning and end, and hence the sentence gets ambiguous. In Russian, on the other hand, the commas are necessary, so that the intended reading would be clear.³ Therefore, no specific translation ambiguity can be identified here — which would be the case if, for instance, the *daily...* phrase were morphologically marked to attach to *visit* rather than to *buried*.

On the menu of a Swiss restaurant:
Our wines leave you with nothing to hope for

The closest German origin that we can imagine is *Unsere Weine lassen nichts zu wünschen übrig*, where the idiomatic phrase *nichts zu wünschen übrig lassen* denotes ‘being perfect’ and in a word-by-word translation reads ‘leave nothing to desire’. Now, ‘to desire’ and ‘to hope for’ are indeed quite similar, and thus the overall mistranslation can indeed be traced back to an infelicitous treatment of an idiom — a rich source of wrong translations, as pointed out in the previous section.

¹Uta Weis, personal communication

²Daniel Marcu, personal communication

³Leo Wanner, personal communication

In a Zurich hotel:
Because of the impropriety of entertaining guests of the opposite sex in your room, we suggest that the lobby be used for such purposes.

If this example is regarded funny, then it can only be due to the ambiguity involving the display of “very traditional” moral attitudes; it has nothing to do with translation.

In a Tokyo bar:
Special cocktails for the ladies with nuts

In English, the word order strongly suggests an incorrect PP attachment, whereas in Japanese, word order is much less constrained. Instead, the activation of world knowledge or ‘common expectations’ plays a very important role, so that utterances are typically much more underspecified than in English.⁴ Furthermore, Japanese lacks determiners indicating (in-) definiteness; when translating into English, they have to be filled in. In our example, the newly introduced determiner adds to perception that *the ladies with nuts* forms one coherent phrase. One could thus argue that the example stems from the translation process being forced to add information to the TU which is only implicit in the SU, which here went wrong and produced translation humour.

In summary, closer inspection reveals that some of the popular “translation jokes” are in fact no instances of translation humour in a narrow sense — the sentences do not bear a specific relationship to their source language. The same holds, in general, for the infamous cases of instruction manuals that got badly translated from Asian languages; if they are funny at all, then only because the translators did not know their target language very well. The manuals typically do not contain jokes arising from false yet interesting disambiguations.

3.3 The genre of translation humour

We see an utterance as an instance of *translation humour*, when it is humorous *only* in connection with a source language utterance from which it got translated. In some cases, though, the mere information that the utterance was translated from another language can be a sufficient setting for translation humour. Such underspecified occurrences become more successful if the source language is specified, thereby leading the reader to an attempt at reconstructing the source utterance or some of its fragments. Our analysis in the next two sections assumes a complete translation humor setting, and it is assumed that partial settings (the kind of which appear in section 3.2) can be easily tailored to fit the same analysis by applying minor changes to the scenario (i.e. rather than explicitly reading the SU as the first step, the reader first encounters the TU which implicitly entails a reconstructed SU).

As for distinguishing bad translations (*machine translation humour*) from funny translations (*translation hu-*

⁴Michiko Miyano and Dietrich Bollmann, personal communication

mour), we should first note that both are based on ‘incorrect’ disambiguation(s) in the SU. The results of our small scale case study indicate that the grammatical correctness of the TU, as well as the feature of semantic sensibility is important for making the distinction. The syntactic requirement could be somewhat softened due to correction mechanisms that joke recipients seem to automatically apply to the perceived text, disregarding minor inaccuracies when the joke’s essence is detected. In order for that to be possible, the TU should be *approximately* grammatical. TUs that are utter garbage yield machine translation humor, if anything.

How, then, does translation humour work? A standard analysis of the linear structure of jokes (cf. (Attardo, 1994, p. 95)) states: First, a sense S1 is established, until the recipient encounters an element that causes the passage from the first sense to a second sense S2, which is antagonistic to the first one. The passage from S1 to S2 must be “unexpected” on the one hand, and “immediate” on the other.

This maps to translation humour as follows: The SU introduces a sense S1 based on individual lexeme senses for the content words W1 ... Wn. The reader expects to find the same S1 in the TU, but encounters an incompatible sense S2; this derives from unexpected translations (and hence senses) of some of the W1 ... Wn, thereby demonstrating ambiguities in the SU that the reader had no reason to notice earlier. For the humour to arise, a suitable incompatibility between the senses is needed.

As is indeed reflected in the discriminating role which feature (6) seems to play in our case study (section 3.1), we follow Raskin (1985) in assuming a criterion of script opposition to hold for humorous translations; thus, if the two domains of SU and TU are opposed to one another, more successful humour is to be expected. Of the various opposition types that are described by (Raskin, 1985, p.127) the one which immediately lends itself to the translation humour scenario is the *normal/abnormal* dichotomy. The initial S1 sense of the SU is rendered by scripts that are invoked by the more standard readings of W1 ... Wn in their context, thus representing the “normal” reading of the SU. The alternative sense S2 is based on scripts that are invoked by less standard readings of W1 ... Wn, thus representing an “abnormal” interpretation of the SU.

The *normal/abnormal* script opposition could be seen, in that sense, as a conceptual framework underlying the complete set of translation humour instances. It seems reasonable to assume, however, that other script oppositions that co-exist in particular instances can act as humour-intensifiers. Maximizing the opposition between S1 and S2 seems a reasonable initial rule of thumb for guiding the search for suitable partial translations in the procedure we will introduce below. Developing a more specific account of ‘opposition’ is a crucial aspect of our ongoing work.

4 Humourizing machine translation

We now turn to the question of providing MT systems with the ability to produce humorous translations. Of the MT paradigms in use today, the more recent statistics-driven or example-based systems do not seem suitable candidates for installing a humour parameter; even if jokes could be produced by mere statistical means, one would not be able to obtain a theory of how humour works. For in-depth investigations, an account of creativity or humour needs to have explicit knowledge of ambiguity and its mechanisms, and thus of meaning: fine-grained lexical knowledge as well as world knowledge, including information about cultural stereotypes etc., is required.⁵ Hence, we assume a framework of knowledge-based MT (see, e.g., Nirenburg et al. (1992)), which might either be interlingua- or transfer-based. In the following, we assume a transfer architecture.

In a nutshell, a transfer-based MT system first activates a parser that produces a syntactic and (possibly in a subsequent step) semantic representation of the source sentence; then the transfer module maps the representation into a corresponding syntactic/semantic representation of the target language; finally, a generator transforms that representation into a target language utterance. This ‘standard’ sequence of the translation process is depicted in the upper part of Figure 1.

Since the humour treatment relies on ambiguity, a necessary first step in adapting such a system for our purposes is spelling out the different readings of the SU (in contrast to a standard parser that is content with the one ‘correct’ reading). This ‘correct’ reading and its semantic interpretation need to be determined, too, in order to enable semantic comparisons with possible TL expressions. In addition, though, the whole range of alternative SU analyses is required.

A promising implementation strategy is using a chart parser and having it produce a full chart with all partial analyses, including all different (target language) readings of individual lexemes, and all decompositions of idioms, i.e., separate entries for each word that (for a correct translation) participates in a phrasal expression. Syntactic ambiguities such as PP attachment, on the other hand, might be resolved right away if we trust our preliminary finding that syntactic ambiguity does not seem to play a decisive role for humorous translation. The ‘humour’-version of the parser thus produces a maximal chart holding all partial analyses of the SU — see the lower part of Figure 1.

The chart now constitutes the search space for the humour processor, which has to decide what to transfer to the TL and then activate generation. Since the TU, as pointed out above, should make sense, the choice obviously requires some knowledge of the target language. A hypothetical, brute-force solution would be to generate

⁵Cf. the detailed analyses of analogies in Ambrose Bierce jokes by Shelley et al. (1996), or the semantic database used by Raskin (1996).




Figure 1: Proposed system architecture

all possible translations and then have an understanding component select the (most) humorous one. This is neither practical nor elegant, and hence the humour processor should be integrated into the transfer step and control the search. A possible strategy is to inspect all the TL readings of the content words in the SU (as they appear in the chart), determine their domains (or, scripts) and compare them to the domain of the ‘correct’ SU analysis; using a script opposition measure, the TL lexemes can then be ranked according to suitability. This ranking then is a measure for the desirability of using the words in the TU.

Now, the problem is that we want to produce a coherent TU, while the chart holds no knowledge of the TL other than the individual readings of lexemes. Thus, we envisage to split the overall search task in two stages. In the first one, partial analyses of the SU chart that include the “desired” words are transferred individually into the TL, which loosely corresponds to translating individual constituents of sentences. The TL chunks are also arranged in a chart, such that a second search can take place in conjunction with TL generation: Given the weights associated with the chart entries (which are copied from the SL chart), this process tries to construct a well-formed and coherent TU from (some of) the chart entries such that the overall weight is maximized.

Splitting the overall search into these two components appears to be a way to solve the problem of being faithful to the SU on the one hand (the translation, even if bad, must still be recognizable as a *translation*) and respecting the syntactic and semantic well-formedness of the TL on the other hand (the translation should not be ill-formed or nonsensical).

An important prerequisite for establishing the rankings of options in the transfer and selection steps is the availability of fine-grained lexical knowledge and world knowledge. As for lexical knowledge reflecting nuances of meaning, DiMarco et al. (1993) made a proposal for gathering it from dictionary usage notes and representing it in a description logic framework; Edmonds (1999) describes a system that utilizes such knowledge for fine-grained lexical choice in language generation.

The implementation of our proposed system is currently under way. Its basis is the MT system KIT-FAST (Weisweber, 1994), which realizes a transfer architecture as described above and translates between English and German. Following the syntactic chart parsing, FAST con-

structs up a functor-argument structure and then maps it to a conceptual representation in a description logic; on this level, reasoning with world knowledge can be performed. Due to its modular architecture, integrating a humour processor into the finished system is feasible. The additional lexical and world knowledge will be represented along the lines of the frameworks just mentioned.

5 Example

Finally, to illustrate the proposed translation scheme, we provide a description of a sample translation. Our example is taken from the ‘Orchy homepage’ (<http://www.orchy.com/home/english/whispers.htm>), which gives the Australian National Anthem together with a German version that was produced by the online MT system ‘Systan’ (<http://www.systransoft.com>). One line runs as follows:

*With courage let us all combine
To Advance Australia Fair*

And the corresponding German output is:

*Mit Mut lassen Sie uns aller
Mähdrescher Australienmesse vorrücken*

This can be glossed in English as:

*With courage let us move forward the Aus-
tralian fairs of combines*

The translation failure results mainly from two mix-ups. First, the rhetorical figure of adjective–noun reversal (hyperbaton) in *Australia Fair* is not recognized; instead, *fair* is analyzed in its noun reading and then correctly translates to the German compound noun *Australienmesse*. Second, *combine* is analyzed not as a verb but as a noun denoting a harvest machine, yielding its German equivalent *Mähdrescher*. In this way, the sense communicated by the utterance (and glossed above) arises, even though the German syntax is not exactly correct; but readers are able to fill in the plural morphology quite easily, as we found.

Crucial for the humour is the extreme unsuitability of the word *Mähdrescher* in the context of a national anthem. It is a compound noun that literally means *mower–thrasher*. *Dreschen* in German is also used with the prefix *ver-* and

then means ‘to beat up’ (somewhat similar to the English *to thrash*); therefore, *Mähdrescher* has strong connotations of violence. As such it combines nicely with the introductory phrase *mit Mut* (‘with courage’) — only the kind of courage is not the one intended in the original anthem.

To demonstrate how the German translation can arise from the English input, Figure 2 shows an excerpt of the parse chart yielded by syntactic analysis. Note that we are not interested in syntactic details here and use rather coarse labels to denote nodes in the parse forest, and we do not address the issue of representing the English *let us ...* construction correctly. Within the dashed box, the different syntactic readings of the lexemes are given, as taken from an online Webster’s dictionary. Of course, the individual readings typically correspond to multiple German words — the translation ambiguity adds to the mere syntactic one. For illustration, a few German translations are shown below the box, linked to their corresponding words by dotted lines.

For reasons of space, the parse forest, too, is incomplete and serves merely to illustrate the idea. For some of the nodes, we have shown the translations that would result from transferring the particular sub-tree to German. Notice, for instance, that *us* can be mis-read as an abbreviation of *United States*, which combines with one verb sense of *let* yielding *die USA vermieten* (‘to rent out the United States’).

The complete chart represents the search space for the transfer component enhanced by the humour processor. The first step is attaching weights to the word readings, the second step is transferring high-ranking words to the TL chart, such that the entire input is covered. To determine weights, in this case we need to know that *Mähdrescher* has connotations of violence, which are in conflict with the script underlying a national anthem (at least the Australian one). Furthermore, our lexical and conceptual knowledge should tell us that *with courage* meshes well with the violent connotations, and that a combine, since it is a vehicle, can move forward (*vorrücken*), thereby creating an overall coherent target sentence.

A more complex algorithm thus could perform the steps of weight assignment and transfer not only for individual word readings, but for arbitrary chart entries; this would account for effects stemming from composing TL words into phrases. In this case, we construct a TL chart with high-ranking words and phrases that cover the entire SU. Now, a promising strategy for generating the TL sentence, which we are currently exploring, is the ‘bag generation’ proposed as part of ‘shake and bake machine translation’ by Popovich (1996).

6 Summary

We have argued that machine translation provides an effective framework for developing and testing computational theories of humour. The genre of ‘translation hu-

mour’ can be viewed as jokes consisting of two sentences in different languages; while the source sentence as such is not funny, a wrong translation illuminates ambiguities that would otherwise typically not be perceived. When the overall sense of the incorrect translation stands in an opposition relation to that of the source sentence, a humorous note arises — in the absence of such an opposition, a merely bad translation might be an instance of ‘machine translation humour’ that people find funny because machines are so stupid. We have not developed a more explicit notion of ‘opposition’ yet; rather, this is a central task of our ongoing work. At any rate, good translation humour presupposes knowledge of the two languages involved, and the better the reader knows their subtleties, the better are the chances for achieving good humour.

We proposed a system architecture that extends the standard transfer-based machine translation paradigm with a humour processor. It inspects the syntactic chart of the source sentence and uses a script–opposition measure to rank partial analyses according to their suitability for a humorous translation. High-ranking parts are transferred to the target language, where they again form a chart, from which generation commences. We are currently implementing our approach, using the MT system ‘KIT-FAST’.

Other sources of ambiguity such as near-homophones (the infamous ‘false friends’) can be integrated into such a framework: The humour processor would enlarge the search space by employing phonological knowledge in addition to standard lexicons. Furthermore, when extending the approach to spoken–language translation, one would envisage that a system could on purpose craft translations such as the following, which was produced by ‘Verbmo-bil’ inadvertently:

Speaker: *Hello, good afternoon*
Speech recognizer delivered: “hell the afternoon”
Translation output: *Hölle dem Nachmittag*
(‘To hell with the afternoon’)

Acknowledgements

We thank the anonymous reviewers for their suggestions on improving an earlier version of this paper.

References

- S. Attardo. *Linguistic Theories of Humor*. Mouton de Gruyter, Berlin, 1994.
- C. DiMarco, G. Hirst, and M. Stede. The semantic and stylistic differentiation of synonyms and near-synonyms. In *Working notes of the AAAI Spring Symposium on Building Lexicons for Machine Translation*. Stanford University, 1993.

- P. Edmonds. *Semantic Representations of Near-Synonyms for Automatic Lexical Choice*. PhD thesis, Dept. of Computer Science, University of Toronto, 1999.
- S. Nirenburg, J. Carbonell, M. Tomita, and K. Goodman. *Machine Translation: A knowledge-based approach*. Morgan Kaufmann, San Mateo, 1992.
- F. Popovich. A chart generator for shake and bake machine translation. In *Proceedings of AI '96 – 11th Canadian Conference on Artificial Intelligence*. Toronto, 1996.
- V. Raskin. *Semantic mechanisms of humor*. Reidel, Dordrecht, 1985.
- V. Raskin. Computer implementation of the general theory of verbal humour. In *Proceedings of the International Workshop on Computational Humor (IWCH '96)*. Twente, 1996.
- C. Shelley, T. Donaldson, and K. Parsons. Humorous analogy: modelling ‘the devil’s dictionary’. In *Proceedings of the International Workshop on Computational Humor (IWCH '96)*. Twente, 1996.
- W. Wahlster. Verbmobil: Translation of face-to-face dialogues. In *Proceedings of the Third European Conference on Speech Communication and Technology*. Berlin, 1993.
- W. Weisweber. The experimental mt system system of the project kit-fast. In *Proceedings of the International Conference Machine Translation: Ten Years On*, pages 12.1–12.19. Cranfield, 1994.

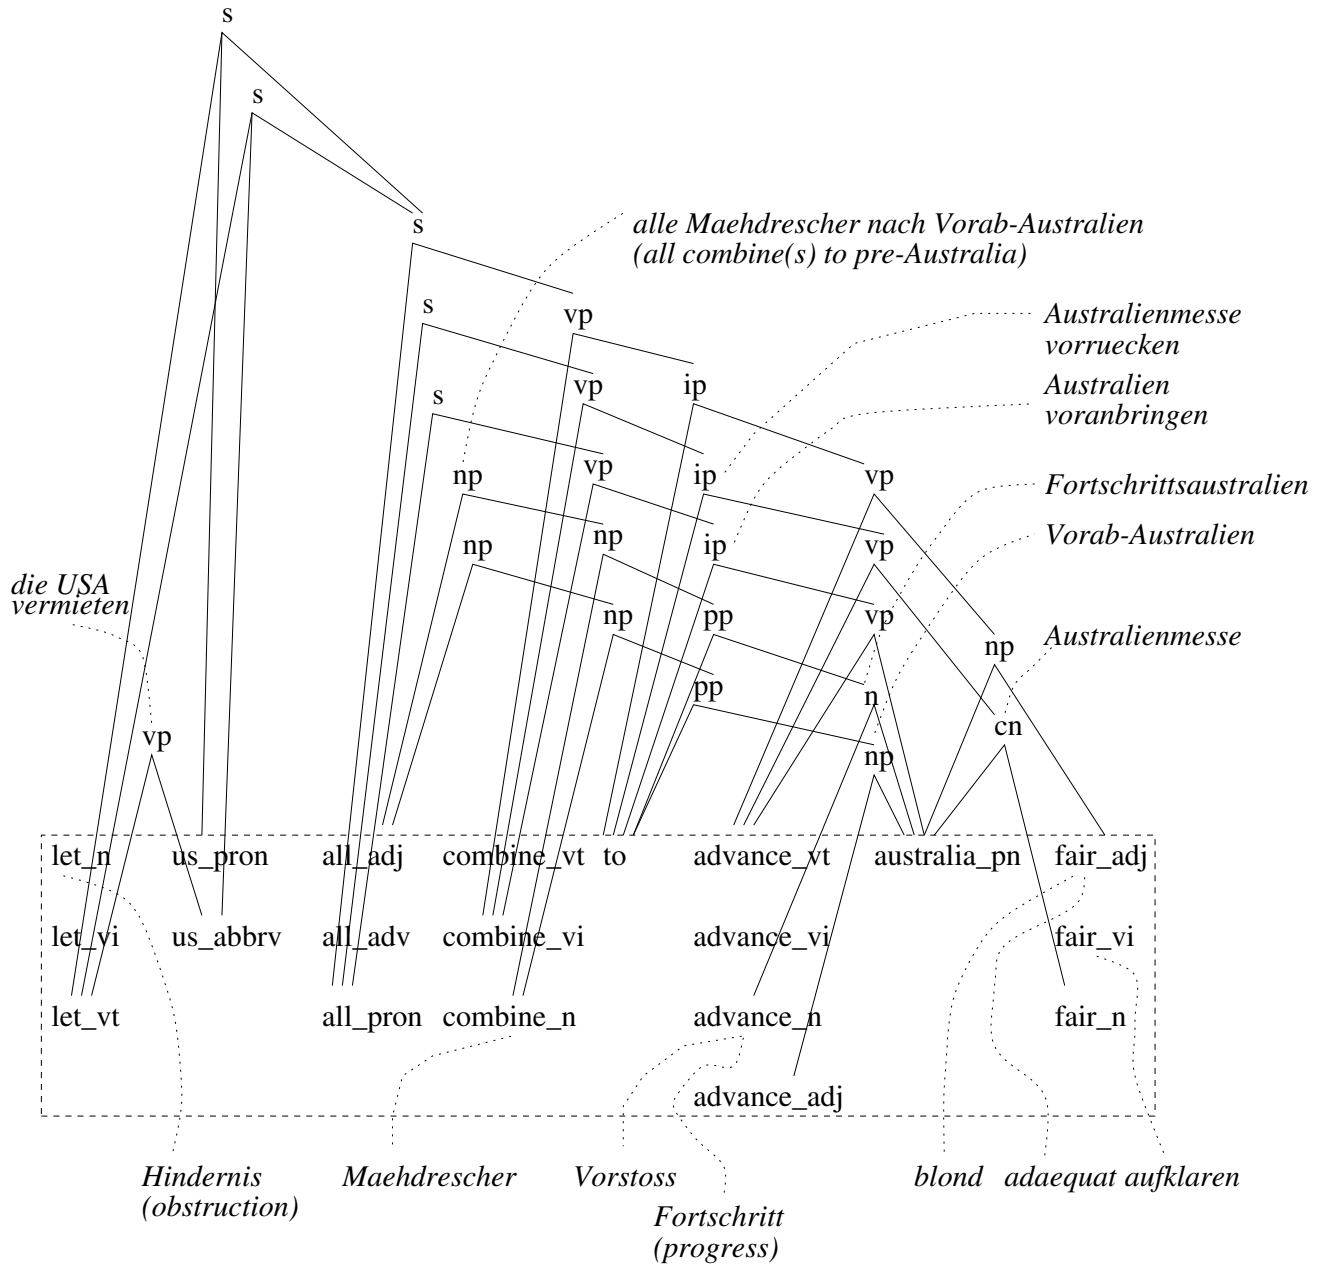


Figure 2: Excerpt from parse chart and translations