

Wordform- and class-based prediction of the components of German nominal compounds in an AAC system

Marco Baroni* Johannes Matiasek* Harald Trost†

February 22, 2002

Abstract

In word prediction systems for augmentative and alternative communication (AAC), productive word-formation processes such as compounding pose a serious problem. We present a model that predicts German nominal compounds by splitting them into their modifier and head components, instead of trying to predict them as a whole. The model is improved further by the use of class-based modifier-head bigrams constructed using semantic classes automatically extracted from a corpus. The evaluation shows that the split compound model with class bigrams leads to an improvement in keystroke savings of more than 8% over a no split compound baseline model. We also present preliminary results obtained with a word prediction model integrating compound and simple word prediction.

1 Introduction

N-gram language modeling techniques have been successfully embedded in a number of natural language processing applications, including word predictors for augmentative and alternative communication (AAC). *N*-gram based techniques rely crucially on the assumption that the large majority of words to be predicted have also occurred in the corpus used to train the models.

Productive word-formation by compounding in languages such as German, Dutch, the Scandinavian languages and Greek, where compounds are commonly written as single orthographic words, is problematic for this assumption.

Productive compounding implies that a sizeable number of new words will constantly be added to the language. Such words cannot, in principle, be contained in any

*Austrian Research Institute for Artificial Intelligence, Vienna

†Department of Medical Cybernetics and Artificial Intelligence, University of Vienna

already existing training corpus, no matter how large. Moreover, the training corpus itself is likely to contain a sizeable number of newly formed compounds that, as such, will have an extremely low frequency, causing data sparseness problems.

New compounds, however, differ from other types of new/rare words in that, while they are rare, they can typically be decomposed into more common smaller units (the words that were put together to form them). For example, in the corpus we analyzed, *Abend* 'evening' and *Sitzung* 'session', the two components of the German compound *Abendsitzung* 'evening session', are much more frequent words than the latter. Thus, a natural way to handle productively formed compounds is to treat them not as primitive units, but as the concatenation of their components.

A model of this sort will be able to predict newly formed compounds that never occurred in the training corpus, as long as they can be analyzed as the concatenation of constituents that did occur in the training corpus. Moreover, a model of this sort avoids the specific type of data sparseness problems caused by newly formed compounds in the training corpus, since it collects statistics based on their (typically more frequent) components.

Building upon previous work (e.g., [14], [5], [9], and [12]), [1] reported encouraging results obtained with a model in which two-element nominal German compounds are predicted by treating them as the concatenation of a *modifier* (left element) and a *head* (right element).

Here, we report of further improvements to this model that we obtained by adding a new, class-based bigram term to head prediction. As far as we know, this is the first time that semantic classes automatically extracted from the training corpus have been used to enhance compound prediction, independently of the domain of application of the prediction model.

Moreover, we present the results of preliminary experiments we conducted in the integration of compound predictions and simple word predictions within the AAC word prediction task.

The remainder of this paper is organized as follows. In section 2, we describe the AAC word prediction task. In section 3, we describe the basic properties of German compounds. In section 4, we present our split compound prediction model, focusing on the new class-based head prediction component. In section 5, we report the results of simulations run with the enhanced compound prediction model. In section 6, we report about our preliminary experiments with the integration of compound and simple word prediction. Finally, in section 7, we summarize the main results we obtained and indicate directions for further work.

2 Word prediction for AAC

Word prediction systems based on n -gram statistics are an important component of AAC devices, i.e., software and possibly hardware typing aids for disabled users (see for example [4], [6]).

Word predictors provide the user with a *prediction window*, i.e. a menu that, at any time, lists the most likely next word candidates, given the input that the user has typed until the current character. If the word that the user intends to type next is in the prediction window, the user can select it from there. Otherwise, the user will keep typing letters, until the target word appears in the prediction window (or until she finishes typing the word).

The (percentage) *keystroke savings rate* (*ksr*) is a standard measure used in AAC research to evaluate word predictors. The *ksr* can be thought of as the percentage of keystrokes that a “perfect” user would save by employing the relevant word predictor to type the test set, over the total number of keystrokes that are needed to type the test set without using the word predictor.

Usually, the *ksr* is defined by

$$ksr = \left(1 - \frac{k_i + k_s}{k_n}\right) * 100 \quad (1)$$

where: k_i is the number of input characters actually typed, k_s is the number of keystrokes needed to select among the predictions presented by the model and k_n is the number of keystrokes that would be needed if the whole text was typed without any prediction aid. Typically, the user will need one keystroke to select among the predictions, and thus we assume that k_s equals 1.¹

The *ksr* is influenced not only by the quality of the prediction model but also by the size of the prediction window. In our simulations, we use a 7 word prediction window.

Ksr is not a function of perplexity, but it is generally true that there is an inverse correlation between *ksr* and perplexity (see [4] and references quoted there).

3 Compounding in German

Compounding is an extremely common and productive mean to form words in German.

In an analysis of the APA newswire corpus (a corpus of over 28 million words),² it was found that almost half (47%) of the word types were compounds. However, the compounds accounted for a small portion of the overall token count (7%). This suggests that, as expected, many of them are productively formed *hapax legomena* or very rare words (83% of the compounds had a corpus frequency of 5 or lower).

By far the most common type of German compound is the N+N type, i.e., a sequence of two nouns (62% of the compounds in our corpus have this shape). Thus, we decided to limit ourselves, for now, to handling compounds of this shape.

¹In the split compound model, the user needs one keystroke to select the modifier and one keystroke to select the head.

²This corpus was kindly made available to us by the Austria Presse Agentur.

In German, nominal compounds, including the N+N type, are *right-headed*, i.e., the rightmost element of the compound determines its basic semantic and morphosyntactic properties.

Thus, the context of a compound is often more informative about its right element (the *head*) than about its left element (the *modifier*).

In modifier context, nouns are sometimes followed by a *linking suffix* ([11],[7]), or they take other special inflectional shapes.

As a consequence of the presence of linking suffixes and related patterns, the forms that nouns take in modifier position are sometimes specific to this position only, i.e., they are bound forms that do not occur as independent words.

We did not parse special modifier forms in order to reconstruct their independent nominal forms. Thus, we treat all inflected modifier forms, including bound forms, as unanalyzed primitive nominal wordforms.

4 The split compound prediction model

In [1], we present and evaluate a split compound model in which N+N compounds are predicted by treating them as the sequence of a modifier and a head.

Modifiers are predicted on the basis of weighed probabilities deriving from the following three terms: the unigram and bigram training corpus frequency of nominal wordforms as modifiers or independent words, and the training corpus type frequency of nominal wordforms as modifiers:³

$$P_{mod}(w) = \lambda_1 P(w) + \lambda_2 P(w|c) + \lambda_3 P_{ismod}(w) \quad (2)$$

The type frequency of nouns as modifiers is determined by the number of distinct compounds in which a noun form occurs as modifier.

Heads are predicted on the basis of weighted probabilities deriving from three terms analogous to the ones used for modifiers: the unigram and bigram frequency of nouns as heads or independent words, and the type frequency of nouns as heads:

$$P_{head}(w) = \lambda_1 P(w) + \lambda_2 P(w|c) + \lambda_3 P_{ishead}(w) \quad (3)$$

The type frequency of nouns as heads is determined by the number of distinct compounds in which a noun form occurs as head.

Given that compound heads determine the syntactic properties of compounds, bigrams for head prediction are collected by considering not the immediate left context of heads (i.e., their modifiers), but the word preceding the compound (e.g., *die Abendsitzung* is counted as an instance of the bigram *die Sitzung*).

³Here and below, *c* stands for the last word in the left context of *w*; *w* is the suffix of the word to be predicted minus the (possibly empty) prefix typed by the user up to the current point.

For reasons of size and efficiency, single n -gram count lists are used for modifiers and heads.⁴ For the same reasons, and to minimize the chances of over-fitting to the training corpus, all n -gram/frequency tables are trimmed by removing elements with a frequency of 5 or lower.

We currently use a simple interpolation model, in which all terms are assigned equal weight.

4.1 Improving head prediction

While we obtained encouraging results with it (see [1]), we feel that a particularly unsatisfactory aspect of the model described in the previous section is that information on the modifier is not exploited when trying to predict the head of a compound. Intuitively, knowing what the modifier is should help us in guessing the head of a compound. However, constructing a plausible head-prediction term based on modifier-head dependencies is not straightforward.

The word-form-based compound-bigram frequency of a head, i.e., the number of times a specific head occurs after a specific modifier, is not a very useful measure: Counting how often a modifier-head pair occurs in the training corpus is equivalent to collecting statistics on unanalyzed compounds, and it will not help us to generalize beyond the compounds encountered in the training corpus. Moreover, if a specific modifier-head bigram is frequent, i.e., the corresponding compound is a frequent word, it is probably better to treat the whole compound as an unanalyzed lexical unit anyway.

POS-based head-modifier bigrams are not going to be of any help either, since we are considering only N+N compounds, and thus we would collect a single POS bigram (N N) with probability 1.⁵

We decided instead to try to exploit a semantically-driven route. It seems plausible that modifiers that are semantically related will tend to co-occur with heads that are, in turn, semantically related. Think for example to the relationship between the class of fruits and the class of sweets in English compounds. It is easy to think of compounds in which a member of the class of fruits (bananas, cherries, apricots...) modifies a member of the class of sweets (pies, cakes, muffins...). Thus, if you have to predict the head of a compound given a fruit modifier, it would be reasonable, all else being equal, to guess some kind of sweet.

4.1.1 Class-based modifier-head bigrams

While semantically-driven prediction makes sense in principle, clustering nouns into semantic classes is certainly not a trivial job, and, if a large input lexicon must be

⁴This has a distorting effect on the bigram counts (words occurring before compounds are counted twice, once as the left context of the modifier and once as the left context of the head). However, preliminary experiments indicated that the empirical effect of this distortion is minimal.

⁵Even if the model handled other compound types, very few POS combinations are attested within compounds.

partitioned, it is not a task that could be accomplished by a human expert. Drawing inspiration from [3], we constructed instead semantic classes using a clustering algorithm extracting them from a corpus, on the basis of the *average mutual information* (MI) between pairs of words ([13]).⁶

MI values were computed using Adam Berger’s trigger toolkit ([2]).⁷ The same training corpus of about 25.5M words (and with N+N compounds split) that we describe below was used to collect MI values for noun pairs. All modifiers and heads of N+N compounds and all corpus words that were parsed as nouns by the Xerox morphological analyzer ([10]) were counted as nouns for this purpose.

MI was computed only for pairs that co-occurred at least three times in the corpus. Valid co-occurrences were bound by a maximal distance between elements of 500 words, and a minimal distance of 2 words (to avoid lexicalized phrases, such as proper names or phrasal loanwords).

Having obtained a list of pairs from the toolkit, the next step was to cluster them into classes, by grouping together nouns with a high MI. For space reasons, we do not describe our clustering algorithm in detail here (we describe and analyze the algorithm in a paper currently in preparation).

In short, the algorithm starts by building classes out of nouns that occur with very few other nouns in the MI pair list, and thus their assignment to classes is relatively unambiguous, and it then adds progressively more ambiguous nouns (ambiguous in the sense that they occur in a progressively larger number of MI pairs, and thus it becomes harder to determine with which other nouns they should be clustered).

This is a heuristic clustering procedure and there is no guarantee that it will construct classes that maximize MI. However, a cursory inspection of the output list indicates that most classes constructed by our algorithm are intuitively reasonable. Table 1 reports a list of ten randomly selected classes that were constructed using this procedure.

The algorithm generated 3744 classes, containing a total of 14059 nouns (about one third of the nouns in the training corpus).

Class-based modifier-head bigrams were then collected by labeling all the modifiers and heads in the training corpus with their semantic classes, and counting how often each combination of modifier and head class occurred.

In the simulations reported below, class-based bigrams were trimmed by removing elements with a frequency of 5 or lower, but comparable results were also obtained using different trimming thresholds.

⁶We are aware of the fact that other measures of lexical association have been proposed (see, among others, [8]) and are sometimes claimed to be more reliable than MI, and we are planning to run our clustering algorithm using alternative measures.

⁷The trigger toolkit returns directional MI values (i.e., separate MI values for the pairs N_1N_2 and N_2N_1). Since we were not interested in directional information, we merged pairs containing identical nouns by summing their MI. We realize that this is not mathematically equivalent to computing symmetric MI values, but it is a practical approximation that allowed us to use the trigger toolkit for our purposes.

Alleinstehende, Singles, Alben, Platten, Platte, Sound, Hits, Hit, Live, Songs, Single, Album, Pop, Studio, Rock, Fans, Band
Atrophie, Hartung, Neurologe
Magische, Magie
Bilgen, Tivoli, Baur, Scharrer, Streiter, Winkel, Pfeffer, Schmid, M
Effizienz, Transparenz
Harm, Radar, Jets, Flugzeugen, Typs, Abwehr, Raketen, Maschinen, Angriffen, Flugzeuge, Kampf
Relegation, Birmingham, Stephen
Partnerschafts, Partnerschaft, Kooperation, Bereichen, Aktivitäten
Importeure, Zölle
Labyrinths, Labyrinth

Table 1: Randomly selected noun classes

4.1.2 The class-based head prediction model

We compute the class-based probability of a compound head given its modifier in the following way:

$$P_{class}(h|m) = P(Cl(h)|Cl(m))P(h|Cl(h)) \quad (4)$$

where

$$P(Cl(h)|Cl(m)) = \frac{\text{count}(Cl(m), Cl(h))}{\text{count}(Cl(m))} \quad (5)$$

and

$$P(h|Cl(h)) = \frac{1}{|Cl(h)|} \quad (6)$$

The latter term assigns equal probability to all members of a class, but lower probability to members of larger classes.

Class-based probability is added to the wordform-based terms of equation 3 obtaining the following formula to compute head probability:

$$P_{head}(w) = \lambda_1 P(w) + \lambda_2 P(w|c) + \lambda_3 P_{ishead}(w) + \lambda_4 P_{class}(w|m) \quad (7)$$

5 Evaluation

The new split compound model and a baseline model with no compound processing were evaluated in a series of simulations, using the APA newswire articles from Jan-

uary to September 1999 (containing 25,466,500 words) as the training corpus, and the 123,025 compounds in the newswire articles from October 1999 as the testing targets.

In order to train and test the split compound model, all words in both sets were run through the morphological analyzer, and all N+N compounds were split into their modifier and head surface forms.

We first ran simulations in which compound heads were predicted using each of the terms in equation 7 separately. The results are reported in table 2.

model	$P(w)$	$P(w c)$	P_{ishead}	$P_{class}(w m)$
head <i>ksr</i>	49.4	39.8	52.5	45.0

Table 2: Predicting heads with single term models

As an independent predictor, the class-based term does not perform as good as unigram- and head-type-frequency-based prediction, but it outperforms wordform-based bigram prediction.

We then simulated head and compound prediction using the head prediction model of equation 7. The results of this simulation are reported in table 3, together with the results of a simulation in which class-based prediction was not used, and the results of the baseline no-split-compound model.

Model	split w/ classes	split no classes	no split
head <i>ksr</i>	61.4	57.7	N/A
compound <i>ksr</i>	59.8	57.9	51.5

Table 3: Predicting heads and compounds

When used in conjunction with the other terms, class bigrams lead to an improvement in head prediction of almost 4% over the split compound model without class-based prediction. This translates into an improvement of about 2% in the prediction of whole compounds. Overall, the split compound model with class bigrams leads to an improvement of more than 8% over the baseline model.

Clearly, there is still much room for improvement of our split-compound model, but the positive results reported here confirm our hypothesis that semantically related heads often follow semantically related modifiers, and they indicate that our approach to constructing semantic class-based bigrams is on the right track.

6 Preliminary experiments in integration

In a working word prediction system, compounds are obviously not the only type of words that the user needs to type. Thus, the predictions provided by the compound

model must be integrated with predictions of simple words. In this section, we report preliminary results we obtained with a model limited to the integration of N+N compound prediction with simple noun prediction.

Our results indicate that the integrated model is currently only approximating the performance of a model without compound prediction. This is encouraging when our experiments are seen as a first attempt at integration, but it clearly suggests that further work is needed on improving the integrated model, given that the whole point of adding compound prediction to an AAC system should be to make the system achieve a higher *ksr* rate.⁸

In our approach to compound/simple prediction integration, candidate modifiers are presented together and in competition with simple word solutions as soon as the user starts typing a new word. The user can distinguish modifiers from simple words in the prediction window because the former are suffixed with a special symbol (for example an underscore). If the user selects a modifier, the head prediction model is activated, and the user can start typing the prefix of the desired compound head, while the system suggests completions based on the head prediction model.

For example, if the user has typed *Abe*, the prediction window could contain, among other things, the candidates *Abend* and *Abend_*. If the user selects the latter, possible head completions for a compound having *Abend* as its modifier are presented.

Modifier candidates are proposed on the basis of $P_{mod}(w)$ computed as in equation 2 above. Simple noun candidates are proposed on the basis of their unigram and bigram probabilities (interpolated with equal weights).

We experimented with two versions of the integrated model.

In one, modifier and simple noun candidates are ranked directly on the basis of their probabilities. This risks to lead to over-prediction of modifier candidates (recall that, from the point of view of token frequency, compounds are much rarer than simple words; thus, the prediction window should not be cluttered by too many modifier candidates when, most of the time, users will want to type simple words).

Thus, we constructed a second version of the integrated model in which $P_{mod}(w)$ is multiplied by a penalty term. This term discounts the probability of modifier candidates built from nominal wordforms that occur more frequently in the training corpus as independent nouns than as modifiers (forms that are more – or as – frequent in modifier position are not affected by the penalty).

The same training corpus and procedures described in section 5 above were used to train the two versions of the integrated model, and the baseline model that does not use compound prediction.

These models were tested by treating *all* the nouns in the test corpus as prediction targets. The integrated test set contained 123,025 N+N tokens and 704,314 more

⁸The performance of the model without compound prediction must be assessed with the important *caveat* that our current training and test corpora display a large degree of overlap (see discussion in [1]). Thus, the simple model can perform relatively good compound prediction simply in virtue of the fact that most of the compounds in the test set also occur in the training corpus. Clearly, we must very soon run simulations with a test set further removed from the training corpus.

nouns. The results of the simulations are reported in table 4.

Model	integrated no penalty	integrated w/ penalty	simple pred only
compound <i>ksr</i>	57.7	56.0	51.5
simple n <i>ksr</i>	52.3	54.5	57.3
combined <i>ksr</i>	53.6	54.9	55.9

Table 4: Integrated prediction

Because of the simple noun predictions getting in the way, the integrated models perform compound prediction worse than the non-integrated split compound model of table 3. However the integrated models still perform compound prediction considerably better than the baseline model.

The integrated model with modifier penalties performs worse than the model without penalties when predicting compounds. This is expected, since the modifier penalties make this model more conservative in proposing modifier candidates.

However, the model with penalties outperforms the model without penalties in simple noun prediction. Given that in our test set (and, we expect, in most German texts) simple nouns greatly outnumber compounds, this results in an overall better performance of the model with penalties.

The integrated model with penalties achieves an overall *ksr* that is exactly 1% lower than that achieved by the baseline model.

Thus, these preliminary experiments indicate that an approach to integrating compound and simple word predictions along the lines sketched at the beginning of this section, and in particular the version of the model in which modifier predictions are penalized, is feasible. However, the model is clearly in need of further refinement, given that a close approximation of the performance of a baseline model without compound prediction is obviously not what we are aiming for.

7 Conclusion

The main result concerning German compound prediction that was reported in this paper pertains to the introduction of class-based modifier-head bigrams to enhance head prediction.

We presented a procedure to cluster nominal wordforms into semantic classes and to extract class-based modifier-head bigrams, and then a model to calculate the class-based probability of candidate heads using these bigrams.

The addition of class-based head prediction to the split compound of [1] leads to a considerable improvement in head prediction (from a *ksr* of 57.7% to a *ksr* of 61.4%). This translates into an improvement of about 2% in whole compound prediction (from 57.9% to 59.8%). Overall, the split compound model with class bigrams led to an improvement of more than 8% over a no split compound baseline model.

This result was presented in the context of the AAC word prediction task, but we believe that the class-based prediction model we proposed could be extended to any other domain in which n -gram-based compound prediction must be performed.

Future work should bring further enhancements to the semantic-class-based head prediction model and to compound prediction in general.

For example, we plan to test in the near future whether measuring the degree of relatedness of words using measures other than average mutual information will result in improved semantic classes.

Moreover, we hope to improve the overall performance of the compound predictor by resorting to a better interpolation strategy than the "equal weight everywhere" model we are currently adopting.

We also reported results obtained with a preliminary model in which split compound prediction is integrated with simple noun prediction. This model is currently only able to closely approximate the performance of a baseline model without compound prediction. Clearly, further work in this area is also necessary.

Acknowledgements

This work was supported by the European Union in the framework of the IST programme, project FASTY (IST-2000-25420). Financial support for ÖFAI is provided by the Austrian Federal Ministry of Education, Science and Culture.

References

- [1] M. Baroni, J. Matiassek, H. Trost, *Predicting the components of German nominal compounds*, ÖFAI Technical Report TR-2002-09, 2002.
- [2] A. Berger: Trigger Toolkit, publicly available software, 1997.
<http://www-2.cs.cmu.edu/aberger/software/TTK.html>
- [3] P. Brown, V. Della Pietra, P. DeSouza, J. Lai, and R. Mercer, 'Class-based n -gram models of natural language', *Computational Linguistics* 18(4), pp.467-479, 1990.
- [4] J. Carlberger, *Design and Implementation of a Probabilistic Word Prediction Program*, Royal Institute of Technology (KTH), 1998.
- [5] D. Carter, J. Kaja, L. Neumeyer, M. Rayner, F. Weng, and M. Wirèn, 'Handling Compounds in a Swedish Speech-Understanding System', *Proc. ICSLP-96*.
- [6] A Copestake, 'Augmented and alternative NLP techniques for augmentative and alternative communication', *Proceedings of the ACL workshop on Natural Language Processing for Communication Aids*, 1997.
- [7] W. Dressler, G. Libben, J. Stark, C. Pons, and G. Jarema, 'The processing of interfixed German compounds', *Yearbook of Morphology 1999*, pp. 185-220, 2001.

- [8] S. Evert and B. Krenn: Methods for the Qualitative Evaluation of Lexical Association Measures in *Proceedings of the 39th Annual Meeting of the Association for Computational Linguistics* Toulouse, France, 2001.
- [9] P. Fetter, *Detection and Transcription of OOV Words*, Verbmobil Report 231, 1998.
- [10] L. Karttunen, K. Gal, and A. Kempe, *Xerox Finite-State Tool*, Xerox Research Centre Europe, Grenoble, 1997.
- [11] A. Krott, *Analogy in Morphology*, Max Planck Institute for Psycholinguistics, Nijmegen, 2001.
- [12] M. Larson, D. Willett, J. Kohler, and G. Rigoll, 'Compound splitting and lexical unit recombination for improved performance of a speech recognition system for German parliamentary speeches', in *Proceedings of the 6th International Conference of Spoken Language Processing (ICSLP-2000)*, October 16-20., Peking, China, 2000.
- [13] R. Rosenfeld. A Maximum Entropy Approach to Adaptive Statistical Language Modeling. *Computer Speech and Language* 10, 187–228, 1996.
- [14] M. Spies, 'A Language Model for Compound Words', *Proc. Eurospeech '95*, pp.1767-1779, 1995.