Association Norms of German Noun Compounds

Sabine Schulte im Walde
Institut für Maschinelle Sprachverarbeitung
Universität Stuttgart
schulte@ims.uni-stuttgart.de

Susanne Borgwaldt
Institut für Germanistik
TU Braunschweig
s.borgwaldt@tu-bs.de

Ronny Jauch
Institut für Maschinelle Sprachverarbeitung
Universität Stuttgart
jauchry@ims.uni-stuttgart.de

Abstract

This paper introduces association norms of German noun compounds as a lexical-semantic resource for cognitive and computational linguistics research on compositionality. Based on an existing database of German noun compounds, we collected human associations to the compounds and their constituents within a web experiment. The current study describes the collection process and a part-of-speech analysis of the association resource. In addition, we demonstrate that the associations provide insight into the semantic properties of the compounds, and perform a case study that predicts the degree of compositionality of the experiment compound nouns, as relying on the norms. Applying a comparatively simple measure of association overlap, we reach a Spearman rank correlation coefficient of $r_s = 0.5228, p < .000001$, when comparing our predictions with human judgements.

Keywords: German noun compounds, association norms, compositionality

1. Introduction

This paper introduces a new resource containing association norms of German noun compounds. Association norms have a long tradition in psycholinguistic research, where the implicit notion that associates reflect meaning components of words has been used for more than 30 years to investigate semantic memory. In the last decade, association norms have also found their way into lexical-semantic research in computational linguistics. For example, Rapp (2002) developed corpus-based approaches to predict paradigmatic and syntagmatic associations; de Deyne and Storms (2008a) created semantic networks from Dutch associations; and Schulte im Walde (2008) used associations to German verbs to select features for semantic classification.

We collected associations to German noun compounds because we believe that the associations are a valuable resource for cognitive and computational linguistics research on compositionality. Based on an existing collection of German noun compounds, we therefore gathered associations for the compounds and also for their constituents (e.g., *Ahornblatt/Ahorn/Blatt*). The data were collected within a web experiment. The association norms can be used as lexical semantic resource concerning the target stimuli, i.e., the compound nouns and their constituents. The data should be relevant for research on the lexical semantic properties of the stimulus words, the (semantic) relations between stimuli and associations, and the semantic relatedness between the compounds and their constituents.

In order to demonstrate that the associations provide insight into semantic properties of the compounds (and their constituents) that should be useful for computational models of compound compositionality, we perform a case study that predicts the degree of compositionality of the experiment compound nouns. We plan to work on more elaborate predictions of compound compositionality, and consider the current case study as a baseline to future computational experiments.

In the remainder of the paper, Section 2 provides an introduction to the compound data; Section 3 describes the collection of the associations, and Section 4 the resulting norms; finally, Section 5 illustrates the case study on compositionality.

2. German Noun Compounds

Compounds are combinations of two or more simplex words. Traditionally, a number of criteria (such as compounds being syntactically inseparable, and that compounds have a specific stress patterns) have been defined, in order to establish a border between compounds and non-compounds. However, Lieber and Stekauer (2009a) demonstrate that none of these tests are universally reliable to distinguish compounds from phrases and other types of derived words.

Compounds have thus been a recurrent focus of attention within theoretical, cognitive, and in the last decade also within computational linguistics. Recent evidence of the strong interest are the Handbook of Compounding (Lieber and Stekauer, 2009b) on theoretical perspectives, and a series of workshops[1] and special journal issues (Journal of Computer Speech and Language, 2005; Language Resources and Evaluation, 2010; ACM Transactions on Speech and Language Processing; to appear) with respect to the computational perspective.

Our focus of interest is on German noun compounds (cf. Fleischer and Barz (2012) for a detailed overview), such as *Ahornblatt* 'maple leaf', *Feuerwerk* 'fireworks', *Nähmaschine* 'sewing machine', *Obstkuchen* 'fruit cake', and *Rotkohl* 'red cabbage', where the head (as the rightmost constituent) is a noun, and the modifier can be from a set of various parts-of-speech.

More specifically, we are interested in the degrees of compositionality of German noun compounds, i.e., the relation between the meaning of the whole compound (e.g., *butter-*

[1] www.multiword.sourceforge.net

fly) and the meaning of its parts (e.g., *butter, fly*), which has been studied intensively by psycholinguists (mostly with respect to the term *semantic transparency*), in order to find out how compound words are processed and represented in the mental lexicon. There is an ongoing debate about whether morphologically complex words are stored as single units (Butterworth's full listing approach (1983)), whether they are decomposed into their morphemes (Taft, 2004; Taft and Forster, 1975), or whether they can be accessed both ways: as whole forms and componentially, via their constituent morphemes (dual route models, cf. Caramazza et al. (1988) and Baayen and Schreuder (1999)).

One variable that might be important for the processing of compounds is their compositionality, with some researchers (e.g., Longtin et al. (2003), and Marslen-Wilson et al. (1994)) arguing that morphological decomposition happens only in semantically transparent words, not in semantically opaque ones. In a similar argument, studies by Sandra (1990) and Zwitserlood (1994) showed that the meanings of the constituents of semantically transparent compounds (e.g. *dog* and *house* in *doghouse*) were activated during processing, whereas the meanings of the constituents of opaque compounds (e.g. *butter* and *fly* in *butterfly*) were not activated.

While the long-term goal of our work is to provide a computational model for the degree of compositionality of noun compounds, the current study presents a first step to obtain insight into salient properties of German noun compounds: A database of associations has been collected both on a selection of German noun compounds as well as on their constituents. Association norms have been used for decades to investigate semantic memory (cf. Section 4 for an overview), and we believe that the norms provide a useful starting point for computational features regarding the compositionality of the compounds.

Our work is based on a selection of noun compounds by von der Heide and Borgwaldt (2009), who were interested in semantic transparency ratings of noun compounds, and created a set of 450 concrete, depictable German noun compounds. As part of their work, they collected judgements on the compositionality of the 450 compounds. The compounds were distributed over 5 lists, and 270 participants in a paper experiment judged the degree of compositionality of the compounds with respect to their first as well as their second constituent, on a scale between 1 (opaque) and 7 (strong compositionality). Then for each compound noun the mean compositionality value was calculated. For example, the mean compositionality value for *Fliegenpilz* 'fly agaric' with respect to its first constituent *Fliege* 'fly/bow tie' is 1.93, and with respect to its second constituent *Pilz* 'mushroom' is 6.55.

We rely on von der Heide and Borgwaldt (2009) in two ways: (a) We use the noun compounds and their constituents as stimuli within the association experiment in Section 3, and (b) we use the compositionality judgements to evaluate the degrees of compositionality we predict in Section 5.

3. Web Experiment

This section introduces our method for collecting the associations for the German noun compounds. The common way to obtain associations is by presenting *target stimuli* to the participants in an experiment, who then provide *associate responses*, i.e., words that are spontaneously called to mind by the stimulus words. The quantification of the resulting target–association pairs (i.e., how often a certain association is provided for a certain target) is called *association norms*.

Our experiment used as target stimuli 442 of the compound nouns and all of their constituents as described in Section 2. In total, our material comprised 996 target stimuli.[2] The stimuli were divided randomly into 12 separate experimental lists of 83 nouns each.

The experiment was administered over the Internet. When participants loaded the experimental page, they were first asked for their biographical information, such as linguistic expertise, age and regional dialect. Next, the participant was presented with the written instructions for the experiment and an example item with potential responses. In the actual experiment, each trial consisted of a single word presented in a box at the top of the screen. The word was either one of the noun compounds, or one of the constituents. If a compound constituent was not a noun, the base form was nominalised by starting it with a capital letter. The order of the target words was random for each data set and each participant. Below the target were three data input lines where participants could type their associations. They were instructed to type at most one word per line and, following German grammar, to distinguish nouns from other parts-of-speech with capitalisation. Below the three input lines was a box that participants were asked to check if they did not know the word.

268 participants took part in the experiment, between 14 and 28 for each data set.Because the participants could provide between none and three associations per target, the actual number of associations per stimulus varies between 6 and 74.

In total, we collected 46,989 associations from 17,906 trials, an average of 2.6 associations per trial. The 46,989 association tokens are distributed over 29,221 association types. Considering only the first responses to the target stimuli, the total number of associations per stimulus is between 2 and 26, and in total there are 17,045 association tokens distributed over 11,038 association types. In 861 trials, the participants did not provide any association, out of which 327 targets were explicitely checked as not known by the participants.

4. Association Norms

Association norms have a long tradition in psycholinguistic research. They have been used for more than 30 years to investigate semantic memory, making use of the implicit notion that associates reflect meaning components of words. In this section, we will first provide an overview of existing norms and analyses of associations norms (Section 4.1),

[2]The total number of target stimuli is less than 442×3 because some compounds share constituents.

before we introduce our association norms of the German noun compounds (Section 4.2) and the morpho-syntactic analysis of the norms (Section 4.3).

4.1. Collections and Analyses of Association Norms

One of the first collections of word association norms was done by Palermo and Jenkins (1964), comprising associations for 200 words. The *Edinburgh Association Thesaurus* (Kiss et al., 1973) was a first attempt to collect association norms on a larger scale, and also to create a network of stimuli and associates, starting from a small set of stimuli derived from the Palermo and Jenkins norms. A similar motivation underlies the association norms from the University of South Florida (Nelson et al., 1998), who grew a stimulus-associate network over more than 20 years, from 1973. More than 6,000 participants produced nearly three-quarters of a million responses to 5,019 stimulus words. In another long-term project, Simon de Deyne and Gert Storms are collecting associations to Dutch words, cf. `www.smallworldofwords.com`. Previously, they performed a three-year collection for associations to 1,424 Dutch words (de Deyne and Storms, 2008b).

Smaller sets of association norms have also been collected for example for German (Russell and Meseck, 1959; Russell, 1970), Dutch (Lauteslager et al., 1986), French (Ferrand and Alario, 1998) and Spanish (Fernández et al., 2004) as well as for different populations of speakers, such as adults vs. children (Hirsh and Tree, 2001). Association norms have been used extensively in experimental psychology to conduct studies using the variations on the semantic priming technique to investigate, among other things, word recognition, knowledge representation and semantic processes (see McNamara (2005) for a review of methods, issues, and findings). Last but not least, association norms comparable to the ones presented here have been collected in earlier work by the authors and colleagues (Borgwaldt et al., 2005; Melinger and Weber, 2006; Schulte im Walde et al., 2008; von der Heide and Borgwaldt, 2009).

In parallel to the interest in collecting association norms, researchers have analysed association data in order to get insight into semantic memory and – more specifically – issues concerning semantic relatedness. The following paragraphs provide an overview of these analyses, starting with theoretical considerations on relationships between stimuli and responses in association norms (not actually based on collected data), and progressing towards analyses of collected norms.

Clark (1971) identified relations between stimulus words and their associations on a theoretical basis, not with respect to collected association norms. He categorised stimulus-association relations into sub-categories of paradigmatic and syntagmatic relations, such as synonymy and antonymy, selectional preferences, etc. Heringer (1986) concentrated on syntagmatic associations to a small selection of 20 German verbs. He asked his subjects to provide question words as associations (e.g., *wer* 'who', *warum* 'why'), and used the responses to investigate the valency behaviour of the verbs. Spence and Owens (1990) showed that associative strength and word co-occurrence are correlated. Their investigation was based on 47 pairs of semantically related concrete nouns, as taken from the *Word Association Norms* (Palermo and Jenkins, 1964), and their co-occurrence counts in a window of 250 characters in the 1-million-word Brown corpus. Church and Hanks (1989) were the first to apply information-theoretic measures to corpus data in order to predict word association norms. However, they did not rely on or evaluate against existing association data, but rather concentrated on the usage of the measure for lexicographic purposes. Rapp (2002) brought together research questions and methods from the above previous work: He developed corpus-based approaches to predict paradigmatic and syntagmatic associations, relying on the 100-million word BNC corpus. Rapp's work used the *Edinburgh Association Thesaurus* as association database.

Work by Fellbaum in the 1990s focused on human judgements concerning the semantic relationships between verbs. Fellbaum and Chaffin (1990) asked participants in an experiment to provide associations to verbs. The resulting verb-verb pairs were manually classified into five pre-defined semantic relations. Fellbaum (1995) investigated the relatedness between antonymous verbs and nouns and their co-occurrence behaviour. Within that work, she searched the Brown corpus for antonymous word pairs in the same sentence, and found that regardless of the syntactic category, antonyms occur in the same sentence with much higher-than-chance frequencies.

Another class of work addressed and partly classified the semantic relationships between stimuli and associates within association norms. Examples of this kind are Schlaghecken and Bölte (1998), Schulte im Walde et al. (2008), and von der Heide and Borgwaldt (2009) for German associations, and de Deyne and Storms (2008a) for Dutch.

Our own work, of course, is also closely related to this paper: Next to the analyses of German noun and verb associations at the syntax-semantics interface by Schulte im Walde et al. (2008), Schulte im Walde and Melinger (2008) performed a more in-depth analysis of the co-occurrence distributions of stimulus-response pairs. Guida (2007) can be considered as a first piece of cross-linguistic work. She replicated most of our analyses on verb association norms for Italian verbs. Finally, Melinger et al. (2006) took the noun associations as input to a soft clustering approach, in order to predict noun ambiguity, and to discriminate the various noun senses of ambiguous stimulus nouns.

4.2. Noun Compound Association Norms

For the current work, we created the association norms based on our experiment data as follows. For each stimulus, we quantified over all responses in the experiment, in two different modes: (i) considering only the first response in each trial, and (ii) considering all responses, and disregarding the order of the associates.

The reason for distinguishing between the first and all responses is the following: Association experiments differ with respect to collecting only the first associate response vs. several associate responses. A reason towards a discrete procedure is partly due to concerns about association chain effects, i.e., that the nth response is associated to the

$(n-1)$th response rather than the stimulus, and that association chaining would contaminate the later responses (McEvoy and Nelson, 1982). For example, given a target word *tree*, a first response could be *leaf* and a second response could be *to float*, which is arguably more related to *leaf* than it is to the target word *tree*. On the other hand, there are cases where several responses provide a more complete picture of the stimulus meaning than a single response. For example, in most cases the first associate response to the target word *blood* is *red*. In order to obtain an extended description of the target word, several responses are required, cf. de Deyne and Storms (2008b). Furthermore, Schulte im Walde et al. (2008) provided a series of analyses where they distinguished first responses vs. several responses. They showed that the general picture provided by the analyses was similar; often, the results obtained with the first responses were stronger.

Tables 1 and 2 provide an example of our association norms, and list the 10 most frequent responses for the compound noun *Ahornblatt* 'maple leaf' and its constituents *Ahorn* 'maple' and *Blatt* 'leaf'. The two tables differ with respect to quantifying over only the first vs. all responses. If there are less than 10 different responses (as it is the case in Table 1), then only those responses that were provided are listed. Table 3 provides a second example over all association responses, for the compound noun *Fliegenpilz* 'fly agaric' and its constituents *Fliege* 'fly/bow tie' and *Pilz* 'mushroom'. Note that *Fliegenpilz* is less transparent than *Ahornblatt* (at least with respect to its modifier), so that the associations of the compound and the modifier differ more strongly. Note also that for the polysemous noun *Fliege*, we receive associations to both the animal sense 'fly' as well as the clothes sense 'bow tie'. Table 4 provides a third example over all association responses, for the compound noun *Schlittenhund* 'sledge dog' and its constituents *Schlitten* 'sledge' and *Hund* 'dog'. This example is interesting because even though the compound seems rather transparent with respect to both constituents, the associations provided for the head noun *Hund* are very different to those of the compound, as they largely refer to prototypical properties of a general domestic dog, while the associations to *Schlittenhund* rather point to specific properties and conditions of sledge dogs.

4.3. Morpho-Syntactic Analysis

The description of the association norms is completed by a morpho-syntactic analysis of the response tokens that distinguishes and quantifies the part-of-speech categories of the associate responses: Each response to the stimuli was assigned its – possibly ambiguous – part-of-speech (*pos*). As resource for the pos assignment we relied on a lemmatised and pos-tagged frequency list of the *sdeWaC* corpus (Faaß et al., 2010), a cleaned version of the German web corpus *deWaC* created by *WaCky* (Baroni and Kilgarriff, 2006). We disregarded fine-grained distinctions such as case, number and gender features and considered only the major categories verb (V), noun (N), and adjective (ADJ). A fourth category 'OTHER' comprises all other part-of-speech categories such as adverbs, prepositions, particles, interjections, conjunctions, etc. Ambiguities between the categories arose e.g. when the experiment participant could have been referring either to a (capitalised) adjective or a noun, such as *Fett* 'fat'.

Having assigned part-of-speech tags to the responses, we were able to distinguish and quantify the morpho-syntactic categories of the responses. In non-ambiguous situations, the unique part-of-speech received the total stimulus-response frequency. For example, *Sofa* 'sofa' was provided as response to the stimulus *Wohnzimmer* 'living room' by 11 participants. Our pos resource saw *Sofa* in the corpus only as a noun. So *Wohnzimmer* received a contribution of all 11 mentions for noun pos. In ambiguous situations, the stimulus-response frequency was split over the possible part-of-speech tags according to the pos proportions. For example, *Fett* 'fat' was provided as response to the stimulus *Bratpfanne* 'frying pan' by four participants. Our pos resource saw *Fett* 10,780 in the corpus as a noun, and 493 times as an adjective. So *Bratpfanne* received a contribution of $\frac{4 \times 10{,}780}{10{,}780+493}$ nouns, and of $\frac{4 \times 493}{10{,}780+493}$ adjectives.

The output of this analysis is frequency distributions of the part-of-speech tags for each stimulus individually, and also as a sum over all stimuli. Table 5 presents the total numbers both over only the first and all associate responses, and for specific examples. Participants provided noun associates in the clear majority of token instances, 74%/71% with respect to first/all responses; adjectives were given in 12%/14% of the responses, and verbs in 10%/11%. The table also shows that the pos distributions vary across specific stimuli. For example, *Dose* 'can' and *Notenschlüssel* 'clef' both received very large proportions of nouns. In the first case, this was mainly due to associations referring to either the material (of the can), e.g., *Metall* 'metal', and *Plastik* 'plastic'; or the content, e.g., *Suppe* 'soup', *Katzenfutter* 'cat food', and *Mais* 'maize'. In the latter case, this was mainly due to participants providing associations referring to *Musik* 'music', or to an instrument, e.g., *Klavier* 'piano', and *Gitarre* 'guitar'. *Faden* 'thread' and *Türklinke* 'door handle' are two example stimuli that received very large proportions of verbs. In both cases, this was mainly due to associations referring to typical functions, such as *nähen* 'sew', and *binden* 'tie' for *Faden*, and *drücken* 'press', and *öffnen* 'open' for *Türklinke* (where the former is actually metonymic because it refers to the door rather than the door handle). *Zitrone* 'lemon' and *Wollschal* 'woollen scarf' are two example stimuli that received very large proportions of adjectives. In both cases, this was mainly due to associations referring to typical properties, such as *sauer* 'sour' and *gelb* 'yellow' for *Zitrone*, and such as *warm* 'warm' and *weich* 'soft' for *Wollschal*.

5. Case Study: Compositionality of German Noun Compounds

Addressing the compositionality of multi-word expressions is a crucial ingredient for lexicography and NLP applications, to know whether the expression should be treated as a whole, or through its constituents, and what the expression means. Only recently work on exploring distributional models on compositionality have emerged, such as Mitchell and Lapata (2010), Reddy et al. (2011), Baroni et al. (2012).

Ahornblatt ‘maple leaf’			*Ahorn* ‘maple’			*Blatt* ‘leaf’		
Kanada	‘Canada’	4	*Baum*	‘tree’	7	*Papier*	‘paper’	4
Baum	‘tree’	4	*Sirup*	‘syrup’	5	*Baum*	‘tree’	2
Herbst	‘autumn’	4	*Blatt*	‘leaf’	4	*grün*	‘green’	2
rot	‘red’	2	*Kanada*	‘Canada’	3	*schreiben*	‘write’	1
Klee	‘clover’	1	*Biologie*	‘biology’	1	*Blüte*	‘blossom’	1
Verfärbung	‘discolouring’	1	*Samen*	‘seed’	1	*Wald*	‘forest’	1
zackig	‘jagged’	1	*Ahornsirup*	‘maple syrup’	1	*Käfer*	‘bug’	1
						fallen	‘fall’	1
						weiß	‘white’	1

Table 1: Association frequencies for example compound *Ahornblatt* and its constituents (first associations).

Ahornblatt ‘maple leaf’			*Ahorn* ‘maple’			*Blatt* ‘leaf’		
Kanada	‘Canada’	8	*Baum*	‘tree’	14	*Baum*	‘tree’	10
Baum	‘tree’	7	*Sirup*	‘syrup’	11	*Papier*	‘paper’	8
Herbst	‘autumn’	7	*Kanada*	‘Canada’	9	*schreiben*	‘write’	4
Sirup	‘syrup’	4	*Blatt*	‘leaf’	6	*grün*	‘green’	4
rot	‘red’	3	*Blätter*	‘leaves’	4	*Herbst*	‘autumn’	2
Wald	‘forest’	2	*rot*	‘red’	2	*Blume*	‘flower’	2
Form	‘shape’	2	*Ahornsirup*	‘maple syrup’	2	*Käfer*	‘bug’	1
bunt	‘colourful’	2	*grün*	‘green’	2	*fallen*	‘fall’	1
zackig	‘jagged’	2	*Herbst*	‘autumn’	1	*Blattadern*	‘leaf veins’	1
Ahornsirup	‘maple syrup’	2	*Kindheit*	‘childhood’	1	*weiß*	‘white’	1

Table 2: Association frequencies for example compound *Ahornblatt* and its constituents (all associations).

Fliegenpilz ‘fly agaric’			*Fliege* ‘fly/bow tie’			*Pilz* ‘mushroom’		
giftig	‘poisonous’	12	*nervig*	‘annoying’	4	*Wald*	‘forest’	13
rot	‘red’	7	*summen*	‘buzz’	2	*Fliegenpilz*	‘fly agaric’	4
Wald	‘forest’	5	*lästig*	‘annoying’	2	*sammeln*	‘collect’	3
Gift	‘poison’	2	*Insekt*	‘bug’	2	*giftig*	‘poisonous’	3
Hut	‘cap’	1	*Tier*	‘animal’	2	*Schimmel*	‘mould’	2
Glück	‘fortune’	1	*Fliegenklatsche*	‘fly flap’	2	*Suche*	‘search’	2
Kinderbuch	‘children’s book’	1	*Krawatte*	‘tie’	2	*Hut*	‘cap’	2
Pflanze	‘plant’	1	*Sommer*	‘summer’	2	*Pilzpfanne*	‘mushroom pan’	2
Muster	‘pattern’	1	*Anzug*	‘suit’	1	*essbar*	‘eatable’	1
weiß	‘white’	1	*fangen*	‘catch’	1	*Suppe*	‘soup’	1

Table 3: Association frequencies for example compound *Fliegenpilz* and its constituents (all associations).

Schlittenhund ‘sledge dog’			*Schlitten* ‘sledge’			*Hund* ‘dog’		
Schnee	‘snow’	10	*Schnee*	‘snow’	12	*bellen*	‘bark’	9
Husky	‘husky’	8	*Winter*	‘winter’	11	*Katze*	‘cat’	7
Winter	‘winter’	5	*rodeln*	‘luge’	4	*Halsband*	‘collar’	2
schnell	‘fast’	3	*kalt*	‘cold’	3	*Leine*	‘leash’	2
Schlitten	‘sledge’	3	*fahren*	‘drive’	3	*freundlich*	‘friendly’	1
kalt	‘cold’	3	*rutschen*	‘slide’	2	*Haustier*	‘domestic animal’	1
treu	‘faithful’	1	*Spaß*	‘fun’	2	*Schwanz*	‘tail’	1
rennen	‘run’	1	*Glöckchen*	‘little bell’	1	*Hundehütte*	‘doghouse’	1
Glöckchen	‘little bell’	1	*Hunde*	‘dogs’	1	*wedeln*	‘wag’	1
Skandinavien	‘Scandinavia’	1	*Schlittenfahren*	‘tobogganing’	1	*Freund*	‘friend’	1

Table 4: Association frequencies for example compound *Schlittenhund* and its constituents (all associations).

	N	ADJ	V	OTHER
FIRST ASSOCIATIONS:				
TOTAL FREQ	12,659	2,124	1,726	624
TOTAL PROB	74%	12%	10%	4%
ALL ASSOCIATIONS:				
TOTAL FREQ	33,322	6,835	5,264	1,827
TOTAL PROB	71%	14%	11%	4%
Dose 'can'	92%	3%	5%	0%
Notenschlüssel 'clef'	96%	2%	0%	3%
Faden 'thread'	43%	10%	44%	3%
Türklinke 'door handle'	54%	3%	38%	5%
Zitrone 'lemon'	20%	74%	3%	3%
Wollschal 'woollen scarf'	37%	49%	13%	1%

Table 5: Part-of-speech tag distributions of responses.

Compound	Modifier	Head	Modifier Scores		Head Scores	
			system	human	system	human
Ahornblatt 'maple leaf'	*Ahorn* 'maple'	*Blatt* 'leaf'	.69	5.63	.35	5.70
Badeanzug 'bathing costume'	*baden* 'bath'	*Anzug* 'suit'	.68	6.13	.00	3.03
Blockflöte 'flute'	*Block* 'block'	*Flöte* 'flute'	.14	1.73	.65	5.38
Buntstifte 'crayons'	*bunt* 'colourful'	*Stifte* 'pen'	.26	4.80	.68	5.93
Erdnuss 'peanut'	*Erde* 'earth/soil'	*Nuss* 'nut'	.00	2.87	.32	6.57
Feuerwerk 'fireworks'	*Feuer* 'fire'	*Werk* 'opus'	.02	4.20	.02	2.80
Fliegenpilz 'fly agaric'	*Fliege* 'fly/bow tie'	*Pilz* 'mushroom'	.00	1.93	.47	6.55
Kochtopf 'cooking pot'	*kochen* 'cook'	*Topf* 'pot'	.49	5.03	.61	6.52
Lesebrille 'reading glasses'	*lesen* 'read'	*Brille* 'glasses'	.02	4.93	.32	5.97
Mülleimer 'dustbin'	*Müll* 'rubbish'	*Eimer* 'bucket'	.34	5.50	.18	5.23
Nähmaschine 'sewing machine'	*nähen* 'sew'	*Maschine* 'machine'	.16	6.03	.00	4.93
Obstkuchen 'fruit cake'	*Obst* 'fruit'	*Kuchen* 'cake'	.15	4.80	.29	5.93
Rotkohl 'red cabbage'	*rot* 'red'	*Kohl* 'cabbage'	.03	2.70	.30	5.83
Schlagsahne 'whipping cream'	*schlagen* 'beat'	*Sahne* 'cream'	.00	3.67	.68	6.77
Schlittenhund 'sledge dog'	*Schlitten* 'sledge'	*Hund* 'dog'	.44	4.30	.02	5.33
Weihnachtsbaum 'christmas tree'	*Weihnachten* 'christmas'	*Baum* 'tree'	.16	5.70	.13	5.10

Table 6: Compositionality scores for modifiers and heads (system vs. human).

Concerning German, little effort has been done to address compound compositionality by computational means. One of the few existing approaches is Zinsmeister and Heid (2004) who suggested a distributional model that determined the compositionality of German noun compounds by comparing the verbs that subcategorise the noun compound vs. the head noun as direct objects.

The goal of our case study is to demonstrate that associations provide insight into properties of compounds (and their constituents) that should be useful for computational models of compound compositionality. For that purpose, the case study relies on a simple association overlap measure to predict the degree of compositionality of the experiment compound nouns: We use the proportion of shared associations of the compound and a constituent with respect to the total number of associations of the compound. The degree of compositionality of a compound noun is calculated with respect to each constituent of the compound. We plan to work on more elaborate predictions of compound compositionality, and consider the current case study as a baseline to future computational experiments.

As an example of the calculation, when considering the 10 most frequent responses of the compound noun *Ahornblatt* 'maple leaf' and its constituents, as provided in Table 2, the compound noun received a total of 39 associations, out of which it shares 31 with the first constituent *Ahorn* 'maple', and 14 with the second constituent *Blatt* 'leaf'. Thus, the predicted degrees of compositionality are $\frac{31}{39} = 0.79$ for *Ahornblatt–Ahorn*, and $\frac{14}{39} = 0.36$ for *Ahornblatt–Blatt*. The predicted degrees of compositionality are compared against the mean compositionality judgements as collected by von der Heide and Borgwaldt (2009), using the Spearman rank-order correlation coefficient. This correlation is a non-parametric statistical test that measures the association between two variables that are ranked in two ordered series. The resulting correlation is $r_s = 0.5228, p < .000001$, which we consider a surprisingly successful result concerning our simple measure.

Table 6 presents a number of example compounds and their constituents, accompanied by the predicted degrees of compositionality as well as the mean human judgement scores, with respect to both constituents. We picked a mixture of examples to illustrate good matches of system and human scores (e.g., *Blockflöte*), bad matches (e.g., *Nähmaschine*) as well as moderate matches (e.g., *Ahornblatt/Blatt*, and *Obstkuchen/Obst*).

6. Summary

This paper introduced association norms of German noun compounds as a lexical-semantic resource for cognitive and computational linguistics research on compositionality. Based on an existing database of German noun compounds, we collected human associations to the compounds and their constituents within a web experiment. In addition to describing the collection, quantification and part-of-speech analysis of the association norm, we provided a case study on the compositionality of the noun compounds, as relying on the norms, and reached a Spearman rank correlation coefficient of $r_s = 0.5228, p < .000001$ to predict the degree of compositionality of the compounds.

7. Acknowledgements

Sabine Schulte im Walde is funded by the DFG Heisenberg Fellowship SCHU-2580/1-1.

8. References

H. Baayen and R. Schreuder. 1999. War and peace: Morphemes and full forms in a noninteractive activation parallel dual–route model. *Brain and Language*, 68:27–32.

Marco Baroni and Adam Kilgarriff. 2006. Large Linguistically-processed Web Corpora for Multiple Languages. In *Proceedings of the 11th Conference of the European Chapter of the Association for Computational Linguistics*, Trento, Italy.

Marco Baroni, Raffaella Bernardi, Ngoc-Quynh Do, and Chung chieh Shan. 2012. Entailment above the Word Level in Distributional Semantics. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, Avignon, Francs.

Susanne Borgwaldt, Catherine-Marie Longtin, Rachel Kemps, and Gary Libben. 2005. Semantic Transparency Ratings and Associations to English Compounds. Unpublished raw data.

Alfonso Caramazza, Alessandro Laudanna, and Cristina Romani. 1988. Lexical Access and Inflectional Morphology. *Cognition*, 28:297–332.

Kenneth W. Church and Patrick Hanks. 1989. Word Association Norms, Mutual Information, and Lexicography. In *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics*, pages 76–83, Vancouver, Canada.

Herbert H. Clark. 1971. Word Associations and Linguistic Theory. In John Lyons, editor, *New Horizon in Linguistics*, chapter 15, pages 271–286. Penguin.

Simon de Deyne and Gert Storms. 2008a. Word associations: Network and semantic properties. *Behavior Research Methods*, 40(1):213–231.

Simon de Deyne and Gert Storms. 2008b. Word associations: Norms for 1,424 dutch words in a continuous task. *Behavior Research Methods*, 40(1):198–205.

Gertrud Faaß Ulrich Heid, and Helmut Schmid. 2010. Design and Application of a Gold Standard for Morphological Analysis: SMOR in Validation. In *Proceedings of the 7th International Conference on Language Resources and Evaluation*, pages 803–810, Valletta, Malta.

Christiane Fellbaum and Roger Chaffin. 1990. Some Principles of the Organization of Verbs in the Mental Lexicon. In *Proceedings of the 12th Annual Conference of the Cognitive Science Society of America*.

Christiane Fellbaum. 1995. Co-Occurrence and Antonymy. *Lexicography*, 8(4):281–303.

Ana Fernández, Emiliano Diez, María Angeles Alonso, and María Soledad Beato. 2004. Free-Association Norms for the Spanisch Names of the Snodgrass and Vanderwart Pictures. *Behavior Research Methods, Instruments and Computers*, 36(3):577–583.

Ludovic Ferrand and F.-Xavier Alario. 1998. French Word Association Norms for 366 Names of Objects. *L'Annee Psychologique*, 98(4):659–709.

Wolfgang Fleischer and Irmhild Barz. 2012. *Wortbildung der deutschen Gegenwartssprache*. de Gruyter.

Annamaria Guida. 2007. The Representation of Verb Meaning within Lexical Semantic Memory: Evidence from Word Associations. Master's thesis, Universit degli studi di Pisa.

Hans Jürgen Heringer. 1986. The Verb and its Semantic Power: Association as the Basis for Valence. *Journal of Semantics*, 4:79–99.

Katherine W. Hirsh and Jeremy Tree. 2001. Word Association Norms for two Cohorts of British Adults. *Journal of Neurolinguistics*, 14(1):1–44.

George R. Kiss, Christine Armstrong, Robert Milroy, and James Piper. 1973. An Associative Thesaurus of English and its Computer Analysis. In *The Computer and Literary Studies*. Edinburgh University Press.

Max Lauteslager, Theo Schaap, and Dick Schievels. 1986. *Schriftelijke Woordassociatienormen voor 549 Nederlandse Zelfstandige Naamworden*. Swets and Zeitlinger.

Rochelle Lieber and Pavol Stekauer. 2009a. Introduction: Status and Definition of Compounding. In *The Oxford Handbook on Compounding* (Lieber and Stekauer, 2009b), chapter 1, pages 3–18.

Rochelle Lieber and Pavol Stekauer, editors. 2009b. *The Oxford Handbook of Compounding*. Oxford University Press.

Catherine-Marie Longtin, Juan Segui, and Pierre A. Halle. 2003. Morphological Priming without Morphological Relationship. *Language and Cognitive Processes*, 18:313–334.

William D. Marslen-Wilson, Lorraine K. Tyler, Rachelle Waksler, and Lianne Older. 1994. Morphology and Meaning in the English Mental Lexicon. *Psychological Review*, 101:3–33.

Cathy L. McEvoy and Douglas L. Nelson. 1982. Category Name and Instance Norms for 106 Categories of Various Sizes. *American Journal of Psychology*, 95:581–634.

Timothy P. McNamara. 2005. *Semantic Priming: Perspectives from Memory and Word Recognition*. Psychology Press, New York.

Alissa Melinger and Andrea Weber. 2006. Database of Noun Associations for German. URL: www.coli.uni-saarland.de/projects/nag/.

Alissa Melinger, Sabine Schulte im Walde, and Andrea Weber. 2006. Characterizing Response Types and Revealing Noun Ambiguity in German Association Norms. In *Proceedings of the EACL Workshop 'Making Sense of Sense: Bringing Computational Linguistics and Psycholinguistics together'*, pages 41–48, Trento, Italy.

Jeff Mitchell and Mirella Lapata. 2010. Composition in Distributional Models of Semantics. *Cognitive Science*, 34:1388–1429.

Douglas L. Nelson, Cathy L. McEvoy, and Thomas A. Schreiber. 1998. The University of South Florida Word Association, Rhyme, and Word Fragment Norms.

D.S. Palermo and James J. Jenkins. 1964. *Word Association Norms: Grade School through College*. University of Minnesota Press, Minneapolis.

Reinhard Rapp. 2002. The Computation of Word Associations: Comparing Syntagmatic and Paradigmatic Approaches. In *Proceedings of the 19th International Conference on Computational Linguistics*, Taipei, Taiwan.

Siva Reddy, Ioannis P. Klapaftis, Diana McCarthy, and Suresh Manandhar. 2011. Dynamic and Static Prototype Vectors for Semantic Composition. In *Proceedings of the 5th International Joint Conference on Natural Language Processing*, pages 705–713, Chiang Mai, Thailand.

Wallace A. Russell and O.R. Meseck. 1959. Der Einfluss der Assoziation auf das Erinnern von Worten in der deutschen, französischen und englischen Sprache. *Zeitschrift für Experimentelle und Angewandte Psychologie*, 6:191–211.

Wallace A. Russell. 1970. The complete German Language Norms for Responses to 100 Words from the Kent-Rosanoff Word Association Test. In Leo Postman and Geoffrey Keppel, editors, *Norms of Word Association*, pages 53–94. Academic Press, New York.

Dominiek Sandra. 1990. On the Representation and Processing of Compound Words: Automatic Access to Constituent Morphemes does not occur. *The Quarterly Journal of Experimental Psychology*, 42A:529–567.

Friederike Schlaghecken and Jens Bölte. 1998. Assoziationsklassen bei Basis- und Oberbegriffen. *Sprache und Kognition*, 17(1–2):21–40.

Sabine Schulte im Walde and Alissa Melinger. 2008. An In-Depth Look into the Co-Occurrence Distribution of Semantic Associates. *Italian Journal of Linguistics.*, 20(1):89–128.

Sabine Schulte im Walde, Alissa Melinger, Michael Roth, and Andrea Weber. 2008. An Empirical Characterisation of Response Types in German Association Norms. *Research on Language and Computation*, 6(2):205–238. DOI 10.1007/s11168-008-9048-4.

Sabine Schulte im Walde. 2008. Human Associations and the Choice of Features for Semantic Verb Classification. *Research on Language and Computation*, 6(1):79–111. DOI 10.1007/s11168-008-9044-8.

Donald P. Spence and Kimberly C. Owens. 1990. Lexical Co-Occurrence and Association Strength. *Journal of Psycholinguistic Research*, 19:317–330.

Marcus Taft and Kenneth I. Forster. 1975. Lexical Storage and Retrieval of Prefixed Words. *Journal of Verbal Learning and Verbal Behavior*, 14:638–648.

Marcus Taft. 2004. Morphological Decomposition and the Reverse Base Frequency Effect. *The Quarterly Journal of Experimental Psychology*, 57:745–765.

Claudia von der Heide and Susanne Borgwaldt. 2009. Assoziationen zu Unter-, Basis- und Oberbegriffen. Eine explorative Studie. In *Proceedings of the 9th Norddeutsches Linguistisches Kolloquium*, pages 51–74.

Heike Zinsmeister and Ulrich Heid. 2004. Collocations of Complex Nouns: Evidence for Lexicalisation. In *Proceedings of Konvens*, Vienna, Austria.

Pienie Zwitserlood. 1994. The Role of Semantic Transparency in the Processing and Representation of Dutch Compounds. *Language and Cognitive Processes*, 9:341–368.