

Modeling morphological priming in German with naive discriminative learning

R. Harald Baayen¹ & Eva Smolka²

^a *Department of Linguistics, University of Tübingen, harald.baayen@uni-tuebingen.de*

^b *Department of Linguistics, University of Konstanz, eva.smolka@uni-konstanz.de*

Abstract

Both localist and connectionist models, based on experimental results obtained for English and French, assume that the degree of semantic compositionality of a morphologically complex word is reflected in how it is processed. Since priming experiments using English and French morphologically related prime-target pairs reveal stronger priming when complex words are semantically transparent (e.g., *refill–fill*) compared to semantically more opaque pairs (e.g., *restrain–strain*), localist models set up connections between complex words and their stems only for semantically transparent pairs. Connectionist models have argued that the effect of transparency should arise as an epiphenomenon in PDP networks. However, for German, a series of studies has revealed equivalent priming for both transparent and opaque prime-target pairs, which suggests mediation of lexical access by the stem, independent of degrees of semantic compositionality.

This study reports a priming experiment that replicates equivalent priming for transparent and opaque pairs. We show that these behavioral results can be straightforwardly modeled by a computational implementation of Word and Paradigm Morphology (WPM), Naive Discriminative Learning (NDL). Just as WPM, NDL eschews the theoretical construct of the morpheme. NDL succeeds in modeling the German priming data by inspecting the extent to which a discrimination network pre-activates the target lexeme from the orthographic properties of the prime. Measures derived from an NDL network, complemented with a semantic similarity measure derived from distributional semantics, predict lexical decision latencies with somewhat improved precision compared to classical measures such as word frequency, prime type, and human association ratings. We discuss both the methodological implications of our results, as well as their implications for models of the mental lexicon.

1 Introduction

Current mainstream models of lexical processing assume that complex words such as *developmental* comprise several morphemic constituents, *de-*, *develop*, *-ment*, and *-al*, that recur in the language in many other words. Since early research in the seventies (e.g., Taft and Forster, 1975), it has been argued that the recognition of morphologically complex words is mediated by such morphemic units (for a review of models of morphological processing see Milin et al., 2017b).

One of the issues under investigation in this line of research is whether visual input is automatically decomposed into morphemes before semantics is accessed. Several studies have argued in favor of early morpho-orthographic decomposition (Longtin et al., 2003; Rastle et al., 2004; Rastle and Davis, 2008), but others argue that semantics is involved from the start (Feldman et al., 2009), that the effect is task dependent and is limited to the lexical decision task (Dunabeitia et al., 2011; Marelli et al., 2013; Norris and Kinoshita, 2008), or fail to replicate experimental results central to decompositional accounts (Milin et al., 2017a).

Another issue that is still unresolved is whether complex words are potentially accessed through two routes operating in parallel, one involving decomposition and the other whole-form based retrieval (Marslen-Wilson et al., 1994b; Frauenfelder and Schreuder, 1992; Baayen et al., 1997, 2003).

Recent investigations that make use of survival analysis actually suggest that whole-word based effects precede in time constituent-based effects (Schmidtke et al., 2017).

A third issue concerns the role of semantic transparency. Priming studies conducted on English and French prefixed derivations that are semantically transparent, such as *distrust*, have reported facilitation of the recognition of their stems (*trust*), as well as other prefixed or suffixed derivations such as *entrust* or *trustful*. The same holds for suffixed derivations that are semantically transparent like *production* and *productivité* in French or *confession* and *confessor* in English, which prime each other and their stem (*confess*). The critical condition in this discussion, however, concerns semantically opaque (i.e., noncompositional) derivations such as *successor*, which appear not to facilitate the recognition of stems like *success*. This latter finding was replicated under auditory prime presentations or visual priming at long exposure durations at 230 ms or 250 ms (e.g., Feldman and Prostko, 2001; Feldman et al., 2004; Lavric et al., 2011; Meunier and Longtin, 2007; Pastizzo and Feldman, 2002; Rastle et al., 2000). Localist accounts take these findings to indicate that only semantically transparent complex words are processed decompositionally, via their stem, while semantically opaque words are processed as whole word units. Although on different grounds, also distributed connectionist approaches assume that the facilitation between complex words and their stem depends on their meaning relation. In a series of cross-modal priming experiments, Gonnerman et al. (2007) showed for English that morphological effects vary according to the gradual overlap of form and meaning between word pairs. Indeed, word pairs with a strong phonological and semantic relation like *preheat-heat* induced stronger priming than words with a moderate phonological and semantic relation like *midstream-stream*, and words holding a low semantic relation like *rehearse-hearse* induced no priming at all. According to connectionist accounts of lexical processing, this result arises as the consequence of the extent to which orthographic, phonological, and semantic codes converge.

However, these findings for English and French contrast with results repeatedly obtained for German, where morphological priming appears to be unaffected by semantic transparency (Smolka et al., 2015, 2009, 2019, 2014). Under auditory or overt visual prime presentations, morphologically related complex verbs facilitated the recognition of their stem regardless of whether they were semantically transparent (*aufstehen-stehen*, ‘stand up’-‘stand’) or opaque (*verstehen-stehen*, ‘understand’-‘stand’). Smolka and colleagues interpreted these findings to indicate that a German native speaker processes a complex verb like *verstehen* by accessing the stem *stehen* irrespective of the whole-word meaning, and argued that morphological structure overrides meaning in the lexical processing of German complex words. To account for such stem effects without effects of semantic transparency, they hypothesized a model in which the frequency of stem is the critical factor, such that stems of complex words are accessed and activated, independent of the meaning composition of the complex word.

These findings for German receive support from experiments on Dutch — a closely related language with a highly similar system of verbal prefixes, separable particles, and non-separable particles. Work by Schreuder and colleagues (Schreuder et al., 1990), using an intramodal visual short SOA partial priming technique to study Dutch particle verbs, revealed morphological effects without modulation by semantic transparency. Experiments addressing speech production in Dutch (Roelofs, 1997a,b; Roelofs et al., 2002) likewise observed, using the implicit priming task, that priming effects were equivalent for transparent and opaque prime-target pairs. Morphological priming without effects of semantic transparency have recently been replicated in Dutch under overt prime presentations (De Grauwe et al., 2019; Cremeer et al., 2019). Unprimed and primed visual lexical decision experiments on Dutch low-frequency suffixed words with high-frequency base words revealed that the semantics of opaque complex words were equally quickly available as the semantics of transparent complex words (Schreuder et al., 2003), contradicting the original prediction of

this study that transparent words would show a processing advantage compared to their opaque counterparts.

Importantly, there are some studies in English, e.g., [Gomerman et al. \(2007, Exp. 4\)](#) and [Marslen-Wilson et al. \(1994a, Exp. 5\)](#), that applied a similar cross-modal priming paradigm with auditory primes and visual targets, and with similar prefixed stimuli as in the abovementioned studies by Smolka and collaborators, but found no priming for semantically opaque pairs like *rehearse-hearse* (for similar ERP-results in English see [Kielar and Joanisse, 2011](#)). Thus, results for German and results for English appear at present to be genuinely irreconcilable.¹

In what follows, we first present an overt visual priming experiment that provides further evidence for the equivalent facilitation effects seen for German transparent and opaque prime-target pairs. The behavioral results are consistent with localist models in which connections between stems and derived words are hand-wired into a network, as argued by, e.g., [Smolka et al. \(2007, 2009, 2014, 2015\)](#); [Smolka and Eulitz \(2018\)](#). However, this localist model is a post hoc description of the experimental findings, and a computational implementation for this high-level theory is not available.

In this study, we proceed to show that the observed stem priming effects can be straightforwardly modeled by naive discriminative learning (NDL, [Baayen et al., 2011](#); [Milin et al., 2017a](#); [Baayen et al., 2016b,a](#); [Divjak et al., 2017](#); [Sering et al., 2018b](#); [Arnold et al., 2017](#); [Tomaschek et al., 2019](#)) without reference to stems or other morphological units, and without requiring hand-crafting of connections between such units. In fact, measures derived from an NDL network, complemented with a semantic similarity measure derived from distributional semantics, turn out to predict lexical decision latencies with greater precision compared to classical measures such as word frequency, prime type, and semantic association ratings. Importantly, the NDL model predicts the effects of stem priming without a concomitant effect of semantic compositionality. According to the NDL model, the crucial predictor is the extent to which the target is pre-activated by the sublexical form features of the prime. In the final section, we discuss both the methodological implications of our results, as well as their implications for models of the mental lexicon.

2 The priming experiment

2.1 Previous results for German complex verbs

German complex verbs present a very useful means to study the effects of morphological structure with or without meaning relatedness to the same base verb. German complex verbs are very productive and frequently used in standard German. The linguistic literature ([Eisenberg, 2004](#); [Fleischer and Barz, 1992](#)) distinguishes two word formations: prefix verbs and particle verbs. Both consist of a verbal root and either a verbal prefix or a particle.

In spite of some prosodic and morphosyntactic differences (see [Smolka et al., 2019](#)), prefix and particle verbs share many similar semantic properties. Both may differ in the degree of semantic transparency with respect to the meaning of their base. For example, the particle *an* ('at') only slightly alters the meaning of the base *führen* ('guide') in the derivation *anführen* ('lead'), but radically does so with respect to the base *schicken* ('send') in the opaque derivation *anschicken* ('get ready'). Similarly, the prefix *ver-* produces the transparent derivation *verschicken* ('mail') as well as the opaque derivation *verführen* ('seduce'). Prefix and particle verbs are thus a particularly useful means by which the effects of meaning relatedness to the same base verb can be studied. For instance, derivations of the base *tragen* ('carry'), such as *hintragen* ('carry to'), *forttragen*

¹ See [Günther et al. \(2018\)](#) and [Smolka et al. \(2014\)](#) for possible explanations.

(‘carry away’), *zurücktragen* (‘carry back’), *abtragen* (‘carry off’), *auftragen* (‘apply’), *vertragen* (‘get along’), *ertragen* (‘suffer’), alter the meaning relatedness from fully transparent to fully opaque with respect to the base. It is important to note that in general complex verbs in German are true etymological derivations of their base, regardless of the degree of semantic transparency they share with it. Because morphological effects of prefix and particle verbs are alike in German ((see [Smolka and Eulitz, 2018](#); [Smolka et al., 2019](#)) and Dutch ([Schriefers et al., 1991](#)), henceforth, we refer to them as ‘complex verbs’ or ‘derived verbs’.

Previous findings on complex verbs in German have shown that these verbs strongly facilitate the recognition of their stem, without any effect of semantic transparency ([Smolka et al., 2015, 2009](#); [Smolka and Eulitz, 2018](#); [Smolka et al., 2014, 2019](#)). That is, semantically opaque verbs such as *verstehen* (‘understand’) primed their base *stehen* (‘stand’) to the same extent as did transparent verbs such as *aufstehen* (‘stand up’). Further, the priming by both types of morphological primes was stronger than that by either purely semantically related primes like *aufspringen* (‘jump up’) or purely form-related primes like *bestehlen* (‘steal’). The morphological effects remained unaffected by semantic transparency under conditions that were sensitive to detecting semantic and form similarity, that is, when semantic controls like *verlangen–fordern* (‘require’–‘demand’) and *Biene–Honig* (‘bee’–‘honey’) induced semantic facilitation or when form-controls with embedded stems, as in *bekleiden–leiden* (‘dress’–‘suffer’) and *Bordell–Bord* (‘brothel’–‘board’) induced form inhibition (see Exp. 3 in [Smolka et al., 2014](#)). This offered assurance that the lack of a semantic transparency effect between semantically transparent and opaque complex verbs was not a null effect but rather indicated that morphological relatedness overrides both semantic and form relatedness.

Further studies explored the circumstances of stem facilitation in more detail. For example, in spite of several differences in the phonological and morpho-syntactic properties of prefix and particle verbs, prefix verbs showed processing patterns that were substantially the same as those for particle verbs and, crucially, were uninfluenced by semantic transparency ([Smolka et al., 2019](#)). Furthermore, stem access occurs regardless of the directionality of prime and target *entwerfen–werfen* vs. *werfen–entwerfen* vs. *entwerfen–bewerfen* ([Smolka and Eulitz, 2011](#)).

Stem access is modality independent, as it occurs under both intra-modal (visual-visual) and cross-modal (auditory-visual) priming conditions ([Smolka et al., 2014, 2019](#)). Finally, event-related brain potentials revealed wide-spread N400 brain potentials in response to semantically transparent and opaque verbs without effects of semantic transparency — N400 brain potentials that are generally taken to be characteristic to indicate expectancy and (semantic) meaning integration. Most importantly, these brain potentials revealed that stem facilitation in German occurs without an overt behavioral response and is stronger than the activation by purely semantically related verbs or form-related verbs ([Smolka et al., 2015](#)).

The present experiment was closely modeled after previous experiments addressing priming effects for German verb pairs (e.g., [Smolka et al., 2009, 2014](#)).

2.2 Design

We compared the differential effects of semantic, form, or morphological relatedness between complex verbs and a base verb in four priming conditions: (a) semantic condition, where the complex verb was a synonym of the target verb, (b) morphological transparent condition, where the complex verb was a semantically transparent derivation of the target verb, (c) morphological opaque condition, where the complex verb was a semantically opaque derivation of the target verb, and (d) form condition, where the base of the complex verb was form-related with the base of the target. We measured lexical decision latencies to the target verbs and calculated priming relative to an unrelated control condition. In addition to the (unrelated minus related) priming effects, the

influence of the stem should surface in the comparison of the conditions (a) and (b), where both types of primes are synonyms of the base verb – the former holding a different stem as the target, the latter holding the same stem as the target; the influence of the degree of semantic transparency should surface in the comparison between conditions (b) and (c), where both types of primes are true morphological derivations of the base target. The influence of form similarity should surface in the comparison between conditions (b) and (d), where both types of primes have stems that are form-similar with the target.

As in our previous experiments, we used overt visual priming conditions to tap into lexical processing. We used only verbs as materials to avoid word category effects, and inserted a large number of fillers to prevent expectancy or strategic effects. Different from our previous experiments, though, we applied a between-subject and between-target design.

In summary, the primes in all conditions were complex verbs with the same morphological structure and were thus (a) of the same word category, and (b) closely matched on distributional variables like lemma frequency, number of syllables and letters. They differed only with respect to the morphological, semantic, or form-relatedness with the target. Prime conditions are exemplified in Table 1; all critical items are listed in the Appendix. Our prediction is that both semantically transparent and opaque complex verbs will induce the same amount of priming to their base, and that this priming will be stronger than the priming by either semantically related or form-related verbs.

2.3 Method

Participants Fifty students of the University of Konstanz participated in the experiment (14 males; mean age = 22.69, range 19 – 32). All were native speakers of German, were not dyslexic, and had normal or corrected-to-normal vision. They were paid for their participation.

Materials As critical stimuli, eighty-eight prime-target pairs with complex verbs as primes and base verbs as targets were selected from the CELEX German lexical database (Baayen et al., 1993), 22 pairs in each of four conditions (see also Table 1): (a) morphologically unrelated synonyms of the base (e.g., *vorangehen*–*führen*, ‘antecede’–‘guide’), (b) morphologically related synonyms of the base, these were semantically transparent derivations of the base, (e.g., *anführen*–*führen*, ‘head’–‘guide’), (c) semantically opaque derivations of the base (e.g., *verführen*–*führen*, ‘seduce’–‘guide’), and (d) morphologically unrelated form controls that kept the onset of the target’s base and changed a letter in the rime (e.g., *befühlen*–*führen*, ‘palpate’–‘guide’). Complex verbs in conditions (a) and (b) were synonyms of the target base and were selected by means of the online synonym dictionaries <http://www.canoo.net/> and <http://synonyme.woxikon.de/>.

For each of the 88 related primes, we selected an unrelated control that served as baseline and (a) was morphologically, semantically, and orthographically unrelated to the target and (b) matched the related prime in word class, morphological complexity (i.e., it was a complex verb), number of letters and syllables. In addition, control primes were pair-wise matched to the related primes on lemma frequency according to CELEX. Furthermore, primes across conditions were matched on lemma frequency according to CELEX.

The critical set of 88 prime-target pairs was selected from a pool of verb pairs that had been subjected to semantic association tests, in which participants rated the meaning relatedness between the verbs of each prime-target pair on a 7-point scale from completely unrelated (1) to highly related (7) (for a detailed description of the database see Smolka and Eulitz (2018)). The following criteria determined whether a verb pair was included in the critical set: The mean ratings for a semantically-related pair (in the synonym and semantically transparent conditions) had to be

higher than 4, and those for a semantically unrelated pair (in the semantically opaque and form-related conditions) lower than 3. The final set had mean ratings of 5.5 (range 4.3 – 6.7) for synonyms, 5.7 (range 4.78 – 6.56) for semantically transparent derivations, 2.13 (range 1.5 – 2.88) for semantically opaque derivations, and 1.7 (range 1.0 – 2.89) for form-related pairs. Table 1 provides the prime characteristics (lemma frequency, number of letters and syllables, semantic association); the Appendix lists all stimuli.

In order to prevent strategic effects, a total of 140 prime-target pairs was added as fillers. All had complex verbs as primes, 48 had verbs and 92 had pseudoverbs as targets. With respect to the former, 18 of the 48 prime-verb fillers comprised related prime-target pairs of the other lists. These types were included to assure that participants would not detect a certain type of prime-target relatedness in a list. For example, list A held six items of list B, six of list C, and six of list D as fillers. The other 30 prime-target pairs were semantically, morphologically and orthographically unrelated.

Regarding the prime-pseudoverb fillers, 44 of the pseudoverb targets were closely matched to the critical verb targets by keeping the onset of the verbs' first syllable (e.g. *binden* – *binken*). To further ensure that participants did not respond with 'word' decisions for any trial where prime and target were orthographically similar, eleven pseudoverbs were preceded by a form-related prime (e.g., *umwerben* – *wersen*) to mimic the form condition. All pseudoverbs were constructed by exchanging one or two letters in real verbs, while preserving the phonotactic constraints of German.

Overall, the large amount of fillers reduced the proportion of (a) critical prime-target pairs to 24% per list. All filler items differed from the critical items. Throughout the experiment, all primes and targets were presented in the infinitive (stem/-en), which is also the citation form in German.

Apparatus Stimuli were presented on a 18.1" monitor, connected to an IBM-compatible AMD Athlon 1.4 GHz personal computer. Stimulus presentation and data collection were controlled by the Presentation software developed by Neurobehavioral Systems (<http://nbs.neuro-bs.com/>). Response latencies were recorded from the left and right buttons of a push-button box.

Procedure Each participant saw only one list. Each list was divided into four blocks, each block containing the same amount of stimuli per condition. The critical prime-target pairs were rotated over the four blocks according to a Latin Square design in such a way that the related and unrelated primes of the same target were separated by a block. The related fillers (form-related prime-pseudoverb pairs, related prime-verb pairs) and unrelated filler pairs were evenly allocated to the blocks.

In total, an experimental session comprised 184 prime-target pairs, with 66 pairs per block. Within blocks, prime-target pairs were randomized separately for each participant. Twenty additional prime-target pairs were used as practice trials. Participants were tested individually in a dimly lit room, seated at a viewing distance of about 60 cm from the screen. Stimuli were presented in Sans-Serif letters on a black background. To ensure that primes and targets were perceived as physically distinct stimuli, primes were presented in uppercase letters, point 32, in light blue (RGB: 0-255-255), 20 points above the center of the screen. Targets were presented centrally in lowercase letters, point 36, in yellow (RGB: 255-255-35).

Each trial started with a fixation cross in the center of the screen for 300 ms. This was followed by the presentation of the prime for 400 ms, followed by an offset (i.e. a blank screen) for 100 ms, resulting in a stimulus onset asynchrony (SOA) of 500 ms. After the offset, the target immediately followed and remained on the screen until a participant's response. The intertrial interval was 1500 ms. Participants were instructed to make lexical decisions to the targets, as fast and as accurately as possible. 'Word' responses were given with the index finger of the dominant hand, 'pseudoword'

Table 1: Stimulus Characteristics of Related Primes and Their Matched Unrelated Controls in the Semantic Synonym List, Semantically Transparent List, Semantically Opaque List, and Form Control List. Statistics are given for the total set of stimuli: mean (SD; range); sample stimuli of prime-target pairs are italicized. Frequencies are from the CELEX database (Baayen et al., 1993), count is per million. Lists were between-subjects.

List	Relatedness	Lemma Frequency	Word Length	Syllable Length	Age of Acquisition	Relatedness Score
Semantic						
<i>vorangehen-führen</i>	Related	4.9	9.2	3.0	5.4	5.5
(‘antecede’-‘guide’)		(6.3; 0-24)	(1.2; 7-12)	(0.2; 3-4)	(1.3; 3.4-7.4)	(0.6; 4.3-6.6)
	Unrelated	4.7	9.1	3.0	7.1	
		(6.1; 0-24)	(1.1; 7-11)	(0; 3-3)	(1.7; 4.3-9.8)	
Transparent						
<i>anführen-führen</i>	Related	5.5	9.3	3.0	5.2	5.7
(‘head’-‘guide’)		(5.6; 0-17)	(1.5; 7-11)	(0, 3-3)	(1.2; 3.1-8.3)	(0.5; 4.8-6.6)
	Unrelated	5.6	9.2	3.0	6.5	
		(5.8; 0-20)	(1.1; 7-11)	(0; 3-3)	(1.7; 3.7-9.2)	
Opaque						
<i>verführen-führen</i>	Related	8.8	9.6	3.1	8.0	2.1
(‘seduce’-‘guide’)		(9.5; 0-34)	(1.4; 8-13)	(0.3; 3-4)	(1.7; 4.4-9.8)	(0.5; 1.1-2.9)
	Unrelated	9.1	9.1	3.0	6.6	
		(10.7; 0-47)	(1.1; 7-11)	(0.2; 3-4)	(1.6; 3.7-9.2)	
Form						
<i>befühlen-führen</i>	Related	5.8	9.7	3.0	7.7	1.3
(‘palpate’-‘guide’)		(9.2; 0-36)	(1.2; 8-12)	(0; 3-3)	(1.8; 4.7-10.3)	(0.5; 1-2.9)
	Unrelated	6.4	9.2	3.0	6.5	
		(10.9; 0-47)	(1.2; 7-11)	(0.2; 3-4)	(1.6; 3.4-9.2)	

responses with the subordinate hand. Feedback was given on both correct (‘richtig’) and incorrect (‘falsch’) responses during the practice session, and on incorrect responses during the experimental session. The experiment lasted for about 12 minutes, during which participants self-administered the breaks between blocks.

2.4 Results

A generalized additive mixed model (Wood, 2017) was fitted to the inverse-transformed reaction times with predictors Prime Type (using treatment coding, with the unrelated condition as reference level) and log target frequency. Random intercepts were included for target and prime, and a factor smooth for the interaction of subject by trial number (see Baayen et al., 2017, for detailed discussion of this nonlinear counterpart to what in a linear mixed model would be obtained with by-subject random intercepts and by-subject random slopes for trial). Table 2 presents the model summary. Prime-target pairs in the semantic condition were responded to slightly more quickly than prime-target pairs in the unrelated condition. Prime-target pairs in the transparent and opaque conditions showed substantially larger facilitation of equal magnitude. Prime-target pairs in the form condition elicited reaction times that did not differ from those seen in the control condition.

To obtain further insight into the effects of the predictors not only for the median, but across the distribution of reaction times, we fitted quantile GAMs to the deciles 0.1, 0.2, . . . , 0.9, using the `qgam` package (Fasiolo et al., 2017). For the median, the quantile GAM also complements the Gaussian GAMM reported in Table 2. The Gaussian GAMM could have been expanded with further random effects for the interaction of subject by priming effect, but such models run the risk of overspecification (Bates et al., 2015). More importantly, the distribution of the residuals showed clear deviation from normality that resisted correction. As quantile GAMS are distribution free, simple main effects can be studied without having to bring complex random effects into the model as a safeguard against anti-conservative p -values.

A. parametric coefficients	Estimate	Std. Error	t-value	p-value
Intercept (PrimeType=Unrelated)	-1.8791	0.0298	-63.0083	< 0.0001
PrimeType=Semantic	-0.0761	0.0240	-3.1709	0.0015
PrimeType=Transparent	-0.2514	0.0248	-10.1222	< 0.0001
PrimeType=Opaque	-0.2519	0.0249	-10.1305	< 0.0001
PrimeType=Form	0.0016	0.0242	0.0669	0.9467
B. smooth terms	edf	Ref.df	F-value	p-value
TPRS smooth LogTargetFreq	1.0001	1.0001	9.1823	0.0025
factor smooths for Trial \times Subject	113.0046	449.0000	2.6135	< 0.0001
random intercepts Prime	16.7699	116.0000	0.1907	0.0585
random intercepts Target	21.4123	39.0000	1.6359	< 0.0001

Table 2: Generalized additive mixed model fitted to inverse-transformed primed lexical decision latencies.

Figure 1 presents, from top left to bottom right, the effects of Prime Type for the deciles 0.1, 0.2, . . . , 0.9. The p -values above the bars concern the contrasts with the unrelated condition (the reference level). Across the distribution, the form condition was never significantly different from the unrelated condition. The small effect of the semantic condition hardly varied in magnitude across deciles, but was no longer significant at the last decile. The magnitude of the effects for the transparent and opaque conditions was significantly different from that for the unrelated condition across all deciles, and increased in especially the last three deciles. Across the deciles, the transpar-

ent condition showed a growing increase in facilitation compared to the opaque condition, but as indicated by the p -values in red, the difference between these two conditions was never significant.

Figure 2 visualizes the effect of target frequency. From the second decile onwards, target frequency was significant, with greater frequencies affording shorter reaction times, as expected. The magnitude of the frequency effect, as well as its confidence interval, increased across deciles.

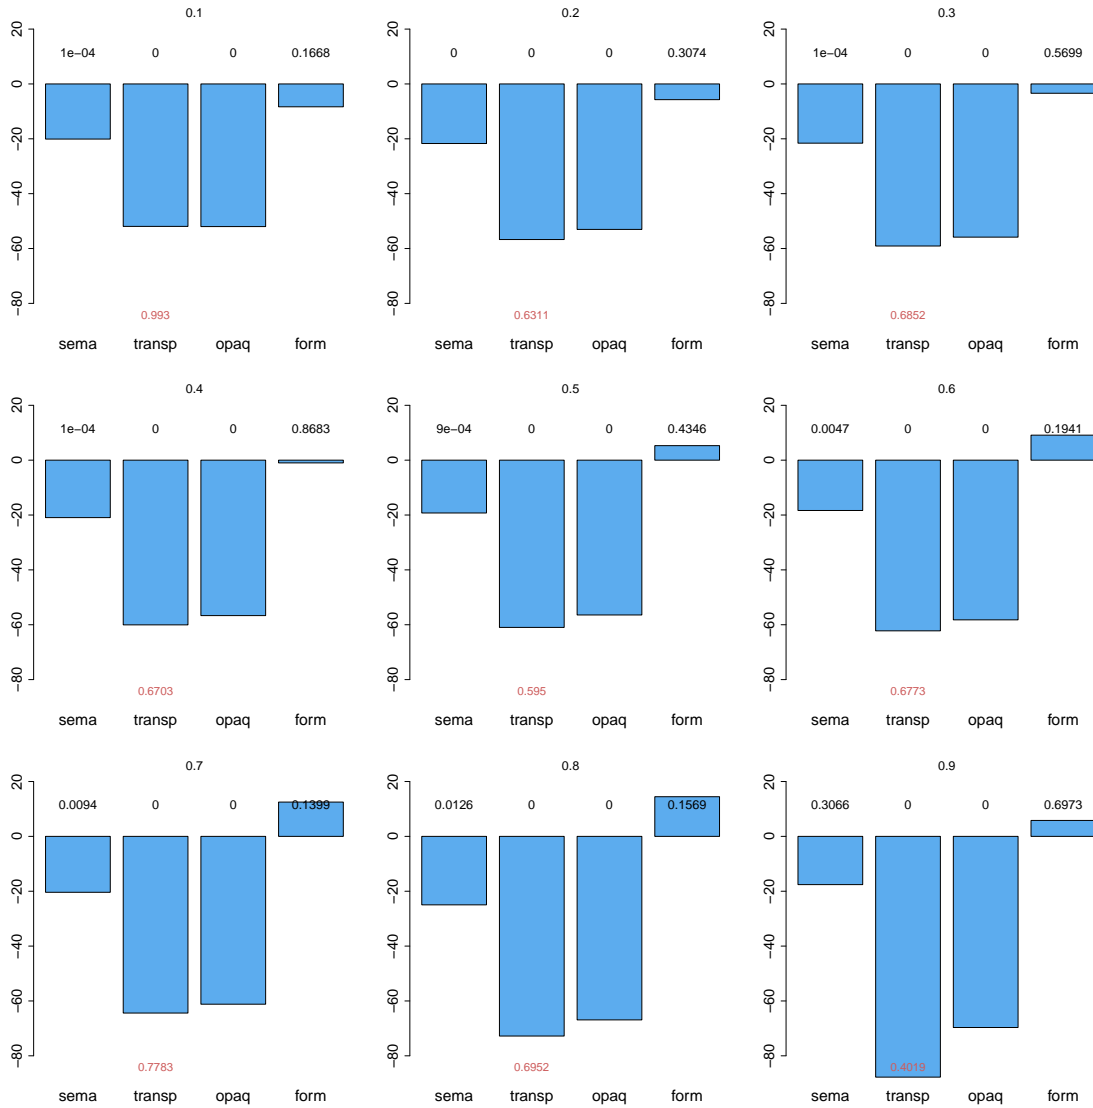


Figure 1: Effects of Prime Type (with unrelated as reference level) in a Quantile GAM fitted to primed lexical decision latencies, for deciles 0.1, 0.2, ..., 0.9.




Figure 2: Partial effect of log target frequency in a Quantile GAM fitted to primed lexical decision latencies, for deciles 0.1, 0.2, \dots , 0.9.

In summary, replicating earlier studies, morphologically related primes elicited a substantial priming effect that did not vary with semantic transparency. The priming effect tended to be somewhat stronger at larger deciles, for which the effect of target frequency was also somewhat larger.

Considered jointly, the behavioral results of ‘pure morphological priming’ without semantic transparency effects in German, as well as the older results obtained for speech production in Dutch, appear to indicate a fundamental role in lexical processing for morphemic units such as the stem. However, perhaps surprisingly, developments in current linguistic morphology indicate that the theoretical construct of the morpheme is in many ways problematic. In what follows, we show that the present results can be explained within the framework of naive discriminative learning, even though this theory eschews morphemic units altogether.

3 Computational modeling with naive discriminative learning

Before introducing naive discriminative learning (NDL), we first provide a brief overview of developments in theoretical morphology over the last decades that motivated the development of NDL.

3.1 Developments in theoretical morphology

The concept of the morpheme, as the minimal linguistic sign combining form and meaning, traces its history to the American structuralists that sought to further systematize the work of Leonard Bloomfield (see [Blevins, 2016](#), for detailed discussion). The morpheme as minimal sign has made it into many introductory textbooks (e.g., [Plag, 2003](#); [Butz and Kutter, 2016](#)). The hypothesis that semantically transparent complex words are processed compositionally, whereas semantically opaque words are processed as units, is itself motivated by the belief that morphemes are linguistic signs. For semantically opaque words, the link between form and meaning is broken, the morpheme is no longer a true sign, and hence the rules operating over true signs in comprehension and production are no longer relevant.

The theoretical construct of the morpheme as smallest sign of the language system has met with substantial criticism (see, e.g., [Matthews, 1974](#); [Beard, 1977](#); [Aronoff, 1994](#); [Stump, 2001](#); [Blevins, 2016](#)). Whereas the morpheme-as-sign appears a reasonably useful construct for agglutinating languages such as Turkish, as well as for morphologically simple languages such as English (but see [Blevins, 2003](#)), it fails to provide much insight for typologically very dissimilar languages such as Latin, Estonian, or Navajo (see, e.g., [Baayen et al., 2019, 2018](#); [Chuang et al., 2019](#), for detailed discussion). One important insight from theoretical morphology is that systematicities in form are not coupled in a straightforward one-to-one way with systematicities in meaning. Realizational theories of morphology (e.g., [Stump, 2001](#)) therefore focus on how sets of semantic features are expressed in phonological form, without seeking to find atomic form features that line up with atomic semantic features. Interestingly, as pointed out by [Beard \(1977\)](#), form and meaning are subject to their own laws of historical change or resistance to change.

Within realizational theories, two main approaches have been developed, *Realizational Morphology* and *Word and Paradigm Morphology*. Realizational Morphology formalizes how bundles of semantic (typically inflectional) features are realized in phonological form by making use of units for stems, stem variants, and the morphs (now named *exponents*) that realize (or express) sets of inflectional or derivational features (see, e.g., [Stump, 2001](#)). Realizational Morphology is to some extent compatible with localist models in psychology, in that the stems and exponents of realizational morphology can be seen as corresponding to the ‘morphemes’ (now understood strictly as

form-only units, henceforth ‘morphs’) in localist networks. The compatibility is only partial, however, as current localist models typically remain underspecified as to how, in comprehension, the pertinent semantic feature bundles are activated once the proper exponents have been identified. For instance, in the localist interactive activation model of Veríssimo (2018), the exponent *-er* that is activated by the form *teacher* has a connection to a lemma node for ER as deverbal nominalization, but no link is given from the *-er* exponent to an inflectional function that in English is also realized with *-er*, namely, the comparative. Furthermore, even the node for deverbal *-ER* is semantically underspecified, as *-ER* realizes a range of semantic functions, including AGENT, INSTRUMENT, CAUSER, and PATIENT (Booij, 1986; Bauer et al., 2015).

A further, empirical, problem for decompositional theories that take the first step in lexical processing to be driven by units for morphs are experiments indicating that quantitative measures tied to properties of whole words, rather than their component morphs, are predictive much earlier in time than expected. For eye-tracking studies on Dutch and Finnish, see Kuperman et al. (2010, 2009, 2008) and for reaction times analysed with survival analysis, see Schmidtke et al. (2017). These authors consistently find that measures linked to whole words are predictive for shorter response times, and that measures linked to morphs are predictive for longer response times. This strongly suggests that properties of whole words determine early processing and properties of morphs arise later in processing.

There is a more general problem specifically with models that make use of localist networks and the mechanism of interactive activation to implement lexical access. First of all, interactive activation is a very expensive mechanism, as inhibitory connections between morpheme nodes grow quadratically with the number of nodes, and access times increase polynomially or even exponentially. Furthermore, interactive activation as a method for candidate selection in a what amounts to a straightforward classification task is unattractive as it would have to be implemented separately for each classification task that the brain has to carry out. Redgrave et al. (1999) and Gurney et al. (2001) therefore propose a central single mechanism, supposed to be carried out by the basal ganglia, that receives a probability distribution of alternatives as input from any system requiring response selection, and returns the best-supported candidate (see Stewart et al., 2012, for an implementation of their algorithm with spiking neurons).

The second main approach within morpheme-free theories, *Word and Paradigm Morphology*, rejects the psychological reality of stems and exponents, and calls upon proportional analogies between words to explain how words are produced and comprehended (Matthews, 1974; Blevins, 2006, 2003, 2016). Although attractive at a high level of abstraction, without computational implementation, supposed proportional analogies within paradigms do not generate quantitative predictions that can be tested experimentally. As discussed in detail by Baayen et al. (2019), discrimination learning provides a computational formalization of Word and Paradigm Morphology that does generate testable and falsifiable predictions.

In what follows, we will use NDL to estimate a distribution of activations (proxies for probabilities) over the set of possible word meanings given the visual input. Specifically, we investigate whether we can predict how prior presentation of a prime word affects the activation of the target meaning.

3.2 Morphological processing without morphemes

Naive discriminative learning is not the first cognitive computational model that seeks to move away from morphemes. The explanatory adequacy of the morpheme for understanding lexical processing has also been questioned within psychology by the parallel distributed processing programme (McClelland and Rumelhart, 1986; Rumelhart and McClelland, 1986). As mentioned previously,

the triangle model (Harm and Seidenberg, 2004) has been argued to explain the effects of semantic transparency observed for English derived words as reflecting the convergence of phonological and semantic codes (Plaut and Gonnerman, 2000; Gonnerman and Anderson, 2001; Gonnerman et al., 2007). It is noteworthy, however, that to our knowledge, actual simulation studies demonstrating this have not been forthcoming. Importantly, if indeed the triangle model makes correct predictions for English, then one would expect its predictions for German to be wrong, because it would predict semantic transparency effects and no priming for semantically opaque word pairs.

Like the PDP programme, the twin theories of Naive Discriminative Learning (NDL Baayen et al., 2011, 2016b; Milin et al., 2017a; Sering et al., 2018b) and Linear Discriminative Learning (Baayen et al., 2018, 2019), eschew the construct of the morpheme. But instead of using backpropagation multi-layer networks, NDL and LDL build on simple networks with input units that are fully connected to all output units.




Figure 3: The two-layer network corresponding to \mathbf{W} . The outer vertices in blue represent the input nodes, the inner vertices in red represent the output nodes. All input nodes are connected to each of the three output nodes.

An NDL network is defined by its weight matrix \mathbf{W} . By way of example, consider the following weight matrix,

$$\mathbf{W} = \begin{array}{c} \text{QAID} \quad \text{SAID} \quad \text{HID} \\ \begin{array}{l} \#q \\ qa \\ ai \\ id \\ d\# \\ \#s \\ sa \\ \#h \\ hi \end{array} \left(\begin{array}{ccc} 0.33 & -0.17 & -0.08 \\ 0.33 & -0.17 & -0.08 \\ 0.17 & 0.17 & -0.17 \\ 0.08 & 0.08 & 0.17 \\ 0.08 & 0.08 & 0.17 \\ -0.17 & 0.33 & -0.08 \\ -0.17 & 0.33 & -0.08 \\ -0.08 & -0.08 & 0.33 \\ -0.08 & -0.08 & 0.33 \end{array} \right), \end{array}$$

which is visualized in Figure 3. The output nodes are on the inner circle in red, and the input nodes in the outer circle in blue. A star layout was chosen in order to guarantee readability of the connection weights. The network corresponding to this weight matrix comprises nine sublexical input units, shown in the left margin of the matrix. We refer to these units, here the letter bigrams of the words *qa*id, *sa*id and *hi*d, as cues; the # symbol (a + in Figure 3) represents the space character). There are three output units, the outcomes, shown in the upper margin of the matrix. The entries in the matrix present the connection strengths of the digraphs to the lexical outcomes. The digraph *qa* provides strong support (0.33) for QAID (‘tribal chieftain’), and *sa* provides strong support (0.33) for SAID. Conversely, *ai*, which is a valid cue for two words, QAID and SAID, has connection strengths to these lexemes of only 0.17. The weights from *hi* and *sa* to QAID are negative, -0.08 and -0.17 respectively. For QAID, the cue that best discriminates this word from the other two words is *qa*. Conversely, *sa* is a (somewhat less strong) discriminative cue arguing against QAID. Informally, we can say that the model concludes the outcome must be QAID given *qa*, and that the outcome cannot be QAID given *sa*.

In the present example, form cues are letter pairs, but other features have been found to be effective as well. Depending on the language and its writing conventions, larger letter or phone n -grams (typically with $1 < n \leq 4$) may outperform letter bigrams. For auditory comprehension, low-level acoustic features have been developed for modeling auditory comprehension (Arnold et al., 2017; Shafaei Bajestan and Baayen, 2018; Baayen et al., 2019). For visual word recognition, low-level visual ‘histograms of gradient orientation’ features have been applied successfully in (Linke et al., 2017).

The total support a_j for an outcome j given the set of cues \mathcal{C} in the visual input to the model, henceforth its activation, is obtained by summing the weights on the connections from these cues to that outcome:

$$a_j = \sum_{i \in \mathcal{C}} w_{ij}.$$

For QAID ($j = 1$), the total evidence a_1 given the cues *#q*, *qa*, *ai*, *id*, and *d#* is $0.33 + 0.33 + 0.17 + 0.08 + 0.08 = 1$.

The values of the weights are straightforward to estimate. We represent the digraph cues of the words by a matrix \mathbf{C} , with a 1 representing the presence of a cue in the word, and a 0 its absence:

$$\mathbf{C} = \begin{array}{c} \text{QAID} \\ \text{SAID} \\ \text{HID} \end{array} \begin{array}{c} \#q \quad qa \quad ai \quad id \quad d\# \quad \#s \quad sa \quad \#h \quad hi \\ \left(\begin{array}{ccccccccc} 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \end{array} \right). \end{array}$$

We also represent the outcomes using a matrix, again using binary coding:

$$\mathbf{T} = \begin{matrix} & \text{QAID} & \text{SAID} & \text{HID} \\ \text{QAID} & \left(\begin{array}{ccc} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{array} \right) \\ \text{SAID} & & & \\ \text{HID} & & & \end{matrix}.$$

The vectors representing the outcomes in an NDL network are orthogonal: each pair of row vectors of \mathbf{T} is uncorrelated. The weight matrix \mathbf{W} follows by solving²

$$\mathbf{CW} = \mathbf{T}.$$

In other words, \mathbf{W} projects words' forms, represented by vectors in a form space $\{\mathbf{C}\}$, onto words' meanings, represented by vectors in a semantic space $\{\mathbf{T}\}$.³

The outcomes of an NDL network represent lexical meanings that are discriminated in a language. Milin et al. (2017a) refer to these outcomes as lexomes, which they interpret as pointers to (or identifiers of) locations (or vectors) in some high-dimensional space as familiar from distributional semantics (see, e.g., Landauer and Dumais, 1997; Mikolov et al., 2013). However, as illustrated above with the \mathbf{T} matrix, NDL's lexomes can themselves be represented as high-dimensional vectors, the length k of which is equal to the total number of lexomes. The vector for a given lexome has one bit on and all other bits off (cf. Sering et al., 2018b). Thus, the lexomes jointly define a k -dimensional orthonormal space.

However, the orthonormality of the outcome space does not do justice to the fact that some lexomes are more similar to each other than others. Within the general framework of NDL, such similarities can be taken into account, but to do so, measures gauging semantic similarity have to be calculated from a separate semantic space that constructs lexomes' semantic vectors (known as word embeddings in computational linguistics) from a corpus. A technical complication is that, because many words share semantic similarities, the dimensionality of NDL's semantic space, k , is much higher than it need be. As a consequence, the classification accuracy of the model is lower than it could be (see Baayen et al., 2019, for detailed discussion).

The twin model of NDL, LDL, therefore replaces the one-hot encoded semantic vectors as exemplified by \mathbf{T} by real-valued vectors. For the present example, this amounts to replacing \mathbf{T} by a matrix such as \mathbf{S} :

$$\mathbf{S} = \begin{matrix} & \text{QAID} & \text{SAID} & \text{HID} \\ \text{QAID} & \left(\begin{array}{ccc} 0.4 & -0.2 & 0.3 \\ -0.2 & -0.2 & -0.3 \\ -0.1 & 0.3 & 0.3 \end{array} \right) \\ \text{SAID} & & & \\ \text{HID} & & & \end{matrix}.$$

Actual corpus-based semantic vectors are much longer than this simple example suggests, with hundreds or even thousands of elements. The method implemented in Baayen et al. (2019) produces vectors the values of which represent a given lexome's collocational strengths with all the other lexomes in the corpus.

Model accuracy is evaluated by examining how close a predicted semantic vector $\hat{\mathbf{s}}$ is to the targeted semantic vector \mathbf{s} , a row vector of \mathbf{S} . In the case of NDL, this evaluation is straightforward: The lexome that is best supported by the form features in the input, and that thus receives the

² In R: `W = ginv(C) %*% T`, see Baayen et al. (2018) and Baayen et al. (2019) for further details on linear transformations from form to meaning (and from meaning to form).

³ The weight matrix \mathbf{W} is identical to the weight matrix obtained by applying the equilibrium equations of Danks (2003) for the Rescorla-Wagner learning rule (Rescorla and Wagner, 1972) that was used by Baayen et al. (2011), see Sering et al. (2018b) for detailed discussion.

highest activation, is selected. In the case of LDL, that word ω is selected as recognized for which the predicted semantic vector \hat{s} is most strongly correlated with the targeted semantic vector s_ω .

As the dimensionality of the row vectors of \mathbf{S} , \mathbf{T} , and \mathbf{C} can be large, with thousands or tens of thousands of values, we refer to the network \mathbf{W} as a ‘wide learning’ network, as opposed to ‘deep learning’ networks which have multiple layers but usually much smaller numbers of units on these layers.

Of specific relevance to the present study is how NDL and LDL deal with morphologically complex words. With respect to the forms of complex words, exactly the same encoding scheme is used as for simple words, with either n-grams or low-level modality-specific features used as descriptor sets. No attempt is made to find morpheme boundaries, stems, affixes, or allomorphs.

At the semantic level, both NDL and LDL are analytical. NDL couples inflected words such as *walked* and *swam* with the lexomes WALK and PAST, and SWIM and PAST, respectively. In the example worked out in Table 3, the word form *walk* has Lx7 as identifier; the lexome for *past* is indexed by Lx4. The form *walked* is linked with both Lx7 and Lx4. For clarity of exposition, instead of using indices, we refer to lexomes using small caps: WALK and PAST. An NDL network is thus trained to predict, for morphologically complex words, on the basis of the form features in the input, the simultaneous presence of two (or more) lexomes. Mathematically, as illustrated in the top half of Table 3, this amounts to predicting the sum of the one-hot encoded vectors for the stem (WALK) and the inflectional function (PAST). Thus, NDL treats the recognition of complex words as a multi-label classification problem (Sering et al., 2018b).

LDL proceeds in exactly the same way, as illustrated in the bottom half of Table 3. Again, the semantic vector of the content lexome and the semantic vector of the inflection are added. The columns now label semantic dimensions. In the model of Baayen et al. (2019), these dimensions quantify collocational strengths with — in the present example — 10 well-discriminated lexomes. Regular past tense forms such as *walked* and irregular past tense forms such as *swam* are treated identically at the semantic level. It is left to the mapping \mathbf{W} (the network taking form vectors as input and producing semantic vectors as output) to ensure that the different forms of regular and irregular verbs are properly mapped on the pertinent semantic vectors.

NDL										
	Lx1	Lx2	Lx3	Lx4	Lx5	Lx6	Lx7	Lx8	Lx9	Lx10
<i>walk</i>	0	0	0	0	0	1	0	0	0	0
<i>past</i>	0	0	0	1	0	0	0	0	0	0
<i>walked</i>	0	0	0	1	0	1	0	0	0	0
LDL										
	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10
<i>walk</i>	-0.16	0.25	-0.07	-0.04	0.13	0.03	-0.04	0.10	0.22	-0.25
<i>past</i>	0.16	0.13	-0.36	0.01	-0.07	-0.04	-0.32	0.07	0.45	0.14
<i>walked</i>	0.00	0.38	-0.43	-0.03	0.06	-0.01	-0.36	0.17	0.67	-0.11

Table 3: Semantic vector representations for inflected words in NDL (top) and LDL (bottom).

The NDL model as laid out by Baayen et al. (2011) treats transparent derived words in the same way as inflections, but assigns opaque derived words their own lexomes. For opaque words in which the semantics of the affix are present, even though there is no clear contribution from the semantics of the base word, a lexome for the affix is also included (e.g., *employer*: EMPLOY + ER; *cryptic*: CRYPTIC + IC).

The LDL model, by contrast, takes the idea seriously that derivation serves word formation,

in the onomasiological sense. Notably, derived words are almost always characterized by semantic idiosyncracies, the exception being inflection-like derivation such as adverbial *-ly* in English.⁴ For instance, the English word *worker* denotes not just ‘someone who works’, but “one that works especially at manual or industrial labor or with a particular material”, a “factory worker”, “a member of the working class”, or “any of the sexually underdeveloped and usually sterile members of a colony of social ants, bees, wasps, or termites that perform most of the labor and protective duties of the colony” (<https://www.merriam-webster.com/dictionary/worker>, s.v.). Given these semantic idiosyncracies, when constructing semantic vectors from a corpus, LDL assigns each derived word its own lexome. However, in order to allow the model to assign an approximate interpretation to unseen derived words, each occurrence of a derived word is also coupled with a lexome for the semantic function of the affix. For instance, *worker* (in the sense of the bee) is associated with the lexomes WORKER and AGENT, and *amplifier* with the lexomes AMPLIFIER and INSTRUMENT. In this way, semantic vectors are created for derivational functions, along with semantic vectors for the derived words themselves (see Baayen et al., 2019, for detailed discussion and computational and empirical evaluation).⁵

To put LDL and NDL in perspective, consider the substantial advances made in recent years in machine learning and its applications in natural language engineering. Computational linguistics initially worked with deterministic systems applying symbolic units and formal grammars defined over these units. It then became apparent that considerable improvement in performance could be obtained by making these systems probabilistic. The revolution in machine learning that has unfolded over the last decade has made clear that yet another substantial step forward can be made by moving away from hand-crafted systems building on rules and representations, and to make use instead of deep learning networks such as autoencoders, LSTM networks for sequence to sequence modeling, and deep convolutional networks, outperforming almost all classical symbolic algorithms on tasks as diverse as playing *Go* (AlphaGo, Silver et al., 2016) and speech recognition (deep speech, Hannun et al., 2014). How far current natural language processing technology has moved away from concepts in classical (psycho)linguistics theory is exemplified by Hannun et al. (2014), announcing in their abstract that they “...do not need a phoneme dictionary, nor even the concept of a ‘phoneme’ ” (p. 1).

The downside of the algorithmic revolution in machine learning is that what exactly the new networks are doing often remains a black box. What is clear, however, is that these networks are sensitive to what in regression models would be higher-order nonlinear interactions between predictors (Cheng et al., 2018). Crucially, such complex interactions are impossible to reason through analytically. As a consequence, models for lexical processing that are constructed analytically by hand-crafting lexical representations for stems and exponents, and hand-crafting inhibitory or excitatory connections between these representations, as in standard interactive activation models, are unable to generate sufficiently accurate estimates for predicting with precision aspects of human lexical processing.

We note here that NDL and LDL provide high-level functional formalizations of lexical processing. They should not be taken as models for actual neural processing: biological neural networks involve cells that fire stochastically, with connections that are stochastic (Kappel et al., 2015, 2017) as well. Furthermore, most neural computations involve ensembles of spiking neurons (Eliasmith et al., 2012).

NDL and LDL are developed to provide a linguistically fully interpretable model using mathe-

⁴ The reason that adverbial *-ly* is generally treated as derivational is that the word category of words with *-ly* is not identical to the word category of its base word.

⁵ For a compositional approach to the semantics of complex words using distributional semantics, see Lazaridou et al. (2013); Marelli and Baroni (2015).

matically well-understood networks that, even though very simple, are powerful enough to capture important aspects of the interactional complexities in language, and to generate predictions that are sufficiently precise to be pitted against experimental data. Although NDL and LDL make use of the simplest possible networks, these networks can, in combination with carefully chosen input features, be surprisingly effective. For instance, for auditory word recognition, an NDL model trained on the audio of individual words extracted from 20 hours of German free conversation correctly recognized around 20% of the words, an accuracy that was subsequently found to be within the range of human recognition accuracy (Arnold et al. (2017)). Furthermore, Shafaei Bajestan and Baayen (2018) observed that NDL outperforms deep speech networks by a factor 2 on isolated word recognition. With respect to visual word recognition, Linke et al. (2017) showed, using low-level visual features, that NDL outperforms deep convolutional networks (Hannagan et al., 2014) on the task of predicting word learning in baboons (Grainger et al., 2012).

3.3 Computational modeling with naive discriminative learning

In the present study, we model our experiment with NDL, rather than LDL, for two reasons. First, it turns out that NDL, the simpler model, is adequate. Second, work is in progress to derive corpus-based semantic vectors for German along the lines of Baayen et al. (2019), which will include semantic vectors for inflectional and derivational semantic functions, but these vectors are not yet available to us.

The steps in modeling with NDL are the following. First, the data on which the network is to be trained have to be prepared. Next, the weights on the connections from the form features to the lexemes are estimated. Once the network has been trained, it can be used to generate predictions for the magnitude of the priming effect. In the present study, we generate these predictions by inspecting the extent to which the form features of the prime support the lexeme of the target.

3.3.1 Data preparation

The data on which we trained our NDL network comprised 18,411 lemmas taken from the CELEX database, under the restrictions that (i) they contained no more than two morphemes according to the CELEX parses, (ii) that the word was not a compound, and (iii) that it either had a nonzero CELEX frequency or occurred as a stimulus in the experiment. One stimulus word, *betraten*, was not listed in CELEX, and hence this form was not included in the simulation study. For each lemma, its phonological representation and its frequency were retrieved from CELEX. As form cues, we used triphones (for the importance of the phonological route in reading, see Baayen et al., 2019, and references cited there).

Each lemma was assigned its own lexeme (but homophones were collapsed). The decision to assign each lemma its own lexeme follows Baayen et al. (2019) and departs from Baayen et al. (2011). This similar treatment of transparent and opaque verbs is motivated by several theoretical considerations. First, there is no binary distinction between transparent and opaque. The meanings of particle verbs lie on a continuum between relatively semantically compositional and relatively semantically opaque. Second, even the compositional interpretation of a supposedly transparent verb such as *aufstehen* ('stand up') is not straightforward in the absence of situational experience — the particle *auf* (roughly meaning *on* or *onto*) may express a wide range of meanings, depending on cotext and context. In what follows, we therefore assume that even transparent complex words possess somewhat idiosyncratic meanings, and hence should receive their own lexemes in the NDL network.

The resulting input to the model was a file with 4,492,525 rows and two columns, one column spelling out a word's triphones, and the other column listing its lexeme. Each word appeared in

the file with a number of tokens equal to its frequency in CELEX. The order of the words in the file was randomized.

3.3.2 Training the network

An NDL network with 10,180 input nodes (triphones) and 18,404 output nodes (lexomes) was trained on the input list, with incremental updating of the weights on the connections from features in the input to the lexomes, using the learning rule of Rescorla and Wagner (1972) ($\lambda = 1, \alpha = 0.001, \beta = 1$; i.e., with a learning rate of 0.001). As there were 4,492,525 learning events in the input file, the total number of times that weights were updated was 4,492,525.⁶

3.4 Modeling priming

To model the effect of priming, we presented the triphones of the prime to the network, and summed the weights on the connections from these triphones to the pertinent target to obtain a measure of the extent to which the prime pre-activates its target (henceforth `PrimeToTargetPreActivation`). Figure 4, upper left panel, presents a boxplot for `PrimeToTargetPreActivation` as a function of `PrimeType`. Interestingly, the opaque and transparent prime types comprise prime-target pairs for which the prime provides substantial and roughly the same amount of pre-activation for the target. For the other prime types, pre-activation is close to zero. Form-related prime-target pairs show some pre-activation, but this pre-activation is much reduced compared to the prime-target pairs in the opaque and transparent conditions.

The upper right panel of Figure 4 presents the results obtained when the empirical frequencies with which words were presented to the NDL network are replaced by uniformly distributed frequencies. This type-based simulation generates predictions that are very similar to those of the token-based simulation. This result shows that imprecisions in the frequency counts underlying the token-based analysis are not responsible for the model’s predictions.

Above, we called attention to the finding of Smolka and Eulitz (2011) that very similar priming effects are seen when the order of prime and target is reversed. We therefore also ran a simulation in which we reversed the order of prime and target, and investigated the extent to which the current targets (now primes) co-activate the current primes (now targets). The distributions of the predicted pre-activations are presented in the lower left panel of Figure 4 (`target-to-prime pre-activation`). Apart from one extreme outlier for the opaque condition, the pattern of results is qualitatively the same as for the Prime-to-Target Pre-Activation presented in the upper panel. For both simulations, there is no significant difference in the mean for the opaque and transparent conditions, whereas these two conditions have means that are significantly larger than those of the other three condition (Wilcoxon-tests with Bonferroni correction). In summary, our NDL model generates the correct prediction that the priming effect does not depend on the order of prime and target.

Reaction times are expected to be inversely proportional to `PrimeToTargetPreActivation`. We therefore ran a linear model on the stimuli, and used the reciprocal of `PrimeToTargetPreActivation` as response variable, based on the simulation in which the model was presented with the empirical word frequencies. As the resulting distribution is highly skewed, the response variable was transformed to $\log(1/(\text{PrimeToTargetPreActivation} + 0.14))$.⁷ The opaque and transparent priming

⁶ Optimized software, e.g., Sering et al. (2018a), makes it possible to harness multiple cores in parallel. Using 6 cores, training the network takes less than 10 minutes. Incremental learning is much faster than weight estimation by means of the Danks equilibrium equations, which were used by Baayen et al. (2011).

⁷ The shift 0.14 is slightly larger than the absolute value of the most negative pre-activation. This shift thus ensures that all pre-activation values are positive, so that a log-transform becomes possible.

conditions were supported as having significantly shorter simulated reaction times compared to the unrelated condition (both $p \ll 0.0001$), in contrast to the other two conditions (both $p > 0.5$).

Recall that the outcome vectors of NDL are orthogonal, and that hence the present NDL models all make predictions that are driven purely by form similarity. The model is blind to potential semantic similarities between primes and targets, not only for the primes and targets in the transparent and opaque conditions, but also to semantic similarities present for the other prime types. To understand to what extent semantic similarities might be at issue in addition to form similarities, we therefore inspected prime and target’s semantic similarity using distributional semantics.

3.5 Semantic vectors from tweets

As LDL-based semantic vectors for German are currently under construction, we fell back on the word embeddings (semantic vectors) provided at <http://www.spinningbytes.com/resources/wordembeddings/> (Cieliebak et al., 2017; Deriu et al., 2017). These embeddings (obtained with `word2vec` Mikolov et al., 2013) are 300-dimensional vectors derived from a 50 million word corpus of German tweets. Tweets are relatively short text messages that reflect spontaneous and rather emotional conversation. Tweets from facebook have been shown to outperform frequencies from standard text corpora in predicting lexical decision latencies (Herdağdelen and Marelli, 2017).

Cieliebak et al. (2017) and Deriu et al. (2017) provide separate semantic vectors for words’ inflected variants. For instance, the particle verb *vorwerfen* (‘accuse’) occurs in their database in the forms *vorwerfen* (infinitive and 1st or 3rd person plural present), *vorwerfe* (1st person singular present), *vorwirfst* (2nd person singular present), *vorwirft* (3rd person singular present), *vorwerft* (2nd person plural present), *vorgeworfen* (past participle), and *vorzuwerfen* (infinitive construction with *zu*). As we can expect for tweets, not all inflected forms, in particular the more formal ones, appear in the database. Importantly, the semantic vectors are probably obtained without taking into account that the particle of a particle verb can appear separated from its verb, sometimes at a considerable distance (see Schreuder, 1990, for discussion of the cognitive consequences of this separation), as in the sentence ‘Sie *wirft* ihm seinen Leichtsinn *vor*’, ‘She accuses him of his thoughtlessness’. Given the computational complexity of identifying particle-verb combinations when the particle appears at a distance, it is highly likely that for split particle verbs, the base verb of the verb-particle combination is processed as if it were a simple verb (e.g., *werfe*, *wirfst*, *wirft*, *werfen*, and *werft*, 1st, 2nd, and 3rd person singular and plural present, respectively). As a consequence, the semantic similarity of simple verbs and particle verbs computed from the word embeddings provided by Cieliebak et al. (2017); Deriu et al. (2017) is in all likelihood larger than it should be.

Not all words in the experiment are in this database; but for six words, we were able to replace the infinitive by a related form (*einpassen* → *reinpassen*, *verqualmen* → *verqualmt*, *fortlaufen* → *fortlaufend*, *bestürzen* → *bestürzend*, *verfinstern* → *verfinstert*, *beschneien* → *beschneites*).

For each prime-target pair for which we had data, we calculated the cosine similarity of the semantic vectors of prime and target, henceforth **Prime-to-Target Cosine Similarity**. Figure 4, lower right panel, shows that the transparent pairs have the greatest semantic similarity, followed by the opaque pairs, then the semantic pairs and the form pairs, and the least semantic similarity by the unrelated pairs.

Surprisingly, the semantic controls have a rather low semantic similarity, substantially less than that of the opaque pairs. A linear model with the semantic primes as reference level clarifies that the semantic pairs are on a par with the form controls, more similar than the unrelated pairs, but less similar than both the opaque and transparent pairs (Table 4).

There is a striking discrepancy between the assessment of semantic similarity across prime types

based on the cosine similarity of the semantic vectors on the one hand, and an assessment based on the ratings for semantic relatedness between word pairs, as documented in Table 1. In the former, semantic pairs pattern with form controls and differ from transparent ones, while in the latter, semantic pairs pattern with transparent pairs (5.5 and 5.7, respectively), and opaque pairs with form-related ones (2.1 and 1.7, respectively).

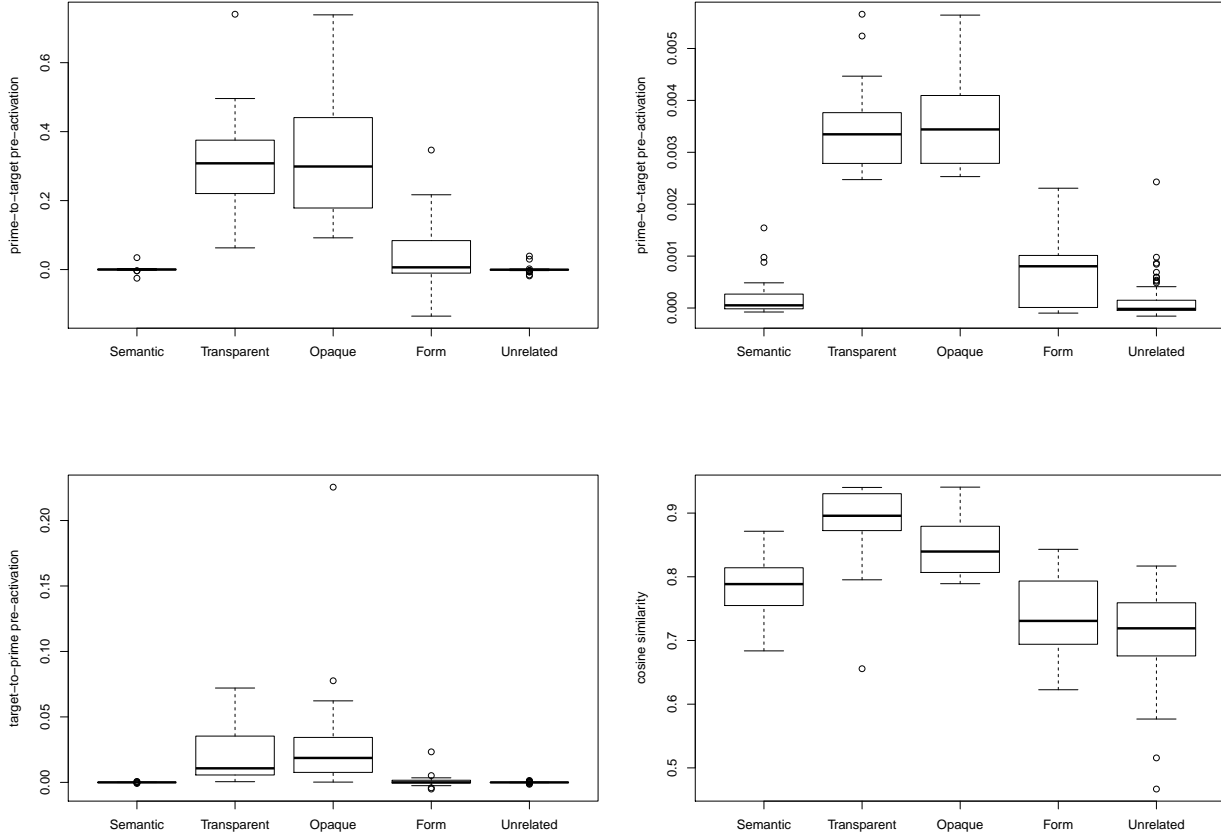


Figure 4: Predicted NDL Prime-to-Target Pre-Activation using empirical word frequencies (upper left panel), Predicted NDL Prime-to-Target Pre-Activation using uniformly distributed frequencies (upper right panel), Target-to-Prime Pre-Activation using empirical frequencies (lower left panel), and Prime-target Cosine Similarity (lower right panel), broken down by prime type.

	Estimate	Std. Error	<i>t</i> -value	Pr(> <i>t</i>)
Intercept (Semantic)	0.7778	0.0134	58.19	0.0000
PrimeType=Form	-0.0404	0.0197	-2.05	0.0418
PrimeType=Opaque	0.0696	0.0191	3.64	0.0004
PrimeType=Transparent	0.1083	0.0191	5.66	0.0000
PrimeType=Unrelated	-0.0700	0.0150	-4.67	0.0000

Table 4: Effect of PrimeType in a linear model predicting cosine similarity, using treatment coding with the semantic condition as reference level.

Most important to our study is that the opaque pairs show significantly less semantic similarity than the transparent ones ($p < 0.0047$, Wilcoxon test): The analysis of word embeddings confirms that there is a true difference in semantic transparency between the transparent and opaque prime-target pairs. And yet, this difference is not reflected in our reaction times.

Given the strong track record of semantic vectors in both psychology and computational linguistics, the question arises of whether the prime-target cosine similarities are predictive for reaction times, and how the magnitude of their predictivity compares to that of the NDL Prime-to-Target Pre-Activation.

3.6 Putting it all together: predicting reaction times

To address these questions, we fitted a new GAMM to the inverse-transformed reaction times of our experiment, replacing the factorial predictor `PrimeType` with the model-based predictor `Prime-to-Target Pre-Activation`. We also replaced target frequency by the activation that the target word receives from its own triphones (henceforth `TargetActivation`, see Baayen et al. (2011) for detailed analyses using this measure). Target activation is proportional to frequency, and hence larger values of target activation are expected to indicate shorter response times.⁸

Of the experimental dataset, about 7% of the observations was lost due to 7 words not being available in CELEX or in the dataset of word embeddings. To set a baseline for model comparison, we refitted the GAMM discussed above to the 1999 datapoints of the reduced dataset. The fREML score for this model was 360.76. A main effects model replacing `Target Frequency` by `Target Activation`, `PrimeType` by `Prime-to-Target Pre-Activation`, and as additional predictor the `Prime-to-Target Cosine Similarity` had a slightly higher fREML score, 370.97. An improved model was obtained by allowing the three new covariates to interact, using a tensor product smooth. The fREML score of this model, summarized in Table 5, was 354.16. A chi-squared test for model comparison (implemented in the `compareML` function of the `itsadug` package (van Rij et al., 2017)) suggests that the investment of 4 additional effective degrees of freedom is significant ($p < 0.0001$). As the models are not nested and the increase in goodness of fit is not that big (the evidence ratio, i.e., the relative probability that the NDL-based GAMM is more likely to provide a more precise model of the data than the model with classical predictors, is only 27), we conclude that — to obtain a model that is at least equally good — it is possible to replace the classical predictors such as `Frequency` and `PrimeType` by model-based predictors without loss of prediction accuracy.⁹

The three-way interaction involving `Target Frequency`, `Prime-To-Target Pre-Activation`, and `Prime-to-Target Cosine Similarity` is visualized in Figure 5. The five panels show the joint effect of the first two predictors for selected quantiles of `Prime-to-Target Cosine Similarity`: From top left, to bottom right, `Prime-to-Target Cosine Similarity` is set to its 0.1, 0.3, 0.5, 0.7, and 0.9 deciles. Darker shades of blue indicate shorter reaction times, and darker shades of yellow longer RTs.

As can be seen in the upper left panel (for the first decile of `Cosine Similarity`), reaction times decrease slightly as `Target Activation` is increased, but only when there is little `Prime-To-Target Pre-Activation`. A clear effect of `Prime-to-Target Pre-Activation` is present for the larger values of `Target Activation`.

⁸ Both activation measures were log-transformed after adding a small number, 0.14 for `Prime-to-Target Pre-Activation` and 0.01 for `Target Activation`, to ensure that all pertinent numbers were positive before taking logarithms.

⁹ Including an interaction of `PrimeType` by `Target Frequency` in the GAMM with classical predictors led to an increase in the fREML score, indicating overfitting and increased model complexity without increased prediction accuracy.

A. parametric coefficients	Estimate	Std. Error	<i>t</i> -value	<i>p</i> -value
Intercept	-1.9404	0.0294	-66.0214	< 0.0001
B. smooth terms	edf	Ref.df	<i>F</i> -value	<i>p</i> -value
te(Target Activation × Prime-To-Target Pre-Activation × Prime-to-Target Cosine Similarity)	14.9326	17.4660	10.3959	< 0.0001
random intercepts Prime	15.4237	108.0000	0.1888	0.0547
factor smooths for Trial × Subject	101.6786	449.0000	2.4620	< 0.0001
random intercepts Target	19.9620	40.0000	1.2681	< 0.0001

Table 5: GAMM fitted to inverse transformed reaction times using model-based predictors; te(): tensor product smooth.

Recall that, as shown in Figure 4, transparent and opaque prime-target pairs have the same mean pre-activation, whereas the mean cosine similarity is greater for transparent prime-target pairs compared to opaque pairs. If both pre-activation and cosine similarity would have independent effects, one would expect a difference in the mean reaction times for these two prime types, contrary to fact. The interaction of pre-activation by target activation by cosine similarity resolves this issue by decreasing the effect of pre-activation as cosine similarity increases. When prime and target are more similar semantically, the effect of pre-activation is reduced, and reaction times are longer than would otherwise have been the case. This increase in RTs may reflect the cognitive system slowing down to deal with two signals for very similar meanings being presented in quick succession in a way that is extremely rare in natural language.

We checked whether the association ratings that were used for stimulus preparation were predictive for the reaction times. This turned out not to be the case, not for the model with classical predictors, nor when the association ratings were added to the model with discrimination-based predictors.

Finally, Table 6 presents the fREML scores for the full model, and the three models obtained when one covariate is removed at a time. Since smaller fREML scores indicate a better fit, Table 6 clarifies that **Prime-to-Target Pre-Activation** is the most important covariate, as its exclusion results in the worst model fit (386.11). The variable importances of **Prime-to-Target Cosine Similarity** is also substantial (383.49) whereas removing **Target Activation** from the model specification reduces model fit only slightly (361.62). The model that best predicts (AIC 354.16) pure morphological priming includes all three factors: **Prime-to-Target Pre-Activation** (capturing Prime Type), **Prime-to-Target Cosine Similarity** (capturing minor semantic effects), and **Target Activation** (capturing the frequency effect).

model	fREML score
full model	354.16
te(Target Activation × Prime-to-Target Pre-Activation)	383.49
te(Target Activation × Prime-to-Target Cosine Similarity)	386.11
te(Prime-to-Target Pre-Activation × Prime-to-Target Cosine Similarity)	361.62

Table 6: fREML scores for 4 models: the full model, a model without Prime-to-Target Cosine Similarity, a model without Prime-to-Target Pre-Activation, and a model without Target Activation. A smaller fREML score indicates a better fit. te(): tensor product smooth.

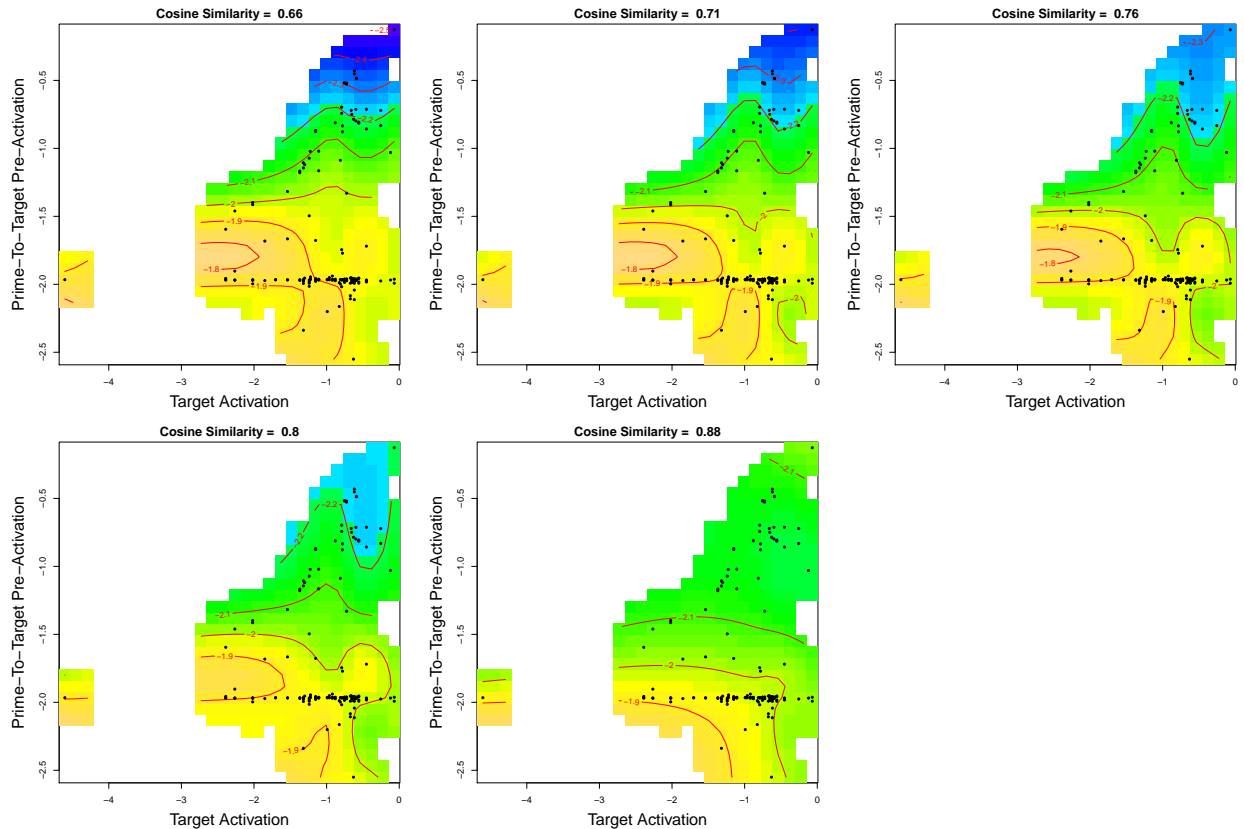


Figure 5: Interaction of Target Activation by Prime-to-Target Pre-Activation by Prime-to-Target Cosine Similarity in a GAMM fitted to the inverse-transformed reaction times. Darker shades of blue indicate shorter reaction times, and darker shades of yellow longer RTs. For visualization, one extreme outlier value for Prime-to-Target Pre-Activation was removed (13 datapoints).

4 Discussion

We presented an overt primed visual lexical decision experiment that replicated earlier results for German complex verbs: Priming effects were large and equivalent for semantically transparent and semantically opaque prime-target pairs. These findings add to the cumulative evidence of “pure morphological priming” patterns that suggest stem access independent of semantic compositionality in German, in contrast to English and French, where semantic compositionality has been reported to co-determine word processing. Because neither localist nor connectionist models of lexical processing are able to account for the German findings, Smolka et al. (2009, 2014, 2019) proposed a stem-based frequency account, according to which stems constitute the crucial morphological units regulating lexical access in German, irrespective of semantic transparency.

In the present study, we took the next step and modelled the German stem priming patterns using naive discriminative learning (NDL). This morpheme-free computational model clarifies that the observed priming effects across all prime types may follow straightforwardly from basic principles of discrimination learning. The extent to which sublexical features of the prime (letter triphones) pre-activate the lexeme of the target is the strongest predictor for the reaction times. A substantially smaller effect emerged for the activation of the target (comparable to a frequency effect). The semantic similarity of prime and target as gauged by the cosine similarity measure also had a solid

effect.

The semantic vectors (word embeddings) used for calculating the cosine similarity between primes and their targets were taken from a database of German tweets. It is noteworthy that the cosine similarity measure provided good support for the transparent prime-target pairs being on average more semantically similar than the opaque prime-target pairs. However, this difference between the two prime types was not reflected in the corresponding mean reaction times.

A three-way interaction between prime-to-target pre-activation, target activation, and prime-target cosine similarity detected by a generalized additive mixed model fitted to the reaction times clarified that as the semantic similarity of primes and targets increases, the facilitatory effect of pre-activation decreases. Apparently, when primes and targets are more similar, pre-activation by the prime forces the cognitive system to slow down in order to resolve the near simultaneous activation of two very similar, but conflicting, meanings.

Interestingly, even though across prime types stimuli were matched for association ratings, these ratings were not predictive for reaction times. Stimuli were not matched across prime types for the cosines of the angles between primes' and targets' tweet-based semantic vectors, yet, surprisingly, these were predictive for reaction times. This finding is particularly surprising for the present data on German, as in previous work semantic similarity measures (not only human but also vector-based measures like LSA (Landauer and Dumais, 1997) and HAL (Lund and Burgess, 1996)) were observed not to be predictive of reaction times. It is conceivable that the present semantic measure is superior to LSA and HAL, due to it being calculated from a large volume of tweets — Herdağdelen and Marelli (2017) point out that measures based on distributional semantics calculated from corpora of social media provide excellent predictivity for lexical processing.

A caveat is in order, though, with respect to the cosine similarity measure, as in all likelihood particle verbs and their simple counterparts are estimated to be somewhat more similar than they should be. Particles can be separated by several words from their stems, and these stems will therefore be treated as simple verbs by the algorithm constructing semantic vectors (especially when the particle falls outside word2vec's 5-word moving window). This, however, implies that the cosine similarity measure must be less sensitive than it could have been: The vector for the simple verb is artifactually shifted in the direction of the vector of its particle verb, with the extent of this shift depending on the frequency of the simple verb, and the frequencies of the separated and non-separated derived particle verbs. As simple verbs typically occur with several particles, the semantic vectors for these simple verbs are likely to have shifted somewhat in the direction of the centroid of the vectors of its particle verbs.

Nevertheless, the present tweet-based semantic vectors contribute to the prediction of reaction times. Importantly, there is no a-priori reason for assuming that the rate at which particles occur separated from their stem would differ across prime types. As a consequence, the partial confounding of particle verbs and simple verbs by the distributional semantics algorithm generating the semantic vectors cannot be the main cause of the different distributions of cosine similarities across the different prime types. (In this context, it is worth noting that any skewing in frequency counts of complex and simple verbs does not affect the qualitative pattern of results for the predictor with the greatest effect size, the prime-to-target activation measure, as shown by the simulation in which all words are presented to the network with exactly the same frequency.) We leave answering the question of why this particular frequency measure appears to be effective as predictor for the present reaction times to further research.

An important result is that a generalized additive mixed model fitted to the reaction times provides a fit that is at least as good, if not better, when the activation, pre-activation, and cosine similarity measures are used, compared to when prime type and frequency of occurrence are used as predictors (evidence ratio: 27).

The theoretical contribution of this study is that it challenges the general localist interactive activation framework that dominates the current discourse on morphological processing. Stems and morphemes are assumed to be psychologically real (see, e.g., [Zwitserslood, 2018](#)), and to excite or inhibit each other. Furthermore, these high-level concepts are apparently understood to be sufficient for explaining the effects of experimental manipulations. NDL, by contrast, provides a framework within which quantitative measures can be derived that can be pitted against experimental response variables. NDL (and LDL) make use of the simplest possible network, the mathematics of which are well understood — in essence, NDL is nothing more (or less) than incremental multivariate multiple (logistic) regression. An NDL model is essentially parameter-free,¹⁰ and driven completely by the distributional properties of the words in the corpus it is trained on. NDL, just as the interactive activation framework, requires the analyst to make decisions on input and output nodes, but unlike the interactive activation framework, no hand-crafting of connections is required, and no search is required for finding a set of parameters that make the model behave in the way desired. As a consequence, the measures derived from an NDL network can be used simply as a statistical tool for assessing how well a word’s meaning can be ‘classified’ or ‘discriminated’ given its form features. As expected, reaction times become shorter when target meanings are better discriminated, i.e., when target activation is higher and the probability of the target being correctly classified is greater. Furthermore, when a morphologically related prime is presented to the network, a distribution of activations over the lexemes ensues in which the activation of the target is greater compared to trials with primes that are not morphologically related. This pre-activation of the target by reading the prime apparently carries over to the reading of the target.¹¹

Although NDL-based measures can be used in the same way as measures such as word frequency and neighborhood density, the linguistic theory underlying NDL holds that morphemes (in the sense of minimal signs) as well as sublexical form units such as stems and exponents are not necessary. At the same time, this theory is analytical at the semantic level. What [Baayen et al. \(2019\)](#) and [Baayen et al. \(2018\)](#) have shown for LDL is that accurate mappings between form vectors and distributional semantic vectors can be set up with linear transformations, i.e., with simple two-layer networks (and no hidden layers). The more comprehensive model of the mental lexicon developed in [Baayen et al. \(2019\)](#) and [Chuang et al. \(2019\)](#) makes use of multiple such networks to generate quantitative predictions for auditory comprehension (with audio as input), visual comprehension, and speech production. A proof of concept that inflected forms of rich paradigms can be predicted from their corresponding semantic vectors without requiring sublexical form units such as stems and exponents is provided by [Baayen et al. \(2018\)](#). It is within this wider context that the present computational modeling of overt primed visual lexical decision comes into its own.

It is important to note that the design decision to assign complex verbs their own lexeme, irrespective of semantic transparency, following [Baayen et al. \(2019\)](#), is crucial for enabling NDL to simulate the German behavioral priming data. This design decision is well motivated, as it is widely recognized that in word formation, in contrast to inflection, complex forms almost always

¹⁰ The learning rule of Rescorla and Wagner has several parameters that were introduced specifically to model differences in the salience of input features and the importance to the animal of different outcomes. In our implementation, we always set λ (representing the maximum amount of learning) to 1, and use a fixed learning rate (the product of the α and β parameters). In [Baayen et al. \(2011\)](#), the learning rate was 0.01, but subsequent work showed optimal performance when the learning rate is set to 0.001. For the simulation reported here, these values were used, and no simulations were run with different values.

¹¹ The NDL framework currently does not provide mechanisms that account for how the cognitive system reaches a response on the basis of lexeme activations. Analyses with the generalized additive model indicate non-linear interactions for these mechanisms. At present, all we can do is bring these nonlinearities to light with GAMs, in the hope that across experiments consistent patterns will emerge that then can be the subject of further computational modeling.

have their idiosyncratic shades of meaning, even when classified as ‘transparent’.

We have shown that the effect of morphological priming can be modeled precisely with a simple network that eschews morphemes and sublexical units such as stems, affixes, and exponents, a result that is consistent with Word and Paradigm Morphology (Blevins, 2016) and the model of the mental lexicon proposed in Baayen et al. (2019). Do these findings imply that morphemes or morphs do not have any psychological or cognitive reality? Answering this question is not straightforward.

First, it is logically possible that morphemes are actually cognitively real and crucially involved in the lexical processing of German verbs. In this case, NDL is no more than a machine learning algorithm that generates correct predictions, but for the wrong reasons. To give substance to this argumentation, it will be essential for alternative computational models to be put forward that make use of units representing morphemes.

Second, the sublexical cues (triphones) that are shared by prime and target drive the prime-to-target pre-activation. For the prime-target pair ANSTOSSEN, STOSSEN, the set of shared trigrams is **Sto**, **tos**, **os@**, **s@n**, **s@n**, **@n#**, and the set of trigrams unique to only one of the two words is **#an**, **anS**, **nSt**, and **#St**. Figure 6 presents the connection strengths of these cues to the lexome STOSSEN (upper panel) and ANSTOSSEN (lower panel). The cues that occur in both lexomes (the blue dots with inner red dots) support STOSSEN (upper panel), or ANSTOSSEN (lower panel) to exactly the same extent — as expected, as the cues are sublexical features that by definition must provide the same support for STOSSEN irrespective of whether they are embedded in the word form **Stos@n** or **anStos@n**.

Figure 6 brings to the fore two important points. First, the two central triphones of the stem that are shared by both the simple and the complex verb, **Sto** and **tos**, provide substantial contributions to the (pre-)activation of both STOSSEN and ANSTOSSEN, as expected. This observation fits well with the stem-frequency hypothesis, according to which the stem is the crucial unit mediating lexical access. Second, however, triphones at the boundary of the stem can have even greater strengths than these central triphones. For STOSSEN, this is the case for **os@**, and for ANSTOSSEN, this happens for both **nSt** and **os@**. The boundary triphone **s@n** also makes a non-negligible, albeit much smaller, contribution to the activation of these lexomes. Crucially, it is exactly here that NDL moves beyond the stem-frequency hypothesis. Triphones at the boundary of the stem often carry substantial discriminatory potential. The boundary triphone **os@** is important for discriminating (AN)STOSSEN from nominal lexemes such as TROSS, STOSS, SCHLOSS and adjectival lexemes such as GROSS, and the boundary triphone **nSt** is important for discriminating ANSTOSSEN from STOSSEN (note its high positive value as cue for ANSTOSSEN and its slightly negative value as cue for STOSSEN).

In other words, the present NDL model can be viewed as a refinement of the stem-frequency model that, by taking into account not only central sublexical features of the stem, but also the discriminatory potential of features at the stem boundary, achieves superior predictivity.¹²

Third, NDL (and LDL) target implicit learning, the continuous recalibration of the lexicon that goes on without conscious thought and attention, similar to the way that object recognition is continually recalibrated (Marsolek, 2008). The finding of Smolka et al. (2015) that morphological priming effects are visible in evoked response potentials even in the absence of behavioral correlates is consistent with NDL capturing subliminal lexical processing. However, we also reflect on language, we enjoy word play, we have poetry, we teach grammar in schools, and in second language teaching we instruct learners how to put words together from their parts. Patients suffering from a stroke may benefit from explicit instruction about how inflected words can be put together (Nault, 2010). This knowledge about language is cognitively real, and it may also affect lexicality decisions in

¹² For detailed discussion of the importance of sublexical features at the boundary of morphemic units from a discrimination learning perspective, see Baayen et al. (2016b). For frailty at stem boundaries in speech production, see Baayen et al. (2019).

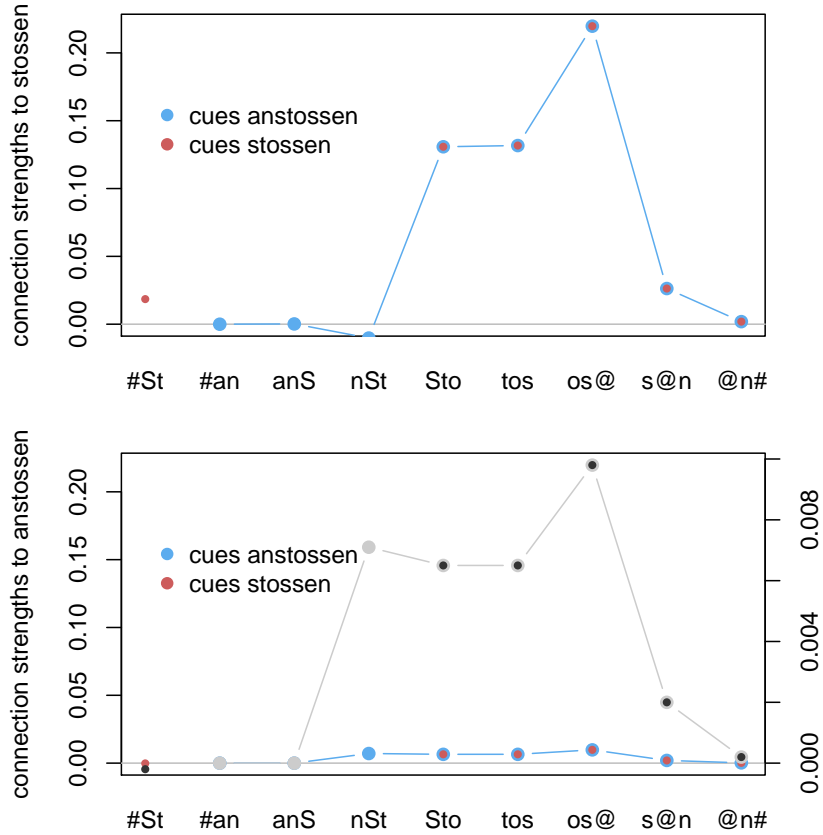


Figure 6: Connection weights from triphones to the lexeme STOSSEN (upper panel), and to the lexeme ANSTOSSEN (lower panel). As the primes are of much lower frequency of occurrence than the targets, the magnitude of weights in the lower panel is substantially reduced compared to the upper panel. In the lower panel, the blue and red disks are presented on the same scale as the upper panel. The disks in gray and black represent the same datapoints, but now on an enlarged scale, depicted at the right-hand side of the plot.

meta-linguistic tasks such as primed lexicality decision making. NDL and LDL, however, are blind to this higher-order knowledge.

Fourth, in LDL, assemblies of low-level units such as triphones may form, depending on a language’s morphological system, clusters in a two-dimensional space when predominantly used to realize specific inflectional or derivational functions (see Baayen et al., 2018, for detailed discussion). Such clusters (see also Chersi et al., 2014; Marzi et al., 2018) can be understood, from a functional perspective, as morphemes. Crucially, such ‘morphemes’ may emerge under constraints of topological self-organization, and do not require hand-crafting of discrete form units and connections between these units, as in the interactive activation framework. In other words, morphemes can be emergent properties of a complex dynamic system.

Fifth, NDL and LDL are based on wide learning networks, simple two-layer networks with large numbers of units that can be trained very quickly, the mathematics of which are well-understood, and that can perform with surprising accuracy given well-chosen input and output representations (see, e.g., Linke et al., 2017; Shafaei Bajestan and Baayen, 2018). Deep learning networks offer architectures in which units on hidden layers have the potential to become sensitive to, and in

some sense ‘represent’, morpheme-like units. Such networks are powerful statistical classifiers, but require decisions about the number of hidden layers, the number of units on these layers, and where to position convolutional and/or recurrent layers. Unfortunately, deep learning networks are widely recognized to have a ‘black box’ nature, although progress is being made towards understanding why they work (see, e.g., [Cheng et al., 2018](#); [Daniel and Yeung, 2019](#)). NDL and LDL are specifically designed to provide both interpretational transparency and accurate and falsifiable predictions.

Importantly, what sets both wide learning and deep learning apart from the interactive activation framework is, first, that the former models are dynamically learning classifiers whereas the latter approach builds on the idea of static classifier with a large number of parameters that have to be set manually, and second, that the former are end-to-end models whereas the latter solves only a partial task. The interactive activation framework is set up to select one word form and suppress all others, given visual input, but it remains silent about the semantics to which this form is supposed to provide access. This approach is still chained to the metaphor of the paper dictionary, in which form entries have to be located that, once found, provide access to meaning (see also [Elman, 2009](#), for detailed criticism of the dictionary metaphor). By contrast, wide learning, following many practical applications in computational linguistics that make use of deep learning, is set up to predict the ultimate true goal of comprehension: the semantics targeted by the input signal. The results obtained with NDL and LDL obtained thus far suggest that this goal can be reached without mediation by form units such as stems, affixes, or exponents. It is likely that future versions of the general LDL model of [Baayen et al. \(2019\)](#) will incorporate deep learning for some of its components, sacrificing interpretational transparency for increased accuracy. If morph-like units arise in such versions of the model, these units will not be part of a classical morphological calculus with symbolic representations and rules operating on these representations such as proposed by [Chomsky and Halle \(1968\)](#) and [Pinker \(1999\)](#). Their function would be to statistically integrate high-dimensional evidence for semantics in interaction with large numbers of other such units.

In the light of these considerations, it is clear that the present study cannot provide a full answer to the question of whether morphemes are, or are not, cognitively real. Clearly, the behavioral findings that have been interpreted as evidence for stem-driven lexical access are real. What the present study adds to this is that there actually is another possible interpretation for the observed priming effects, namely, prime-to-target pre-activation in a discriminative lexicon. The discriminative lexicon provides detailed quantitative predictions within a larger conceptual framework that is informed by recent developments in linguistic morphology, using a computational algorithm that is completely driven by the distributional properties of the lexicon, and that does not require (nor allow) tuning of free parameters to bring the model’s performance in line with the observed data. The present modeling results thus challenge proponents of morpheme-driven, decompositional, lexicons to demonstrate that their high-level conceptual theories can actually be made to work algorithmically. This in turn will make it possible to pit against each other the detailed quantitative predictions of morpheme-based and morpheme-free models.

References

- Arnold, D., Tomaschek, F., Lopez, F., Sering, T., and Baayen, R. H. (2017). Words from spontaneous conversational speech can be recognized with human-like accuracy by an error-driven learning algorithm that discriminates between meanings straight from smart acoustic features, bypassing the phoneme as recognition unit. *PLOS ONE*, 12(4):e0174623.
- Aronoff, M. (1994). *Morphology by Itself: Stems and Inflectional Classes*. The MIT Press, Cambridge, Mass.

- Baayen, R. H., Chuang, Y., and Blevins, J. P. (2018). Inflectional morphology with linear mappings. *The Mental Lexicon*, 13(2):232–270.
- Baayen, R. H., Chuang, Y.-Y., Shafaei-Bajestan, E., and Blevins, J. (2019). The discriminative lexicon: A unified computational model for the lexicon and lexical processing in comprehension and production grounded not in (de)composition but in linear discriminative learning. *Complexity*.
- Baayen, R. H., Dijkstra, T., and Schreuder, R. (1997). Singulars and plurals in Dutch: Evidence for a parallel dual route model. *Journal of Memory and Language*, 36:94–117.
- Baayen, R. H., McQueen, J., Dijkstra, T., and Schreuder, R. (2003). Frequency effects in regular inflectional morphology: Revisiting Dutch plurals. In Baayen, R. H. and Schreuder, R., editors, *Morphological structure in language processing*, pages 355–390. Mouton de Gruyter, Berlin.
- Baayen, R. H., Milin, P., Filipović Durdević, D., Hendrix, P., and Marelli, M. (2011). An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. *Psychological Review*, 118:438–482.
- Baayen, R. H., Milin, P., and Ramscar, M. (2016a). Frequency in lexical processing. *Aphasiology*, 30(11):1174–1220.
- Baayen, R. H., Piepenbrock, R., and van Rijn, H. (1993). *The CELEX lexical database (CD-ROM)*. Linguistic Data Consortium, University of Pennsylvania, Philadelphia, PA.
- Baayen, R. H., Shaoul, C., Willits, J., and Ramscar, M. (2016b). Comprehension without segmentation: A proof of concept with naive discriminative learning. *Language, Cognition, and Neuroscience*, 31(1):106–128.
- Baayen, R. H., Vasishth, S., Bates, D., and Kliegl, R. (2017). The cave of shadows. Addressing the human factor with generalized additive mixed models. *Journal of Memory and Language*, 94:206–234.
- Bates, D. M., Kliegl, R., Vasishth, S., and Baayen, R. H. (2015). Parsimonious mixed models.
- Bauer, L., Lieber, R., and Plag, I. (2015). *The Oxford reference guide to English morphology*. Oxford University Press.
- Beard, R. (1977). On the extent and nature of irregularity in the lexicon. *Lingua*, 42:305–341.
- Blevins, J. P. (2003). Stems and paradigms. *Language*, 79:737–767.
- Blevins, J. P. (2006). Word-based morphology. *Journal of Linguistics*, 42(03):531–573.
- Blevins, J. P. (2016). *Word and paradigm morphology*. Oxford University Press.
- Booij, G. E. (1986). Form and meaning in morphology: the case of Dutch àgent nouns. *Linguistics*, 24:503–517.
- Butz, M. V. and Kutter, E. F. (2016). *How the mind comes into being: Introducing cognitive science from a functional and computational perspective*. Oxford University Press.
- Cheng, X., Khomtchouk, B., Matloff, N., and Mohanty, P. (2018). Polynomial regression as an alternative to neural nets. *arXiv preprint arXiv:1806.06850*.

- Chersi, F., Ferro, M., Pezzulo, G., and Pirrelli, V. (2014). Topological self-organization and prediction learning support both action and lexical chains in the brain. *Topics in cognitive science*, 6(3):476–491.
- Chomsky, N. and Halle, M. (1968). *The sound pattern of English*. Harper and Row, New York.
- Chuang, Y., Vollmer, M.-L., Shafaei-Bajestan, E., Gahl, S., Hendrix, P., and Baayen, R. H. (2019). The processing of nonword form and meaning in production and comprehension: A computational modeling approach using linear discriminative learning. Manuscript, University of Tübingen.
- Cieliebak, M., Dertu, J., Uzdilli, F., and Egger, D. (2017). A twitter corpus and benchmark resources for german sentiment analysis. In *Proceedings of the 4th International Workshop on Natural Language Processing for Social Media (SocialNLP 2017)*, Valencia, Spain.
- Cremer, A., Goodwin Davies, A., Wilder, R. J., Tamminga, M., and Embick, D. (2019). Opacity, transparency, and morphological priming: A study of prefixed verbs in Dutch. Manuscript submitted for publication.
- Daniel, W. B. and Yeung, E. (2019). A constructive approach for one-shot training of neural networks using hypercube-based topological coverings. *arXiv preprint arXiv:1901.02878*.
- Danks, D. (2003). Equilibria of the Rescorla-Wagner model. *Journal of Mathematical Psychology*, 47(2):109–121.
- De Grauwe, S., Lemhöfer, K., and Schriefers, H. (2019). Processing derived verbs: the role of motor-relatedness and type of morphological priming. *Language, Cognition and Neuroscience*, pages 1–18.
- Deriu, J., Lucchi, A., De Luca, V., Severyn, A., Mller, S., Cieliebak, M., Hoffmann, T., and Jaggi, M. (2017). Leveraging large amounts of weakly supervised data for multi-language sentiment classification. In *Proceedings of the 26th International World Wide Web Conference (WWW-2017)*, Perth, Australia.
- Divjak, D., Milin, P., and Baayen, R. H. (2017). A learning perspective on individual differences in skilled reading: Exploring and exploiting orthographic and semantic discrimination cues. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 43(11):1730–1751.
- Dunabeitia, J., Kinoshita, S., Carreiras, M., and Norris, D. (2011). Is morpho-orthographic decomposition purely orthographic? evidence from masked priming in the same–different task. *Language and Cognitive Processes*, 26(4-6):509–529.
- Eisenberg, P. (2004). *Grundriss der deutschen Grammatik: Das Wort (2nd ed. Vol. 1)*. J. B. Metzler, Stuttgart-Weimar.
- Eliasmith, C., Stewart, T. C., Choo, X., Bekolay, T., DeWolf, T., Tang, Y., and Rasmussen, D. (2012). A large-scale model of the functioning brain. *science*, 338(6111):1202–1205.
- Elman, J. L. (2009). On the meaning of words and dinosaur bones: Lexical knowledge without a lexicon. *Cognitive science*, 33(4):547–582.
- Fasiolo, M., Goude, Y., Nedellec, R., and Wood, S. (2017). Fast calibrated additive quantile regression. <https://github.com/fasiolo/qgam>.

- Feldman, L. and Prostko, B. (2001). Graded aspects of morphological processing: Task and processing time. *Brain and Language*, 81:1–16.
- Feldman, L. B., O’Connor, P. A., and Moscoso del Prado Martin, F. (2009). Early morphological processing is morpho-semantic and not simply morpho-orthographic: evidence from the masked priming paradigm. *Psychonomic Bulletin & Review*, 16(4):684–691.
- Feldman, L. B., Soltano, E. G., Pastizzo, M. J., and Francis, S. E. (2004). What do graded effects of semantic transparency reveal about morphological processing? *Brain and Language*, 90:17–30.
- Fleischer, W. and Barz, I. (1992). Wortbildung der deutschen gegenwartssprache. tübingen.
- Frauenfelder, U. H. and Schreuder, R. (1992). Constraining psycholinguistic models of morphological processing and representation: The role of productivity. In Booij, G. E. and Marle, J. v., editors, *Yearbook of Morphology 1991*, pages 165–183. Kluwer Academic Publishers, Dordrecht.
- Gonnerman, L. and Anderson, E. (2001). Graded semantic and phonological similarity effects in morphologically complex words. In Bendjaballah, S., Dressler, W., Pfeiffer, O. E., and Voeikova, M. D., editors, *Morphology 2000: Selected papers from the 9th Morphology meeting*, pages 137–148. John Benjamins, Amsterdam.
- Gonnerman, L. M., Seidenberg, M. S., and Andersen, E. S. (2007). Graded semantic and phonological similarity effects in priming: Evidence for a distributed connectionist approach to morphology. *Journal of experimental psychology: General*, 136(2):323.
- Grainger, J., Dufau, S., Montant, M., Ziegler, J. C., and Fagot, J. (2012). Orthographic processing in baboons (papio papio). *Science*, 336(6078):245–248.
- Günther, F., Smolka, E., and Marelli, M. (2018). understandingdiffers between english and german: Capturing systematic language differences of complex words. *Cortex*.
- Gurney, K., Prescott, T., and Redgrave, P. (2001). A computational model of action selection in the basal ganglia. *Biol. Cybern.*, 84:401–423.
- Hannagan, T., Ziegler, J. C., Dufau, S., Fagot, J., and Grainger, J. (2014). Deep learning of orthographic representations in baboons. *PLOS-one*, 9:e84843.
- Hannun, A., Case, C., Casper, J., Catanzaro, B., Diamos, G., Elsen, E., Prenger, R., Satheesh, S., Sengupta, S., Coates, A., et al. (2014). Deep speech: Scaling up end-to-end speech recognition. *arXiv preprint arXiv:1412.5567*.
- Harm, M. W. and Seidenberg, M. S. (2004). Computing the meanings of words in reading: Cooperative division of labor between visual and phonological processes. *Psychological Review*, 111:662–720.
- Herdağdelen, A. and Marelli, M. (2017). Social media and language processing: How facebook and twitter provide the best frequency estimates for studying word recognition. *Cognitive science*, 41(4):976–995.
- Kappel, D., Habenschuss, S., Legenstein, R., and Maass, W. (2015). Network plasticity as bayesian inference. *PLoS computational biology*, 11(11):e1004485.
- Kappel, D., Legenstein, R., Habenschuss, S., Hsieh, M., and Maass, W. (2017). Reward-based stochastic self-configuration of neural circuits. *arXiv preprint 1704.04238*.

- Kielar, A. and Joanisse, M. F. (2011). The role of semantic and phonological factors in word recognition: An erp cross-modal priming study of derivational morphology. *Neuropsychologia*, 49(2):161–177.
- Kuperman, V., Bertram, R., and Baayen, R. H. (2008). Morphological dynamics in compound processing. *Language and Cognitive Processes*, 23:1089–1132.
- Kuperman, V., Bertram, R., and Baayen, R. H. (2010). Processing trade-offs in the reading of Dutch derived words. *Journal of Memory and Language*, 62:83–97.
- Kuperman, V., Schreuder, R., Bertram, R., and Baayen, R. H. (2009). Reading of multimorphemic Dutch compounds: Towards a multiple route model of lexical processing. *Journal of Experimental Psychology: HPP*, 35:876–895.
- Landauer, T. and Dumais, S. (1997). A solution to Plato’s problem: The latent semantic analysis theory of acquisition, induction and representation of knowledge. *Psychological Review*, 104(2):211–240.
- Lavric, A., Rastle, K., and Clapp, A. (2011). What do fully visible primes and brain potentials reveal about morphological decomposition? *Psychophysiology*, 48(5):676–686.
- Lazaridou, A., Marelli, M., Zamparelli, R., and Baroni, M. (2013). Compositionally derived representations of morphologically complex words in distributional semantics. In *ACL (1)*, pages 1517–1526.
- Linke, M., Broeker, F., Ramscar, M., and Baayen, R. H. (2017). Are baboons learning “orthographic” representations? probably not. *PLOS-ONE*, 12(8):e0183876.
- Longtin, C., Segui, J., and Hallé, P. (2003). Morphological priming without morphological relationship. *Language and Cognitive Processes*, in press:0.
- Lund, K. and Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behaviour Research Methods, Instruments, and Computers*, 28(2):203–208.
- Marelli, M., Amenta, S., Morone, E. A., and Crepaldi, D. (2013). Meaning is in the beholder’s eye: Morpho-semantic effects in masked priming. *Psychonomic bulletin & review*, 20(3):534–541.
- Marelli, M. and Baroni, M. (2015). Affixation in semantic space: Modeling morpheme meanings with compositional distributional semantics. *Psychological Review*, 122(3):485.
- Marslen-Wilson, W., Older, L., Janssen, D., and Waksler, S. (1994a). Stems and affixes: local and non-local attachment. *Paper presented to the Experimental Psychology Society, London, January 1994*.
- Marslen-Wilson, W., Tyler, L. K., Waksler, R., and Older, L. (1994b). Morphology and meaning in the English mental lexicon. *Psychological Review*, 101:3–33.
- Marsolek, C. J. (2008). What antipriming reveals about priming. *Trends in Cognitive Science*, 12(5):176–181.
- Marzi, C., Ferro, M., and Pirrelli, V. (2018). Is inflectional irregularity dysfunctional to human processing? In Kuperman, V., editor, *Abstract booklet, The Mental Lexicon 2018*, page 60. University of Alberta, Edmonton.

- Matthews, P. H. (1974). *Morphology. An Introduction to the Theory of Word Structure*. Cambridge University Press, Cambridge.
- McClelland, J. L. and Rumelhart, D. E., editors (1986). *Parallel Distributed Processing. Explorations in the Microstructure of Cognition. Vol. 2: Psychological and Biological Models*. MIT Press, Cambridge, Mass.
- Meunier, F. and Longtin, C. (2007). Morphological decomposition and semantic integration in word processing. *Journal of Memory and Language*, 56(4):457–471.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Milin, P., Feldman, L. B., Ramscar, M., Hendrix, P., and Baayen, R. H. (2017a). Discrimination in lexical decision. *PLoS-one*, 12(2):e0171935.
- Milin, P., Smolka, E., and Feldman, L. B. (2017b). Models 11 of lexical access and morphological processing. *The handbook of psycholinguistics*, page 240.
- Nault, K. (2010). *Morphological Therapy Protocol*. Ph. D. thesis, University of Alberta, Edmonton.
- Norris, D. and Kinoshita, S. (2008). Perception as evidence accumulation and bayesian inference: Insights from masked priming. *Journal of Experimental Psychology*, 137(3):434–455.
- Pastizzo, M. J. and Feldman, L. B. (2002). Does prime modality influence morphological processing? *Brian and Language*, 81:28–41.
- Pinker, S. (1999). *Words and Rules: The Ingredients of Language*. Weidenfeld and Nicolson, London.
- Plag, I. (2003). *Word-formation in English*. Cambridge University Press, Cambridge.
- Plaut, D. C. and Gonnerman, L. M. (2000). Are non-semantic morphological effects incompatible with a distributed connectionist approach to lexical processing? *Language and Cognitive Processes*, 15(4/5):445–485.
- Rastle, K. and Davis, M. (2008). Morphological decomposition based on the analysis of orthography. *Language and Cognitive Processes*, 23(7-8):942–971.
- Rastle, K., Davis, M. H., Marslen-Wilson, W., and Tyler, L. K. (2000). Morphological and semantic effects in visual word recognition: A time-course study. *Language and Cognitive Processes*, 15:507–537.
- Rastle, K., Davis, M. H., and New, B. (2004). The broth in my brother’s brothel: Morpho-orthographic segmentation in visual word recognition. *Psychonomic Bulletin & Review*, 11:1090–1098.
- Redgrave, P., Prescott, T., and Gurney, K. (1999). The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience*, 89:1009–1023.
- Rescorla, R. A. and Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In Black, A. H. and Prokasy, W. F., editors, *Classical conditioning II: Current research and theory*, pages 64–99. Appleton Century Crofts, New York.

- Roelofs, A. (1997a). A case for nondecomposition in conceptually driven word retrieval. *Journal of Psycholinguistic Research*, 26:33–67.
- Roelofs, A. (1997b). Morpheme frequency in speech production: Testing WEAVER. In Booij, G. E. and Van Marle, J., editors, *Yearbook of Morphology 1996*, pages 135–154. Kluwer, Dordrecht.
- Roelofs, A., Baayen, R. H., and van den Brink, D. (2002). Morphology by itself in planning the production of spoken words. *Psychonomic bulletin and review*, 9:132–138.
- Rumelhart, D. E. and McClelland, J. L. (1986). On learning the past tenses of English verbs. In McClelland, J. L. and Rumelhart, D. E., editors, *Parallel Distributed Processing. Explorations in the Microstructure of Cognition. Vol. 2: Psychological and Biological Models*, pages 216–271. The MIT Press, Cambridge, Mass.
- Schmidtke, D., Matsuki, K., and Kuperman, V. (2017). Surviving blind decomposition: A distributional analysis of the time-course of complex word recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43(11):1793.
- Schreuder, R. (1990). Lexical processing of verbs with separable particles. *Yearbook of Morphology*, 3:65–79.
- Schreuder, R., Burani, C., and Baayen, R. H. (2003). Parsing and semantic opacity. In Assink, E. and Sandra, D., editors, *Reading Complex Words. Cross-language studies*, pages 159–189. Kluwer, Dordrecht.
- Schreuder, R., Grendel, M., Poulisse, N., Roelofs, A., and Voort, M. v. d. (1990). Lexical processing, morphological complexity and reading. In Balotta, D. A., Flores d’Arcais, G. B., and Rayner, K., editors, *Comprehension Processes in Reading*. Lawrence Erlbaum Associates, London.
- Schriefers, H., Zwitserlood, P., and Roelofs, A. (1991). The identification of morphological complex spoken words: Continuous processing or decomposition? *Journal of Memory and Language*, 30:26–47.
- Sering, K., Weitz, M., Kuenstle, D., and Schneider, L. (2018a). Pyndl: Naive discriminative learning in python.
- Sering, T., Milin, P., and Baayen, R. H. (2018b). Language comprehension as a multiple label classification problem. *Statistica Neerlandica*, pages 1–15.
- Shafaei Bajestan, E. and Baayen, R. H. (2018). Wide learning for auditory comprehension. In *Interspeech*.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. (2016). Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484.
- Smolka, E. and Eulitz, C. (2011). Asymmetric meaning assembly for semantically transparent and opaque complex verbs in german. Paper presented at the Seventh International Morphological Processing Conference (MOPROC), San Sebastian.
- Smolka, E. and Eulitz, C. (2018). Psycholinguistic measures for german verb pairs: Semantic transparency, semantic relatedness, verb family size, and age of reading acquisition. *Behavior Research Methods*, 50(4):1540–1562.

- Smolka, E., Gondan, M., and Rösler, F. (2015). Take a stand on understanding: Electrophysiological evidence for stem access in German complex verbs. *Frontiers in Human Neuroscience*, 9(62).
- Smolka, E., Komlosi, S., and Rösler, F. (2009). When semantics means less than morphology: The processing of German prefixed verbs. *Language and Cognitive Processes*, 24(3):337–375.
- Smolka, E., Libben, G., and Dressler, W. U. (2019). When morphological structure overrides meaning: evidence from German prefix and particle verbs. *Language, Cognition and Neuroscience*, 34(5):599–614.
- Smolka, E., Preller, K. H., and Eulitz, C. (2014). ‘verstehen’(‘understand’) primes ‘stehen’(‘stand’): Morphological structure overrides semantic compositionality in the lexical representation of German complex verbs. *Journal of Memory and Language*, 72:16–36.
- Smolka, E., Zwitserlood, P., and Rösler, F. (2007). Stem access in regular and irregular inflection: Evidence from German participles. *Journal of Memory and Language*, 57(3):325–347.
- Stewart, T. C., Bekolay, T., and Eliasmith, C. (2012). Learning to select actions with spiking neurons in the basal ganglia. *Frontiers in Neuroscience*, 6.
- Stump, G. (2001). *Inflectional Morphology: A Theory of Paradigm Structure*. Cambridge University Press.
- Taft, M. and Forster, K. I. (1975). Lexical storage and retrieval of prefixed words. *Journal of Verbal Learning and Verbal Behavior*, 14:638–647.
- Tomaschek, F., Plag, I., Ernestus, M., and Baayen, R. H. (2019). Modeling the duration of word-final s in English with naive discriminative learning. *Journal of Linguistics*, page in press.
- van Rij, J., Wieling, M., Baayen, R. H., and van Rijn, H. (2017). *itsadug*: Interpreting time series and autocorrelated data using GAMMs. R package version 2.3.
- Veríssimo, J. (2018). Taking it a level higher: The leia model of complex word recognition.
- Wood, S. N. (2017). *Generalized Additive Models*. Chapman & Hall/CRC, New York.
- Zwitserlood, P. (2018). Processing and representation of morphological complexity in native language comprehension and production. In Booij, G. E., editor, *The construction of words. Advances in construction morphology*, pages 583–602. Springer.

Appendix

Legenda:

Frequency = Lemma Frequency per one million from the CELEX database; AoR = Mean Age of Acquisition; AoR = Mean Age of Reading; Association = Mean Semantic Association between prime and target (on a scale from 1 “strongly related” to 7 “unrelated”); the asterisks indicate targets that occur in both the transparent and opaque lists.

Table 7: Targets, semantically transparent primes, and controls. The asterisks indicate targets that occur in both the transparent and opaque lists.

Target	Frequency	Related	Frequency	Association	Unrelated	Frequency
binden*	58	zubinden	1	5.67	verrenken	1
bleiben*	799	hierbleiben	2	6.31	entziffern	2
fliegen	106	fortfliegen	0	5.71	zuschnüren	0
führen*	423	anführen	17	5.50	austauschen	15
holen*	120	abholen	16	5.44	verkürzen	14
hören*	383	anhören	17	5.94	vermissen	20
laufen*	217	fortlaufen	3	4.78	verzieren	3
lesen*	159	vorlesen	6	6.05	anfassen	6
leuchten*	24	beleuchten	11	5.67	zerbrechen	13
malen	21	abmalen	1	5.83	verqualmen	1
mischen*	18	vermischen	6	5.50	beneiden	5
reißen	72	zerreißen	14	5.24	bedanken	15
schenken*	63	beschenken	3	6.36	auslachen	2
schicken	58	verschicken	2	6.33	abblättern	2
schießen	75	beschiessen	4	5.39	abdrucken	4
schlafen	67	ausschlafen	1	5.82	beschmutzen	2
stossen*	101	anstossen	3	5.24	verstauben	3
stürzen*	59	einstürzen	1	4.83	wegfegen	1
täuschen*	27	vortäuschen	4	6.00	beeilen	4
tragen	333	hintragen	1	6.14	ankleben	1
wechseln*	42	auswechseln	6	6.56	abzeichnen	6
werfen*	152	bewerfen	3	5.59	aufräumen	3

Table 8: Targets, semantically opaque primes, and controls

Target	Frequency	Related	Frequency	Association	Unrelated	Frequency
binden*	58	entbinden	5	2.33	anlehnen	4
bleiben*	799	unterbleiben	5	1.50	beeilen	4
brechen	61	verbrechen	5	2.00	auffinden	6
führen*	423	verführen	9	1.67	einordnen	8
holen*	120	ausholen	2	1.56	auslachen	2
hören*	383	aufhören	33	1.06	vermissen	20
laufen*	217	auslaufen	6	2.64	anfassen	6
lesen*	159	auslesen	0	1.58	abschminken	0
leuchten*	24	einleuchten	5	2.81	beneiden	5
mischen*	18	einmischen	5	2.69	abzeichnen	6
passen	54	aufpassen	8	2.00	vorangehen	10
schenken*	63	einschenken	1	1.92	auftanken	1
schneiden	30	überschneiden	2	2.88	beschmutzen	2
stechen	12	bestechen	5	1.50	abdrucken	4
steigen	209	umsteigen	2	2.39	verstauben	3
stossen*	101	verstossen	14	2.11	verkürzen	14
stürzen*	59	bestürzen	3	1.71	verzieren	3
täuschen*	27	enttäuschen	34	2.73	bedrohen	47
tragen	333	ertragen	16	2.67	bedanken	15
treffen*	256	auftreffen	1	2.61	verrenken	1
wechseln*	42	verwechseln	13	2.67	zerschlagen	15
werfen*	152	entwerfen	20	1.75	verbrennen	24

Table 9: Targets, synonym primes, and controls

Target	Frequency	Prime	Frequency	Association	Prime	Frequency
binden	58	zuschnüren	0	5.44	abholzen	0
fangen	72	erwischen	17	4.92	ausbreiten	13
führen	423	vorangehen	10	5.88	belehren	10
greifen	104	anfassen	6	5.44	bescheren	7
holen	120	herbringen	1	6.17	verrenken	1
laufen	217	nachrennen	1	4.93	abtrocknen	1
lehnen	67	abstützen	2	4.29	abblättern	2
malen	21	abzeichnen	6	4.58	beneiden	5
ordnen	24	aufräumen	3	5.33	verzieren	3
schicken	58	zusenden	3	6.62	verstauben	3
schlafen	67	einnicken	0	5.76	abschminken	0
schlagen	124	verprügeln	5	6.33	beeilen	4
schliessen	212	absperren	4	5.92	abdrucken	4
schneiden	30	zerteilen	1	5.43	wegfegen	1
schreien	72	anbrüllen	1	5.73	ankleben	1
schützen	58	behüten	2	6.65	abfeuern	3
stossen	101	anrempeeln	1	5.07	verqualmen	1
stürzen	59	umfallen	2	4.71	entziffern	2
täuschen	27	belügen	3	5.25	auslachen	2
treffen	256	versammeln	24	5.83	verbrennen	24
wechseln	42	austauschen	15	6.06	verwalten	16
werfen	152	wegschmeissen	0	5	vordrängeln	0

Table 10: Targets, form-related primes, and controls

Target	Frequency	Prime	Frequency	Association	Prime	Frequency
binden	58	abbilden	2	1.08	auslachen	2
brechen	61	verbrennen	24	1.11	versammeln	24
dienen	135	addieren	4	1.00	beeilen	4
finden	787	verfinstern	1	1.00	ankleben	1
fliegen	106	zerfliessen	1	1.31	zerteilen	1
führen	423	befühlen	3	1.19	verstauben	3
lehnen	67	belehren	10	NA	verkürzen	14
mischen	18	vermissen	20	1.00	bedanken	15
schenken	63	ausschelten	2	NA	abblättern	2
schiessen	75	anschieben	0	1.00	vordrängeln	0
schneiden	30	beschneien	0	1.78	abschminken	0
schreien	72	abschreiben	8	1.00	vorangehen	10
schützen	58	verschütten	6	1.21	abzeichnen	6
schwingen	19	entschwinden	2	1.14	beschmutzen	2
stechen	12	einstecken	2	1.55	auftanken	1
steigen	209	versteifen	2	1.19	umfallen	2
stossen	101	verstopfen	3	NA	beneiden	5
stürzen	59	bestürmen	1	2.89	abtrocknen	1
tragen	333	betraten	36	NA	bedrohen	47
treffen	256	zertrennen	0	1.00	abholzen	0
wechseln	42	aufwecken	1	1.25	verqualmen	1
zwingen	112	zuzwinkern	0	1.13	abwiegen	0