



Modeling Morphology With Linear Discriminative Learning: Considerations and Design Choices

Maria Heitmeier*, Yu-Ying Chuang and R. Harald Baayen

Department of Linguistics, Eberhard-Karls Universität, Tübingen, Germany

This study addresses a series of methodological questions that arise when modeling inflectional morphology with Linear Discriminative Learning. Taking the semi-productive German noun system as example, we illustrate how decisions made about the representation of form and meaning influence model performance. We clarify that for modeling frequency effects in learning, it is essential to make use of incremental learning rather than the end-state of learning. We also discuss how the model can be set up to approximate the learning of inflected words in context. In addition, we illustrate how in this approach the wug task can be modeled. The model provides an excellent memory for known words, but appropriately shows more limited performance for unseen data, in line with the semi-productivity of German noun inflection and generalization performance of native German speakers.

OPEN ACCESS

Edited by:

Juhani Järviö,
University of Alberta, Canada

Reviewed by:

Xiaowei Zhao,
Emmanuel College, United States
Antti Arppe,
University of Alberta, Canada

*Correspondence:

Maria Heitmeier
maria.heitmeier@uni-tuebingen.de

Specialty section:

This article was submitted to
Language Sciences,
a section of the journal
Frontiers in Psychology

Received: 04 June 2021

Accepted: 04 October 2021

Published: 15 November 2021

Citation:

Heitmeier M, Chuang Y-Y and
Baayen RH (2021) Modeling
Morphology With Linear Discriminative
Learning: Considerations and Design
Choices. *Front. Psychol.* 12:720713.
doi: 10.3389/fpsyg.2021.720713

Keywords: German nouns, linear discriminative learning, semi-productivity, multivariate multiple regression, Widrow-Hoff learning, frequency of occurrence, semantic roles, wug task

1. INTRODUCTION

Computational models of morphology fall into two broad classes. The first class addresses the question of how to produce a morphologically complex word given a morphologically related form (often a stem, or an identifier of a stem or lexeme) and a set of inflectional or derivational features. We refer to these models as form-oriented models. The second class comprises models seeking to understand the relation between words' forms and their meanings. We refer to these models as meaning-oriented models.

Prominent form-oriented models comprise Analogical Modeling of Language (AML; Skousen, 1989, 2002) and Memory Based Learning (MBL; Daelemans and Van den Bosch, 2005), which are nearest-neighbor classifiers. Input to these models are tables with observations (words) in rows, and factorial predictors and a factorial response in columns. The response specifies an observation's outcome class (e.g., an allomorph), and the model is given the task to predict the outcome classes from the other predictor variables (for allomorphy, specifications of words' phonological make-up). Predictions are based on sets of nearest neighbors, serving as constrained exemplar sets for generalization. These models have clarified morphological phenomena ranging from the allomorphy of the Dutch diminutive (Daelemans et al., 1995) to stress assignment in English (Arndt-Lappe, 2011).

Ernestus and Baayen (2003) compared the performance of the MBL, AML, and Generalized Linear Models (GLM), as well as a recursive partitioning tree (Breiman et al., 1984), on the task of predicting whether word-final obstruents in Dutch alternate with respect to their voicing. They observed similar performance across all models, with the best performance, surprisingly,

for the only parameter-free model, AML. Their results suggest that the quantitative structure of morphological datasets may be straightforward to discover for any reasonably decent classifier. The model proposed by Belth et al. (2021) is a recent example of a classifier based on recursive partitioning.

Minimum Generalization Learning (MGL; Albright and Hayes, 2003) offers an algorithm for rule induction (for comparison with nearest neighbor methods, see Keuleers et al., 2007). The model finds rules by an iterative process of minimal generalization that combines specific rules into ever more general rules. Each rule comes with a measure of prediction accuracy, and the rule with the highest accuracy is selected for predicting a word's form.

All models discussed thus far are exemplar-based, in the sense that the input to any of these models consists of a table with exemplars, exemplar features selected on the basis of domain knowledge, and a categorical response variable specifying targeted morphological form changes. In other words, all these models are classifiers that absolve the analyst from hand-engineering lexical entries, rules or constraints operating on these lexical entries, and theoretical constructs such as inflectional classes. In this respect, they differ fundamentally from the second group of the following computational methods.

The DATR language (Evans and Gazdar, 1996) defines non-monotonic inheritance networks for knowledge representation. This language is optimized for removing redundancy from lexical descriptions. A DATR model requires the analyst to set up lexical entries that specify information about, for instance, inflectional class, gender, the forms of exponents, and various kinds of phonological information. The lexicon is designed in such a way that the network is kept as small as possible, while still allowing the model, through its mechanism of inheritance, to correctly predict all inflected variants. Realizational morphology (RM; Stump, 2001) sets up rules for realizing bundles of inflectional and lexical features in phonological form. This theory can also be defined as a formal language (a finite-state transducer) that provides mappings from underlying representations onto their corresponding surface forms and vice versa (Karttunen, 2003). The Gradual Learning Algorithm (GLA; Boersma, 1998; Boersma and Hayes, 2001) works within the framework of optimality theory (Prince and Smolensky, 2008). The algorithm is initialized with a set of constraints and gradually learns an optimal constraint ranking by incrementally moving through the training data, and upgrading or downgrading constraints.

The third group of form-oriented computational models comprises connectionist models. The past-tense model of Rumelhart and McClelland (1986) was trained to produce English past-tense forms given the corresponding present-tense form. An early enhancement of this model was proposed by MacWhinney and Leinbach (1991), for an overview of the many follow-up models, see Kirov and Cotterell (2018). Kirov and Cotterell proposed a sequence-to-sequence deep learning network, the Encoder-Decoder (ED) learner, that they argue does not suffer from the drawbacks noted by Pinker and Prince (1988) for the original paste-tense model. Malouf (2017) introduced a recurrent deep learning model trained to predict upcoming

segments, showing that this model has high accuracy for predicting paradigm forms given the lexeme and the inflectional specifications of the desired paradigm cell.

In summary, the class of form-oriented models comprises three subsets: statistical classifiers (AML, MBL, GLM, recursive partitioning), generators based on linguistic knowledge engineering (DATR, RM, GLA), and connectionist models (paste-tense model, ED learner). The models just referenced presuppose that when speakers use a morphologically complex form, this form is derived on the fly from its underlying form. The sole exception is the model of Malouf (2017), which takes the lexeme and its inflectional features as point of departure. As pointed out by Blevins (2016), the focus on how to create one form from another has its origin in pedagogical grammars, which face the task of clarifying to a second language learner how to create inflected variants. Unsurprisingly, applications within natural language processing also have need of systems that can generate inflected and derived words.

However, it is far from self-evident that native speakers of English would create past-tense forms from present-tense forms. Meaning-oriented models argue that in comprehension, the listener or reader can go straight from the auditory or visual input to the intended meaning, without having to go through a pipeline requiring initial identification of underlying forms and exponents. Likewise, speakers are argued to start from meaning, and realize this meaning directly in written or spoken form.

The class of meaning-oriented models comprises both symbolic and subsymbolic models. The symbolic models of Dell (1986) and Levelt et al. (1999) implement a form of realizational morphology. Concepts and inflectional features activate stems and exponents, which are subsequently combined into syllables. Both models hold that the production of morphologically complex words is a compositional process in which units are assembled together and ordered for articulation at various hierarchically ordered levels. These models have been worked out only for English, and to our knowledge have not been applied to languages with richer morphological systems.

The subsymbolic model of Harm and Seidenberg (2004) sets up multi-layer networks between orthographic, phonological, and semantic units. No attempt is made to define morphemes, stems, or exponents. To the extent that such units have any reality, they are assumed to arise, statistically, at the hidden layers. Mirković et al. (2005) argue for Serbian that gender is an emergent property of the network that arises from statistical regularities governing both words' forms and their meanings (see Corbett, 1991, for discussion of semantic motivations for gender systems). The model for auditory comprehension of Gaskell and Marslen-Wilson (1997) uses a three-layer recurrent network to map speech input onto distributed semantic representations, again without attempting to isolate units such as phonemes or morphemes.

The naive discrimination learning (NDL) model proposed by Baayen et al. (2011) represents words' forms sub-symbolically, but words' meanings symbolically. The modeling set-up that we discuss in the remainder of this study, that of linear discriminative learning (LDL, Baayen et al., 2019), replaces the symbolic representation of word meaning in

NDL by sub-symbolic representations building on distributional semantics (Landauer and Dumais, 1997; Mikolov et al., 2013b).

LDL is an implementation of Word and Paradigm Morphology (Matthews, 1974; Blevins, 2016). Sublexical units such as stems and exponents play no role. Semantic representations in LDL, however, are analytical: the semantic vector (word embedding, i.e., a distributed representation of meaning) of an inflected word is constructed from the semantic vector of the lexeme and the semantic vectors of the pertinent inflectional functions. Both NDL and LDL make use of the simplest possible networks: networks with only input and output layers, and no hidden layers.

At this point, the distinction made by Breiman (2001) between statistical models and machine learning is relevant. Statistical models aim to provide insight into the mechanisms that generate the data. Machine learning, on the other hand, aims to optimize prediction accuracy, and it is not an issue whether or not the algorithms are interpretable. LDL is much closer to statistical modeling than to the black boxes of machine learning. All input and output representations can be set up in a theoretically transparent way (Baayen et al., 2019). Furthermore, because LDL implements multivariate multiple regression, its mathematical properties are well-understood. Importantly, modeling results do not depend on the choice of hyper-parameters (e.g., the numbers of LSTM layers and LSTM units), instead, they are completely determined by the representations chosen by the analyst.

The goals of this study are, first, to clarify how such choices of representation affect LDL model performance; second, to illustrate what can be achieved simply with multivariate multiple regression; and third, to call attention to the kind of problems that are encountered when word meaning is integrated into morphology. Our working example is the comprehension and production of German nouns. In what follows, we first introduce the German noun system, and review models that have been proposed for German nouns. We then introduce LDL, after which we present a systematic overview of modeling choices, covering the representation of form, the representation of meaning, and the learning algorithm (incremental learning vs. the regression “end-state of learning” solution).

2. GERMAN NOUN MORPHOLOGY

The German noun system is both highly irregular and semi-productive, featuring three different genders, two numbers and four cases. In this section, we will give an overview over this system, show where irregularity and semi-productivity arise, and which (non-computational) models have been employed to account for it.

Plural forms are marked with one of four suffixes *-(e)n*, *-er*, *-e*, *-s* or without adding a suffix [∅; a “zero” morpheme (Köpcke, 1988, p. 306)], three of which can pair with stem vowel fronting [e.g., *a* (/a/) → *ä* (/ε/)] (e.g., Köpcke, 1988) (Table 1). There are additional suffixes which usually apply to words with foreign origin, such as *-i* (e.g., *Cello* → *Celli*, “cellos”) (Cahill and Gazdar, 1999). Cahill and Gazdar (1999) sub-categorize the nouns into 11 classes, based on whether singulars have a

different suffix than plurals (*Album* → *Alben*, “albums”). Nakisa and Hahn (1996) distinguish between no less than 60 inflection classes. No plural class is prevalent overall (Köpcke, 1988), and it is impossible to fully predict plural class from gender, syntax, phonology or semantics (Köpcke, 1988; Cahill and Gazdar, 1999; Trommer, 2021). Further complications arise when case is taken into account. German has four cases: *nominative*, *genitive*, *dative*, and *accusative*, which are marked with two exponents (applied additional to the plural markers): *-(e)n* and *-(e)s* (Schulz and Griesbach, 1981). Case forms are also not fully predictable from gender, phonology or meaning. Since many forms do not receive a separate marker, the system has been described as “degenerate” (Bierwisch, 2018, p. 245) (see Table 2). German speakers do, however, get additional disambiguating information from the definite and indefinite articles which accompany nouns and likewise encode gender, number, and case. Table 2 shows the definite articles for all genders. Additionally, there are indefinite articles available for singular forms which also express case in their endings (e.g., Gen. sg. m./n./f. *eines*, Dat. sg. m./n. *einem*, Dat. sg. f. *einer*).

Unsurprisingly, it has been the subject of a long-standing debate whether a distinction between regular and irregular nouns is useful for German (the debate has mostly focused on the formation of the nominative plural which we accordingly also

TABLE 1 | Plural classes of German nouns (relative frequencies from Gaeta, 2008).

Plural class	Example	Type frequency (%)
<i>-(e)n</i>	Tasse → Tassen “cup(s)”	56.5
<i>(uml+)-e</i>	Tag → Tage “day(s)”	
	Topf → Töpfe “pot(s)”	23.9
<i>(uml+)-er</i>	Brett → Bretter “board(s)”	
	Glas → Gläser “glass(es)”	2.3
<i>(uml+)∅</i>	Daumen → Daumen “thumb(s)”	
	Apfel → Äpfel “apple(s)”	13.3
<i>-s</i>	Kamera → Kameras “camera(s)”	2.6

Most of the classes can appear with both masculine and neuter nouns. Feminine nouns belong mostly to the *-(e)n* class (97%; Gaeta, 2008).

TABLE 2 | German noun declension.

Case and number	Masculin I	Masculin II	Neutral	Feminin
Nom. sg.	der Freund	der Mensch	das Kind	die Mutter
Gen. sg.	des Freundes	des Menschen	des Kindes	der Mutter
Dat. sg.	dem Freund	dem Menschen	dem Kind	der Mutter
Acc. sg.	den Freund	den Menschen	das Kind	die Mutter
Nom. pl.	die Freunde	die Menschen	die Kinder	die Mütter
Gen. pl.	der Freunde	der Menschen	der Kinder	der Mütter
Dat. pl.	den Freunden	den Menschen	den Kindern	den Müttern
Acc. pl.	die Freunde	die Menschen	die Kinder	die Mütter

Plural endings vary with declension class. Table adapted from Schulz and Griesbach (1981, p. 105).

focus on here). It is also unsurprising that the system shows limited productivity. Several so-called “wug” studies, where participants are asked to inflect nonce words, have clarified that German native speakers struggle with predicting unseen plurals. Köpcke (1988), Zaretsky et al. (2013), and McCurdy et al. (2020) reported high variability across speakers with respect to the plural forms produced. Köpcke (1988) took this as evidence for a “modified schema model” of German noun inflection, arguing that plural forms are generated based not only on a speaker’s experience with the German noun system, but also on the “cue validity” of the plural markers. For example, *-(e)n* is a good cue for plurality, as it does not occur with many singular forms. By contrast, *-er* has low cue validity for plurality, as it occurs with many singulars.

Köpcke (1988) also observed that *-s* is used slightly more in his wug experiments than would be expected from corpus data. Marcus et al. (1995) and Clahsen (1999) therefore argued that *-s* serves as the regular default plural marker in German, contrasting with all other plural markers that are described as irregular. Others, however, have argued that an *-s* default rule does not provide any additional explanatory value (Nakisa and Hahn, 1996; Behrens and Tomasello, 1999; Indefrey, 1999; Zaretsky and Lange, 2015).

Despite the irregularity and variability in the system, some sub-regularities within the German noun system have also been pointed out (Wiese, 1999; Wunderlich, 1999). For instance, Wunderlich (1999, p. 7f.) reports a set of rules that German nouns adhere to, which can be overridden on an item-by-item basis through “lexical storage.” For example, he notes that

- a. Masculines ending in schwa are weakly inflected (and thus also have n-plurals).
- b. Non-umlauting feminine have an n-plural.
- c. Non-feminines ending in a consonant have a ∂ -plural. [...]
- e. All atypical nouns have an s-plural. [...]

He also allows for semantics to co-determine class membership. For instance, masculine animate nouns show a tendency to belong to the *-n* plural class (see also Gaeta, 2008). A further remarkable aspect of the German noun system, especially for second language learners, is that whereas it is remarkably difficult to learn to produce the proper case-inflected forms, understanding these forms in context is straightforward.

In the light of these considerations, the challenges for computational modeling of German noun inflection, specifically from a cognitive perspective, are the following:

1. To construct a memory for a highly irregular, “degenerate,” semi-productive system,
2. To ensure that this memory shows some moderate productivity for novel forms, but with all the uncertainties that characterize the generalization capacities of German native speakers,
3. To furthermore ensure that the performance of the mappings from form to meaning, and from meaning to form, within the

framework of the discriminative lexicon (Baayen et al., 2019), are properly asymmetric with respect to comprehension and production accuracy (see also Chuang et al., 2020a).

2.1. Computational Models for German Nouns

The complexity of the German declension system has inspired many computational models. The DATR model of Cahill and Gazdar (1999) belongs to the class of generating models based on linguistic knowledge engineering. It assigns lexemes to carefully designed hierarchically ordered declension classes. Each class inherits the properties from classes further up in the hierarchy, but will override some of these properties. This model provides a successful and succinct formal model for German noun declension. Other models from this class include GERTWOL which is based on finite-state operations (Haapalainen and Majorin, 1994), as well as the model of Trommer (2021) which draws on Optimality Theory (OT) and likewise requires careful hand-crafting and constraint ranking (but does not currently have a computational implementation).

Belth et al. (2021) propose a statistical classifier based on recursive partitioning, with as response variable the morphological change required to transform a singular into a plural, and as predictors the final segments of the lexeme, number, and case. At each node, nouns are divided by their features, with one branch comprising the most frequent plural ending (which will inevitably include some nouns with a different plural ending, labeled as exceptions), and with the other branch including the remainder of the nouns. Each leaf node of the resulting tree is said to be productive if a criterion for node homogeneity is met. An older model, also a classifier building up rules inductively, was developed 20 years earlier by Albright and Hayes (2003).

Connectionist models for the German noun system include a model using a simple recurrent network (Goebel and Indefrey, 2000), and a deep learning model implementing a sequence-to-sequence encoder-decoder (McCurdy et al., 2020). The latter model takes letter-based representations of German nouns in their nominative singular form as input, together with information on the grammatical gender of the noun. The model is given the task to produce the corresponding nominative plural form. The model learned the task with high accuracy on held out data (close to 90%), but was more locked in on the “correct” forms compared to native speakers, who in a wug task showed substantially more variability in their choices.

The models discussed above also differ with respect to how they generate predictions for novel nouns. The sequence-to-sequence deep learning model of (McCurdy et al., 2020) can do so relatively easily, straight from a word’s form and its gender specification but its inner workings are not immediately interpretable (though recent work has started to gain some insights, see e.g., Linzen and Baroni, 2021, for syntactic structure in deep learning). By contrast, the linguistically more transparent DATR model can only generate a novel word’s inflectional

variants once this word has been assigned to an inflectional class. This may to some extent be possible given its principal parts (Finkel and Stump, 2007), but clearly requires additional mechanisms to be in place.

In what follows, we introduce the LDL model. LDL is a model of human lexical processing, with all its limitations and constraints, rather than an optimized computational system for generating (or understanding) morphologically complex words. It implements a simple linear mapping between form and meaning, where form is represented as a binary vector of sublexical cues, and meaning is represented in a distributed fashion.

By applying LDL to the modeling of the German noun system (including its case forms), we address a question that has thus far not been addressed computationally, namely the incorporation of semantics. Semantic subregularities in the German noun system have been noted by several authors (e.g., Wunderlich, 1999; Gaeta, 2008), and although deep learning models can be set up that incorporate semantics (see e.g., Malouf, 2017), LDL by design must take semantics into account.

The next section introduces the LDL model. The following sections proceed with an overview of the many modeling decisions that have to be made. An important part of this overview is devoted to moving beyond the modeling of isolated words, as words come into their own only in context (Elman, 2009), and case labels do not correspond to contentful semantics, but instead are summary devices for syntactic distribution classes (Blevins, 2016; Baayen et al., 2019).

3. LINEAR DISCRIMINATIVE LEARNING

LDL is the computational engine of the discriminative lexicon model (DLM) proposed by Baayen et al. (2019). The DLM implements mappings between form and meaning for both reading and listening, and mappings from meaning to form for production. It also allows for multiple routes operating in parallel. For reading in English, for instance, it sets up a direct route from form to meaning, in combination with an indirect route from visual input to a phonological representation that in turn is mapped onto the semantics (cf. Coltheart et al., 1993). In what follows, we restrict ourselves to the mappings from form onto meaning (comprehension) and from meaning onto form (production). Mappings can be obtained either with trial-to-trial learning, or by estimating the end-state of learning. In the former case, the model implements incremental regression using the learning rule of Widrow and Hoff (1960); in the latter case, it implements multivariate multiple linear regression, which is mathematically equivalent to a simple network with input units, output units, no hidden layers, and simple summation of incoming activation without using thresholding or squashing functions.

Each word form of interest is represented by a set of cues. For example, *wordform1* might feature the cues *cue1*, *cue2*, and *cue3*, while *wordform2* could be marked by *cue1*, *cue4*, and *cue5*. We can thus express a word form as a binary vector, where

1 denotes presence and 0 absence. This information is coded in the cue matrix **C**:

$$\mathbf{C} = \begin{matrix} & \text{cue1} & \text{cue2} & \text{cue3} & \text{cue4} & \text{cue5} \\ \text{wordform1} & \begin{pmatrix} 1 & 1 & 1 & 0 & 0 \end{pmatrix} \\ \text{wordform2} & \begin{pmatrix} 1 & 0 & 0 & 1 & 1 \end{pmatrix} \end{matrix}$$

Words' meanings are also represented by numeric vectors. The dimensions of these vectors can have a discrete interpretation, or have a latent interpretation (see section 4.2 below for detailed discussion). In the following example, *wordform1* has strong negative support for semantic dimensions *S3* and *S5*, while *wordform2* has strong positive support for *S4* and *S5*. This information is brought together in a semantic matrix **S**:

$$\mathbf{S} = \begin{matrix} & \text{S1} & \text{S2} & \text{S3} & \text{S4} & \text{S5} \\ \text{wordform1} & \begin{pmatrix} 0.1 & 0.004 & -1.95 & 0.03 & -0.54 \end{pmatrix} \\ \text{wordform2} & \begin{pmatrix} -0.49 & -0.32 & 0.03 & 1.06 & 0.98 \end{pmatrix} \end{matrix}$$

Comprehension and production in LDL are modeled by means of simple linear mappings from the form matrix **C** to the semantic matrix **S**, and vice versa. These mappings specify how strongly input nodes are associated with output nodes. The weight matrix for a given mapping can be obtained in two ways. First, using the mathematics of multivariate multiple regression, a comprehension weight matrix **F** is obtained by solving **F** from

$$\mathbf{S} = \mathbf{C} \cdot \mathbf{F},$$

and a production weight matrix **G** is obtained by solving **G** from

$$\mathbf{C} = \mathbf{S} \cdot \mathbf{G}.$$

As for linear regression modeling, the predicted row vectors are approximate. Borrowing notation from statistics, we write

$$\hat{\mathbf{S}} = \mathbf{C} \cdot \mathbf{F}$$

for predicted semantic vectors (row vectors of $\hat{\mathbf{S}}$), and

$$\hat{\mathbf{C}} = \mathbf{S} \cdot \mathbf{G}$$

for predicted form vectors (row vectors of $\hat{\mathbf{C}}$).

These equations amount to estimating multiple outcomes from multiple variables, which in statistics is referred to as multivariate multiple regression. In simple linear regression, a single value y is estimated from a value x via an intercept β_0 and a weighing of x with scalar β_1 :

$$\hat{y} = \beta_0 + \beta_1 x \quad (1)$$

which can easily be expanded to estimating y from a vector \mathbf{x} (multiple linear regression), using a vector of beta coefficients $\beta_i \in \boldsymbol{\beta}$ to weigh each value $x_i \in \mathbf{x}$:

$$\hat{y} = \beta_0 + x_1\beta_1 + x_2\beta_2 + \dots + x_n\beta_n \quad (2)$$

Finally, to estimate a vector \mathbf{y} from a vector \mathbf{x} (multivariate multiple regression), we need an entire matrix of beta coefficients $\beta_{ij} \in \mathbf{B}$. A single value $y_i \in \mathbf{y}$ is then estimated via

$$\hat{y}_i = \beta_{0i} + x_1\beta_{1i} + x_2\beta_{2i} + \dots + x_p\beta_{pi} \quad (3)$$

Thus, estimating the mappings \mathbf{F} and \mathbf{G} in LDL amounts to computing the coefficients matrix \mathbf{B} for mappings from \mathbf{C} to \mathbf{S} and vice versa. As such, each value in a predicted semantic vector $\hat{\mathbf{s}}$ (form vector $\hat{\mathbf{c}}$) is a linear combination (i.e., weighted sum) of the values in the corresponding form vector \mathbf{c} (semantic vector \mathbf{s}) it is predicted from. This means that LDL is mathematically highly constrained: it cannot handle non-linearities that even shallow connectionist models (e.g., Goldsmith and O'Brien, 2006) can take in their stride. Nevertheless, we have found that these simple linear mappings result in high accuracies (e.g., Baayen et al., 2018, 2019) suggesting that morphological systems are surprisingly simple. Cases where model predictions are less precise due to the limitations of linearity become indicative of learning bottlenecks.

Furthermore, note that estimating the mappings \mathbf{F} and \mathbf{G} using the matrix algebra of multivariate multiple regression provides optimal estimates, in the least squares sense, of the connection weights (or equivalently, beta coefficients) for datasets that are type-based, in the sense that each pair of row vectors \mathbf{c} of \mathbf{C} and \mathbf{s} of \mathbf{S} is unique. Having multiple instances of the same pair of row vectors in the dataset does not make sense, as it renders the input completely singular and does not add any further information. Thus, models based on the regression estimates of \mathbf{F} and \mathbf{G} are comparable to type-based models such as AML, MBL, MGL, and models using recursive partitioning.

Making the estimates of the mappings sensitive to frequency of use requires incremental learning, updating weights after each word token that is presented for learning. Incremental learning is implemented using the learning rule of Widrow and Hoff (1960) and Milin et al. (2020), which defines the matrix \mathbf{W}^{t+1} with updated weights at time $t + 1$ as the weight matrix \mathbf{W}^t at time t , modified as follows:

$$\mathbf{W}^{t+1} = \mathbf{W}^t + \mathbf{c} \cdot (\mathbf{o}^T - \mathbf{c}^T \cdot \mathbf{W}^t) \cdot \eta,$$

where \mathbf{c} is the current cue (vector), \mathbf{o} the current outcome vector, and η the learning rate. Conceptually, this means that after each newly encountered word token, the weight matrix is changed such that the next time that the same cue vector has to be mapped onto its associated outcome vector, it will be slightly closer to the target outcome vector than it was before. The learning rule of Widrow-Hoff implements incremental regression. As the number of times that a model is trained again and again on a training set increases (training epochs), the network's weights will converge to the matrix of beta coefficients obtained by approaching the estimation problem with multivariate multiple regression (see e.g., Evert and Arppe, 2015; Chuang et al., 2020a; Shafaei-Bajestan et al., 2021). As a consequence, the regression-based estimates pertain to the "end-state of learning," at which the data have been worked through infinitely many times. Unsurprisingly, effects of frequency and order of learning

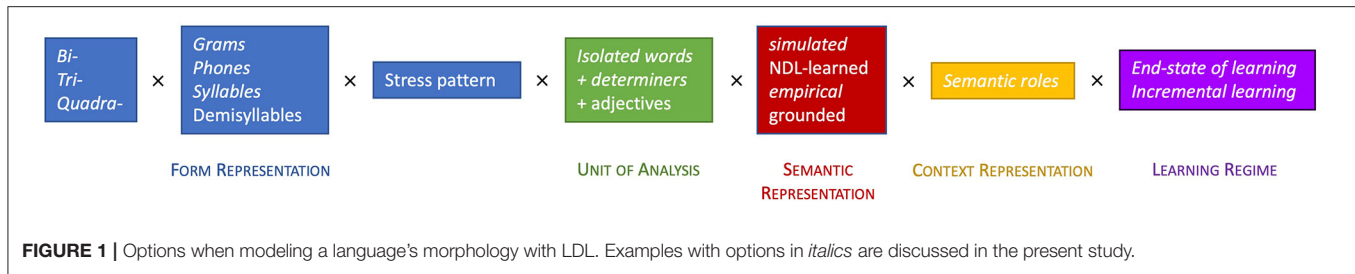
are not reflected in model predictions based on the regression estimates. Such effects do emerge with incremental learning (see section 4.5).

Comprehension accuracy for a given word ω is assessed by comparing its predicted semantic vector $\hat{\mathbf{s}}_\omega$ with all gold standard semantic vectors in \mathbf{S} (the creation of gold standard semantic vectors will be described in subsequent sections), using either the cosine similarity measure or the Pearson correlation r . In what follows, we use r , and select as the meaning that is recognized that gold standard row vector \mathbf{s}_{\max} of \mathbf{S} that shows the highest correlation with $\hat{\mathbf{s}}_\omega$. If \mathbf{s}_{\max} is the targeted semantic vector, the model's prediction is classified as correct, otherwise, it is taken to be incorrect.

For the modeling of production, a supplementary algorithm is required for constructing actual word forms. The predicted vectors $\hat{\mathbf{c}}$ provide information about the amount of support that form cues receive from the semantics. However, information about the amount of support received by the full set of cues does not provide information about the order in which a subset of these cues have to be woven together into actual words. Algorithms that construct words from form cues make use of the insight that when form cues are defined as n -grams ($n > 1$), the cues contain implicit information about order. For instance, for digraph cues, cues ab and bc can be combined into the string abc , whereas cues ab and cd cannot be merged. Therefore, when n -grams are used as cues, directed edges can be set up in a graph with n -grams as vertices, for any pair of n -grams that properly overlap. A word form is uniquely defined by a path in such a graph starting with an initial n -gram (starting with an initial word edge symbol, typically a $\#$ is used) and ending at a final n -gram (ending with $\#$). This raises the question of how to find word paths in the graph. This is accomplished by first discarding n -grams with low support from the semantics below a threshold θ^1 , then calculating all possible remaining paths, and finally selecting for articulation that path for which the corresponding predicted semantic vector (obtained by mapping its corresponding cue vector \mathbf{c} onto \mathbf{s} using comprehension matrix \mathbf{F}) best matches the semantic vector that is the target for articulation. This implements "synthesis by analysis," see Baayen et al. (2018, 2019) for further details and theoretical motivation. For a discussion of the cognitive plausibility of this method, see Chuang et al. (2020b).

The first algorithm that was used to enumerate possible paths made use of a shortest-paths algorithm from graph theory. This works well for small datasets, but becomes prohibitively expensive for large datasets. The **JudiLing** package (Luo et al., 2021) offers a new algorithm that scales up better. This algorithm is first trained to predict, from either the $\hat{\mathbf{C}}$ or the \mathbf{S} matrix, for each possible word position, which cues are best supported at that position. All possible paths with the top k best-supported cues are then calculated, and subjected to synthesis by analysis. Details about this algorithm, implemented in julia in the **JudiLing** package as the function `learn_paths` can be found in Luo (2021). The `learn_paths` function is used throughout the

¹This is a simple cut-off point for n -grams with low support, not to be confused with thresholds as often used in deep learning.



remainder of this study. A word form is judged to be produced correctly when it exactly matches the targeted word form.

4. MODELING CONSIDERATIONS

When modeling a language's morphology within the framework of the DLM, the analyst is faced with a range of choices, illustrated in **Figure 1**. From left to right, choices are listed for representing form, for the unit of analysis, for the representation of semantics, for the handling of context, and for the learning regime.

With respect to form representations, the kind of n-gram has to be selected, setting n , deciding on phonological or orthographic grams, and specifying how stress or lexical tone are represented. With respect to the unit of analysis, the analyst has to decide whether to model isolated words, or words in phrasal contexts. A third set of choices concerns what semantic representations to use: simulated representations, or word embeddings such as *word2vec* (Mikolov et al., 2013b), or grounded vectors (Shahmohammadi et al., in press). A further set of choices for languages with case concerns how to handle case labels, as these typically refer to syntactic distribution classes rather than contentful inflectional features (Blevins, 2016). Finally, a selection needs to be made with respect to whether incremental learning is used, or instead the end-state of learning using regression-based estimation. In what follows, we illustrate several of these choice points using examples addressing the German noun system.

The dataset on German noun inflection that we use for our worked examples was compiled as follows. First, we extracted all monomorphemic nouns and their inflections with a frequency of at least 1 from CELEX (Baayen et al., 1995), resulting in a dataset of about 6,000 word forms. Of these we retained the 5,486 word forms for which we could retrieve grammatical gender from Wiktionary, thus including word forms of 2,732 different lemmas. The resulting data was expanded such that each attested word form was listed once for each possible paradigm cell it could belong to. For instance, *Aal* ("eel") is listed once as singular nominative, once as dative and once as accusative (**Table 3**). This resulted in a dataset with 18,147 entries, with word form frequencies ranging from 1 to 5,828, (M log frequency 2.56, SD 1.77). Word forms are represented in their DISC notation, which represents German phones with single characters². **Table 3** clarifies that there are many homophones. As a consequence, the

TABLE 3 | Representation of the paradigm for *Aal* "eel" in our dataset.

Word form	Pronunciation	Lemma	Case	Number	Frequency	Gender
<i>Aal</i>	a1	<i>Aal</i>	Nominative	Singular	29	M
<i>Aal</i>	a1	<i>Aal</i>	Dative	Singular	29	M
<i>Aal</i>	a1	<i>Aal</i>	Accusative	Singular	29	M
<i>Aale</i>	a1@	<i>Aal</i>	Nominative	Plural	34	M
<i>Aale</i>	a1@	<i>Aal</i>	Genitive	Plural	34	M
<i>Aalen</i>	a1@n	<i>Aal</i>	Dative	Plural	17	M
<i>Aalen</i>	a1@n	<i>Aal</i>	Accusative	Plural	17	M

Genitive singular (Aals) is not included as it has a frequency of 0 in CELEX.

actual number of distinct word forms in our dataset is only 5,486, which amounts to on average about two word forms per lemma.

There are many ways in which model performance can be evaluated. First, we may be interested in how well the model performs as a memory. How well does the model learn to understand and produce words it has encountered before? Note that because the model is not a list of forms, this is not a trivial question. For evaluation of the model as a memory, we consider its performance on the training data (henceforth *train*). Second, we may be interested in the extent to which the memory is productive. Does it generalize so that new forms can be understood or produced? Above, we observed that the German noun system is semi-regular, and that German native speakers are unsure about what the proper plural is of words they have not encountered before (McCurdy et al., 2020). If our modeling approach properly mirrors human limitations on generalization from data with only partial regularities, evaluation on unseen, held-out data of German should not be perfect. At this point, however, several issues arise that require careful thought.

For one, from the perspective of the linguistic system, it seems unreasonable to assume that any held-out form can be properly produced (or understood) if some of the principal parts (Finkel and Stump, 2007) of the lexeme are missing in the training data. In what follows, we will make the simplifying assumption that under cross-validation with sufficient training data, this situation will not arise.

A further question that arises is how to evaluate held-out words that have homophones in the training data. Such homophones present novel combinations of a form vector (shared with another data point in the training data) and a semantic vector (not attested for this form in the training data). We may want to impose a strict evaluation criterion requiring

²Data and code are available in the **Supplementary Materials** at <https://osf.io/zrw2v/>

TABLE 4 | Types of model evaluation.

Evaluation type			
Simple	Blind evaluation of all held-out data		val_all
Nuanced	Evaluation on novel forms only		val_newform
	Evaluation on homophones	Strict	val_strict
		Lenient	val_lenient

that the model gets the semantic vector exactly right. However, when presented with a homophone in isolation, a human listener cannot predict which of a potentially large set of paradigm cells is the targeted one (the problem of modeling words in isolation). We may therefore want to use a lenient evaluation criterion for comprehension according to which comprehension is judged to be accurate when the predicted semantic vector \hat{s} is associated with one of a homophonic word's possible semantic interpretations. Yet a further possible evaluation metric is to see how well the model performs on words with forms that have not been encountered in the training data. These possibilities are summarized in **Table 4**. Below, in section 4.3.1, we consider further complications that can arise in the context of testing the model on unseen forms.

For evaluating the productivity of the model, we split the full dataset into 80% training data and 20% validation data, with 14,518 and 3,629 word forms, respectively. In the validation data, 3,309 forms are also present in the training data (i.e., homophones), and 320 are new forms. Among the 320 new forms, 8 have novel lemmas that are absent in the training data. Since it is unrealistic to expect the model to understand or produce inflected forms of completely new words, these 8 words are excluded from the validation dataset for new forms, although they are taken into consideration when calculating the overall accuracy for the validation data. The same training and validation data are used for all the simulations reported below, unless indicated otherwise.

4.1. Representing Words' Forms

Decisions about how to represent words' forms depend on the modality that is to be modeled. For auditory comprehension, Arnold et al. (2017) and Shafaei-Bajestan et al. (2021) explore ways in which form vectors can be derived from the audio signal. Instead of using low-level audio features, one can also use more abstract symbolic representations such as phone n-grams³. For visual word recognition, one may use letter n-grams, or, as lower-level visual cues, for instance, features derived from histograms of oriented gradients (Dalal and Triggs, 2005; Linke et al., 2017). In what follows, we use binary vectors indicating the presence or absence of phonological phone or syllable n-grams.

4.1.1. Phone-Based Representations

Sublexical phone cues can be of different granularity, such as biphones and triphones. For the word *Aale* (pronunciation a1@),

³Other work (e.g., Joannis and Seidenberg, 1999) has used slot-coding for representing phonology, but we do not think that this representation is optimal, since, for example, we are not sure how prefixation is to be modeled without hand-engineering (details in Heitmeier and Baayen, 2020).

the biphone cues are #a, a1, 1@, and @#, and the triphone cues are #a1, a1@, and 1@#. The number of unique cues (and hence the dimensionality of the form vectors) increases as granularity decreases. For the present dataset, there are 931 unique biphone cues, but 4,656 triphone cues. For quadrphones, there are no less than 9,068 unique cues. Although model performance tends to become better with more unique cues, we also run the risk of overfitting. That is, the model does not generalize and thus performs worse on validation data. The choice of granularity therefore determines the balance of having a precise memory on the one hand and a productive memory on the other hand. In the simulation examples with n-phones that follow, we made use of simulated semantic vectors. Details on the different kinds of semantic vectors that can be used are presented in section 4.2.1.

Accuracy for n-phones is presented in the first three rows of **Table 5**. For the training data, comprehension accuracy is high with both triphones and quadrphones. For biphones, the small number of unique cues clearly does not offer sufficient discriminatory power to distinguish word meanings. Under strict evaluation, unsurprisingly given the large number of homophones in German noun paradigms, comprehension accuracy plummets substantially to 8, 33, and 35% for biphone, triphone, and quadrphone models, respectively. Given that there is no way to tell the meanings of homophones apart without further contextual information, we do not provide further details for strict evaluation. However, in section 4.1.1 we will address the problem of homophony by incorporating further contextual information into the model.

With regards to model accuracy for validation data, we see that overall accuracy (val_all) is quite low for biphones, while it remains high for both triphones and quadrphones. Closer inspection reveals that this high accuracy is mainly contributed by homophones (val_lenient). Since these forms are already present in the training data, a high comprehension accuracy under lenient evaluation is unsurprising. As for unseen forms (i.e., val_newform), quadrphones perform slightly better than triphones.

Production accuracy, presented in the right half of **Table 5**, is highly sensitive to the threshold θ used by the learn_paths algorithm. Given that usually only a relatively small number of cues receive strong support from a given meaning, we therefore set the threshold such that the algorithm does not need to take into account large numbers of irrelevant cues. Depending on the form and meaning representations selected, some fine-tuning is generally required to obtain a threshold value that optimally balances both accuracy and computation time. Once the threshold is fine-tuned for the training data, the same threshold is used for the validation data.

Production accuracy is similar to comprehension accuracy, albeit systematically slightly lower. Triphones and quadrphones again outperform biphones by a large margin. For the training data, triphones are somewhat less accurate than quadrphones. Interestingly, in order to predict new forms in the validation data, triphones outperform quadrphones. Clearly, triphones offer better generalizability compared to quadrphones, suggesting that we are overfitting when modeling with quadrphones as cues. Accuracy under the val_newform criterion is quite low,

TABLE 5 | Comprehension and production accuracy for train and validation datasets, with biphones, triphones, quadraphones, and bisyllables as cues.

	Comprehension				Production			
	Train (%)	val_all (%)	val_lenient (%)	val_newform (%)	Train (%)	val_all (%)	val_lenient (%)	val_newform (%)
Biphone	22	16	17	8	48	31	33	12
Triphone	93	88	92	51	84	64	68	21
QuadrAPHone	97	93	97	53	91	67	73	11
Bisyllable	99	93	99	20	95	63	69	0.3
word2vec	87	72	79	0.3	97	88	94	25

For the first four rows, we used simulated semantic vectors. For the last row, cues are triphones, and semantic vectors are word2vec embeddings (discussed in section 4.2.2). For the *learn_paths* algorithm, the threshold θ was set to 0.05, 0.008, 0.005, 0.005, and 0.008, respectively.

which is perhaps not unexpected given the uncertainty that characterizes native speakers' intuitions about the forms of novel words (McCurdy et al., 2020). In section 4.3.2, we return to this low accuracy, and consider in further detail generated novel forms and the best supported top candidates.

4.1.2. Syllable-Based Representations

Instead of using n-phones, the unit of analysis can be a combination of n syllables. The motivation for using syllables is that some suprasegmental features, such as lexical stress in German, are bound to syllables. Although stress information is not considered in the current simulation experiments, suprasegmental cues can be incorporated (see Chuang et al., 2020a, for an implementation).

As for n-phones, when using n-syllables, we have to choose a value for the unit size n . For the word *Aale*, the bi-syllable cues are #-a, a-l@, and l@-#, with "-" indicating syllable boundary. When unit size equals two, there are in total 8,401 unique bi-syllable cues. For tri-syllables, the total number of unique cues increases to 10,482. Above, we observed that the model was already overfitting with 9,068 unique quadrAPHone cues. We therefore do not consider tri-syllable cues, and only present modeling results for bi-syllable cues⁴.

As shown in the fourth row of **Table 5**, comprehension accuracy (for bi-syllables) for the training data is almost error-free, 99%, the highest among all the cue representations. For the validation data, the overall accuracy is also high, 93%. This is again due to the high accuracy for the seen forms ($\text{val_lenient} = 99\%$). Only one fifth of the unseen forms, however, is recognized successfully ($\text{val_newform} = 20\%$). Production accuracies for the training and validation data are 95 and 63%, respectively. The model again performs well for homophones ($\text{val_lenient} = 69\%$) but fails to produce unseen forms ($\text{val_newform} = 0.3\%$). This extremely low accuracy is in part due to the large number of cues that appear only in the validation dataset (325 for bisyllables, but only 23 for triphones). Since such novel cues do not receive any training, words with such cues are less likely to be produced correctly. We

⁴Even though the number of bi-syllables is close to that of quadrAPHones, the fact that quadrAPHones still outnumber bi-syllables suggests that quadrAPHones have captured within-syllable phone collocations that are not available in bi-syllable cues. These further fine-grained cues might include, for example, consonant clusters, as in *Sprache* "language."

will come back to the issue of novel cues in section 4.3.1. For now, we conclude that triphone-based form vectors are a good choice as they show a good balance of comprehension and production accuracy on training and validation data.

4.2. Semantic Representation

There are many ways in which words' meanings can be represented numerically. The simplest method is to use one-hot encoding (i.e., a binary vector where a single value/bit is set to one), as implemented in the naive discriminative learning model proposed by Baayen et al. (2011). One-hot encoding, however, misses out on the semantic similarities between lemmas: all lemmas receive meaning representations that are orthogonal. Instead of using one-hot encoding, semantic vectors can also be derived by turning words' taxonomies in WordNet into binary vectors with multiple bits on (details in Chuang et al., 2020a). In what follows, however, we work with real-valued semantic vectors, known as "word embeddings" in natural language processing. Semantic vectors can either be simulated, or derived from corpora using methods from distributional semantics (see e.g., Landauer and Dumais, 1997; Mikolov et al., 2013b).

4.2.1. Simulated Semantic Vectors

When corpus-based semantic vectors are unavailable, semantic vectors can be simulated. The **JudiLing** package enables the user to simulate such vectors using normally distributed random numbers for content lexemes and for inflectional functions. By default, the dimension of the semantic vectors is set to be identical to that of the form vectors.

The semantic vector for an inflected word is obtained by summing the vector of its lexeme and the vectors of all the pertinent inflectional functions. As a consequence, all vectors sharing a certain inflectional feature are shifted in the same direction in semantic space. By way of example, consider the German plural dative of *Aal* "eel," *Aalen*. We compute its semantic vector by adding the semantic vector for PLURAL and DATIVE to the lemma vector \vec{Aal}_{lemma} :

$$\vec{Aalen}_{dat.pl} = \vec{Aal}_{lemma} + \vec{PLURAL} + \vec{DATIVE}$$

The corresponding singular dative *Aal* can be coded as:

$$\vec{Aal}_{dat.sg} = \vec{Aal}_{lemma} + \vec{SINGULAR} + \vec{DATIVE}$$

Alternatively, the singular form could be coded as unmarked, following a privative opposition approach:

$$\vec{Aal}_{dat.sg.} = \vec{Aal}_{lemma} + \vec{DATIVE}$$

For the remainder of the paper, we treat number as an equipollent opposition. Finally, a small amount of random noise is added to each semantic vector ($M 0, SD 1$; compare this to $M 0, SD 4$ for lexeme and inflectional vectors), as an approximation of further semantic differences in word use other than number and case [see Sinclair (1991, e.g., p.44ff.)⁵, Tognini-Bonelli (2001) and further discussion below]. The results reported thus far were all obtained with simulated vectors.

It is worth noting that when working with simulated semantic vectors, the meanings of lexemes will still be orthogonal, and that as a consequence, all similarities between semantic vectors originate exclusively from the semantic structure that comes from the inflectional system.

4.2.2. Empirical Semantic Vectors

A second possibility for obtaining semantic vectors is to derive them from corpora. Baayen et al. (2019) constructed semantic vectors from the TASA corpus (Ivens and Koslin, 1991), in such a way that semantic vectors were obtained not only for lexemes but also for inflectional functions. With their semantic vectors, the semantic vector of *Aalen* can be straightforwardly constructed from the semantic vectors of *Aal*, PLURAL, and DATIVE.

However, semantic vectors that are created with standard methods from machine learning, such as *word2vec* (Mikolov et al., 2013a), *fasttext* (Bojanowski et al., 2017), or *GloVe* (Pennington et al., 2014), can also be used (albeit without semantic vectors for inflectional features; see below). In what follows, we illustrate this for 300-dimensional vectors generated with *word2vec*, trained on the German Wikipedia (Yamada et al., 2020). For representing words' forms, we used triphones.

The model in general performs well for the training data (Table 5). For the validation data, while the homophones are easy to recognize and produce, the unseen forms are again prohibitively difficult. Interestingly, if we compare the current results with the results of simulated vectors (cf. second row, Table 5), we observe that while the *train* and *val_all* accuracies are fairly comparable for the two vector types, their *val_newform* accuracies nonetheless differ. Specifically, understanding new forms is substantially more accurate with simulated vectors (51 vs. 0.3%), whereas *word2vec* embeddings yield slightly better results for producing new forms (21 vs. 25%).

To understand why these differences arise, we note, first, that lexemes are more similar to each other than is the case for simulated vectors (in which case lexemes are orthogonal), and second, that *word2vec* semantic vectors are exactly the same for each set of homophones within a paradigm, so that inflectional structure is much less precisely represented. This lack

⁵Our approach of adding small semantic differences to individual word forms does probably not do justice to Sinclair (1991)'s view that word forms can have completely idiosyncratic meanings, since we still assume commonalities across word forms such as e.g., a shared meaning of plurality. We hope to be able to address this issue in future research.

of inflectional structure may underlie the inability of the model to understand novel inflected forms correctly. Furthermore, the lack of differentiation between homophones simplifies the mapping from meaning to form, leading to more support from the semantics for the relevant triphones, which in turn facilitates synthesis by analysis.

In addition, we took the *word2vec* vectors, and reconstructed from these vectors the vectors of the lexemes and of the inflectional functions. For a given lexeme, we created its lexeme vector by averaging over the vectors of its inflectional variants⁶. For plurality, we averaged over all vectors of forms that can be plural forms. Using these new vectors, we constructed semantic vectors for a given paradigm cell by adding the semantic vector of the lexeme and the semantic vectors for its number and case values. The mean correlation between the new "analytical" *word2vec* vectors and the original empirical vectors was 0.79 ($sd = 0.076$). Apparently, there is considerable variability in how German inflected words are actually used in texts, a finding that has also emerged from corpus linguistics (Sinclair, 1991; Tognini-Bonelli, 2001). The idiosyncrasies in the use of individual inflected forms renders the comprehension of an unseen, but nevertheless also idiosyncratic, inflected word form extremely difficult. From this we conclude that the small amount of noise that we added to the simulated semantic vectors is likely to be unrealistically small compared to real language use.

Interestingly, semantic similarity facilitates the production of unseen forms. A Linear Discriminant Analysis (LDA) predicting nine plural classes (the eight sub-classes presented in Table 1 plus one "other" class) from the *word2vec* semantic vectors has a prediction accuracy of 62.7% (50.5% under leave-one-out cross validation). Conducting 10-fold cross-validation with Support Vector Machine (SVM) yields an average accuracy of 56.7%, considerably higher than the percentage of the majority choice (the -n plural class, 35.6%). Apparently, semantically similar words tend to inflect similarly. When a novel meaning is encountered in the validation set, it is therefore possible to predict to some extent its general form class. Given the similarities between LDA and regression, the same kind of information is likely captured by LDL.

4.3. Missing Forms and Missing Semantics

Evaluation on held-out data is a means for assessing the productivity of the network. However, it often happens during testing that the model is confronted with novel, unseen cues, or with novel, unseen semantics. Here, linguistically and cognitively motivated choices are required.

4.3.1. Novel Cues

For the cross-validation results presented thus far, the validation data comprise a random selection of words. As a consequence, there often are novel cues in the validation data that the model has never encountered during training. The presence of such novel cues is especially harmful for production. As mentioned in section 4.1.2, the model with bi-syllables as cues fails to produce

⁶Note that these vectors are not sense-disambiguated, so that they can cover homophonous forms from various paradigm cells.

TABLE 6 | Comprehension and production accuracy for train and validation datasets, which are split in such a way that no novel cues are present in the validation set.

	Comprehension				Production			
	Train (%)	val_all (%)	val_lenient (%)	val_newform (%)	Train (%)	val_all (%)	val_lenient (%)	val_newform (%)
Triphone	93	88	91	52	85	63	67	17
Bisyllable	99	99	99	61	95	52	52	12

Both the triphone and bisyllable models make use of simulated semantic vectors.

unseen forms, due to the large number of novel cues in the validation data.

What is the theoretical status of novel cues? To answer this question, first consider that actual speakers rarely encounter new phones or new phone combinations in their native languages. Furthermore, novel sounds encountered in loan words are typically assimilated into the speaker's native phonology⁷. Also, many cues that are novel for the model actually occur not only in the held-out nouns, but also in verbs, adjectives, and compounds that the model has no experience with. Thus, the presence of novel cues is in part a consequence of modeling only part of the German lexicon.

Since novel cues have zero weights on their efferent connections (or, equivalently, zero beta coefficients), they are completely inert for prediction. One way to address this issue is to select the held-out data with care. Instead of randomly holding out words, we make sure that in the validation data all cues are already present in the training data. We therefore split the dataset into 80% training and 20% validation data, but now making sure that there are no novel triphone cues in the validation dataset. Among the 3,629 validation words, 3,331 are homophones, and 298 are unseen forms. We note that changing the kind of cues used typically has consequences for how many datapoints are available for validation. When bi-syllables are used instead of triphones, due to the sparsity of bi-syllable cues, we have to increase the percentage of validation data to include sufficient numbers of unseen forms. Even for 65% training data and 35% validation data, the majority of validation data are homophones (98.5%), and only 76 cases represent unseen forms (with only known cues).

For the triphone model (top row, **Table 6**), for both comprehension and production, the `train`, `val_all`, and `val_lenient` accuracies are similar to the results presented previously (**Table 5**). For the evaluation of unseen forms (`val_newform`), there is only a slight improvement for comprehension (from 51 to 52%); for other datasets, the improvement can be larger. However, for production, `val_newform` becomes worse (decreasing from 21 to 17%). The reason is that even though all triphone cues of the validation words are present in the training data, they obtain insufficient

support from the semantics. The solution here is to allow a small number of triphone cues with weak support (below the threshold θ) to be taken into account by the algorithm that orders triphones into words. This requires turning on the `tolerance` mode in the `learn_paths` function of the **JudiLing** package. By allowing at most two weakly supported triphones to be taken into account, production accuracy for unseen forms increases to 57%.

The bi-syllable model benefits more from the removal of novel cues in the validation data. Especially for comprehension, the accuracy of unseen forms reaches 61%, compared to 20% with random selection. For production, we observe a non-negligible improvement as well (from 0.3 to 12%). Further improvements are expected when tolerance mode is used, but given the large number of bi-syllables, this comes at considerable computation costs. In other words, bi-syllables provide a model that is an excellent memory, but a memory with very limited productivity specifically for production.

4.3.2. Unseen Semantics

In real language, speakers seldom encounter words that are completely devoid of meaning: even novel words are typically encountered in contexts which narrow down their interpretation. In the wug task, by contrast, participants are often confronted with novel words presented without any indication of their meaning, as, for instance, in the experiment on German nouns reported by McCurdy et al. (2020). Within the framework of the discriminative lexicon, this raises the question of how to model the wug task, as the model has no way to produce inflected variants without semantics.

For modeling the wug task, and comparing model performance with that of German native speakers, we begin with observing that the comprehension system generates meanings for non-words. Chuang et al. (2020c) showed that measures derived from the semantic vectors of non-words were predictive for both reaction times in an auditory lexical decision task and for non-words' acoustic durations in a reading task. In order to model the wug task, we therefore proceeded as follows:

1. We first simulated a speaker's lexical knowledge prior to the experiment by training a comprehension matrix using all the words described in section 4. Here, we made use of simulated semantic vectors.
2. We then used the resulting comprehension network to obtain semantic vectors $s_{\text{nom.sg}}$ for the nominative singular forms of the non-words by mapping their cue vectors into the semantic space, resulting in semantic vectors $s_{\text{nom.sg}}$.

⁷Note that such assimilation effects could be modeled using real acoustic input (i.e., audio files) with LDL-AURIS (Shafaei-Bajestan et al., 2021). Here, unseen sounds would presumably be assimilated to the closest seen sounds, similar to human performance. Of course, given sufficient training data, such a model would over time also be able to acquire the new sounds. We have, however, restricted ourselves to modeling using letter/phone representations.

3. Next, we created the production mapping from meaning to form, using not only all real words but also the non-words (known only in their nominative singular form).
4. Then, we created the semantic vectors for the plurals ($s_{\text{nom.pl}}$) of the non-words by adding the plural vector to their nominative singular vectors while subtracting the singular vector.
5. Finally, these plural semantic vectors were mapped onto form vectors ($\hat{c}_{\text{nom.pl}}$) using the production matrix, in combination with the `learn_paths` algorithm that orders triphones for articulation.

We applied these modeling steps to a subset of the experimental materials provided by Marcus et al. (1995) (reused by McCurdy et al., 2020), in order to compare model predictions with the results of McCurdy et al. (2020). The full materials of Marcus et al. (1995) contained non-words that were set up such that only half of them had an existing rhyme in German. We restricted ourselves to the non-words with existing rhymes, first, because non-rhyme words have many cues that are not in the training data; and second, because, as noted by Zaretsky and Lange (2015), many of the non-rhyme words have unusual orthography and thus are strange even for German speakers. Furthermore, many of the non-rhyme non-words share endings and therefore do not provide strong data for testing model predictions.

McCurdy et al. (2020) presented non-words visually and asked participants to write down their plural form. To make our simulation more comparable to their experiment, in the following we made use of letter trigrams rather than triphones. We represented words without their articles, as the wug task implemented by McCurdy et al. (2020) presented the plural article as a prompt for the plural form; participants thus produced bare plural forms. For assessing what forms are potential candidates for production, we examined the set of candidate forms, ranked by how well their internally projected meanings (obtained with the synthesis-by-analysis algorithm, see section 3), correlated with the targeted meaning $s_{\text{nom.pl}}$. We then examined the best supported candidates as possible alternative plural forms.

The model provided a plausible plural form as the best candidate in 7 out of 12 cases. Five of these belonged to the *-en* class. A further plausible candidate was also only provided in 5 of the cases. The lack of diversity as well as the bias for *-en* plurals does not correspond to the responses given by German speakers in McCurdy et al. (2020).

Upon closer inspection, it turns out that a more variegated wug performance can be obtained by changing two parameters. First, we replaced letter trigrams by letter bigrams. This substantially reduces the number of n-grams that are present in the non-words but that do not occur in the training data. Second, we made a small but important change to how semantic vectors were simulated. The default parameter settings provided with the **JudiLing** package generate semantic vectors with the same standard deviation for both content words and inflectional features. Therefore, the magnitudes of the values in semantic vectors is very similar for content words and inflectional features. Since words are inflected for case and number, their semantic vectors are numerically dominated by the inflectional vectors. To enhance the importance of the lexemes, and to reduce the dominance of the inflectional functions, we reduced the standard deviation (by a factor of $\frac{1}{10}$) when generating the semantic vectors for number and case. As a consequence, the mean of the absolute values in the plural vector decreased from 3.25 to 0.32. (Technical details are provided in the **Supplementary Materials**.) With these two changes, the model generated a more diverse set of plural non-word candidates (**Table 7**). Model performance is now much closer to the performance of native speakers as reported by Zaretsky et al. (2013); McCurdy et al. (2020).

The model also produces some implausible plural candidates, all of which however are phonotactically legal; these are marked with an asterisk in **Table 7**. Sometimes a plural marker is interfixed instead of suffixed (e.g., Spand, Span-*en-d*; Pund, Pun-*en-d*). Almost all words have a candidate which shows double plural marking (e.g., Bral, Bral-*en-en*; Nuhl, Nuhl-*er-e*; Pind, Pind-*er-n*; cf. Dutch kind-*er-en*), or a mixture of both (e.g., Span, Span-*en-d-e*; Spert, Spert-*er-t-en*). For Klot, doubling of the *-t* can be observed, as this form is presumably more plausible in German [e.g., Motte (“moth”), Gott (“god”), Schrott (“scrap, rubbish”)]. One plural has been attracted to an existing singular (Spand, Spaten-d). Apparently, by downgrading the strength (or more precisely, the L1-norm) of the semantic vectors of inflectional functions, the model moves in the direction of interfixation-like changes.

The model does not produce a single plural form with an umlaut, even though in corpora umlauted plurals are relatively frequent (see e.g., Gaeta, 2008). Interestingly, the German speakers in McCurdy (2019) also tended to avoid umlauted forms (with the exception of *Kach* → *Kächer*). Interestingly, children at the age of 5 also tend to avoid umlauts when producing plurals

TABLE 7 | First five candidates for the plural forms of non-words.

Bral	Kach	Klot	Mur	Nuhl	Pind	Pisch	Pund	Raun	Spand	Spert	Vag
Bralen	Kachen	Klot	Muren	Nuhlen	Pinden	Pischen	Punden	Raunen	*Spanend	Sperten	Vag
Bral	Kach	*Klotten	Murn	Nuhl	Pind	Pisch	*Punend	Raun	Spand	Sperte	Vagen
*Bralenen	Kacher	*Klotte	Mur	Nuhle	Pinder	Pischer	Pund	*Raunern	*Spanende	Sperter	Vage
*Bralern	Kache	*Klotter	*Murnen	*Nuhlern	Pinde	Pische	Punde	Rauner	*Spanenden	*Spererten	Vager
Braler	*Kachern	*Klieloten	Murer	*Nuhlere	*Pindern	*Pischern	*Pundene	Raune	*Spatend	*Spererte	*Vagern

Forms that are implausible as plurals are marked with an asterisk. Non-words are taken from Marcus et al. (1995).

for German non-words, but usage increases for 7-year-olds and adults (Van de Vijver and Baer-Henney, 2014).

Finally, most non-words have a plural in *-en* as one of the candidates (10 out of 12 cases), with as runners-up the *-e* plural (8 out of 12 cases), and the *-er* plural (8 out of 12). There is not a single instance of an *-s* plural, which fits well with the low prevalence (around 5%) of *-s* plurals in the experiment of McCurdy et al. (2020).

This simulation study shows that it is possible to make considerable headway with respect to modeling the wug task for German. The model is not perfect, unsurprisingly, given that we have worked with simulated semantic vectors and estimates of non-words' meanings. It is intriguing that a strong weight imposed on the stem shifts model performance in the direction of interfixation-like morphology. However, the model has no access to information about words' frequency of use, and hence is blind to an important factor shaping human learning (see section 4.5 for further discussion). Nevertheless, the model does appear to mirror the uncertainties of German speakers fairly well.

4.4. Words in Context

Thus far, we have modeled words in isolation. However, in German, case and number information is to a large extent carried by preceding determiners. In addition, in actual language use, a given grammatical case denotes one of a wide range of different possible semantic roles. The simplifying assumption that an inflectional function can be represented by a single vector, which may be reasonable for grammatical number, is not at all justified for grammatical case. In this section, we therefore explore how context can be taken into account. We first present modeling results of nouns learned together with their articles. Next, we break down grammatical cases into actual semantic functions, and show how we can begin to model the noun declension system with more informed semantic representations.

4.4.1. Articles

We first consider definite articles. Depending on gender and case, a noun can follow one of the six definite articles in German—*der*, *die*, *das*, *dem*, *den*, *des*. We added these articles, transcribed in DISC notation, before the nouns. Although in writing articles and nouns are separated by a space character (e.g., *der Aal*), to model auditory comprehension we removed the space character (e.g., *deral*). By adding the articles to the noun forms, the number of homophones in our dataset was reduced to a substantial extent,

whereas the number of unique word forms more than doubled (from 5,427 to 12,798).

In the first set of simulations we used the same semantic vectors as we did previously for modeling isolated words. That is, the meanings of the definite articles are not taken into account in the semantic vectors, as all forms would be shifted in semantic space in the exactly the same way. After including articles, the validation data now only contained 3,982 homophones, but the number of unseen forms increased to 3,260. Using triphones as cues, we ran two models, one with simulated vectors and the other with *word2vec* semantic vectors. For simulated vectors the results (Table 8) are generally similar to those obtained without articles (Table 5). However, if we look at the evaluation of comprehension with the strict criterion (according to which recognizing a homophone is considered incorrect), without articles *val_strict* is 6%, whereas it is 34% with articles. The generalizability of the model also improves as the number of homophones in the dataset decreases. Even though there are more unseen forms in the current dataset with articles than in the original one without articles, the *val_newform* for comprehension increases by 12% from 51 to 63%.

When using *word2vec* embeddings, adding articles to form representations also improved the comprehension of unseen forms: the *val_newform* astonishingly increased from 0.3 to 58%. Without articles, homophones all shared the same form representations and exactly the same *word2vec* vectors. As a consequence, many triphone cues were superfluous and not well-positioned to discriminate between lemma or inflectional meanings. Now, with the addition of articles, the form space is better discriminated. With an increased number of triphone cues, the model is now able to predict and generalize more accurately for comprehension. However, for production, model performance is generally worse when articles have to be produced. For the training data, for instance, production accuracy drops from 97% (without articles) to 48%. This is of course unsurprising. In the simulation with articles, the semantic representations remain the same, but now identical semantic vectors have to predict more variegated triphone vectors. The learning task has become more challenging, and inevitably resulted in less accurate performance. Replacing the contextually unaware *word2vec* vectors by contextually aware vectors obtained using language models such as BERT (Corbett et al., 2019; Miaschi and Dell'Orletta, 2020) should alleviate this problem.

We can test the model on more challenging data by including indefinite articles (*ein*, *eine*, *einem*, *einen*, *einer*, *eines*), and

TABLE 8 | Comprehension and production accuracy for train and validation datasets with articles.

	Comprehension				Production			
	Train (%)	val_all (%)	val_lenient (%)	val_newform (%)	Train (%)	val_all (%)	val_lenient (%)	val_newform (%)
Simulated	94	76	92	63	81	37	57	19
<i>word2vec</i>	91	69	81	58	48	14	28	1
def + indef	94	80	93	64	82	40	60	15

All three simulations use triphones as cues. The first two rows present results with simulated vectors and *word2vec* embeddings as semantic representations. The simulation presented in the bottom row also makes use of simulated vectors, but includes both definite and indefinite articles.

creating two additional semantic vectors, one for definiteness and one for indefiniteness. This doubles the size of our dataset: half of the words are preceded by definite articles, and the other half by indefinite articles. However, because German indefinite articles are restricted to singular forms, only indefinite singular forms are preceded by indefinite articles. On the meaning side, the $\overrightarrow{\text{DEFINITE}}$ vector is added to the semantic vectors of words preceded by definite articles, and the $\overrightarrow{\text{INDEFINITE}}$ vector is added to vectors for words preceded by either indefinite articles in the singular, or no article in the plural.

The validation data of this dataset confronts the model with in total 3,982 homophones and 3,260 unseen forms. Homophones comprise slightly more words with indefinite articles (57%) whereas unseen forms comprise slightly more definite articles (59%). The results, presented in the bottom row of **Table 8**, are very similar to those with only definite articles (top row). Closer inspection of the results for the validation data shows that for comprehension, accuracies do not differ much across definite and indefinite forms. For production, however, especially for unseen forms, the accuracy for definite articles is twice higher than that for indefinite articles (20 and 9%, averaging out to 15%). This is a straightforward consequence of the much more diverse realizations of indefinite nouns. For definite nouns, the possible triphone cues at the first two positions in the word are always limited to the triphone cues of the six definite articles. For indefiniteness, however, in addition to the six indefinite articles, initial triphone cues also originate from words' stems—indefinite plural forms are realized without articles. The mappings for production are thus faced with a more complex task for indefinites, and the model is therefore more likely to fail on indefinite forms.

4.4.2. Semantic Roles

The simulation studies thus far suggest it is not straightforward to correctly comprehend a novel German word form in isolation, even when articles are provided. This is perhaps not that surprising, as in natural language use, inflected words appear in context, and usually realize not some abstract case ending, but a specific semantic role (also called *thematic role*, see e.g., Harley, 2010). For example, a word in the nominative singular might express a theme, as *der Apfel* in *Der Apfel fällt vom Baum*. (“The apple falls from the tree”), or it might express an agent as *der Junge* in *Der Junge isst den Apfel*. (“The boy eats the apple.”). Exactly the same lemma, used with exactly the same case and number, may still realize very different semantic roles. Consider the two sentences *Ich bin bei der Freundin* (“I’m at the friend’s”) and *Ich gebe der Freundin das Buch* (“I give the book to the friend”). *der Freundin* is dative singular in both cases, but in the first sentence, it expresses a location while in the second it represents the beneficiary or receiver. Semantic roles can also be reflected in a word’s form, independently of case markers. For example, German nouns ending in *-er* are so-called “Nomina Agentis” (Baeskow, 2011). As pointed out by Blevins (2016), case endings are no more (or less) than markers for the intersection of form variation and a distribution class of semantic roles. Since within the framework of the DLM, the aim is to provide mappings between form and meaning, a case label is not

a proper representation of a word’s actual meaning. All it does is specify a range of meanings that the form can have, depending on context. Therefore, even though we can get the mechanics of the model to work with case specifications, doing so clashes with the “discriminative modeling” approach. In what follows, we therefore implement mappings with somewhat more realistic semantic representations of German inflected nouns.

Our starting point is that in German, different cases can realize a wide range of semantic roles. For our simulations, we restrict ourselves to some of the most prominent semantic roles for each case (**Table 9**). Even though these clearly do not reflect the full richness of the semantics of German cases, they suffice for a proof-of-concept simulation.

In order to obtain a dataset with variegated semantic roles, we expanded the previous dataset, with each word form (including its article) appearing with a specification of its semantic role, according to the probabilities presented in **Table 9**. The resulting dataset had 45,605 entries, which we randomly split into 80% training data and 20% validation data. For generating the semantic matrix, we again used number, but instead of a case label, we provided the *semantic role* as inflectional feature. Comprehension accuracy on this data is comparable to the previous simulations: 89% for the training data *train*, and 85% *val_lenient*. Comprehension accuracy on the validation set drops dramatically when we use strict evaluation (4% accuracy). This is unsurprising given that it is impossible for the model to know which semantic role is intended when only being exposed to the word form and its article in isolation, without syntactic context. Production accuracy is likewise comparable to previous simulations with *train* at 78% and *val_lenient* at 61% (*val_newform* 25%). This simple result clarifies that in order to properly model German nouns, it is necessary to take the syntactic context in which a noun occurs into account. Future research will also have to face the challenge of integrating words’ individual usage profiles into the model (see also section 4.2.1 above).

4.5. Incremental Learning vs. the End-State of Learning

In the simulation studies presented thus far, we made use of the regression method to estimate the mappings between form and meaning. The regression method is strictly type based: the data on which a model is trained and evaluated consists of all unique combinations of form vectors \mathbf{c} and semantic vectors \mathbf{s} . In this

TABLE 9 | Probabilities of semantic roles by cases in the German noun system.

Case	Semantic roles
Nominative	Agent (50%), theme (40%), patient (10%)
Genitive	Possessive (90%), partitive (10%)
Dative	Beneficiary (50%), location (50%)
Accusative	Patient (40%), motion (30%), experiencer (30%)

Semantic roles are informed by Schulz and Griesbach (1981). Percentages are simulated and do not necessarily reflect corpus-frequencies of the respective semantic role.

respect, the regression method is very similar to models such as AML, MBL, MGL, and to statistical analyses with the GLM or recursive partitioning methods. However, word types (understood as unique sets $\{c, s\}$) are not uniformly distributed in language, and there is ample evidence that the frequencies with which word types occur co-determines lexical processing (see e.g., Baayen et al., 1997, 2007, 2016; Tomaschek et al., 2018). While some formal theorists flatly deny that word frequency effects exist for inflected words (Yang, 2016), others have argued that there is no problem with integrating frequency of use into formal theories of the lexicon (Jackendoff, 1975; Jackendoff and Audring, 2019), and yet others have argued that it is absolutely essential to incorporate frequency into any meaningful account of language in action (Langacker, 1987; Bybee, 2010).

Within the present approach, effects of frequency of occurrence can be incorporated seamlessly by using incremental learning instead of the end-state of learning as defined by the regression equations (see Danks, 2003; Evert and Arppe, 2015; Shafaei-Bajestan et al., 2021, for the convergence over learning time of incremental learning to the regression end-state of learning). We illustrate this for our German nouns dataset with number and semantic role as crucial constructors of simulated semantic vectors.

We begin with noting that word forms usually do not instantiate all possible semantic roles equally frequently. For instance, a word such as *der Doktor* (“doctor”) will presumably occur mostly as *agent* in the nominative singular form, rather than as *theme* or *patient*. If the model is informed about the probability distributions of semantic roles in actual language use (both in language generally, and lexeme-specific), it may be expected to make more informed decisions when coming across new forms, for instance, by opting for the best match given its past experience.

Incremental learning with the learning rule of Widrow-Hoff makes it possible to start approximating human word-to-word learning as a function of experience. As a consequence, the more frequent a word type occurs in language use, the better it can be learned: practice makes perfect. This sets the following simulation study apart from models such as proposed by McCurdy et al. (2020) or Belth et al. (2021), who base their training regimes strictly on types rather than tokens.

In the absence of empirical frequencies with which combinations of semantic roles and German nouns co-occur, we simulated frequencies of use⁸. To do so, we proceeded as follows. First, we collected token frequencies for all our word forms from CELEX. Next, we assigned an equal part $freq_p$ of this frequency count to each case/number cell realizing this word form. Third, for each paradigm cell, we randomly set to zero some semantic roles, drawing from a binomial distribution with $n = 1$, $p = \frac{1}{K}$, with K the number of semantic roles for the paradigm cell (see Table 9). In this way, on average, one semantic role was omitted per paradigm cell. Finally, given a

proportional frequency count $freq_p$, the semantic roles associated with a paradigm cell received frequencies proportional to the percentages given in Table 9. Further details on this procedure are available in the **Supplementary Materials**, a full example can be found in Table 10.

Having obtained simulated frequencies, we proceeded by randomly selecting 274 different lemmas (1,274 distinct word forms with definite articles included), in order to keep the size of the simulation down — simulating with the Widrow-Hoff rule is computationally expensive. The total number of tokens in this study was 4,470. For the form vectors, we used triphones. The dimension of the simulated semantic vectors was identical to that of the cue vectors. As before, the data was split into 80% training and 20% validation data. We followed the same procedure as in the previous experiments, but instead of computing the mapping matrices in their closed form (i.e., end-state) solution, we used incremental learning.

While for comprehension, the implementation of the learning algorithm is relatively straightforward, this is not the case for production. The `learn_paths` algorithm calculates the support for each of the n-grams, for each possible position in a word. In the current implementation of **JudiLing**, the calculation of positional support is not implemented for incremental learning. Therefore, we do not consider incremental learning of production here.

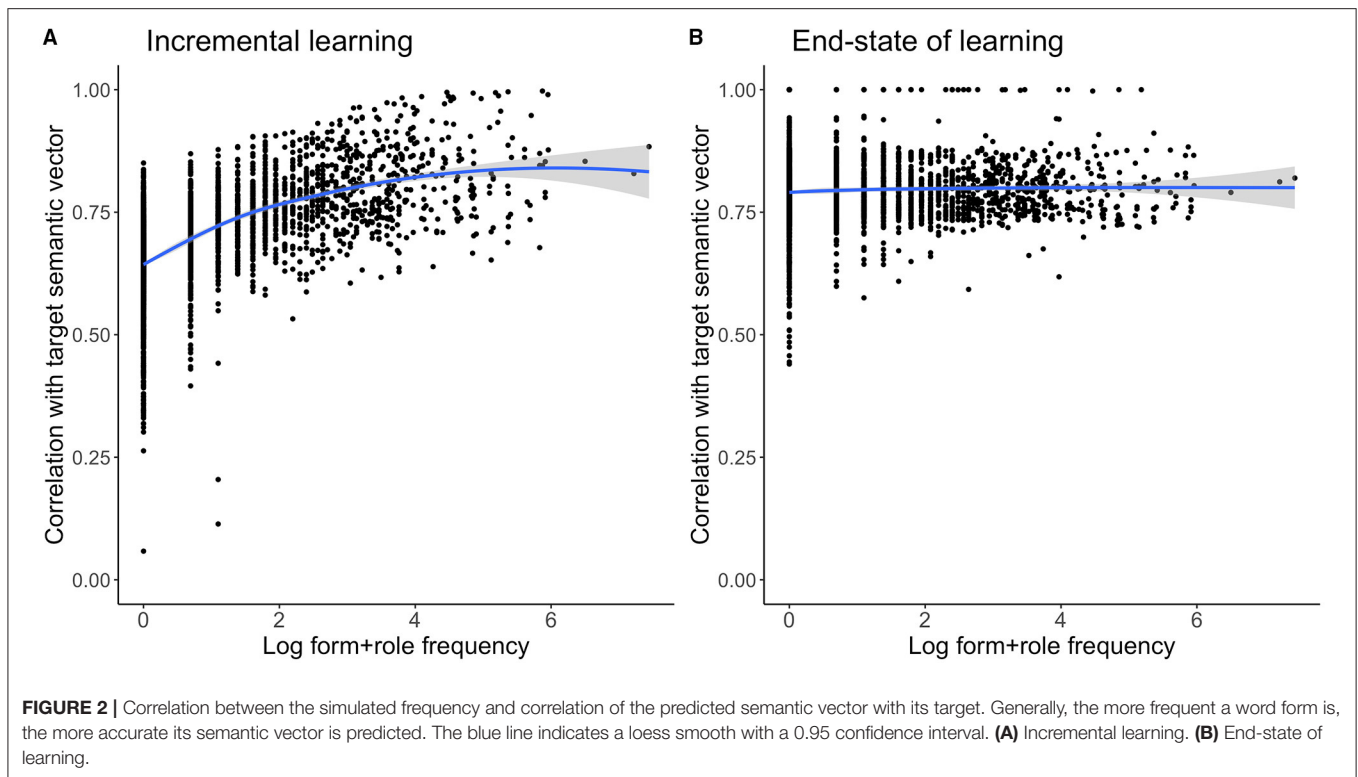
Comprehension accuracy was similar to that observed for previous experiments. Training accuracy when taking into account homophones was 85%, validation accuracy on the full data was 79% (`val_lenient`). Without considering homophones, validation accuracy drops substantially (`val_strict` 7%). This is unsurprising given that from the form alone it is impossible to predict the proper semantic role.

The accuracy of the model's predictions is also closely linked to the frequencies with which words' form+role combinations are encountered in the training data. If a word's form+role combination is very frequent, it is learned better. Figure 2 presents the correlations of words' predicted and targeted semantic vectors against their frequency of occurrence. The left panel presents the results for the incrementally learned

TABLE 10 | Example of simulated frequencies for combinations of case and semantic role for the word form “Adresse.”

Word form	Lemma	Case	Number	Semantic role	Form frequency	Form+role frequency
Adresse	Adresse	Nominative	Singular	Agent	137	20
Adresse	Adresse	Nominative	Singular	Theme	137	16
Adresse	Adresse	Nominative	Singular	Patient	137	0
Adresse	Adresse	Genitive	Singular	Possessive	137	0
Adresse	Adresse	Genitive	Singular	Partitive	137	35
Adresse	Adresse	Dative	Singular	Beneficiary	137	18
Adresse	Adresse	Dative	Singular	Location	137	18
Adresse	Adresse	Accusative	Singular	Patient	137	0
Adresse	Adresse	Accusative	Singular	Motion	137	0
Adresse	Adresse	Accusative	Singular	Experiencer	137	35

⁸Though there are several semantic role labelers available for English [e.g., arising from the CoNLL-2004 and 2005 Shared Tasks (<https://www.cs.upc.edu/~srlconll/home.html>)], there are—to our knowledge—currently no suitable taggers for German.



model, the right panel for the end-state of learning. Clearly, after incremental learning the model predicts the semantics of more frequent form+role combinations more accurately. For the end-state of learning on the other hand, no such effect can be observed. These results clearly illustrate the difference between a token-based model and a typed-based model.

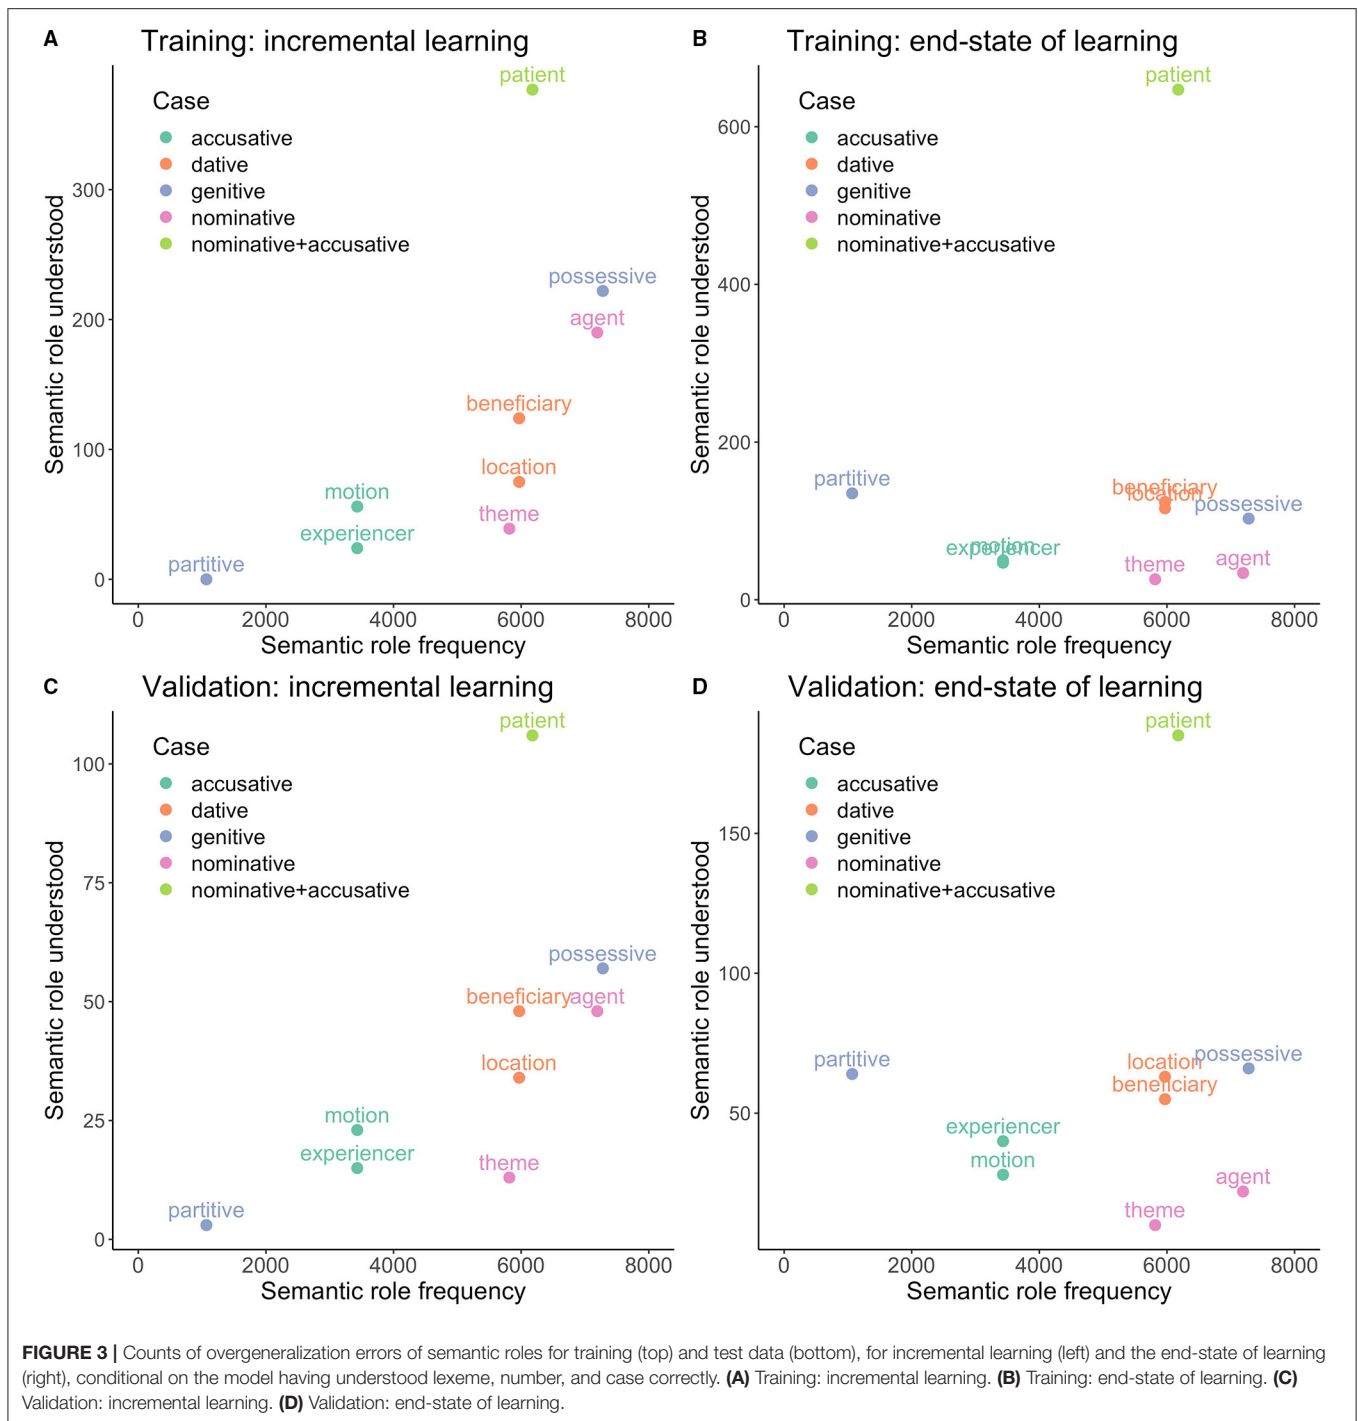
The effect of frequency of use on the kind of errors made by the model is also revealing. We zoom in on those cases where the model was able to correctly identify the lemma and paradigm cell of the word form, but did not get the semantic role right. **Figure 3** provides scatterplots graphing the number of times a semantic role was (incorrectly) understood against the frequency of the form's semantic role, cross-classified by training method (incremental, left panels; end-state of learning, right panels) and by evaluation set (top panels: training data, bottom panels: validation data). For incremental learning, there is a positive correlation between the number of times a semantic role was (incorrectly) identified and the frequency of the semantic role in the training data. Note that the relation is not linear, but curvilinear. A linear relation would have implied that a fixed proportion of word forms would be incorrectly recognized, across all semantic roles. What we see, by contrast, is that greater exposure in language use has an increasingly detrimental effect on learning, with more probable semantic roles being over-identified. Importantly, for the end-state of learning, this curvilinear effect of frequency on learning is absent, with the *PATIENT* role representing an atypical outlier. This outlier status is due to the *patient* semantic role being realized by two cases: *nominative* and *accusative*. As a consequence, it is not only frequent, but it is also predicted by many more different cues

(especially cues from the articles) than is the case for other semantic roles.

In other words, with incremental learning, strong frequency effects emerge, hand in hand with overgeneralization of semantic roles (the study by Ramscar et al., 2013 makes the same point for irregular English noun plurals). By contrast, for the end-state of learning, such effects are absent. Mathematically, this makes sense: as experience (i.e., volume of training data) goes to infinity, all forms are learned an infinite number of times, and frequency is no longer distinctive.

With incremental learning, it is also possible to follow the learning trajectory of the model. **Figure 4** presents this trajectory at 10 evaluation points. Learning proceeds rapidly during the first 15,000 learning events and slows down afterwards. Validation accuracy `val_lenient` closely follows training accuracy, which is a straightforward consequence of the large numbers of homophones. `val_newforms` on the other hand stays relatively low, in accordance with the semi-productivity of the German declension system.

Note that in this simulation we only pass through the data once. If a word form has a form+role frequency of 1, it is only seen a single time during training. As such, it is not possible for the model to reach accuracies as high as at the end-state of learning (indicated as dots in **Figure 4**), which would be reached eventually after an infinite number of passes through the data (Danks, 2003; Evert and Arppe, 2015; Shafaei-Bajestan et al., 2021). This sets our approach apart from deep learning, where models are trained on many iterations through the data set until the loss function reaches a local minimum. Whereas such a procedure makes sense



for language engineering, it does not make sense for human learning: we don't relive the same exposure to data multiple times, and for healthy people, there is no point in learning after which performance degrades. For instance, vocabulary learning is a continuous process straight into old age (Keuleers et al., 2015).

Note finally, that even though incremental learning is certainly superior for modeling realistic frequency effects, there are also cases where the end-state of learning can be the

preferred choice of modeling. Incremental learning is much more computationally expensive which becomes a problem especially if the training set is large and frequencies are high. Moreover, in cases where simulated speakers are expected to have learned a phenomenon well enough, the end-state simulation might be sufficient.

In summary, the present modeling framework offers the possibility to approximate incremental human learning and the consequences of frequency of exposure for learning in a

cognitively motivated way (see also Chuang et al., 2020a, for learning in a multilingual setting).

4.6. Model Complexity

LDL mappings are costly in the number of connection weights, or equivalently, the number of beta coefficients. For example, the mapping matrix F for the dataset discussed in section 4.4.2 has 35 million weights ($5,913 \times 5,913$ dimensions), rendering it much more costly in terms of the number of weights than deep-learning models, models such as AML, MBL, and recursive partitioning methods.

Inspection of the distribution of weights, however, clarifies that many weights are very close to zero. Apparently, many cues have low discriminative value. This suggests their connections can be pruned without seriously affecting model performance. This can be tested by selecting a threshold ϑ and setting all absolute values in the mapping matrix that fall below this

threshold to zero. **Figure 5** shows, for varying ϑ , that up to 40% of the small weights can be pruned without substantially impacting model performance with end-state of learning. As neural pruning is part and parcel of human cortical development (see e.g., Gogtay et al., 2004), an interesting topic for further research is to integrate incremental learning with neural pruning of uninformative connections.

5. DISCUSSION

In this study, we illustrated the methodological consequences of the many different choices that have to be made when modeling morphological systems within the discriminative lexicon framework, using LDL as modeling engine. We illustrated these choices for the German noun system. This system is “degenerate,” as many of its paradigm cells share the same word forms (homophones). This system is also in many ways irregular: a noun’s declension class can often not be fully predicted by its phonology, gender, or semantics (Köpcke, 1988). The results we obtained with LDL reflect this complexity. The model can learn word forms very well, achieving accuracies of more than 90% on both comprehension and production when evaluated on training data. It can also generalize very well to new paradigm cells when it comes to word forms it has already seen, thanks to the ubiquitous homophony that characterizes German noun paradigms. However, it also mirrors the unpredictability of German inflections when it comes to word forms it has not seen before. Accuracies for both comprehension and production suffer. Nevertheless, the model shows some semi-productivity and succeeds in generalizing to many of the sub-regularities found in the German noun system (Wunderlich, 1999), reaching accuracies of 50% on comprehension and 20% on production. Since German speakers encounter similar problems with new German word forms, as has been demonstrated in various wug studies (Zaretsky et al., 2013; McCurdy et al., 2020), our model properly exhibits the limitations that are also characteristic for native speakers.

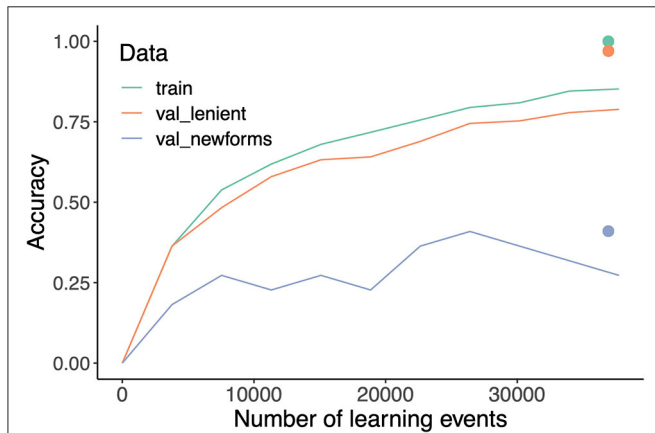


FIGURE 4 | Comprehension accuracy over the course of learning. After a very fast increase in accuracy over the first 15,000 learning events, the amount of learning levels off. Points indicate the accuracy at the end-state of learning which the incremental model would reach eventually after an infinite number of learning events.

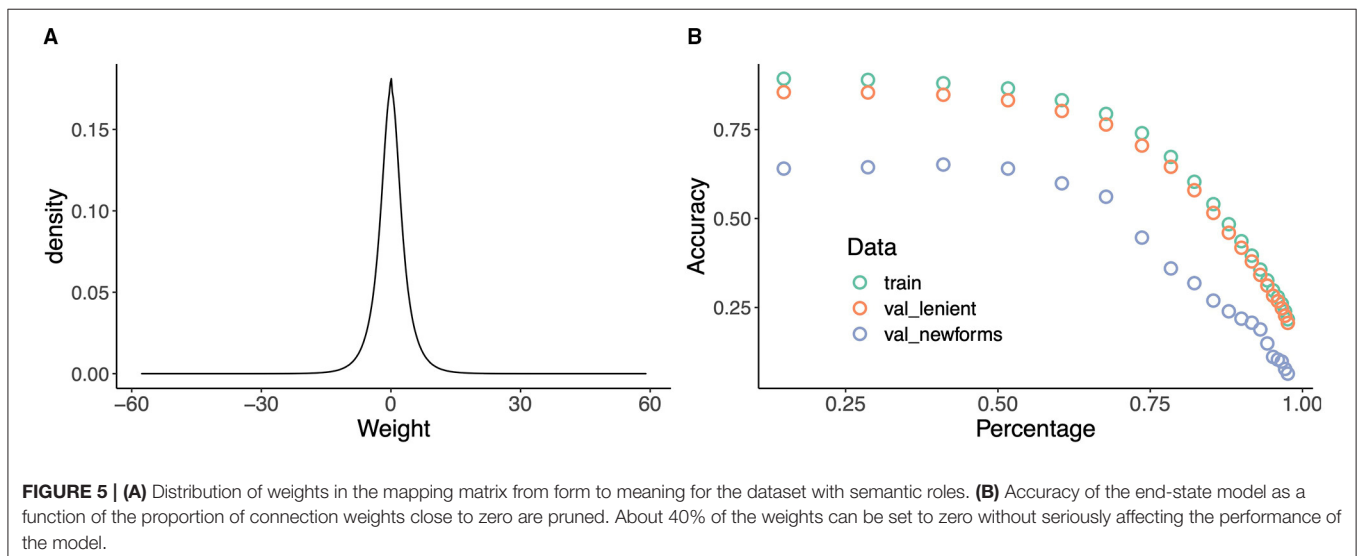


FIGURE 5 | (A) Distribution of weights in the mapping matrix from form to meaning for the dataset with semantic roles. **(B)** Accuracy of the end-state model as a function of the proportion of connection weights close to zero are pruned. About 40% of the weights can be set to zero without seriously affecting the performance of the model.

In this study, we also probed the modeling of German nouns in context. The rampant homophony that characterizes German noun paradigms is a straightforward consequence of considering words in isolation. The amount of homophony can be substantially reduced by including articles, in which case the model still performs well. In context, case-inflected words typically do not realize a specific case meaning, but rather a specific semantic role. As case endings typically do not stand in a one-to-one relation with semantic roles, we also examined to what extent we can make the model more realistic by replacing semantic vectors for cases with semantic vectors for a variety of semantic roles. For the simulated dataset that we constructed, the model again performed well.

For this dataset, we also demonstrated how the consequences of frequency of occurrence can be brought into the model, namely, by moving from the end-state of learning (estimated with regression) to incremental learning using the Widrow-Hoff learning rule.

One limitation of the present approach is that most models have been using very high-level abstract representations. The phone-based representation, for example, involves tremendous simplifications compared to real speech, as variability in pronunciations is enormous (Ernestus et al., 2002; Johnson, 2004; Shafaei-Bajestan et al., 2021). On the meaning side, traditional case labels have no intrinsic semantic content, and although we can replace cases with semantic roles, these too are still too simplistic to be able to capture the full complexity of the semantics of words in context. However, we note that even with the present high-level representations, the model can still generate useful predictions. We note here that various other studies carried out within this framework have successfully modeled a range of aspects of human lexical processing (see Chuang and Baayen, 2021, for further details). In summary, even though the current framework undoubtedly misses out on a great number of nuanced but potentially informative features of forms and meanings in real language use, it can still serve as a useful linguistic tool to explore the strengths and weaknesses of morphological systems.

A question that inevitably arises in the context of computational modeling is how cognitively plausible a model is. In the introduction, we called attention to the distinction made by Breiman (2001) between statistical models and machine learning models. We view LDL primarily as a statistical model that enables us to clarify, at a functional level of analysis, quantitative structure in the lexicon as well as understand the challenges a language processing system faces, without claiming that our model is cognitive reality. However, it is worth noting that LDL helps incorporate biologically and psychologically plausible learning into linguistic theory by making use of the principle of error-driven learning (when training the model incrementally). The very simple learning rules of Widrow-Hoff and Rescorla-Wagner have been shown to excellently explain phenomena from a range of domains in e.g., biology and psychology (see e.g., Rescorla, 1988; Schultz, 1998; Marsolek, 2008; Oppenheim et al., 2010; Trimmer et al., 2012).

It is possible to take the model as point of departure for addressing questions at the level of neural organization in the

brain. For instance, Heitmeier and Baayen (2020) were interested in clarifying whether the framework of the discriminative lexicon properly predicts the dissociations of form and meaning observed for aphasic speakers producing English regular and irregular past-tense forms, following Joanisse and Seidenberg (1999). They took the unordered banks of units of form and meaning (the column dimensions of the *C* and *S* matrices) and projected them onto two-dimensional surfaces approximating, however crudely, cortical maps. This made it possible to lesion the network in a topologically cohesive way, rather than by randomly taking out connections across the whole network. For projection, they made use of an algorithm from physics (<http://www.schmuhl.org/graphopt/>) for displaying graphs, but temporal self-organizing maps (TSOMs, Ferro et al., 2011; Chersi et al., 2014) offer a much more fine-grained and principled way for modeling morphological organization that builds on principles of error-driven learning.

Deep learning algorithms provide the analyst with powerful modeling tools, but it seems they are too powerful (see e.g., McCurdy et al., 2020) for understanding not only the strengths but also the weaknesses and the frailties of human lexical memory and lexical processing. However, linguistic models are in a different way also too powerful on the one hand, and too underspecified on the other hand. Paradigms are typically constructed to accommodate any contrast between forms and inflectional functions, even when a contrast is attested only for a few forms in the language. The result is an overabundance of homophones, which are severely underspecified with respect to their real meanings in actual language use (such as their semantic roles). Furthermore, in actual language use, inflected forms can occur at very different frequencies and some are never encountered at all (Karlsson, 1986; Janda and Tyers, 2018), which in turn has demonstrable consequences for lexical processing (Lõo et al., 2018)⁹. An interesting challenge for further research is to clarify how different degrees of paradigm economy (Ackerman and Malouf, 2013) are reflected in the matrices that define mappings between form and meaning within the framework of the discriminative lexicon.

In this study, we have provided an overview of the many choice points that arise in modeling with LDL, each of which requires knowledge of morphology and morphological theory. The implications of our approach to psycho-computational modeling for morphological theory depends on the specifics of a given specific theory of morphology. Our approach is broadly consistent with usage-based approaches to morphology (Bybee, 1985, 2010), and with Word and Paradigm Morphology (Blevins, 2016). It is less clear whether our modeling approach is informative for theories that are only interested in defining possible words. With this methodological study, we have shed some light on the many questions and issues that do not arise in formal theories of morphology, but that have to be addressed

⁹Note that we do not claim that rare inflected word forms cannot be processed. Generally, the more regular a morphological system, the more easily the model can predict new forms (e.g., in Estonian, Chuang et al., 2020b), while in semi-productive cases such as German or Maltese (Nieder et al., 2021) generalization is much more difficult.

in a linguistically informed way when the goal of one's theory is to better understand, and predict, in all its complexity, human lexical processing across comprehension and production.

DATA AVAILABILITY STATEMENT

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession number(s) can be found below: <https://osf.io/zrw2v/>.

AUTHOR CONTRIBUTIONS

RB, Y-YC, and MH: conception and design of the study and writing. MH and Y-YC: computational implementation and modeling. All authors contributed to the article and approved the submitted version.

REFERENCES

- Ackerman, F., and Malouf, R. (2013). Morphological organization: the low conditional entropy conjecture. *Language* 89, 429–464. doi: 10.1353/lan.2013.0054
- Albright, A., and Hayes, B. (2003). Rules vs. analogy in English past tenses: a computational/experimental study. *Cognition* 90, 119–161. doi: 10.1016/S0010-0277(03)00146-X
- Arndt-Lappe, S. (2011). Towards an exemplar-based model of stress in English noun-noun compounds. *J. Linguist.* 47, 549–585. doi: 10.1017/S0022226711000028
- 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:e0174623. doi: 10.1371/journal.pone.0174623
- Baayen, R. H., Chuang, Y.-Y., and Blevins, J. P. (2018). Inflectional morphology with linear mappings. *Mental Lexicon* 13, 232–270. doi: 10.1075/ml.18010.baa
- 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* 2019:4895891. doi: 10.1155/2019/4895891
- Baayen, R. H., Dijkstra, T., and Schreuder, R. (1997). Singulars and plurals in Dutch: evidence for a parallel dual route model. *J. Mem. Lang.* 36, 94–117. doi: 10.1006/jmla.1997.2509
- Baayen, R. H., Milin, P., Filipović Durđević, D., Hendrix, P., and Marelli, M. (2011). An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. *Psychol. Rev.* 118, 438–482. doi: 10.1037/a0023851
- Baayen, R. H., Milin, P., and Ramscar, M. (2016). Frequency in lexical processing. *Aphasiology* 30, 1174–1220. doi: 10.1080/02687038.2016.1147767
- Baayen, R. H., Piepenbrock, R., and Gulikers, L. (1995). *The CELEX Lexical Database [CD ROM]*. Philadelphia, PA: Linguistic Data Consortium, University of Pennsylvania.
- Baayen, R. H., Wurm, L. H., and Aycok, J. (2007). Lexical dynamics for low-frequency complex words. A regression study across tasks and modalities. *Mental Lexicon* 2, 419–463. doi: 10.1075/ml.2.3.06baa
- Baekow, H. (2011). *Abgeleitete Personenbezeichnungen im Deutschen und Englischen: kontrastive Wortbildungsanalysen im Rahmen des minimalistischen Programms und unter Berücksichtigung sprachhistorischer Aspekte*, Vol. 62. Berlin; Boston, MA: Walter de Gruyter.
- Behrens, H., and Tomasello, M. (1999). And what about the chinese? *Behav. Brain Sci.* 22, 1014–1014. doi: 10.1017/S0140525X99222224

FUNDING

This work was supported by ERC-WIDE (European Research Council—Wide Incremental learning with Discrimination nEtworks), grant number 742545.

ACKNOWLEDGMENTS

We acknowledge support by Open Access Publishing Fund of University of Tübingen. We thank Ruben van de Vijver and the two reviewers for their helpful comments on earlier versions of this manuscript.

SUPPLEMENTARY MATERIAL

Supplementary Materials are available at <https://osf.io/zrw2v/>.

- Belth, C., Payne, S., Beser, D., Kodner, J., and Yang, C. (2021). The greedy and recursive search for morphological productivity. *arXiv [Preprint]*. *arXiv:2105.05790*.
- Bierwisch, M. (2018). *Syntactic Features in Morphology: General Problems of So-Called Pronominal Inflection in German*. Berlin; Boston, MA: De Gruyter Mouton.
- Blevins, J. P. (2016). *Word and Paradigm Morphology*. Oxford: Oxford University Press. doi: 10.1093/acprof:oso/9780199593545.001.0001
- Boersma, P. (1998). *Functional Phonology*. The Hague: Holland Academic Graphics.
- Boersma, P., and Hayes, B. (2001). Empirical tests of the gradual learning algorithm. *Linguist. Inq.* 32, 45–86. doi: 10.1162/002438901554586
- Bojanowski, P., Grave, E., Joulin, A., and Mikolov, T. (2017). Enriching word vectors with subword information. *Trans. Assoc. Comput. Linguist.* 5, 135–146. doi: 10.1162/tacl_a_00051
- Breiman, L. (2001). Statistical modeling: the two cultures (with comments and a rejoinder by the author). *Stat. Sci.* 16, 199–231. doi: 10.1214/ss/1009213726
- Breiman, L., Friedman, J. H., Olshen, R., and Stone, C. J. (1984). *Classification and Regression Trees*. Belmont, CA: Wadsworth International Group.
- Bybee, J. (2010). *Language, Usage and Cognition*. Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511750526
- Bybee, J. L. (1985). *Morphology: A Study of the Relation Between Meaning and Form*. Amsterdam: Benjamins. doi: 10.1075/tsl.9
- Cahill, L., and Gazdar, G. (1999). German noun inflection. *J. Linguist.* 35, 1–42. doi: 10.1017/S0022226798007294
- 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. *Top. Cogn. Sci.* 6, 476–491. doi: 10.1111/tops.12094
- Chuang, Y.-Y., and Baayen, R. H. (2021). “Discriminative learning and the lexicon: NDL and LDL,” in *Oxford Research Encyclopedia of Linguistics*, ed M. Aronoff (Oxford: Oxford University Press).
- Chuang, Y.-Y., Bell, M., Banke, I., and Baayen, R. H. (2020a). Bilingual and multilingual mental lexicon: a modeling study with Linear Discriminative Learning. *Lang. Learn.* 71, 219–292. doi: 10.31234/osf.io/adtyr
- Chuang, Y.-Y., Loo, K., Blevins, J. P., and Baayen, R. H. (2020b). “Estonian case inflection made simple: a case study in word and paradigm morphology with linear discriminative learning,” in *Complex Words Advances in Morphology*, eds L. Körtvélyessy and P. Štekauer (Cambridge: Cambridge University Press), 119–141. doi: 10.1017/9781108780643.008
- Chuang, Y.-Y., Vollmer, M.-L., Shafaei-Bajestan, E., Gahl, S., Hendrix, P., and Baayen, R. H. (2020c). The processing of pseudoword form and meaning in production and comprehension: a computational modeling approach using linear discriminative learning. *Behav. Res. Methods* 53, 945–976. doi: 10.3758/s13428-020-01356-w

- Clahsen, H. (1999). Lexical entries and rules of language: a multidisciplinary study of German inflection. *Behav. Brain Sci.* 22, 991–1013. doi: 10.1017/S0140525X99002228
- Coltheart, M., Curtis, B., Atkins, P., and Haller, M. (1993). Models of reading aloud: dual-route and parallel-distributed-processing approaches. *Psychol. Rev.* 100:589. doi: 10.1037/0033-295X.100.4.589
- Corbett, G. G. (1991). "Introduction," in *Gender* (Cambridge: Cambridge University Press). doi: 10.1017/CBO9781139166119
- Corbett, G. G., Devlin, J., Chang, M. W., Lee, K., and Toutanova, K. (2019). "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1* (Minneapolis, MN: Long and Short Papers), 4171–4186. doi: 10.18653/v1/N19-1423
- Daelemans, W., Berck, P., and Gillis, S. (1995). "Linguistics as data mining: Dutch diminutives," in *CLIN V, Papers from the 5th CLIN Meeting*, eds T. Andernach, M. Moll, and A. Nijholt (Enschede: Parlevink), 59–71.
- Daelemans, W., and Van den Bosch, A. (2005). *Memory-Based Language Processing*. Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511486579
- Dalal, N., and Triggs, B. (2005). "Histograms of oriented gradients for human detection," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)* (Los Alamitos, CA), 886–893. doi: 10.1109/CVPR.2005.177
- Danks, D. (2003). Equilibria of the Rescorla-Wagner model. *J. Math. Psychol.* 47, 109–121. doi: 10.1016/S0022-2496(02)00016-0
- Dell, G. S. (1986). A spreading-activation theory of retrieval in sentence production. *Psychol. Rev.* 93:283. doi: 10.1037/0033-295X.93.3.283
- Elman, J. L. (2009). On the meaning of words and dinosaur bones: lexical knowledge without a lexicon. *Cogn. Sci.* 33, 547–582. doi: 10.1111/j.1551-6709.2009.01023.x
- Ernestus, M., and Baayen, R. H. (2003). Predicting the unpredictable: interpreting neutralized segments in Dutch. *Language* 79, 5–38. doi: 10.1353/lan.2003.0076
- Ernestus, M., Baayen, R. H., and Schreuder, R. (2002). The recognition of reduced word forms. *Brain Lang.* 81, 162–173. doi: 10.1006/brln.200.1.2514
- Evans, R., and Gazdar, G. (1996). DATR: a language for lexical knowledge. *Comput. Linguist.* 22, 167–216.
- Evert, S., and Arppe, A. (2015). "Some theoretical and experimental observations on naive discriminative learning," in *Proceedings of the 6th Conference on Quantitative Investigations in Theoretical Linguistics (QITL-6)* (Tübingen).
- Ferro, M., Marzi, C., and Pirrelli, V. (2011). A self-organizing model of word storage and processing: implications for morphology learning. *Lingue e Linguaggio* 10, 209–226.
- Finkel, R., and Stump, G. (2007). Principal parts and morphological typology. *Morphology* 17, 39–75. doi: 10.1007/s11525-007-9115-9
- Gaeta, L. (2008). Die deutsche Pluralbildung zwischen deskriptiver Angemessenheit und Sprachtheorie. *Z. German. Linguist.* 36, 74–108. doi: 10.1515/ZGL.2008.005
- Gaskell, M. G., and Marslen-Wilson, W. (1997). Integrating form and meaning: a distributed model of speech perception. *Lang. Cogn. Process.* 12, 613–656. doi: 10.1080/016909697386646
- Goebel, R., and Indefrey, P. (2000). "A recurrent network with short-term memory capacity learning the German's plural," in *Models of Language Acquisition: Inductive and Deductive Approaches*, eds P. Broeder and J. Murre (Oxford: Oxford University Press), 177–200.
- Gogtay, N., Giedd, J. N., Lusk, L., Hayashi, K. M., Greenstein, D., Vaituzis, A. C., et al. (2004). Dynamic mapping of human cortical development during childhood through early adulthood. *Proc. Natl. Acad. Sci. U.S.A.* 101, 8174–8179. doi: 10.1073/pnas.0402680101
- Goldsmith, J., and O'Brien, J. (2006). Learning inflectional classes. *Lang. Learn. Dev.* 2, 219–250. doi: 10.1207/s15473341lld0204_1
- Haapalainen, M., and Majorin, A. (1994). *Gertwol: Ein System zur Automatischen Wortformerkennung Deutscher Wörter*. Lingsoft, Inc.
- Harley, H. (2010). "Thematic roles," in *The Cambridge Encyclopedia of the Language Sciences*, ed P. Hogan (Cambridge: Cambridge University Press), 861–862.
- Harm, M. W., and Seidenberg, M. S. (2004). Computing the meanings of words in reading: cooperative division of labor between visual and phonological processes. *Psychol. Rev.* 111, 662–720. doi: 10.1037/0033-295X.111.3.662
- Heitmeier, M., and Baayen, R. H. (2020). Simulating phonological and semantic impairment of English tense inflection with Linear Discriminative Learning. *Mental Lexicon* 15, 385–421. doi: 10.1075/ml.20003.hei
- Indefrey, P. (1999). Some problems with the lexical status of nondefault inflection. *Behav. Brain Sci.* 22:1025. doi: 10.1017/S0140525X99342229
- Ivens, S., and Koslin, B. (1991). *Demands for Reading Literacy Require New Accountability Methods*. Brewster, NY: Touchstone Applied Science Associates.
- Jackendoff, R., and Audring, J. (2019). *The Texture of the Lexicon: Relational Morphology and the Parallel Architecture*. Oxford: Oxford University Press. doi: 10.1093/oso/9780198827900.001.0001
- Jackendoff, R. S. (1975). Morphological and semantic regularities in the lexicon. *Language* 51, 639–671. doi: 10.2307/412891
- Janda, A. L., and Tyers, M. F. (2018). Less is more: why all paradigms are defective, and why that is a good thing. *Corpus Linguist. Linguist. Theory.* 17, 109–141. doi: 10.1515/cllt-2018-0031
- Joanisse, M. F., and Seidenberg, M. S. (1999). Impairments in verb morphology after brain injury: a connectionist model. *Proc. Natl. Acad. Sci. U.S.A.* 96, 7592–7597. doi: 10.1073/pnas.96.13.7592
- Johnson, K. (2004). "Massive reduction in conversational American English," in *Spontaneous Speech: Data and Analysis. Proceedings of the 1st Session of the 10th International Symposium* (Tokyo: The National International Institute for Japanese Language), 29–54.
- Karlssohn, F. (1986). Frequency considerations in morphology. *Zeitschrift Phonetik Sprachwissenschaft Kommunikationsforschung* 39, 19–28.
- Karttunen, L. (2003). "Computing with realizational morphology," in *International Conference on Intelligent Text Processing and Computational Linguistics* (Berlin; Heidelberg: Springer), 203–214. doi: 10.1007/3-540-36456-0_20
- Keuleers, E., Sandra, D., Daelemans, W., Gillis, S., Durieux, G., and Martens, E. (2007). Dutch plural inflection: the exception that proves the analogy. *Cogn. Psychol.* 54, 283–318. doi: 10.1016/j.cogpsych.2006.07.002
- Keuleers, E., Stevens, M., Mandera, P., and Brysbaert, M. (2015). Word knowledge in the crowd: measuring vocabulary size and word prevalence in a massive online experiment. *Q. J. Exp. Psychol.* 8, 1665–1692. doi: 10.1080/17470218.2015.1022560
- Kirov, C., and Cotterell, R. (2018). Recurrent neural networks in linguistic theory: revisiting Pinker and Prince (1988) and the past tense debate. *Trans. Assoc. Comput. Linguist.* 6, 651–665. doi: 10.1162/tacl_a_00247
- Köpcke, K.-M. (1988). Schemas in German plural formation. *Lingua* 74, 303–335. doi: 10.1016/0024-3841(88)90064-2
- Landauer, T., and Dumais, S. (1997). A solution to Plato's problem: the latent semantic analysis theory of acquisition, induction and representation of knowledge. *Psychol. Rev.* 104, 211–240. doi: 10.1037/0033-295X.104.2.211
- Langacker, R. W. (1987). *Foundations of Cognitive Grammar: Theoretical Prerequisites*, Vol. 1. Stanford, CA: Stanford University Press.
- Levelt, W., Roelofs, A., and Meyer, A. S. (1999). A theory of lexical access in speech production. *Behav. Brain Sci.* 22, 1–38. doi: 10.1017/S0140525X99451775
- Linke, M., Broeker, F., Ramscar, M., and Baayen, R. H. (2017). Are baboons learning "orthographic" representations? Probably not. *PLoS ONE* 12:e0183876. doi: 10.1371/journal.pone.0183876
- Linzen, T., and Baroni, M. (2021). Syntactic structure from deep learning. *Annu. Rev. Linguist.* 7, 195–212. doi: 10.1146/annurev-linguistics-032020-051035
- Lõo, K., Jaervikivi, J., Tomaschek, F., Tucker, B., and Baayen, R. (2018). Production of Estonian case-inflected nouns shows whole-word frequency and paradigmatic effects. *Morphology* 1, 71–97. doi: 10.1007/s11525-017-9318-7
- Luo, X. (2021). *JudiLing: an implementation for Linear Discriminative Learning in JudiLing* (unpublished Master's thesis).
- Luo, X., Chuang, Y.-Y., and Baayen, R. H. (2021). Judiling: an implementation in Julia of Linear Discriminative Learning algorithms for language modeling. Available online at: <https://github.com/MegamindHenry/JudiLing.jl>
- MacWhinney, B., and Leinbach, J. (1991). Implementations are not conceptualizations: revising the verb learning model. *Cognition* 40, 121–157. doi: 10.1016/0010-0277(91)90048-9
- Malouf, R. (2017). Abstractive morphological learning with a recurrent neural network. *Morphology* 27, 431–458. doi: 10.1007/s11525-017-9307-x

- Marcus, G. F., Brinkmann, U., Clahsen, H., Wiese, R., and Pinker, S. (1995). German inflection: the exception that proves the rule. *Cogn. Psychol.* 29, 189–256. doi: 10.1006/cogp.1995.1015
- Marsolek, C. J. (2008). What antipriming reveals about priming. *Trends Cogn. Sci.* 12, 176–181. doi: 10.1016/j.tics.2008.02.005
- Matthews, P. H. (1974). *Morphology. An Introduction to the Theory of Word Structure*. Cambridge: Cambridge University Press.
- McCurdy, K. (2019). *Neural networks don't learn default rules for German plurals, but that's okay, neither do Germans* (Master's Thesis). University of Edinburgh, Edinburgh, United Kingdom.
- McCurdy, K., Goldwater, S., and Lopez, A. (2020). "Inflecting when there's no majority: limitations of encoder-decoder neural networks as cognitive models for German plurals," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (Stroudsburg, PA: Association for Computational Linguistics), 1745–1756. doi: 10.18653/v1/2020.acl-main.159
- Miaschi, A., and Dell'Orletta, F. (2020). "Contextual and non-contextual word embeddings: an in-depth linguistic investigation," in *Proceedings of the 5th Workshop on Representation Learning for NLP* (Stroudsburg, PA), 110–119. doi: 10.18653/v1/2020.repl4nlp-1.15
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013a). Efficient estimation of word representations in vector space. *arXiv [Preprint]*. arXiv:1301.3781.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013b). "Distributed representations of words and phrases and their compositionality," in *Advances in Neural Information Processing Systems*, eds C. J. C. Burges and L. Bottou and M. Welling and Z. Ghahramani and K.Q. Weinberger (Red Hook, NY: Curran Associates Inc.), 3111–3119.
- Milin, P., Madabushi, H. T., Croucher, M., and Divjak, D. (2020). Keeping it simple: Implementation and performance of the proto-principle of adaptation and learning in the language sciences. *arXiv [Preprint]*. arXiv:2003.03813.
- Mirković, J., MacDonald, M. C., and Seidenberg, M. S. (2005). Where does gender come from? Evidence from a complex inflectional system. *Lang. Cogn. Process.* 20, 139–167. doi: 10.1080/01690960444000205
- Nakisa, R. C., and Hahn, U. (1996). "Where defaults don't help: the case of the German plural system," in *Proc. 18th Annu. Conf. Cogn. Sci. Soc* (Mahwah, NJ), 177–182.
- Nieder, J., Chuang, Y.-Y., van de Vijver, R., and Baayen, R. (2021). Comprehension, production and processing of maltese plurals in the discriminative lexicon. doi: 10.31234/osf.io/rkath
- Oppenheim, G. M., Dell, G. S., and Schwartz, M. F. (2010). The dark side of incremental learning: a model of cumulative semantic interference during lexical access in speech production. *Cognition* 114, 227–252. doi: 10.1016/j.cognition.2009.09.007
- Pennington, J., Socher, R., and Manning, C. D. (2014). "Glove: global vectors for word representation," in *Empirical Methods in Natural Language Processing (EMNLP)* (Stroudsburg, PA), 1532–1543. doi: 10.3115/v1/D14-1162
- Pinker, S., and Prince, A. (1988). On language and connectionism. *Cognition* 28, 73–193. doi: 10.1016/0010-0277(88)90032-7
- Prince, A., and Smolensky, P. (2008). *Optimality Theory: Constraint Interaction in Generative Grammar*. Malden, MA: John Wiley & Sons.
- Ramsar, M., Dye, M., and McCauley, S. M. (2013). Error and expectation in language learning: the curious absence of mouses in adult speech. *Language* 89, 760–793. doi: 10.1353/lan.2013.0068
- Rescorla, R. A. (1988). Pavlovian conditioning. It's not what you think it is. *Am. Psychol.* 43, 151–160. doi: 10.1037/0003-066X.43.3.151
- Rumelhart, D. E., and McClelland, J. L. (1986). "On learning the past tenses of English verbs," in *Parallel Distributed Processing. Explorations in the Microstructure of Cognition. Vol. 2: Psychological and Biological Models*, eds J. L. McClelland and D. E. Rumelhart (Cambridge, MA: The MIT Press), 216–271. doi: 10.7551/mitpress/5237.001.0001
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *J. Neurophysiol.* 80, 1–27. doi: 10.1152/jn.1998.80.1.1
- Schulz, D., and Griesbach, H. (1981). *Grammatik der deutschen Sprache*. Munich: Max Hueber Verlag. doi: 10.1515/infodaf-1981-080512
- Shafaei-Bajestan, E., Tari, M. M., and Baayen, R. H. (2021). LDL-AURIS: error-driven learning in modeling spoken word recognition. *Lang. Cogn. Neurosci.* 1–28. doi: 10.1080/23273798.2021.1954207
- Shahmohammadi, H., Lensch, H. and Baayen, R. H. (in press). "Learning zero-shot multifaceted visually grounded word embeddings via multi-task training," in *Proceedings of the 25th Conference on Computational Natural Language Learning*.
- Sinclair, J. (1991). *Corpus, Concordance, Collocation*. Oxford: Oxford University Press.
- Skousen, R. (1989). *Analogical Modeling of Language*. Dordrecht: Kluwer. doi: 10.1007/978-94-009-1906-8
- Skousen, R. (2002). *Analogical Modeling*. Amsterdam: Benjamins. doi: 10.1075/hcp.10
- Stump, G. (2001). *Inflectional Morphology: A Theory of Paradigm Structure*. Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511486333
- Tognini-Bonelli, E. (2001). *Corpus Linguistics at Work*, Vol. 6. Amsterdam: John Benjamins Publishing. doi: 10.1075/scl.6
- Tomaschek, F., Tucker, B. V., Fasiolo, M., and Baayen, R. H. (2018). Practice makes perfect: the consequences of lexical proficiency for articulation. *Linguist. Vanguard* 4s2. doi: 10.1515/lingvan-2017-0018
- Trimmer, P. C., McNamara, J. M., Houston, A. I., and Marshall, J. A. R. (2012). Does natural selection favour the Rescorla-Wagner rule? *J. Theoret. Biol.* 302, 39–52. doi: 10.1016/j.jtbi.2012.02.014
- Trommer, J. (2021). The subsegmental structure of German plural allomorphy. *Nat. Lang. Linguist. Theory* 39, 601–656. doi: 10.1007/s11049-020-09479-7
- Van de Vijver, R., and Baer-Henney, D. (2014). Developing biases. *Front. Psychol.* 5:634. doi: 10.3389/fpsyg.2014.00634
- Widrow, B., and Hoff, M. E. (1960). "Adaptive switching circuits," in *1960 WESCON Convention Record Part IV* (New York, NY: IRE), 96–104. doi: 10.21236/AD0241531
- Wiese, R. (1999). On default rules and other rules. *Behav. Brain Sci.* 22, 1043–1044. doi: 10.1017/S0140525X99532226
- Wunderlich, D. (1999). German noun plural reconsidered. *Behav. Brain Sci.* 22, 1044–1045. doi: 10.1017/S0140525X99542222
- Yamada, I., Asai, A., Sakuma, J., Shindo, H., Takeda, H., Takefuji, Y., et al. (2020). "Wikipedia2Vec: an efficient toolkit for learning and visualizing the embeddings of words and entities from Wikipedia," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations* (Stroudsburg, PA: Association for Computational Linguistics), 23–30. doi: 10.18653/v1/2020.emnlp-demos.4
- Yang, C. (2016). *The Price of Linguistic Productivity*. Cambridge, MA: The MIT Press. doi: 10.7551/mitpress/9780262035323.001.0001
- Zaretsky, E., and Lange, B. P. (2015). "No matter how hard we try: still no default plural marker in nonce nouns in modern high German," in *A lend of MaLT: Selected Contributions from the Methods and Linguistic Theories Symposium* (Bamberg), 153–178.
- Zaretsky, E., Lange, B. P., Euler, H. A., and Neumann, K. (2013). Acquisition of German pluralization rules in monolingual and multilingual children. *Stud. Second Lang. Learn. Teach.* 3, 551–580. doi: 10.14746/sslt.2013.3.4.6

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2021 Heitmeier, Chuang and Baayen. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.