COGNITIVE SCIENCE

A Multidisciplinary Journal



Cognitive Science 47 (2023) e13334

© 2023 The Authors. Cognitive Science published by Wiley Periodicals LLC on behalf of Cognitive Science Society (CSS).

ISSN: 1551-6709 online DOI: 10.1111/cogs.13334

Predicting Age of Acquisition for Children's Early Vocabulary in Five Languages Using Language Model Surprisal

Eva Portelance, a,# Yuguang Duan, b,# Michael C. Frank, Gary Lupyan

^aDepartment of Linguistics, McGill University ^bDepartment of Psychology, University of Wisconsin-Madison ^cDepartment of Psychology, Stanford University

Received 3 April 2023; received in revised form 26 July 2023; accepted 14 August 2023

Abstract

What makes a word easy to learn? Early-learned words are frequent and tend to name concrete referents. But words typically do not occur in isolation. Some words are predictable from their contexts; others are less so. Here, we investigate whether predictability relates to when children start producing different words (age of acquisition; AoA). We operationalized predictability in terms of a word's surprisal in child-directed speech, computed using n-gram and long-short-term-memory (LSTM) language models. Predictability derived from LSTMs was generally a better predictor than predictability derived from n-gram models. Across five languages, average surprisal was positively correlated with the AoA of predicates and function words but not nouns. Controlling for concreteness and word frequency, more predictable predicates and function words were learned earlier. Differences in predictability between languages were associated with cross-linguistic differences in AoA: the same word (when it was a predicate) was produced earlier in languages where the word was more predictable.

Keywords: Age of acquisition; Word learning; Language models; Early vocabulary; Child-directed speech; Word surprisal

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

[#]Co-first authors.

Correspondence should be sent to Eva Portelance, Department of Linguistics, McGill University, 1085 Dr. Penfield, Room 111, Montreal, Quebec H3A 1A7, Canada. E-mail: eva.portelance@mcgill.ca

1. Introduction

In the first 2 years of life, children's grammatical knowledge and lexicon grow in tandem (Bates et al., 1994; Brinchmann, Braeken, & Lyster, 2019; Frank, Braginsky, Marchman, & Yurovsky, 2021). In addition, the order in which children acquire their first words shows remarkable consistency (Braginsky, Yurovsky, Marchman, & Frank, 2019; Clark, 1993; Goodman, Dale, & Li, 2008; Tardif et al., 2008). For example, the words "ball," "car," and "nose" are, on average, learned almost a year earlier than words like "drawer," "green," and "animal." This general pattern holds across multiple languages (Braginsky et al., 2019). Modeling when words are acquired can, therefore, help us understand what factors drive language learning more generally.

One way to study these ordering effects is by attempting to predict a word's age of acquisition (AoA) from lexical properties such as its part of speech, frequency, length, and concreteness (Braginsky et al., 2019; Goodman et al., 2008; Kuperman, Stadthagen-Gonzalez, & Brysbaert, 2012). A common way is to estimate AoA is to use retrospective reports from adults of their own AoA for words (e.g., Kuperman et al., 2012). In contrast, here, AoA is determined using objective and timely parental reports of their children's linguistic productions. Though these methods are related, they are not the same and produce only moderately correlated AoA estimates for words (Łuniewska et al., 2016). Using child-based AoA estimates, studies find that nouns tend to be learned before verbs; more frequent words before less frequent words, shorter words before longer words, and words with concrete referents before those referring to more abstract entities. Controlling for part of speech, the frequency tends to be one of the strongest predictors of AoA (Braginsky et al., 2019; Fourtassi, Bian, & Frank, 2020; Goodman et al., 2008; Roy, Frank, DeCamp, Miller, & Roy, 2015). Here, we go beyond using these word-level predictors, by considering the linguistic context in which words appear in speech directed at children.

One such contextual predictor, previously examined by Braginsky et al. (2019), is the mean length of utterance (MLU) in which the word occurs. MLU can be considered as a proxy for syntactic complexity. If word learning is constrained by a child's ability to parse longer utterances (which are often more complex), we should find longer MLUs to be associated with later AoAs. This is precisely what Braginsky et al. (2019) found, though MLU was a significant predictor only for predicates and function words (a result that turns out to be relevant to the current investigation). Others have used contextual diversity—a measure of semantic co-occurrence—as a predictor of AoA (Amatuni & Bergelson, 2017; Fourtassi et al., 2020; Hills, Maouene, Riordan, & Smith, 2010; Roy et al., 2015; Stella, Beckage, & Brede, 2017). Words with higher contextual diversity—those appearing in many different semantic contexts in children's input-tend to be easier to learn and process by children (Hsiao & Nation, 2018; Pagán, Bird, Hsiao, & Nation, 2020; Rosa, Tapia, & Perea, 2017). Like with MLU, this predictor does not show the same effect across lexical categories. Hills et al. (2010) found that higher contextual diversity predicted earlier word acquisition, though this effect appeared to be strongest for predicates and function words when controlling for the effect of word frequency. Such a predictor can be considered a proxy for some semantic factors, approximating the overall semantic richness of children's learning environment. However, contextual diversity abstracts away from the semantic information encoded in specific sequential linguistic contexts—utterances or conversations—experienced by learners (a point we discuss in Section 3.1 of this paper) and, thus, does not directly measure syntactic complexity.

Here, we examine another contextual factor—the predictability of words across the linguistic contexts in which they appear. A word's predictability takes into account its linguistic context as a whole, including both syntactic and semantic information. It allows us to identify words that might be more or less difficult to learn as a function of the word sequences they tend to appear in.

There are multiple reasons for focusing on predictability. First, it is easy to compute, as we detail below. Second, it is psychologically real—we know that statistical factors such as transition probabilities play important roles in just about every aspect of early language learning (Saffran, 2020). Predictability (operationalized in terms of *surprisal*, see below) has also been shown to be a strong predictor of processing difficulty in adult psycholinguistic experiments (Demberg & Keller, 2008; Levy, 2008; Smith & Levy, 2013). More predictable words (i.e., those with lower surprisal) tend to be easier to process.¹

Thus, to help us understand the role of sequential predictability on the complexity and growth of children's vocabulary across languages, we propose to study the relation that exists between word predictability and AoA, treating AoA as a proxy for a word learning difficulty. We frame our experiments around the following two sets of concrete questions:

- Does a word's predictability across linguistic contexts help explain how difficult a
 word is to learn beyond previously known predictors? If so, how much sequential
 context matter? And, is predictability equally important for all words or not and why?
- 2. Second, are the effects of predictability on AoA observable across many languages and linguistic communities or are they isolated to only some? And, do differences in word predictability predict differences in AoA between languages?

To determine a word's predictability, we use computational language models that consider a word's previous sequential linguistic context. Specifically, word predictability can be measured as the average surprisal of a word—its negative log probability in a given context averaged across all contexts in which it appears—given a *language model* as our probability model.

Language models are sequential predictive models trained to generate linguistic output which are commonly used in natural language processing (NLP). They do so by learning probability distributions over strings of words in a corpus, where the length of these word strings may vary.² The additional information contained in these substrings in the form of preceding words is what allows us to consider a word's predictability given previous syntactic and semantic context. In this paper, we consider two types of language models: n-gram models and LSTM language models (Sundermeyer, Schlüter, & Ney, 2012). N-gram models are simple conditional probability models, while LSTM models are more sophisticated neural network models.

There has been growing interest in training or evaluating neural network language models on corpora that resemble the linguistic input of children more than the standard newspaper, *Wikipedia*, or web corpora used in NLP. This interest is motivated by two main goals: first,

determining how neural network models come to learn languages and, second, using these models as tools to better understand language learning in children. In one model-centered investigation, Huebner, Sulem, Cynthia, and Roth (2021) found that training language models on corpora of child-directed utterances can actually help these models perform better on grammatical knowledge evaluations than when trained on similar sized or larger corpora of traditional NLP data, suggesting that child-directed language may help grammar learning. Chang and Bergen (2022) have also proposed to use the average surprisal of words as a proxy for the AoA of words for language models in order to develop a new model evaluation task. They evaluated whether previously known predictors of the AoA of words in children—like frequency, concreteness, MLU, lexical category, and number of characters—also predicted the "AoA of words" in language models and found that there are clear differences between the orders in which children and language models acquire words.

Language models have also been proposed as tools to evaluate children's language development. For example, Sagae (2021) suggest using LSTMs trained on children's utterances to quantify children's syntactic development, finding that they perform as well or better than previous metrics.

All of this work—both model-based and child-focused—has been limited to English data. Here, we expand our analyses to cross-linguistic data, considering models trained in five different languages: English, German, French, Spanish, and Mandarin.

Our approach. In the rest of this paper, we will take the following approach: for each language, we fit a set of language models on a corpus of child-directed utterances and extract the average surprisal of words for which we have AoA estimates in children. We then compare regression models of children's AoA with average surprisal as one key predictor in concert with previous significant predictors, using cross-validation to estimate out-of-sample performance. We present two different modeling methods in two experiments. The first considers the effect of predictability on the AoA of words in each language individually, while the second considers its effect on all languages as a whole. We close this paper by considering how what we have learnt about the relation between word predictability, and AoA can inform our understanding about the role predictability plays on both children's vocabulary growth and the complexity of languages as a whole.

2. Data

We relied on two types of data: (1) *Corpora of child-directed utterances used to train the language models*.³ These were taken from the CHILDES database (MacWhinney, 2000); (2) *AoA estimates*. These were based on parental reports of children's language use, taken from the Wordbank repository (Frank, Braginsky, Yurovsky, & Marchman, 2016). We go into further detail in the following subsections about both these resources, and the data they contain. Importantly, in order for a language to be considered in this cross-linguistic study, there had to be sufficient data in this language in both of these resources. This criterion narrowed down the list of languages we could consider to English, German, French, Spanish, and Mandarin.

Table 1

Amount of data available in CHILDES and Wordbank by language

Language	Child-Directed Tokens (CHILDES)	Dialect	Vocabulary Reports (Wordbank)
English	25,659,263	American	7,955
		British	23,129
		Australian	1,497
German	5,663,294		1,181
French	3,183,037	European	863
		Quebecois	1,364
Spanish	2,267,707	European	1,005
		Mexican	1,934
Mandarin	2,369,896	Beijing	1,938
		Taiwanese	2,654

English was by far the most represented language in CHILDES, but there were still enough data for the other languages to be able to fit our language models (see Table 1 for the amounts of data available in each language for both CHILDES and Wordbank).

2.1. CHILDES and child-directed utterances

The CHILDES database (MacWhinney, 2000) is a repository of child language data, containing text transcripts of child-caregiver interactions as well as video and sound recordings of some of these interactions. The data come from many different studies conducted during the past 60 years, spanning multiple languages and countries. For the most part, the children in these studies range in age between 9 months old and five and a half years old.

For this paper, we only considered text transcription data and no other modalities. For each of the five languages considered (English, German, French, Spanish, Mandarin), we collected all of the available transcripts across all corpora available through the childes-db API (Sanchez et al., 2019) in July 2021. We then removed all utterances spoken by the target child, leaving only the utterances said to the child or around the child. These utterances can be considered as an estimate of the linguistic input the children in these transcripts have access to. We combined all of these child-directed utterances⁴ into a corpus used to fit the language models presented in the next section and to calculate the relative frequency of words for the regression models presented in our experiments. The result was five corpora of child-directed utterances, one for each language. Unlike the Wordbank database (Frank et al., 2016) presented in the next subsection, childes-db does not explicitly distinguish data based on dialectal varieties of each language, so we use the same aggregated data for each language across all varieties when fitting our models.

2.2. Wordbank and age of acquisition estimates

The Wordbank database (Frank et al., 2016) is a repository of parental reports about their children's vocabularies—essentially, a checklist of words where parents can check

off words their child produces or understands. Most of these reports are versions of the MacArthur–Bates Communicative Development Inventories (CDI) (Fenson et al., 1993). The database is a collection of reports originating from different studies that were conducted across the world. These studies and the vocabulary checklists they use are dialect specific, so for each of five languages we consider in this study, we collected data from all available dialects. We did not combine the data from different dialects into single languages as each dialect contains different word lists on their reports, leaving fewer words at their intersection.

Our predictive target is the age at which a word is acquired. Since not all children learn a given word at the same time, we instead follow prior work in quantifying AoA as the age at which 50% of children are reported to produce a word on the CDI (Goodman et al., 2008).⁵ There are a number of methods to estimate this 50% point from a group of binary responses for children of different ages. The simplest method is to determine the youngest age group at which the empirical proportion of children producing the word is > 50%, but this approach has several shortcomings. If words are very hard or very easy to learn, then it is possible that for the covered age range some words never reach the 50% point (e.g., beside) or have already surpassed the 50% point (e.g., Mommy) for even the youngest children. Such words would have to be discarded if we were to use this method. Another issue is that this approach is susceptible to bias AoA estimates towards ages for which more CDI instruments were available since the number of observations at each age is not equal (i.e., there may be more CDI instruments from 24-month-olds than with 20-month-olds in the dataset, but this density should not lead to more words being acquired at exactly 24 months). For these reasons, we used Bayesian generalized linear models predicting acquisition as a function of age to estimate the AoA for each word, following the method suggested in Frank et al. (2021).⁶

From the reports available in each language, we narrowed down the list of items used in our experiments to all single-word items on the forms that were classified as either nouns, predicates (verbs and adjectives), or function words (closed class words like pronouns, prepositions, question words, connectives, determiners). We excluded items that were multiword expressions (e.g. "all gone") or that were classified as being part of the "other" lexical category, which included animal sounds, onomatopoeia, and other non-word expressions. Words were also excluded if they were not in the five thousand most common words in each language in our corpora of child-directed utterances from CHILDES, as this was the vocabulary size used for the language models (described in the next section). Table 2 contains the exact number of items taken from Wordbank that we considered for each language as well as their breakdown by a lexical category.

3. Language models and predictability

To determine the predictability of words in the child-directed utterance corpora described above, we use language models as our probability models. Language models define probability distributions over subsequent words in a given context. Here, we consider the predictability of words solely based on linguistic contextual information. Specifically, we define the overall predictability of a word, w_i , as its average surprisal, or average negative log probability, across

Table 2
Number of items and their breakdown by lexical category in each language

Language	Number of Items	Nouns	Predicates	Function Words
English (American)	563	301 (53%)	165 (29%)	97 (17%)
English (British)	333	196 (59%)	100 (30%)	37 (11%)
English (Australian)	368	230 (63%)	138 (37%)	0 (0%)
German	361	189 (52%)	108 (30%)	64 (18%)
French (European)	394	222 (56%)	119 (30%)	53 (13%)
French (Quebecois)	468	248 (53%)	159 (34%)	61 (13%)
Spanish (European)	430	212 (49%)	121 (28%)	97 (23%)
Spanish (Mexican)	339	181 (53%)	90 (27%)	68 (20%)
Mandarin (Beijing)	500	269 (54%)	185 (37%)	46 (9%)
Mandarin (Taiwanese)	412	257 (62%)	119 (29%)	36(8%)

all its contexts of use, C (Eq. 1).

$$predictability(w_i) = \sum_{C: w_i \in C} -\log P(w_i \mid w_1, \dots, w_{i-1}) \times \frac{1}{|C|}, \tag{1}$$

where w_1, \ldots, w_{i-1} is a sequence of words of bounded length representing the preceding linguistic context.

There are many different types of language models. They vary in terms of the context sizes they consider, in how they represent words, and in how they come to calculate the overall probability of a word. As our definition in Eq. 1 suggests, we only consider language models that take into account preceding linguistic context (and not following context) and do so in an incremental order. Specifically, the experiments that follow will contain average surprisal values obtained from two types of language models: n-gram models, and LSTM models.

3.1. N-gram language models

N-gram models are basic language models that consider contexts of sequence length n, such that a bi-gram model keeps track of two word sequences and the contextual probability of a word is based on a single preceding word, and a tri-gram keeps track of three word sequences and contextual probability of a word is based on the two preceding words. Thus, given our formula in Eq. 1, we simply need to replace i by n to determine the predictability measure or average surprisal, of a word for a given n-gram model.

In an n-gram probability model, the probability of a word w_n in a given context, w_1, \ldots, w_{n-1} , or $P_{n-gram}(w_n \mid w_1, \ldots, w_{n-1})$, is simply its normalized count across all words that follow this context in the corpus. For example, if our probability model was a tri-gram model and we wanted to determine the probability of the word "bird" in the context "is that a bird," we would take $P_{tri-gram}(\text{bird} \mid \text{that a})$, which in practice is equal to count(that a bird)/count(that a), where count() returns the number of instances of an expression in a corpus. In this study, we used four different n-gram models: uni-grams, bigrams, tri-grams, and four-grams. Note that uni-gram models only track single-word contexts,

in other words they represent the normalized frequency counts of words so the average surprisal of a word for a uni-gram model is equivalent to its negative log frequency. Uni-gram surprisal is also in effect equivalent to contextual diversity. Though their definitions are different, when calculated on a large enough corpus, these two metrics converge, as we show in Appendix C in the Supporting Information, reaching a correlation score of -0.97 between uni-gram surprisal and log-transformed contextual diversity. For all intents and purposes, uni-gram surprisal, or the predictability of words irrespective of previous linguistic context, can also be interpreted as contextual diversity in the experiments which follow.⁷

One downside of n-gram models is that the context size is fixed across the whole probability model. This means that though some words may be better predicted from a single preceding word while others may be better predicted by two preceding words, we can only consider one of these context sizes at a time. LSTM models can help us get around this problem.

3.2. LSTM language models

Recurrent neural networks (RNNs) which, use long-short term memory gating unit layers (Hochreiter & Schmidhuber, 1997), commonly known as LSTMs, are neural networks that can be trained on sequential data, such as sentences, up to some bounded maximum length n. These models can be used for language modeling (Sundermeyer et al., 2012) and have become a staple baseline that continues to be used in NLP because of their useful analytical properties, even though more recent model architectures outperform them (Vaswani et al., 2017). Furthermore, regular RNNs have previously been proposed as cognitive models for language learning (Christiansen, Allen, & Seidenberg, 1998; Elman, 1990, 1993); however, these earlier models were computationally limited and could be used only with small schematic datasets; in contrast, LSTMs can be applied to larger datasets. Thus, LSTM language models lend themselves well to our analyses.

LSTM language models process utterances incrementally and make use of nested layers of hidden units to learn abstract representations that can predict sequential dependencies between words across a range of dependency lengths (Linzen, Dupoux, & Goldberg, 2016). LSTM neural units use a gating system that allows them to "forget" some of the previous states while "remembering" others, thus learning to prioritize some dependencies in a sequence over others at each state. So, unlike n-gram models, when determining the predictability, or average surprisal, (Eq. 1) of a word w_i for these models, the preceding contexts, w_1, \ldots, w_{i-1} in C, can vary in length, usually representing all of the previous words in the sentence. Further, the probability of a word in context can weigh the importance of preceding words differently based on the information encoded in the model's different layers. The added richness of these representations may lead to a better probability model overall.

For the experiments that follow, we use a two-layered LSTM language model (Fig. 1). The model has randomly initialized 100-dimensional word embeddings as its input layer, which are updated during learning. Hidden states encode information about the preceding context. At each time step, the current word embedding w_t and the hidden state from the previous timestep h_{t-1}^1 are passed through a transformation function, resulting in a new hidden state h_t^1 . This hidden state h_t^1 and the hidden state from the previous time step in the second hidden layer




Fig. 1. The LSTM model architecture incrementally processing the utterance "is that a bird." At each step, the model generates a conditional probability distribution over the vocabulary W representing the likelihood of being the next word. During training, the model updates its weights to maximize the probability of the actual next word. We can estimate the probability of words in context by retrieving their probability from the model output at each state.

 h_{t-1}^2 are then also passed through a transformation function, resulting in a new hidden state h_t^2 . This final hidden state is then resized through a linear layer to the size of our vocabulary before going through a softmax transformation to produce the output—a distribution over the whole vocabulary W representing a prediction about the upcoming word. We use a vocabulary size of 5,000, representing the most frequent words, because we found that including the 5,000 most common words usually resulted in the inclusion of almost all the words we had on our AoA word lists for all languages. Fig. 1 shows how the probability of a word given its preceding context $P_{LSTM}(w_i \mid w_1, \ldots, w_{i-1})$ can be extracted from the model.

The model's objective during training is to maximize the likelihood of the next word at each step—in other words, the model updates its parameters to minimize the surprisal of words in context. We performed cross-validation tests to find the hyperparameter settings for the LSTM language models that best minimized overall word surprisal. The hyperparameters tested were the vocabulary size (2,000, 4,000, or 5,000), the word embedding size (100, 150), the hidden state dimension size (100, 150), the batch size (128, 256, 512), and the number of epochs (up to 50). We found that the optimal parameter combination was a vocabulary size of 5,000, 100-dimensional word embeddings, a 100 hidden state size, a batch size of 256, and about 20 epochs of training.

We trained the models on all of the child-directed utterances for each language since the models were to be used as probability models and not predictive models. We were therefore not concerned with overfitting to the training data. Utterances were shuffled at each epoch of training. (For further details on the model implementation, see Appendix A in the Supporting Information.)

4. Experiment 1: The role of word predictability beyond log frequency

In this first experiment, we consider how word predictability beyond frequency predicts the AoA of words. Previous work (Braginsky et al., 2019; Fourtassi et al., 2020; Goodman et al., 2008; Kuperman et al., 2012; Roy et al., 2015) has found that log frequency—or, equivalently, uni-gram surprisal—is an important predictor of the AoA of words; here, we evaluate the explanatory power of larger context sizes by using their residualized effect beyond unigram surprisal as predictors. We compare models with different versions of average surprisal, obtained using different language models: bi-gram, tri-gram, four-gram, or LSTM average surprisals. We do so using leave-one-out (LOO) cross-validation.

4.1. Predictors

There are two main types of predictors considered in our models. First, we consider several methods for computing average surprisal using language models conditioned on different sizes and types of previous linguistic context. Second, we include other predictors that have been found to be informative in previous work: concreteness and lexical category (Braginsky et al., 2019; Goodman et al., 2008; Kuperman et al., 2012). All predictors are scaled by centering their mean at zero and dividing by their standard deviation so that their magnitudes can be compared.

uni-gram surprisal. is computed as the negative logarithm of frequency. Log frequency has been found to explain substantial variance in AoA in previous work.

Residualized n-gram surprisal. represents the residual variance left after fitting a linear model which predicts n-gram average surprisal as a function of uni-gram surprisal, n-gram average surprisal \sim uni-gram surprisal. n-gram average surprisal is the predictability of a word given all contexts of size n in which it appears in the nth position. We compare the average surprisal of bi-gram, tri-gram, and four-gram language models. If an item had multiple word forms associated with it on the parental report instrument (e.g., "inside/in"), we used the weighted mean average surprisal across all forms—in other words, if one form was overall more frequent, then it was weighted it accordingly.

Residualized LSTM surprisal. represents the residual variance left after fitting a linear model which predicts LSTM average surprisal as a function of uni-gram surprisal, LSTM average surprisal \sim uni-gram surprisal. LSTM average surprisal is calculated using the LSTM language models described above. We compute the average surprisal of each word across all of the child-directed utterances available in each language. We trained three LSTM language models in each language using different random seeds and then used the mean

average surprisal across these three runs as our measure of word predictability in each language. As with n-gram surprisal, we used the weighted mean average surprisal across all word forms.

Concreteness. is a rating score ranging between 1 and 5 for each word representing some measure along the abstract to concrete scale. These scores are taken from Brysbaert, Warriner, and Kuperman (2014). In order to obtain them for languages other than English, Wordbank data associate each item to a "unilemma," which is an equivalent English concept across all languages for that item in the database. The concreteness score for the equivalent English concept was then used for each word. This practice follows previous work (Braginsky et al., 2019).

Lexical category. is included via contrast coding of three lexical categories: nouns (common nouns), predicates (verbs and adjectives), and function words (closed-class words) following Bates et al. (1994). Word categories were derived from the categories on the CDI forms (e.g., verbs are listed as "action words"). Lexical category serves as interacting variable with all predictors.

4.2. Regression models

Our approach involves a nested model comparison between the model containing previously known predictors, the uni-gram model, and augmented models that additionally contain residualized average surprisal from either LSTM or n-gram language models. This comparison allows us to assess whether adding information from more dynamic context sizes using LSTMs or from fixed n-gram context sizes beyond log frequency helps predict AoA. We consider the uni-gram model as our base model:

1. The uni-gram model : AoA \sim lexical category *(uni-gram surprisal + concreteness)

We then compare it to the following augmented models:

- 2. The residualized n-gram model: AoA \sim lexical category *(uni-gram surprisal + n-gram residual average surprisal + concreteness), where n-grams are either bi-grams, tri-grams, or four-grams.
- 3. The residualized LSTM model: AoA \sim lexical category *(uni-gram surprisal + LSTM residual average surprisal + concreteness)

4.3. Results

We compare the uni-gram models in each language to their augmented versions which additionally contain residualized LSTM or n-gram average surprisals as predictors. We analyze the difference between the base uni-gram models and the augmented models using both LOO cross-validation and an analysis of variance (ANOVA)-nested model comparison across

Table 3
Model comparison results augmenting uni-gram surprisal model with residualized surprisals by language and model. Numbers indicate LOO cross-validation MAD (in months) with 95% confidence intervals

Language	Uni-gram	2-gram	3-gram	4-gram	LSTM
English (American)	2.0[1.86,2.14]	2.02[1.88,2.16]	2.01[1.87,2.15]	2.01[1.87,2.14]	2.01[1.87,2.15]*
English (British)	$2.19_{[2.0,2.39]}$	$2.22_{[2.02,2.41]}$	$2.21_{[2.01,2.4]}$	$2.21_{[2.01,2.4]}$	2.18[1.99,2.37]*
English (Australian)	1.93[1.76,2.1]	1.93[1.76,2.1]	1.94[1.77,2.11]	1.94[1.77,2.11]	1.93[1.76,2.11]
German	2.2 _[1.99,2.41]	$2.22_{[2.02,2.43]}$	$2.22_{[2.01,2.43]}$	$2.21_{[2,2.42]}$	$2.3_{[2.02,2.58]}$
French (European)	2.3 _[2.1,2.5]	$2.33_{[2.12,2.53]}$	$2.31_{[2.11,2.52]}$	$2.3_{[2.1,2.51]}$	$2.46_{[2.12,2.8]}$
French (Quebecois)	2.54 _[2.35,2.73]	$2.56_{[2.37,2.75]}$	2.54 _[2.31,2.73]	$2.55_{[2.36,2.74]}$	2.54 _[2.35,2.73]
Spanish (European)	2.48[2.29,2.67]	$2.49_{[2.3,2.69]}$	$2.5_{[2.31,2.69]}$	$2.51_{[2.31,2.7]}$	$2.49_{[2.3,2.68]}$
Spanish (Mexican)	$2.08_{[1.91,2.25]}$	$2.1_{[1.93,2.27]}$	$2.1_{[1.93,2.28]}$	$2.09_{[1.92,2.27]}$	2.06[1.89,2.23]*
Mandarin (Beijing)	1.89[1.75,2.02]	1.89[1.75,2.02]*	1.86[1.73,2.0] **	1.87[1.74,2.01]*	1.9[1.76,2.04]
Mandarin (Taiwanese)	2.98 _[2.75,3.21]	$3.0_{[2.77,3.23]}$	$3.0_{[2.78,3.23]}$	3.0 _[2.77,3.22] *	2.98 _[2.75,3.2] *

Note: * (p < 0.05) and ** (p < 0.01) indicate that the nested ANOVA is significant.

languages. We report the mean absolute deviation (MAD) across all LOO model fits in each language: the lower the MAD, the better the model fit.

4.3.1. Predictability overall

The overall results are available in Table 3, where we report MAD and 95% confidence intervals across LOO cross-validation folds, as well as *p* values from our ANOVA-nested model comparison. The models with the smallest MAD, or best fits, are given in bold.

The nested ANOVA results suggest that adding residualized LSTM surprisal as a predictor significantly increases model fit in four of the ten datasets. However, given the cross-validation results which show that there is very little difference in MAD values between our base models and augmented models, we may want to be cautious about these results and suggest instead that the overall effects are likely small.

The large majority of items across all languages are nouns. Average surprisal using more linguistic context may not be such a good predictor of the AoA of nouns, at least not as much as simple frequency (Portelance, Degen, & Frank, 2020). For this reason, the fact that we do not see a large difference between our base and augmented models here may be expected. If we consider the interaction between lexical category and residualized LSTM surprisal, we find that the interaction terms for predicates are significant (p < .05) in five of the languages (English (American), English (British), English (Australian), French (Quebecois), Spanish (Mexican)).

4.3.2. Predictability by lexical category

Adding residualized surprisal beyond uni-gram surprisal as a predictor of AoA does not substantially improve overall prediction, but it may make a difference in predicting the AoA of specific words. We next consider how effect sizes may differ by lexical category. Here, we do so by plotting the estimated coefficients by lexical category for the different surprisal predictors. 11




Fig. 2. Coefficient estimates for surprisal values by lexical category in each language.

The plots in Fig. 2 show the estimated coefficients for variables across LOO folds by lexical category. In each of these graphs, a point represents the estimated coefficient of a predictor for one fitted model run from the LOO cross-validation, so, for example, in *English (American)* there are 563 items and therefore 563 folds all of which have been plotted here.

For each language in Fig. 2, we see that residualized bi-gram, tri-gram, four-gram, and LSTM surprisals generally have little to no effect on nouns, but have a positive effect on function words and predicates in most languages. In other words, the harder function words and predicates are to predict in their linguistic contexts, the later they are acquired. The exception to this rule is Mandarin, where both function words and predicates show a negative effect

for all types of residual surprisal, meaning predicates with higher average surprisal, or less predictability, seem to be learned earlier and have a lower AoA.

Several explanations for the disparate results in Mandarin are possible. First, Mandarin has been known to pattern differently to other languages in other early word learning studies. For example, it has been said that Mandarin learners do not show the same "noun bias"—the observations that learners tend to initially learn to produce more nouns before increasing their productions of verbs-that English learners seem to have (Tardif, Gelman, & Xu, 1999), an observation that has been reproduced using the Wordbank parental report data (Frank et al., 2021; Yee, 2020). Instead, Mandarin learners have been found to produce more predicates early on during learning, even though both English- and Mandarin-speaking parents produce relatively more predicates than nouns (Tardif et al., 1999). Another second possible explanation for this difference may, however, lie in some of the Wordbank data itself. As Frank, Braginsky, Marchman, & YurovskyFrank et al. (2019, chap. 11) note, there seem to be some discrepancies with some of the Mandarin data in the repository, specifically forms collected for Mandarin (Beijing) from the Tardif, Fletcher, Liang, and Kaciroti (2009) study seem to show a much stronger predicate bias than other forms available for Mandarin (Beijing) or other languages. This data imbalance may also be contributing to this effect for predicates. On the other hand, the *Mandarin (Taiwanese)* data are not known to have this same issue, yet it shows a similarly negative effect for average surprisal with predicates albeit weaker than that of *Mandarin* (*Beijing*). 12

Finally, uni-gram surprisal has a positive effect on all lexical categories across almost all languages: words that are more predictable based on their frequency are generally learned earlier, here again the exception being Mandarin function words and predicates.

4.4. Interim conclusions

In this first experiment, we addressed our first set of research questions: Does a word's predictability across linguistic contexts help explain how difficult a word is to learn beyond previously known predictors? If so, how much sequential context matters? And, is predictability equally important for all words or not and why? Since average surprisal and log frequency are correlated (see Appendix B in the Supporting Information), we elected to remove any variance explained by frequency by using residualized surprisal values as our predictor. We found that including word predictability as a predictor of AoA helped explain some of the variance unaccounted for by previous predictors like log frequency. However, these effects were not universal, showing differences across lexical categories. Specifically, more predictable predicates and function words were found to generally be acquired earlier than their less predictable counterparts. As for how much previous sequential context mattered when determining word predictability, based on our plots in Fig. 2, residual LSTM surprisal had a notably greater effect size than all other n-gram surprisals, suggesting that dynamic context sizes which vary from word to word may be most useful in measuring predictability.¹³

5. Experiment 2: The role of word predictability across languages

In the previous section, we fit models for different languages using different word lists (Table 2). It was, therefore hard to determine whether the differences in effect sizes we

observed across languages were caused by variation within each individual language word list or by real distinctions in effect sizes. To remedy this issue, we needed to use the same word list for all languages. We achieved this by unifying words with the same concept across languages using their unilemmas and then taking their intersection.

First, we aggregated our data across languages to find the intersection of all unilemmas. There were a total of 89 unilemmas for which we have AoA estimates in all languages. These included 64 nouns and 25 predicates; no function words were left because one language, *English (Australian)*, did not have function words in its word lists. Although we were left with very few unilemmas overall, since we have 10 language groups, we still had 640 noun items and 250 predicate items. Our previous results suggested that the effects of residualized surprisal differed by lexical category, for this reason we split our data into nouns and predicates, testing each category separately. We then fitted mixed-effects models on each category with by-language random effects to compare effect sizes across languages.

5.1. Regression models

The models in this section differ from those used in the previous section as they include an additional random effect term: (1 + uni-gram surprisal + residual average surprisal | language). This term means that the models consider by-language differences in coefficient estimates for intercepts, the effects of uni-gram surprisal, and LSTM or n-gram residual average surprisal, taking those differences as a source of variance in the data. We did not include a separate random effect by language for concreteness ratings, since these were based on the unilemmas for words and were therefore identical in all languages.

- The residualized n-gram mixed-effects model: AoA ~ uni-gram surprisal + n-gram residual average surprisal + concreteness + (1 + uni-gram surprisal + n-gram residual average surprisal | language), where ngrams are either bi-grams, tri-grams, or four-grams.
- 2. The residualized LSTM mixed-effects model: AoA \sim uni-gram surprisal + LSTM residual average surprisal + concreteness + (1 + uni-gram surprisal + LSTM residual average surprisal | language)

5.2. Results

All predictors were scaled by language, centering predictors at zero and setting the standard deviation to one. 14

5.2.1. Predictability overall

As shown in Fig. 3, when we examine words occurring in all languages, we find that the effects of frequency (uni-gram surprisal) and concreteness are the same as before. Words with higher uni-gram surprisal are learned later, that is, more frequent words are learned earlier. Controlling for frequency and concreteness, none of the residual n-gram surprisal predictors showed significant effects overall, but LSTM surprisal did. The reason may be that n-grams only include information from a fixed length of context, while LSTM surprisal




Fig. 3. Mean coefficient estimates for surprisal values by lexical category for mixed-effects model using the intersection of items in all languages.




Fig. 4. Coefficient estimates for surprisal values for nouns and predicates in each language. Only words with data for all languages are included in this analysis.

contains information from more dynamic context sizes in the course of word learning. For predicates, higher LSTM average surprisal predicts later AoA, so less predictable predicates across linguistic contexts are harder to learn and, therefore, learned later. However, for nouns, residual LSTM surprisal instead shows a negative effect, meaning that less predictable nouns tend to be learned earlier.

A possible explanation for the case of nouns is that cross-situational learning for nouns across diverse social and visual contexts in theory may also lead to more diverse linguistic contexts and, therefore, higher surprisal (Hills, 2013; Roy et al., 2015). For predicates, however, children rely heavily on linguistic context to learn these words. Therefore, higher predictability across linguistic contexts may indicate that there are more linguistic cues helping children learn certain predicates earlier than others.

5.2.2. Surprisal and AoA across languages

Fig. 4 plots effect sizes of uni-gram surprisal and residual LSTM surprisal for each individual language by lexical category. ¹⁵ Comparing the predictive power of surprisal across languages, we find that uni-gram surprisal (word frequency) is equally and positively predictive of noun AoA. LSTM surprisal for nouns shows some variation, being somewhat

more negative—greater surprisal is associated with earlier AoA—in Mandarin and Spanish than in other tested languages. In the case of predicates, we see more variation. This result may simply be due to there being fewer items used to fit the predicate mixed-effects model, with only 250 words that were predicates compared to 640 words in the nouns mixed-effects model. The positive effect of both uni-gram surprisal and LSTM surprisal is strongest for Mandarin (Taiwanese and Beijing). This is an interesting result because our analyses from our first experiment using separate models for each language (Section 4) found that Mandarin was the only language to show a negative effect for surprisal when predicting the AoA of predicates. It is likely that this difference in effect polarity is due to items included in the mixed-effects models being a small subsample of those included in the previous experiment, suggesting that the specific words used to fit the models may be introducing a bias towards one or the other effect direction. We test and confirm this hypothesis in Appendix H in the Supporting Information.

What makes the items from Experiment 2 different from the rest in Experiment 1 specifically in Mandarin is unclear (the item lists are provided in the Appendix as well). However, we note that in Experiment 1 uni-gram surprisal also seemed to have negative estimates, albeit smaller, for function words and predicates in this language. Less frequent function words and predicates were supposedly acquired earlier contradicting the pattern seen in all other languages. Given that the discrepancy in average surprisal estimate polarity between Experiments 1 and 2 extends to uni-gram surprisal and not only LSTM average surprisal, the issue is unlikely to be with our predictability measure. Instead, it likely follows from issues with overall data quality in this language specifically. There are two places where data quality may be eroded. First, as we explain in the results section of Experiment 1, the AoA estimates for Mandarin may be questionable. Second, the corpus built from CHILDES data that we used to calculate both log frequency and probability model surprisal values could be the issue. It could simply be noisier and of poorer quality than the data in other languages, for example, because the Mandarin data were much sparser in CHILDES than many of the other languages (see Table 1).

Although children learning different languages follow broadly similar learning trajectories (e.g., learning frequent concrete nouns before less common abstract verbs), there *are* differences between the AoA of words that mean roughly the same thing across languages. Taking just the words available in all the languages, we find that for 45% of the word/language pairs, the mean difference in AoA is greater than 2 months. For 16% of the pairs, it is greater than 4 months. Are these differences predictable by differences in average surprisal for the same word across languages?

To answer this question, we first computed differences in AoA, uni-gram surprisal, and LSTM residual average surprisal (i.e., LSTM average surprisal controlling for uni-gram surprisal) for each lemma attested in each pair of languages (e.g., American English and Spanish). Uni-gram surprisal and LSTM residual average surprisal values were scaled within each language before computing the differences. We then used a mixed-effects model to predict cross-linguistic differences in AoA from differences in uni-gram surprisal and differences in LSTM residual surprisal. Since some pairs of languages belong to the same base language, for example, English (American) and English (Australian), we also added a binary variable




Fig. 5. Differences in frequency (uni-gram surprisal) and LSTM-based word surprisal are associated with differences in AoA. A positive coefficient indicates that if the same word has higher surprisal in language 1 than language 2, it is learned later in language 1 than language 2.

to the model indicating whether the two languages have the same base language. Moreover, we also included by-unilemma, by-the-first-base-language, and by-the-second-base-language random intercepts in the model, that is, (1 | unilemma) + (1 | base language 1) + (1 | base language 2) because those random effects could also be a source of variance in the data. Concreteness was not included as a predictor because we did not have separate concreteness estimates for each language.

The AoA differences mixed-effects model: AoA differences ~ differences in uni-gram surprisal + differences in LSTM residual average surprisal + whether the base languages are the same + (1 | unilemma) + (1 | base language 1) + (1 | base language 2)

From our final analysis (Fig. 5), we see that for both nouns and predicates, if a word occurs more frequently in language 1 than its equivalent¹⁶ in language 2, then its AoA is likely to be earlier in language 1. The same pattern is true with higher order LSTM surprisal for predicates, though for nouns we see an opposite effect, which aligns with our previous finding that LSTM surprisal has opposite effects on predicates and nouns. A greater difference in surprisal between two languages is associated with earlier AoAs for nouns and later AoAs for predicates.

5.3. Interim conclusions

In this second experiment, we tried to answer our second set of questions: Are the effects of predictability on AoA observable across many languages and linguistic communities or are they isolated to only some? And, do differences in word predictability predict differences in AoA between languages? We found that word predictability in context was a good predictor across language, showing little variation in effect size. The effects for nouns and predicates had polarity, however, such that less predictable nouns in linguistic contexts were learnt earlier, while less predictable predicates were learnt later. Additionally, in our secondary analysis we found that for predicates, there was a positive relationship between differences in surprisal and differences in AoA: predicate lemmas that have a greater surprisal in language 1 than language 2 tend to have a later AoA in language 1.

6. Discussion

Are words that are more predictable given their linguistic context (i.e., words with lower surprisal), easier to learn by young children? For predicates and function words, the answer seems to be yes. Although uni-gram surprisal (i.e., word log frequency) was overall a better predictor of AoA than higher order surprisal, surprisal derived from an LSTM neural network predicted AoA beyond uni-gram surprisal and the words' frequency. This effect was largely restricted to predicates and function words—words whose meaning is especially dependent on their context.

One route by which linguistic context is known to affect learning is via *syntactic bootstrapping*. Introduced by Brown (1957) and coined by Gleitman (1990), syntactic bootstrapping refers to the use of syntactic context to help determine meaning, especially for predicates. A wealth of evidence from individual experiments and computational models supports the ability of children to use linguistic information to make inferences about meaning Fisher, Gertner, Scott, and Yuan (2010), Fisher, Jin, and Scott (2020). Although our experiments do not directly test this idea, our results are certainly congruent with syntactic bootstrapping accounts.

A somewhat puzzling finding is that when predicting the AoA of nouns, greater LSTM surprisal was associated with *earlier* AoA—the opposite of what we observed for predicates and function words. One possible explanation for this pattern is that noun learning is more dependent on the perception of concrete referents, the salience of which is often obvious from the extra-linguistic context (Gillette, Gleitman, Gleitman, & Lederer, 1999). Frequent nouns are likely to appear in a broader set of contexts, which would *increase* their higher order surprisal while at the same time making it more semantically interconnected and improved learning (Hills, 2013; Hills et al., 2010; Roy et al., 2015). The results in Appendix C (in the Supporting Information), which show how frequency and predictability relate to contextual diversity, support this hypothesis. Frequency and contextual diversity are strongly correlated, thus, more frequent words also tend to be more contextually diverse. Furthermore, contextual diversity is negatively correlated with predictability; in other words, more predictable words also tend to have higher contextual diversity.¹⁷

We found the effects of surprisal to be mostly consistent across the tested languages. The one exception was Mandarin. In our first experiment using separate models for each language, Mandarin showed the opposite effects for surprisal on predicates and function words compared to other languages. However, in Experiment 2 using a single-mixed effects model over common unilemmas, Mandarin showed the largest effect of LSTM-based surprisal on AoA. We suggested that this discrepancy may follow from data quality issues in Mandarin specifically.

6.1. Broader theoretical implications

Beyond helping us understand the role played by sequential word predictability in determining which words are easier or harder to learn, this research can also offer insights into other theoretical questions.

First, a broad literature has explored the so-called "AoA effect" which refers to the finding that earlier learnt words are both easier to process for adults (Elsherif, Preece, & Catling, 2023) and more robustly accessible in aphasia patients (Brysbaert & Ellis, 2016). Most studies have looked at the effect specifically with referential nouns; a few exceptions have also considered it with predicates and none have looked at function words, that is, Colombo and Burani (2002), Morrison, Hirsh, and Duggan (2003), Bogka et al. (2003), Bonin, Boyer, Méot, Fayol, and Droit (2004), and Boulenger, Décoppet, Roy, Paulignan, and Nazir (2007). The first four mentioned studies found an AoA effect for both nouns and verbs, all however used an image naming task, and AoA effects are known to be strongly modulated by the imageability of words (Elsherif et al., 2023). The most recent study, Boulenger et al. (2007), used a lexical decision task which requires participants to read words and decided whether they are a real or nonce word. The AoA effect was only found with nouns and not verbs. The authors suggest that verbs and nouns may be learnt and processed differently resulting in different AoA effects depending on task demands. Our results confirm that nouns and predicates may not be learnt using the same cues and that their AoAs are dependent on different predictive factors, including predictability, or average surprisal, in the case of predicates. These learning differences could then result in different AoA effects, depending on task demands. Surprisal is known to predict sentence processing difficulty in incremental reading tasks (Hale, 2016). Reading task demands—as opposed to image naming task demands—may be different for words whose AoAs are more dependent on surprisal, like predicates, than those that are not, like nouns. 18

Second, our results shed light on a puzzle of linguistic complexity. Languages spoken by larger communities, those with many nonnative speakers, and those that have undergone substantial language contact (factors that are often, but not always positively correlated), tend to be morphologically simpler (Bentz, Dediu, Verkerk, & Jäger, 2018; Lupyan & Dale, 2010; Winters, Kirby, & Smith, 2015; Trudgill, 2011; Wray & Grace, 2007). On the other hand, languages spoken by smaller groups are more complex and less compressible (Koplenig, 2019; Lupyan, 2019; Lupyan & Dale, 2010). But why? One possibility a prediction made by Lupyan and Dale's (2010) linguistic niche hypothesis—is that there may be a trade-off between optimizing languages for use by larger/more diverse groups and optimizing languages for maximally efficient learning by young children. Greater compressibility is intimately linked to predictability: A more compressible language is one that contains utterances where one part is better predicted from other parts, making them informationally redundant. An intriguing possibility is that the reason smaller languages are more compressible—i.e., have overall higher predictability, or lower average surprisal—is that predictability is especially important to young children learning their first language (Lupyan & Dale, 2010, 2016). Our results show initial support for this idea: at least for predicates, a decrease in surprisal is associated with faster word learning suggesting that redundancy (lower surprisal) may provide more effective learning opportunities for young children. Further tests of this hypothesis would require comparison across a much wider range of languages, however.

7. Conclusion and limitations

We investigated the relationship between word predictability—formulated as average surprisal in linguistic contexts—and AoA—a proxy for how difficult learning a given word is for children. Word predictability seems to be especially important for the acquisition of predicates and function words, rather than nouns. Less predictable verbs and adjectives in linguistic contexts tend to be acquired later by children and are therefore harder to learn. We found this effect to be true in multiple languages, and that differences in surprisal predict differences in AoA. This finding is broadly in line with the prediction that decreasing surprisal (i.e., increasing redundancy) may—all else being equal—facilitate child language learning, particularly of more relational words.

Theories of language learning must explain not just words like "ball" and "dog" but also words whose meaning in context is almost entirely dependent on other words. Sequential models like the LSTM we used here may be a promising avenue to help explain the acquisition of these "hard words."

A major limitation of our analyses is that it is based on the data available in CHILDES. These corpora of child-directed utterances were—by necessity—assembled from many subcorpora from different children and studies and may not be the best representations of the regularity, idiosyncrasies, and contextual diversity found in the language targeted to a single child. Furthermore, our corpora contained utterances directed at children who were older in some cases than those surveyed for the AoA estimates. Ideally, these sentences would span the exact same developmental stages. A second limitation is that the words and AoA estimates we use are based on parents' reports of their children's language. Although validated and highly reliable, these data cannot capture the full richness of individual children's vocabularies and language use. Finally, our sample of language is restricted to languages with a large digital footprint. We hope future studies can begin to overcome these limitations.

Open Research Badges

This article has earned Open Data and Open Materials badges. Data and materials are available at https://github.com/evaportelance/multilingual-aoa-prediction

Notes

- 1 For an overview of empirical evidence supporting the validity of surprisal as a predictor of processing difficulty, see Hale (2016).
- 2 Context size can vary depending on the language model, ranging from a single previous token in n-gram models to a bounded ordered list of all previous or subsequent tokens in an ong-short-term-memory (LSTM) model.
- 3 All of the data, models, and experiment code presented in this paper are publicly available at https://github.com/evaportelance/multilingual-aoa-prediction.

- 4 We will use the term child directed somewhat loosely here such that it may also refer to utterances that were directed to other adults or children present, but that the target child could still hear.
- 5 We chose to use these AoA estimates from parental reports over those of Kuperman et al. (2012), which exist for a much larger vocabulary, because the latter are based on adult estimates of their own AoA, rather than timely reports of children's AoA. Additionally, Kuperman et al. (2012) report word comprehension AoA estimates, while the estimates we use here are word production AoA estimates.
- 6 For a more detailed description of this method, see Appendix E of Frank et al. (2021).
- 7 Though contextual diversity is strongly correlated with uni-gram surprisal, importantly, it is not related to higher order surprisal metrics which take into account contextualized sequential predictability, like LSTM average surprisal, as shown in Appendix C in the Supporting Information.
- 8 We also tried considering different forms as separate items; these results are available in Appendix D in the Supporting Information.
- 9 We used contrast coding for our lexical categories, such that interaction terms between a lexical category and residualized surprisal can be interpreted as main effects, indicating a difference from the overall mean.
- 10 Since lexical categories differ significantly in their concreteness ratings, we also explored whether the by-lexical-category differences in surprisal effects seen in this subsection can be better explained by an interaction between residualized surprisal and concreteness in Appendix F in the Supporting Information.
- 11 In Appendix E in the Supporting Information, we also fit models to items in each lexical category individually and consider their MAD; Table S5 shows that nouns are generally best predicted by simply using the uni-gram base model, while predicates and function words often benefit from the augmented versions which include some form of higher order residualized surprisal, LSTM and tri-gram residualized surprisal generally doing best.
- 12 It is also possible that these differences are simply due to strength and nature of the resources we had at our disposal for Mandarin. For example, it is possible that different corpora in CHILDES data used different character systems or word segmentation norms, since there exists different standards for this language, which in turn could have led to noisier data.
- 13 The importance of context size has previously been studied in the case of contextual diversity (Hills et al., 2010), but as we note in Appendix C in the Supporting Information, contextual diversity and predictability measure different properties and may, therefore, require different types of contexts.
- 14 Doing so for each language individually allows us to compare the effects across languages without worrying that any differences in variation that are due to our predictive n-gram or LSTM models having different surprisal distributions in different languages.
- 15 The effect sizes by language for residualized n-gram surprisal values are available in Appendix G in the Supporting Information.
- 16 Equivalent words refer to word with the same unilemma in Wordbank.

- 17 We note that when controlling for the effect of log frequency, predictability and contextual diversity are no longer as correlated. Additionally, nouns pattern differently from predicates and function words in terms of the relationship between predictability and frequency.
- 18 We note the caveat that Boulenger et al. (2007) used subjective AoA measures based on retrospective adult ratings; Our research used objective ratings based of parental reports of children's productions. As mentioned in the introduction, these measures are related but not always strongly correlated.

References

- Amatuni, A., & Bergelson, E. (2017). Semantic networks generated from early linguistic input. bioRxiv. https://doi.org/10.1101/157701
- Bates, E., Marchman, V., Thal, D., Fenson, L., Dale, P., Reznick, J. S., Reilly, J., & Hartung, J. (1994). Developmental and stylistic variation in the composition of early vocabulary. *Journal of Child Language*, 21(1), 85–123.
- Bentz, C., Dediu, D., Verkerk, A., & Jäger, G. (2018). The evolution of language families is shaped by the environment beyond neutral drift. Nature Human Behaviour, 2(11), 816–821.
- Bogka, N., Masterson, J., Druks, J., Fragkioudaki, M., Chatziprokopiou, E.-S., & Economou, K. (2003). Object and action picture naming in English and Greek. European Journal of Cognitive Psychology, 15(3), 371–403.
- Bonin, P., Boyer, B., Méot, A., Fayol, M., & Droit, S. (2004). Psycholinguistic norms for action photographs in French and their relationships with spoken and written latencies. *Behavior Research Methods, Instruments*, & Computers, 36, 127–139.
- Boulenger, V., Décoppet, N., Roy, A. C., Paulignan, Y., & Nazir, T. A. (2007). Differential effects of age-of-acquisition for concrete nouns and action verbs: Evidence for partly distinct representations? *Cognition*, 103(1), 131–146.
- Braginsky, M., Yurovsky, D., Marchman, V. A., & Frank, M. C. (2019). Consistency and variability in children's word learning across languages. *Open Mind*, 3, 52–67.
- Brinchmann, E. I., Braeken, J., & Lyster, S.-A. H. (2019). Is there a direct relation between the development of vocabulary and grammar? *Developmental Science*, 22(1), e12709.
- Brown, R. W. (1957). Linguistic determinism and the part of speech. The Journal of Abnormal and Social Psychology, 55, 1–5.
- Brysbaert, M., & Ellis, A. W. (2016). Aphasia and age of acquisition: Are early-learned words more resilient? Aphasiology, 30(11), 1240–1263.
- Brysbaert, M., Warriner, A. B., & Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas. Behavior Research Methods, 46(3), 904–911.
- Chang, T. A., & Bergen, B. K. (2022). Word acquisition in neural language models. Transactions of the Association of Computational Linguistics, 10, 1–16. https://arxiv.org/abs/2110.02406
- Christiansen, M. H., Allen, J., & Seidenberg, M. S. (1998). Learning to segment speech using multiple cues: A connectionist model. *Language and Cognitive Processes*, 13(2-3), 221–268.
- Clark, E. V. (1993). Early lexical development. The lexicon in acquisition (pp. 21–42). Cambridge University Press.
- Colombo, L., & Burani, C. (2002). The influence of age of acquisition, root frequency, and context availability in processing nouns and verbs. *Brain and Language*, 81(1-3), 398–411.
- Demberg, V., & Keller, F. (2008). Data from eye-tracking corpora as evidence for theories of syntactic processing complexity. Cognition, 109(2), 193–210.
- Elman, J. L. (1990). Finding structure in time. Cognitive Science, 14(2), 179-211.
- Elman, J. L. (1993). Learning and development in neural networks: The importance of starting small. Cognition, 48(1), 71–99.

- Elsherif, M., Preece, E., & Catling, J. (2023). Age-of-acquisition effects: A literature review. Journal of Experimental Psychology: Learning, Memory, and Cognition, 49(5), 812–847.
- Fenson, L., Dale, P., Reznick, J. S., Thal, D., Bates, E., Hartung, J., Pethick, S., & Reilly, J. (1993). MacArthur communicative inventories: User's guide and technical manual. Baltimore, MD: Brookes Publishing.
- Fisher, C., Gertner, Y., Scott, R. M., & Yuan, S. (2010). Syntactic bootstrapping. Wiley Interdisciplinary Reviews: Cognitive Science, 1(2), 143–149.
- Fisher, C., Jin, K.-s., & Scott, R. M. (2020). The developmental origins of syntactic bootstrapping. Topics in Cognitive Science, 12(1), 48–77.
- Fourtassi, A., Bian, Y., & Frank, M. C. (2020). The growth of children's semantic and phonological networks: Insight from 10 languages. *Cognitive Science*, 44(7), e12847.
- Frank, M. C., Braginsky, M., Marchman, V., & Yurovsky, D. (2021). Variability and consistency in early language learning: The Wordbank Project. Cambridge, MA: MIT Press.
- Frank, M. C., Braginsky, M., Yurovsky, D., & Marchman, V. A. (2016). Wordbank: An open repository for developmental vocabulary data. *Journal of Child Language*, 44(3), 677–694.
- Gillette, J., Gleitman, H., Gleitman, L., & Lederer, A. (1999). Human simulations of vocabulary learning. Cognition, 73(2), 135–176.
- Gleitman, L. (1990). The structural sources of verb meanings. Language Acquisition, 1(1), 3-55.
- Goodman, J. C., Dale, P. S., & Li, P. (2008). Does frequency count? Parental input and the acquisition of vocabulary. *Journal of Child Language*, 35(3), 515–531.
- Hale, J. (2016). Information-theoretical complexity metrics. Language and Linguistics Compass, 10(9), 397–412.
 Hills, T. T. (2013). The company that words keep: Comparing the statistical structure of child-versus adult-directed language. Journal of Child Language, 40(3), 586–604.
- Hills, T. T., Maouene, J., Riordan, B., & Smith, L. B. (2010). The associative structure of language: Contextual diversity in early word learning. *Journal of Memory and Language*, 63(3), 259–273.
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735-1780.
- Hsiao, Y., & Nation, K. (2018). Semantic diversity, frequency and the development of lexical quality in children's word reading. *Journal of Memory and Language*, 103, 114–126. https://www.sciencedirect.com/science/article/ pii/S0749596X18300718https://doi.org/10.1016/j.jml.2018.08.005
- Huebner, P. A., Sulem, E., Cynthia, F., & Roth, D. (2021). BabyBERTa: Learning more grammar with small-scale child-directed language. In *Proceedings of the 25th conference on computational natural language learning* (pp. 624–646). Stroudsburg, PA: Association for Computational Linguistics.
- Koplenig, A. (2019). Language structure is influenced by the number of speakers but seemingly not by the proportion of non-native speakers. Royal Society Open Science, 6(2), 181274. https://royalsocietypublishing.org/doi/full/10.1098/rsos.181274
- Kuperman, V., Stadthagen-Gonzalez, H., & Brysbaert, M. (2012). Age-of-acquisition ratings for 30,000 English words. Behavior Research Methods, 44(4), 978–990.
- Levy, R. (2008). Expectation-based syntactic comprehension. Cognition, 106(3), 1126-1177.
- Linzen, T., Dupoux, E., & Goldberg, Y. (2016). Assessing the ability of LSTMs to learn syntax-sensitive dependencies. Transactions of the Association for Computational Linguistics, 4, 521–535.
- Łuniewska, M., Haman, E., Armon-Lotem, S., Etenkowski, B., Southwood, F., Andelković, D., ... Ünal Logacev, O. (2016). Ratings of age of acquisition of 299 words across 25 languages: Is there a cross-linguistic order of words? *Behavior Research Methods*, 48, 1154–1177.
- Lupyan, G. (2019). Larger languages have higher entropy. https://cle.ppls.ed.ac.uk/index.php/ielc2019/
- Lupyan, G., & Dale, R. (2010). Language structure is partly determined by social structure. PLoS ONE, 5(1), e8559. http://doi.org/10.1371/journal.pone.0008559
- Lupyan, G., & Dale, R. (2016). Why are there different languages? The role of adaptation in linguistic diversity. Trends in Cognitive Sciences, 20(9), 649–660.
- MacWhinney, B. (2000). *The CHILDES Project: Tools for analyzing talk. Third Edition*. Mahwah, NJ: Erlbaum. Morrison, C. M., Hirsh, K. W., & Duggan, G. B. (2003). Age of acquisition, ageing, and verb production: Normative and experimental data. *The Quarterly Journal of Experimental Psychology Section A*, 56(4), 1–26.

- Pagán, A., Bird, M., Hsiao, Y., & Nation, K. (2020). Both semantic diversity and frequency influence children's sentence reading. Scientific Studies of Reading, 24(4), 356–364. https://doi.org/10.1080/10888438.2019. 1670664
- Portelance, E., Degen, J., & Frank, M. C. (2020). Predicting age of acquisition in early word learning using recurrent neural networks. In *Proceedings for the 42nd annual meeting of the Cognitive Science Society* (pp. 1057–1063). Virtual: Cognitive Science Society.
- Rosa, E., Tapia, J. L., & Perea, M. (2017). Contextual diversity facilitates learning new words in the classroom. PLoS ONE, 12(6), e0179004.
- Roy, B. C., Frank, M. C., DeCamp, P., Miller, M., & Roy, D. (2015). Predicting the birth of a spoken word. Proceedings of the National Academy of Sciences, 112(41), 12663–12668.
- Saffran, J. R. (2020). Statistical language learning in infancy. Child Development Perspectives, 14(1), 49–54. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8078161/https://doi.org/10.1111/cdep.12355
- Sagae, K. (2021). Tracking child language development with neural network language models. Frontiers in Psychology, 12, 674402.
- Sanchez, A., Meylan, S., Braginsky, M., MacDonald, K., Yurovsky, D., & Frank, M. C. (2019). childes-db: a flexible and reproducible interface to the Child Language Data Exchange System. *Behavior Research Methods*, 51, 1928–1941.
- Smith, N. J., & Levy, R. (2013). The effect of word predictability on reading time is logarithmic. Cognition, 128, 302–319.
- Stella, M., Beckage, N. M., & Brede, M. (2017). Multiplex lexical networks reveal patterns in early word acquisition in children. Scientific Reports, 7(1), 1–10.
- Sundermeyer, M., Schlüter, R., & Ney, H. (2012). LSTM neural networks for language modeling. Paper presented at the 13th Annual Conference of the International Speech Communication Association.
- Tardif, T., Fletcher, P., Liang, W., & Kaciroti, N. (2009). Early vocabulary development in Mandarin (Putonghua) and Cantonese. *Journal of Child Language*, 36, 1115–1144.
- Tardif, T., Fletcher, P., Liang, W., Zhang, Z., Kaciroti, N., & Marchman, V. A. (2008). Baby's first 10 words. Developmental Psychology, 44(4), 929.
- Tardif, T., Gelman, S. A., & Xu, F. (1999). Putting the "noun bias" in context: A comparison of English and Mandarin. Child Development, 70, 620–635.
- Trudgill, P. (2011). Sociolinguistic typology: Social Determinants of linguistic complexity. New York, NY: Oxford University Press.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017).
 Attention is all you need. In Advances in neural information processing systems (pp. 5998–6008). Cambridge, MA: MIT Press.
- Winters, J., Kirby, S., & Smith, K. (2015). Languages adapt to their contextual niche. Language and Cognition, 7(3), 415–449. https://www.cambridge.org/core/journals/language-and-cognition/article/abs/languages-adapt-to-their-contextual-niche/83E9F516875C340E0A9263B4A7C38F43
- Wray, A., & Grace, G. (2007). The consequences of talking to strangers: Evolutionary corollaries of socio-cultural influences on linguistic form. *Lingua*, 117(3), 543–578.
- Yee, S. (2020). Is noun bias universal? Evidence from Chinese and Korean compared with French and English. In Studies in the linguistic sciences: Illinois Working Papers (pp. 32–44). Urbana-Champaign, IL: Linguistics Department, University of Illinois at Urbana-Champaign

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Table S1: Pearson correlation between LSTM, bigram, tri-gram, four-gram average surprisal and uni-gram surprisal in each language

Table S2: Pearson correlation between frequency, uni-gram surprisal, LSTM average surprisal, contextual diversity (CD), and log transformed contextual diversity (log CD) across all languages

Table S3: Pearson correlation between uni-gram surprisal, residualised LSTM average surprisal, and residualised log transformed contextual diversity by lexical category across all languages

Table S4: Approach 1 model comparison results augmenting uni-gram surprisal model with residualised surprisal by language

Figure S1: Coefficient estimates by lexical category in each language using first approach with all words as separate items

Table S5: Approach 1 model comparison results augmenting uni-gram surprisal model with residualised surprisal by language and lexical category

Figure S2: Concreteness distribution for different lexical categories

Table S6: Approach 1 model comparison results using interaction with concreteness by language

Figure S3: By-language random effects in Residualised n-gram mixed-effects models for second approach

Figure S4: Coefficient estimates for surprisal values using the model from Experiment 1 on the dataset from Experiment 2 for Mandarin