

Frequency, acceptability, and selection: A case study of clause-embedding

Aaron Steven White
University of Rochester
aaron.white@rochester.edu

Kyle Rawlins
Johns Hopkins University
kgr@jhu.edu

Abstract We investigate the relationship between the frequency with which verbs are found in particular subcategorization frames and the acceptability of those verbs in those frames, focusing in particular on subordinate clause-taking verbs, such as *think*, *want*, and *tell*. We show that verbs' subcategorization frame frequency distributions are poor predictors of their acceptability in those frames—explaining, at best, less than $\frac{1}{3}$ of the total information about acceptability across the lexicon—and, further, that common matrix factorization techniques used to model the acquisition of verbs' acceptability in subcategorization frames fare only marginally better.

Keywords: subcategorization, selection, acceptability, frequency

1 Introduction

Knowing a language involves at least being able to judge the well-formedness of strings relative to that language. Accounts differ as to whether well-formedness is itself a single concept or multiple interrelated concepts—e.g. *grammaticality* v. *acceptability* (Chomsky 1965 *et seq*)—whether those concepts are discrete or continuous, and what exactly is required to collect the supporting judgments (Bard et al. 1996; Keller 2000; Sorace & Keller 2005; Sprouse 2007; 2011; Featherston 2005; 2007; Gibson & Fedorenko 2010; 2013; Sprouse & Almeida 2013; Sprouse et al. 2013; Schütze & Sprouse 2014; Lau et al. 2017; Sprouse et al. 2018 among many others). But few doubt that some specification of well-formedness is a crucial component of a theory of linguistic knowledge.

Well-formedness is something that must be learned—at least in part—from linguistic experience. A major question is how direct the relationship between well-formedness and linguistic experience is. Approaches

that directly address this question have tended to focus on complex syntactic effects—e.g. effects that might arise from constraints on syntactic movement, such as island effects (see Hofmeister & Sag 2010; Gibson & Fedorenko 2010; Sprouse et al. 2012; Kush et al. 2018 among others). The abstractness of these constraints makes them useful test cases, since they are prime candidates for knowledge that might arise from inductive biases innate to language learners—i.e. not directly from statistical properties of the input (though see Pearl & Sprouse 2013). However, complex syntactic phenomena are far from the only factor contributing to well-formedness.

Effects on well-formedness that are at least partially a product of lexical knowledge have garnered much less focused attention within theoretical linguistics—possibly, because they seem more directly dependent on statistical properties of the input (though see Bresnan 2007; Bresnan et al. 2007; White 2015; White et al. 2017a).¹ That is, lexical, as opposed to grammatical, constraints on well-formedness might have a direct connection to co-occurrence statistics in language learners' input, and therefore be learnable using relatively simple strategies—e.g. tracking co-occurrence frequencies, which has long been believed to be well within the capabilities of even young children (Saffran et al. 1996b; a; Aslin et al. 1998; Maye et al. 2002). As such, lexical constraints on well-formedness could in principle make a better case for a direct connection between linguistic experience and grammatical knowledge.

Against this background, this paper makes two main contributions: we introduce and validate an empirical method for collecting lexical acceptability patterns at a very large scale, and we use a dataset collected via this method to computationally investigate the connection between linguistic experience (in the form of corpus data) and linguist knowledge (in the form of acceptability patterns). The main finding from this second part is that knowledge of acceptability does *not* provide support for a direct connection between frequency and acceptability, but rather supports the idea that a language learner needs to employ substantial abstraction in order to be able to achieve adult/human-like knowledge of acceptability.

¹ See also much work in the sentence processing literature (Trueswell et al. 1993; Spivey-Knowlton & Sedivy 1995; Garnsey et al. 1997; McRae et al. 1998; Altmann & Kamide 1999; Hale 2001; Levy 2008; Wells et al. 2009; Fine & Jaeger 2013; Linzen & Jaeger 2016 among others) and the language acquisition literature (Landau & Gleitman 1985; Pinker 1984; 1989; Gleitman 1990; Naigles 1990; 1996; Naigles et al. 1993; Fisher et al. 1991; Fisher 1994; Fisher et al. 1994; 2010; Lederer et al. 1995; Gillette et al. 1999; Snedeker & Gleitman 2004; Lidz et al. 2004; Gleitman et al. 2005; Papafragou et al. 2007 among others).

In particular, we investigate effects on well-formedness having to do with verbs' c(ategory)-selection behavior—i.e. what kinds of syntactic structures verbs are acceptable in—showing that this assumption is not justified. We focus specifically on verbs that take subordinate clauses—henceforth, *clause-embedding verbs*—such as *think*, *want*, and *tell*.

- (1) a. Jo {thought, told Mo} that Bo left.
b. Jo {wanted, told} Bo to leave.

Clause-embedding verbs are a useful test case because (i) subordinate clauses can have a wide variety of syntactic structures; (ii) many verbs can take a large subset of these clause types; and (iii) there is high variability in which subset of clause types verbs can take. For instance, *remember* can combine with a wide variety of differently structured clauses, as in (2a)-(2e) as well as noun phrases (2f).

- (2) a. Jo remembered that Bo left.
b. Jo remembered Bo to have left.
c. Jo remembered Bo leaving.
d. Jo remembered to leave.
e. Jo remembered leaving.
f. Jo remembered Bo.

Starting from the assumption that the relevant lexical knowledge must be acquired, in some way or other, from the co-occurrence statistics, our investigation specifically aims to measure how direct a relationship there is between verbs' selectional patterns—as measured by *acceptability judgments*—and said statistics. We take this question to be one about the lexicon as a whole, and therefore test it with a *lexicon-scale* dataset of acceptability judgments for clause-embedding verbs. A lexicon-scale investigation is both possible and crucial: there are many clause-embedding verbs in English—by some counts, at least 1,000 (White & Rawlins 2016)—so testing grammatical knowledge involves testing knowledge across this large set.

Our investigation contrasts with prior work, which investigates only a small set of key verbs (Fisher et al. 1991; Lederer et al. 1995; Bresnan et al. 2007; White et al. 2018; though see Kann et al. 2019). The relatively small size of previous investigations is likely a product of the fact that scaling standard methodologies to a larger set of verbs is infeasible without introducing unwanted biases or insurmountable workload. To overcome this obstacle, we propose a novel method—the *bleaching method*—for automatically scal-

ing standard methods while avoiding the introduction of such biases into the item set.

In Section 3, we report on an experiment validating the bleaching method against a more standard acceptability judgment collection method, focusing on a small set of clause-embedding verbs. In Section 4, we report on an experiment in which we deploy the bleaching method on 1,000 clause-embedding verbs in 50 syntactic frames to create the MegaAcceptability dataset, which was first reported on in White & Rawlins 2016 and is publicly available at megaattitude.io under the auspices of the MegaAttitude Project. In Section 5, we use the MegaAcceptability dataset in conjunction with a very large dataset of verbs' subcategorization frequencies to show that the relationship between acceptability and frequency, when considering this entire sublexicon, is surprisingly weak. This throws into question the assumption that c-selectional behavior can be directly read off frequency distributions.

Nonetheless, adult native speakers are able to judge the acceptability of the items that make up this task. If they are unable to learn the information to do this directly from the frequency—even assuming access to the data that is both ideal and uniform—they must have *some* way to get it. This suggests that some abstraction of the frequency distributions in the input is necessary. In Section 6, we consider a variety of such abstractions, showing (i) that common, shallow factorization methods yield miniscule improvements in the prediction of acceptability over more direct models; but (ii) that methods involving multiple layers of abstraction can predict acceptability quite well. The models we present here are proof-of-concept models rather than attempts at actual learning models, and thus, our results do not fully answer the question of exactly how a learner will gather this information; but we take them to strongly confirm the necessity of substantial abstraction in grammatical theory relative to input frequency. In Section 7, we conclude with remarks on what these findings imply for the acquisition of distributional knowledge about lexical items.

2 Background

We begin with a discussion of previous work relating frequency and acceptability (Section 2.1), acceptability and selection (Section 2.2), and frequency and selection (Section 2.3) and then discuss two hypotheses about the joint relationship among the three (Section 2.4).

2.1 *Frequency and acceptability*

In a substantial body of work, Clark, Lappin, and colleagues argue that knowledge of grammaticality, by way of acceptability, can be modeled, to a large extent, with probabilistic models that involve a direct link between probability and acceptability and directly recognize gradience in acceptability (Clark & Lappin 2011; Clark et al. 2013a; b; Lau et al. 2017; see also Bresnan et al. 2007; Bresnan 2007). These models have two components: (i) a way of estimating the probability of a sentence of a language; and (ii) a way of translating probabilities to acceptabilities. For the most part, the latter is quite straightforward: some variant of log probabilities, normalized for sentence length and unigram frequency effects. The former is where much of the action is.

This body of work considers a range of possibilities for how to model probability, ranging from simple n -gram frequency models to neural language models in Lau et al. 2017 (see also Warstadt et al. 2018). Lau et al. consider two kinds of data: (i) sentences sampled from the BNC corpus, fed through google translate to a range of languages, and then translated back to English (with the goal of obtaining a spread of acceptability), and a sample of sentences from Adger 2003 by way of Sprouse & Almeida's (2013) dataset. They then run a variety of acceptability judgment studies on Mechanical Turk, providing the core data for their experiments. Across the board, the models in this paper show what the authors describe as an "encouraging degree of accuracy" in predicting human judgments of acceptability—well exceeding baseline models but generally falling well short of human performance.

The authors take this to be a signal that probabilistic models are the current best way to incorporate gradience in human judgments about acceptability into a theory of grammatical knowledge. This conclusion is controversial (Sprouse et al. 2018), but we take it as a starting point that current probabilistic models can at least do reasonably well at capturing many—though probably not all—facets of acceptability.

Clark, Lappin, and colleagues' work therefore sets the stage for the central question that we are addressing here: what is the relationship between frequency of use, probability, and acceptability? Though the model generating the probabilities itself may be complicated, on Lau et al.'s view, the relationship between the two is *direct*. That is, there is a simple transformation that, as long as some basic normalization is taken care of, more or less directly predicts acceptability.

Existing work has focused on testing hypotheses like this on data that is extremely varied—e.g. random samples of naturalistic corpus data or broad datasets of grammaticality judgments from linguistic theory (Sprouse et al. 2018). For this reason, it is an independently challenging and interesting problem to estimate the probability of sentences across such data, and as Lau et al. (2017) demonstrate, in at least some domains more sophisticated models from natural language processing (NLP) will lead to better predictions of acceptability—e.g. models which involve more complex relationships between the frequencies they are trained on and the probabilities they output. But this approach to the underlying data leads to an additional problem: conclusions about the probabilistic nature of the grammar rest on the degree to which the $[0, 1]$ interval values that these models are producing are in fact good estimates of the probabilities of particular sentences, also making it rather challenging to know what the driving factor for sentence-level probabilities is.

In the present work, we take a different approach. Rather than a broad and diverse data sample, we pick a single phenomenon where we can obtain acceptability data exhaustively, and estimate probabilities from corpora in a relatively transparent way, producing a direct idealization of linguistic experience.² The particular data we investigate is selectional patterns for clause-embedding predicates, where the main point of variation between items is just the verb and its selectional frame. This also allows us to begin to localize the *kinds* of grammatical knowledge that are likely to be involved in variation in acceptability.

2.2 Acceptability and selection

The selectional behavior of verbs in general—and clause-embedding verbs in particular—is a classic topic in linguistic theory (Chomsky 1965; Gruber 1965; Fillmore 1970; Jackendoff 1972; Chomsky 1973; Grimshaw 1979; 1990; Pesetsky 1982; 1991 among many others). There are broadly accepted to be two crucial descriptive factors leading to variance in whether a verb is acceptable in a particular sentence (other things being equal): (i) semantic constraints imposed by the argument structure of the verb (*s(ematic)-selection*); and (ii) (morpho-)syntactic constraints imposed by a verb on its complements (*c(omplement)-selection*; Grimshaw 1979).

² This is analogous to recent approaches within the NLP literature that aim to probe what linguistic knowledge different models learn from corpus data using an array of focused datasets (White et al. 2017b; Poliak et al. 2018; Conneau et al. 2018; Wang et al. 2018; Kann et al. 2019).

It is a matter of substantial debate and discussion whether some or all of these constraints might be predictable from each other or other factors—e.g. event structure—something we will not try to settle here (see [Levin & Rappaport Hovav 2005](#) and references therein). But this domain is useful for our purposes because (i) it is empirically rich—even if we focus just on selection of clauses; and (ii) it involves both a large set of patterns and a large set of lexical idiosyncracies.

For example, Grimshaw introduced the classic *wonder* v. *think* comparison (her Ex. 1):

- (3) a. John wondered who Bill saw.
- b. *John wondered that Bill saw someone.
- c. John thought that Bill saw someone.
- d. *John thought who Bill saw.

We suggest that this kind of data, when scaled up to the entire lexicon, is a perfect test-bed for questions about acceptability and grammar. In this case, the two verbs differ inversely in whether they license interrogative vs. declarative complements (though see [White 2019](#)), something we expect would be mirrored in corpus frequencies (after controlling for, e.g., the fact that *think* is itself much more frequent than *wonder*).

While there are many more frames than just these two, there are a limited number of possible selectional patterns instantiated in English; and while, as we have suggested, there are *many* more verbs that might potentially participate in patterns like this, this number (about 1000) is still tractable with modern experimental and corpus methods. Further, there is a large body of literature arguing that such patterns involve many subregularities that a learner could find, and the key points of between-sentence variation involved in pairs like this are fairly minimal and (relatively) possible to extract from corpus data. Finally, we would expect some gradience in the judgments—e.g. (3b) is sometimes claimed to be ‘not as bad’ as (3d).

To date, there is not a large body of experimental linguistics work on acceptability and selection. But the role of selectional patterns has been crucial in research on language acquisition—and particularly verb learning since [Landau & Gleitman’s \(1985\)](#) and [Gleitman’s \(1990\)](#) seminal work. This line of work suggests that children use syntactic frame information to infer semantic representation when learning the meanings of verbs. We do not review this literature in detail (see [White 2015](#) for a recent review), but since the acceptability judgment method in [Fisher et al. 1991](#) (experi-

ment 1, part B) is a crucial predecessor to what we develop here, we briefly discuss it.

In Fisher et al.'s (1991) method, a set of verbs and a set of syntactic frames are selected, and the full cartesian product of these sets is constructed. For each pair in the cartesian product one or more sentences are produced by instantiating the syntactic frames with lexical items, then placing the past tense form of the verb in the resulting instantiation—unless the frame is explicitly specified for some other tense/modal. Subjects then do an acceptability judgment task that (across all subjects) exhausts the matrix of verb-frame pairs.

In this study, Fisher et al. used carefully hand-constructed non-idiomatic sentences for each cell of this matrix: 24 verbs \times 39 frames = 936 sentences. For many purposes, this method can work well as long as the verbs and frames are appropriately sampled relative to the questions at hand; but compared to the total scale of the lexicon, it does not even come close to being exhaustive. Moreover, the method does not easily scale much beyond the size of the original Fisher et al. study because of the challenge of hand-constructing items; for 1000 verbs this would require construction 39,000 sentences. While it may be possible to automate this process to some degree, as Fisher et al. (1991) note (p. 347), substantial care and effort is required in selecting the sentences to be judged because of the potential for unintended item effects. Thus, the challenge is to understand how it might be possible to generate sentences on this scale without requiring hand-inspection of every such sentence. We take this challenge on in Sections 3 and 4.

2.3 Frequency and selection

The connection between frequency and acceptability in the domain of clause selection is not in general well-studied—or really, even on the radar of the theoretical linguistics work mentioned above (though see Bresnan 2007; Bresnan et al. 2007 and references in fn. 1). But as for acceptability and selection, frequency of exposure has played a role in discussions of the acquisition of verb meaning. In particular, it has been hypothesized that the frequency with which a verb occurs in a syntactic structure (along with the frequency of that syntactic structure across verbs) plays a role in learning that verb's meaning (Lederer et al. 1995; Alishahi & Stevenson 2008; Barak et al. 2012; White 2015; White et al. 2017a).

An important component of these proposals are the mechanisms they employ for normalizing frequency information across verbs—in particular,

the cooccurrence frequencies for verbs in different syntactic structures—and subsequently abstracting that frequency information. These mechanisms tend to take the form of clustering models (Lederer et al. 1995), mixture models (Alishahi & Stevenson 2008; Barak et al. 2012), or matrix factorization models (White 2015; White et al. 2017a). We defer detailed discussion of these models until Sections 5 and 6.

2.4 Hypotheses

Based on the prior work discussed in this section, we suggest the following hypotheses:

- H1** Verb-frame co-occurrence frequencies predict acceptability.
- H2** Verb-frame co-occurrence frequencies require syntactic/semantic abstractions in order to predict acceptability.

Given the findings of this prior work, it would be quite surprising if H1 turns out to be entirely false: there is likely to be some sense in which these frequencies do predict acceptability. However, the linguistics and acquisition literature on verb representations suggests H2 – some kind of abstraction may be necessary to learn even something like acceptability patterns, where in principle there is the potential to just match frequencies. For the human learner, the situation is made starker by something we are idealizing away from: a child does not receive the kind of massive, balanced corpus input our models will.

The program for testing these hypotheses is, at this point, straightforward: we need a (large) selectional acceptability dataset, an estimation of frequencies, and models that compare the two by varying the amount of abstraction. We now turn to the first of these.

3 The Bleaching Method

A major obstacle to scaling standard acceptability judgment tasks to entire subregions of the lexicon is ensuring that plausibility effects—i.e. effects on acceptability that are driven by how prototypical the situation described by the sentence is, as opposed to effects driven by syntactic well-formedness—are controlled for. We propose a method to control for these effects by ‘semantically bleaching’ all lexical category words besides a word of interest—in our case, the clause-embedding verb. Specifically, we manipulate the

syntactic context that a word appears in while instantiating all NPs in that context with indefinites (*someone, something*) and all verbs in that context (besides the one of interest) with a low content eventive (*happen, do*) or stative (*have*) verb.

We first demonstrate the validity of this method on a small set of verbs by showing that agreement is high among naïve participants' acceptability ratings when responding to 'contentful sentences' v. when responding to 'bleached sentences' that are otherwise matched in terms of structure—effectively comparing our bleaching method against the Fisher et al.'s (1991) more standard method, described above. We use the data reported in White et al. 2018 as a dataset of acceptability judgments to contentful sentences—as it focuses on exactly the phenomena we are interested in—and collect acceptability judgments to bleached sentences ourselves.

3.1 Materials

We follow White et al. (2018) in using the 30 propositional attitude verbs found in (4), which were selected in such a way that they evenly span the verb classes presented in Hacquard & Wellwood 2012.

- | | | | |
|-----|---------------|------------|------------|
| (4) | a. think | k. hear | u. forbid |
| | b. realize | l. feel | v. allow |
| | c. understand | m. tell | w. promise |
| | d. suppose | n. say | x. love |
| | e. guess | o. promise | y. hate |
| | f. expect | p. hope | z. bother |
| | g. imagine | q. worry | aa. amaze |
| | h. remember | r. doubt | bb. demand |
| | i. forget | s. pretend | cc. want |
| | j. see | t. deny | dd. need |

We also follow White et al. in testing frames built from the same combinations of syntactic features—including various tense-aspect combinations within matrix and embedded clauses as well as various forms of NP and PP arguments. They construct 30 subcategorization frames from these combi-

nations of features.³ These frames are given in abstracted form in (5), with our instantiation for each constituent type in (7).

- (5)
- | | |
|-----------------------------|----------------------------------|
| a. NP ___ed | q. NP ___ed for NP to VP |
| b. NP ___ed NP | r. NP ___ed to VP |
| c. NP ___ed NP NP | s. NP ___ed WH to VP |
| d. NP ___ed about NP | t. NP ___ed NP to VP |
| e. NP ___ed NP about NP | u. NP was ___ed to VP |
| f. NP ___ed so | v. NP ___ed there to VP |
| g. NP ___ed to | w. NP ___ed VP _{ing} |
| h. NP ___ed S | x. NP ___ed NP VP _{ing} |
| i. NP ___ed that S | y. NP ___ed NP VP |
| j. NP ___ed if S | z. It ___ed NP that S |
| k. NP ___ed S _{wh} | aa. It ___ed NP S _{wh} |
| l. NP ___ed NP S | bb. It ___ed NP WH to VP |
| m. NP ___ed NP that S | cc. It ___ed NP to VP |
| n. NP was ___ed that S | dd. S, NP ___ed |
| o. NP ___ed it that S | ee. S, I ___ |
| p. NP ___ed to NP that S | ff. S, NP ___ed |

To instantiate these abstract frames, [White et al.](#) instantiate each of the phrases with contentful lexical items (following [Fisher et al. 1991](#)). For instance, (6) shows one of three items that instantiate the pair (*think*, *NP ___ed that S*) in [White et al.](#)'s experiment.

- (6) Gary thought that she fit the part.

As noted above, one potential issue that arises when using contentful lexical items to instantiate a frame is that ratings of the resulting items are susceptible to typicality effects—some verbs might sound less plausible in certain frame instantiations even if that verb is otherwise perfectly fine in some other instantiation of that same frame. [White et al.](#) control for such typical effects by creating three different instantiations for each frame and taking into account possible item variability in their analysis. But if at all possible, it is ideal not to do this, since it increases the number of items and lowers

³ [White et al.](#) note that their experiment in fact included some syntactic frames that involve degree modification, though they do not list these frames. [White et al. 2014](#), which presents a preliminary analysis of the dataset presented in [White et al. 2018](#), lists four such frames (see the Appendix of that paper). We do not include these in our experiment, since the analyses in [White et al. 2018](#) do not include them, and thus our statistics would not be comparable to theirs if we did.

the statistical power of subsequent analyses—thus requiring more ratings to get an accurate estimate of the acceptability of any particular verb-frame pair.

To address this both in the current experiment and in the large-scale experiment reported in Section 4, we instead create only one instantiation of each frame with as little lexical content as possible. All instantiations we use are listed in (7).

- (7)
- a. **NP** Noun phrase (*someone or something*)
 - b. **VP** Verb phrase with verb in bare form (*do something*)
 - c. **VP_{ing}** Verb phrase with verb in present progressive form (*doing something*)
 - d. **S** Full clause without complementizer (*something happened*)
 - e. **S_{wh}** Full embedded interrogative clause⁴ (*which thing happened*)
 - f. **S_[-tense]** Tenseless embedded clause (*something happen*)

For example, (8) gives the item instantiating the pair (*think, NP_{ed} that S*).⁵

- (8) Someone thought that something happened.

From the resulting set of 30 verbs \times 46 frames = 1,380 items, we constructed 23 lists of 60 items, constrained such that each verb occurred exactly twice in each list (always with a distinct frame) and each frame occurred at least once and no more than twice.

3.2 Participants

We recruited 115 unique participants = 23 list \times 5 participants per list through Amazon Mechanical Turk. All participants reported speaking American English as their native language.

3.3 Predictions

To measure interannotator agreement, White et al. compute the Spearman rank correlation between the responses for each pair of participants that did the same list and report a mean correlation of 0.64. This measure is not comparable to ours, however, since our items were specifically designed to

⁴ White et al. only use adjunct questions to avoid free relative readings. We instead opt for D-linked WH questions, though for the same reason.

⁵ The *something* instantiation of NP is only used for NPs in object position of a simple transitive (*NP_{NP}*) and second object position in a double object construction (*NP_{NP} NP*).

block lexical information from being used in the acceptability judgments. Under the assumption that less information would yield higher variability in judgments, we thus expect lower agreement. To assess how much lower we should expect, we attempt to factor out the variability across items that instantiate a particular verb-frame pair. To derive this more comparable measure, we simulate the amount of agreement we would expect, assuming they had used a method like ours.

First, we fit an ordinal (linked logit) mixed effects model to the ratings from [White et al.](#)'s data, with fixed effects for VERB, FRAME, and their interaction and random unconstrained cutpoints for each participant (see [Appendix A](#) for details). We then use this model to simulate how each participant from their experiment would respond to each verb-frame pair in their experiment by (i) using the ordinal model to produce a predicted probability distribution over the ordinal scale ratings for each participant and item; (ii) sampling once from each of those distributions; and (iii) computing the Spearman correlation between responses given by all pairs of simulated participants. We repeat this simulation 999 times, computing the mean agreement each time.⁶ This yields a mean correlation of 0.516 (95% CI: [0.511, 0.521]) across all simulations.

3.4 Results

We are concerned with two sorts of results in this validation: (i) interannotator agreement among participants in the validation compared to interannotator agreement in [White et al.'s \(2018\)](#) data; and (ii) agreement in the aggregated ratings for verb-frame pairs.

To measure interannotator agreement, we compute the Spearman rank correlation between the responses for each pair of participants that did the same list. This yields a mean correlation of 0.528 (95% CI: [0.508, 0.546]). Thus, the level of interannotator agreement we observe is exactly what we expect given data collected under a more standard methodology.

Next, we turn to agreement between the average ratings for each verb-frame pair computed from each dataset. To compute these average ratings, we fit the ordinal mixed model described in the last subsection to each

⁶ Throughout the remainder of the paper, all confidence intervals are computed using non-parametric bootstraps with 999 replicates.




Figure 1: Correlation by verb between mean normalized verb-frame acceptability in [White et al.’s \(2018\)](#) data and our replication.

dataset separately and then compute the predicted real-valued acceptability for each verb-frame pair.⁷

Figure 2 plots the Spearman rank correlation between these normalized, real-valued verb-frame acceptability by frame (across verbs), and Figure 1 plots the same correlation by verb (across frames). In both plots, the dashed line shows the mean interannotator agreement, and error bars show 95% confidence intervals.

In Figure 1, we see that all verbs show correlations above the mean interannotator agreement. This suggests that as a measure of verbs’ syntactic distributions, our data are encoding essentially the same distributional information that [White et al.’s](#) data are.

In Figure 2, we see that most frames show average correlations close to or above the mean interannotator agreement. There are five frames that do

⁷ This procedure is analogous to *z*-scoring the ratings by participant and then computing the average *z*-scored rating for each item. As shown by [White et al. \(2018\)](#), the method used here better models how participants actually make acceptability judgments.




Figure 2: Correlation by frame between mean normalized verb-frame acceptability in *White et al.’s (2018)* data and replication.

not show a correlation that is significantly different from zero: *NP* _ *NP*, *NP* _ *NP NP*, *It* _ *NP that S*, *It* _ *NP WH S*, and *It* _ *NP WH to VP*.

For *NP* _ *NP* and *NP* _ *NP NP*, it seems likely that the disagreement arises because *White et al.’s* instantiations of those frames only ever include object NPs that denote concrete inanimate entities that cannot be straightforwardly associated with propositional content—e.g. cups, tables, bottles. In contrast, the inanimate indefinites we use could denote either contentful inanimates or contentless ones. This is likely the cause of higher acceptabilities observed for predicates like *believe* and *tell*. Compare (9a) and (9b) with (10a) and (10b).

- (9) a. #I believed the table.
b. #I told her the table.⁸
- (10) a. Someone believed something.
b. Someone told someone something.

⁸ Note that these ditransitive frames are licit with verbs such as *allow*, *deny*, and *forbid*.

Since we aim to factor out effects due to lexical items, this is a point in favor of our method.

For *It_ NP that S*, *It_ NP WH S*, and *It_ NP WH to VP*, we suspect that the low agreement stems from inherent variability in the judgments for items with expletive subjects. One reason this may arise—pointed out by White et al.—is that *it* can be read referentially in these frames, and thus participants' judgments might vary widely depending on their interpretation of *it*. This predicts that expletive subject frames should show lower agreement on average, which appears to be the case.

Regardless of its source, the existence of this low agreement for expletive subject frames suggests that we should be wary of including such frames in a large-scale experiment like the one we report on in this paper. Nonetheless, we would still like to capture information about whether a verb allows expletive subjects. We discuss our approach to this below.

3.5 Discussion

The results reported above suggest that the bleaching method is a promising way to avoid item effects in acceptability judgment tasks for selectional patterns. Since it involves an extremely simple generation strategy, it therefore is also a promising way of scaling standard acceptability judgment tasks to entire subregions of the lexicon. But even with bleaching, one must thread the needle between useful referential ambiguities—such as the one between entity and propositional reference introduced by the use of *something*—and plausibly syntactic ambiguities that introduce variability into the judgments—such as the one between referential and expletive *it* pointed out above.

We address the particular issue of expletive *it* in our lexicon-scale annotation by using an alternative set of structures for capturing the acceptability of predicates that occur with expletive *it*: at least those predicates that take experiencer objects. Our approach, which we describe in the next section, is to use passivized version of the expletive object frames; compare (11a) and (11b).

- (11) a. It amazed someone that something happened.
 b. Someone was amazed that something happened.

This approach introduces some amount of ambiguity—we don't know whether a verb that is acceptable in contexts like (11b) takes contentful or expletive

subjects—but this ambiguity is resolvable by looking at the acceptability of that verb in contexts such as (12).

(12) ???Someone amazed someone that something happened.

That is, a verb is licensed in an expletive subject frame if it is licensed in a passive transitive frame, and not licensed in a non-passive ditransitive frame. In our large-scale experiment, we present an expanded set of frames using this manipulation.

4 The MegaAcceptability Dataset

The main goal of collecting our large-scale acceptability judgment dataset is to obtain a single normalized acceptability score for every clause-embedding verb in the English lexicon, along with an estimate of the variability in judgments for that item. We discuss our data collection method and how we derive these estimates here. In Section 5, we describe experiments that attempt to predict these normalized acceptabilities from frequency data.

4.1 Materials

To scale up the materials, we selected a set of frames, a set of verbs, and automated a method of constructing bleached sentences for every member of the cartesian product of the two.⁹

For verb selection, we attempted to exhaustively select every verb in English which could take a clause of some kind. First, we took the union of several lists of clause-embedding verbs collected in previous work (Hacquard & Wellwood 2012; Anand & Hacquard 2013; 2014; Rawlins 2013; White et al. 2014) as a ‘seed’ set. Helpfully, a range of existing lists were already aggregated in Rawlins (2013) and constituted the bulk of the seed. This gave us about 500 verbs.

We then searched in VerbNet (Kipper-Schuler 2005)—a database that is, in large part, directly derived from the verb classes in Levin 1993—to find all verbs in all VerbNet classes that any of the seed verbs were present in, with a hand-filtering pass to remove obvious errors—e.g. cooking verbs.

To pick the set of frames, we first collected a set of five basic syntactic features that are believed to be relevant to selectional patterns, and selected either all or the most frequent values for these features, depending on the

⁹ These materials were first described in White & Rawlins 2016.

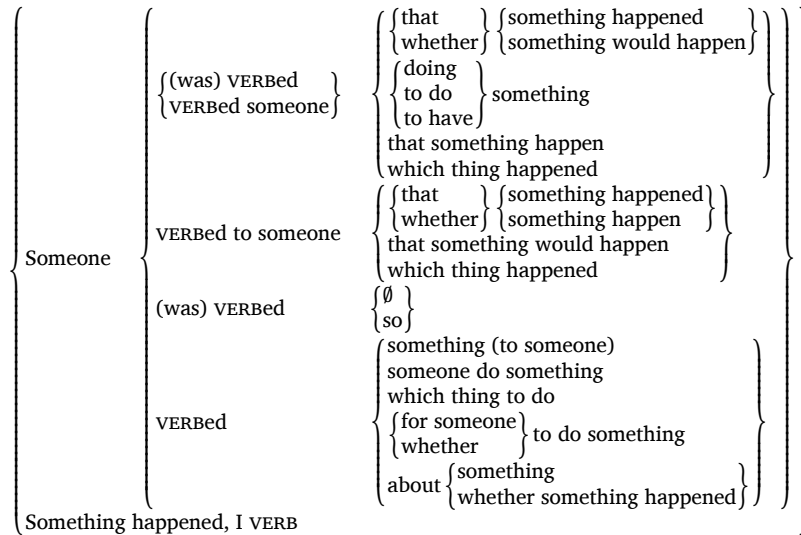


Figure 3: All frame instantiations in the MegaAcceptability dataset..

details. In this case, we did not aim for full exhaustivity, but rather to get as big a sample as possible within constraints imposed by the already large experiment.

For example, for prepositional phrases, we consider only the prepositions *to* and *about*, though many other prepositional markers such as *of* may be relevant to the ultimate question of how to represent selectional patterns. For embedded constituent interrogatives, we chose to use an embedded D-linked WH-phrase (*which thing*) in order to maximize acceptability. To these features, we added passivization—in order to handle expletive subjects as described above—and two more idiosyncratic frame manipulations: declarative slifting (Ross 1973), and the proform *so* (Ross 1972; Hankamer & Sag 1976).

- (13)
- a. **Complementizer:** \emptyset , *that*, *for*, *whether*, *which thing*
 - b. **Embedded tense:** *past*, *future*, *infinitival*, *present participle*, *bare*
 - c. **Matrix object count:** 0, 1, 2
 - (i) **so-frame manipulation:** add *so* to 0-object frames
 - d. **Matrix PP:** \emptyset , *to*, *about*
 - e. **Embedded subject?:** true, false
 - f. **Passivized verb?:** true, false
 - g. **Slifting manipulation:** ‘Something happened, I VERB’

The frame instantiations are shown in Figure 3. As can be seen in this figure, the bleaching manipulation we applied is identical to that applied in the validation experiment reported in Section 3.

4.2 Participants

727 unique participants were recruited through Amazon Mechanical Turk to rate sentences in the 1,000 lists of 50. Participants were allowed to respond to as many unique lists as they liked, but no one participant rated the same list twice. Five ratings were obtained per sentence—i.e. five unique participants rated each list—and each participant responded to a median of four lists (mean: 6.9, min: 1, max: 56).

Four participants reported being native speakers of a language other than English. These participants' responses were removed from the dataset prior to analysis, for a loss of 600 responses total (~0.2% of the data). None of these participants rated the same list.

4.3 Response Normalization

We use a slightly modified form of the ordinal model-based normalization procedure described in Section 3 to produce two pieces of information associated with each verb-frame pair: a real-valued acceptability value (more positive is more acceptable) and the mean (log-)likelihood associated with all acceptability judgments for a particular item. (A full specification of this procedure can be found in Appendix B.) The second score can be viewed as a measure of variability in the judgments: the lower this score is, the higher the variability in ordinal responses to a particular verb-frame.

As an example of what these two measures look like, Figure 4 plots the normalized acceptability scores for the *NP __ed that S* and *NP __ed NP that S* frames. Each point is a verb and only a subset of points are labeled. Smaller labels and grayer points correspond to higher mean variability scores.

4.4 Reliability

In Section 3, we compared the normalized acceptability obtained from White et al.'s (2018) experiment and our replication that used the bleaching method. We cannot derive similar agreement estimates here because we only have a single dataset and because the lists were built in such a way that partici-

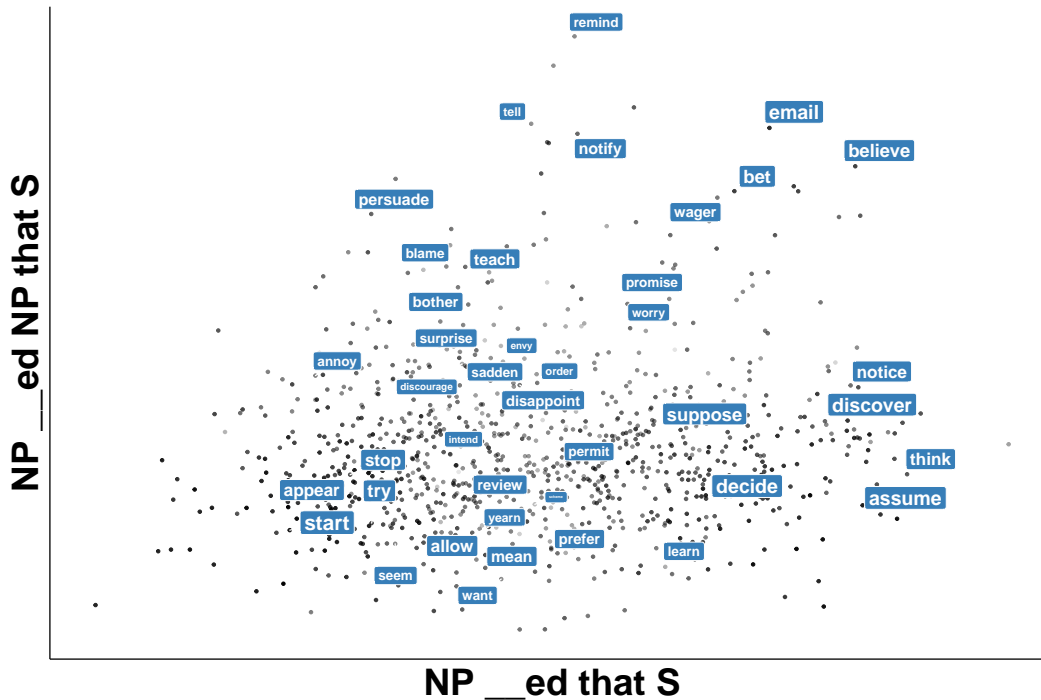


Figure 4: Normalized judgments for two frames. Each point is a verb and only a subset of points are labeled. Smaller labels and grayer points correspond to higher mean variability scores.

pants only saw one verb and one frame and thus it is not possible to assess the correlations within a particular frame.

But because we do have estimates of the variability in judgments for each item. We can get a sense for how much agreement there is within judgments for a particular verb or frame by taking the mean of the above-defined variability scores for each frame across verbs. Figure 5 plots these means in terms of probabilities, where $\frac{1}{7}$ is the lowest possible probability—and thus, the highest possible variability. We see that nearly all frames fall within a narrow band between 0.3 and 0.4, suggesting that no frame shows particularly high disagreement—in contrast to what we saw in Section 3 for the expletive subject and NP direct object frames. Further, our replacements for the expletive subject frames—the passivized frames—do not show systematically lower variability, suggesting that our approach was successful. (See [White & Rawlins \(2016\)](#) for further discussion of how this data might be used or validated in a formal semantics context.)




Figure 5: Mean variability in judgments across verbs for each frame.

5 Relating Frequency and Acceptability

We now turn to the main question of this paper: to what extent can a verb’s subcategorization behavior be predicted directly from the frequency with which it occurs in different syntactic structures in text? To obtain our measure of frequency, we use the VALEX dataset, which is the largest publicly available dataset of subcategorization frame frequencies (Korhonen et al. 2006).¹⁰ VALEX is built from over 900 million words of text and contains 163 subcategorization frame types, described in Briscoe & Carroll 1997, and over 6,000 verbs, 958 of which are shared with the MegaAcceptability dataset. (The verbs that are missing tend to be particle verbs, as these are not treated as separate verbs by VALEX.)

One obstacle to using any frequency dataset for predicting acceptability is that we must determine the importance of any particular verb-frame co-occurrence in the context of its entire distribution. For example, observing a high-frequency verb like *think* once or twice in a ditransitive frame should

¹⁰ VALEX is available at <https://ilexir.co.uk/valex/>. We use the raw counts provided with the data, since all other counts involve some amount of smoothing and/or filtering.

make us less certain that *think* is highly acceptable in that frame than observing a lower frequency verb like *begrudge*. This situation is delicate, though, since it requires the specification of some *frequency normalization model* for processing the raw frequency data and we do not want to introduce too much inductive bias at this stage to ensure that we are in fact testing the predictability of acceptability directly from frequency.

To thread this needle, we consider only normalization models that represent verbs' distributions directly in terms of the original subcategorization frames—as opposed to some set of latent syntactic or semantic factors—while accounting for the importance of a particular observation in the context of those distributions. We then use each of these representations as predictors in a linear model of the normalized acceptability judgments described in the last section.

The use of a linear model for this purpose is important so as not to introduce any further inductive bias—as, e.g., the use of a kernelized support vector machine or multi-layer perceptron might. Further, since linear models learn linear functions from one representation to another and since linear functions are all and only the homomorphisms between those representations, this setup allows us to make stronger conclusion about the character of the relationship—specifically, whether or not the (normalized) frequency distributions and acceptability are homomorphic.

5.1 Normalization Models

We consider two probabilistic models and two information theoretic models of subcategorization frame frequency distributions. Our aim in looking at multiple models is less to compare their relative performance, and more to give the frequency information the best chance at explaining acceptability.

5.1.1 Probabilistic Models

The first probabilistic model we consider models the conditional probability $\mathbb{P}(f \mid v)$ of seeing a particular frame f given a particular verb v as a Categorical distribution with parameters (probabilities) θ_v .

$$\mathbb{P}(f \mid v) = \mathbb{P}(f \mid \theta_v) = \text{Categorical}(f; \theta_v) = \theta_{vf}$$

We use the frequencies c_{vf} for each verb v and frame f to compute the posterior probability $\mathbb{P}(\theta_v \mid c_v)$ of θ_v under the assumptions (i) that the verb-

frame pairs are sampled independently; and (ii) that the prior probability $\mathbb{P}(\boldsymbol{\theta}_v; \boldsymbol{\alpha})$ is given by a Dirichlet distribution with parameter $\boldsymbol{\alpha}$.

$$\mathbb{P}(\boldsymbol{\theta}_v | \mathbf{c}_v; \boldsymbol{\alpha}) \propto \mathbb{P}(\boldsymbol{\theta}_v; \boldsymbol{\alpha}) \prod_f \mathbb{P}(f | \boldsymbol{\theta}_v)^{c_{vf}}$$

We then use the most likely probabilities $\hat{\boldsymbol{\theta}}_v$ —i.e. the *Maximum A Posteriori* (MAP) estimate—for each verb as our representation of a verb’s distribution—i.e. $\hat{\boldsymbol{\theta}}_v$ is what we use to predict a verb’s acceptability. When $\boldsymbol{\alpha}$ is a constant positive vector $(\lambda + 1)\mathbf{1}_{N_F}$ (where n_F is the number of frames), this turns out to be equivalent to standard add- λ smoothing.

$$\hat{\boldsymbol{\theta}}_v = \arg_{\boldsymbol{\theta}_v} \max \mathbb{P}(\boldsymbol{\theta}_v | \mathbf{c}_v; \boldsymbol{\alpha} = (\lambda + 1)\mathbf{1}_{N_F}) = \left[\frac{c_{v1} + \lambda}{\sum_i c_{vi} + \lambda}, \frac{c_{v2} + \lambda}{\sum_i c_{vi} + \lambda}, \dots \right]$$

The second probabilistic model attempts to directly extract acceptability from the frequency data (White 2015; White et al. 2017a). In this model, the conditional probability of seeing a particular frame f with some frequency c_{vf} given a verb v is assumed to have a negative binomial distribution with probability π_{vf} and rate r_v .

$$\mathbb{P}(c_{vf} | v) = \text{NegBin}(c_{vf}; \pi_{vf}, r_v) = \binom{c_{vf} + r_v - 1}{c_{vf}} (1 - \pi_{vf})^{r_v} \pi_{vf}^{c_{vf}}$$

This distribution is a natural choice both (i) because it is known to be a good model of similar kinds of count data (Church & Gale 1995); and (ii) because the parameters themselves have natural interpretations: (a) the parameter π_{vf} can be viewed as a probability of acceptability: when it is close to one, we expect to see more instances of a frame with a verb (though there is a non-zero probability of seeing it rarely); when it is close to zero, we expect to see fewer; and (b) the verb’s rate parameter r_v roughly controls its overall frequency how much variability there is in the frame frequencies for a particular verb. Thus, unlike for the Dirichlet-Categorical model, we straightforwardly separate our knowledge of competence from our knowledge of frequency (for further discussion, see White 2015).

As with the $\boldsymbol{\theta}_v$ parameter of the Categorical-Dirichlet model, we aim to find the most likely pairing $\hat{\pi}_v, \hat{r}_v$ for each verb v , given the counts \mathbf{c}_v . We assume that the prior probability $\mathbb{P}(\pi_{vf}; \beta_1, \beta_2)$ in this case is Beta distributed with parameters $\beta_1 = \beta_2$.¹¹ We assume an improper (uniform) prior for r_v .

¹¹ This is analogous to the Categorical-Dirichlet model: larger values of λ encourage estimates of $\hat{\boldsymbol{\theta}}_v$ nearer $\frac{1}{N_F}$ and larger values of $\beta_1 = \beta_2$ encourage values of π_{vf} nearer $\frac{1}{2}$.

$$(\hat{\pi}_{vf}, \hat{r}_v) = \arg_{\pi_{vf}, r_v} \max \mathbb{P}(\pi_{vf}, r_v \mid \mathbf{c}_v; \beta_1, \beta_2)$$

Unlike for $\hat{\theta}_v$, this MAP estimate for $\hat{\pi}_{vf}, \hat{r}_v$ cannot be computed using a closed form and so we use gradient descent to obtain it.

For both the Dirichlet-Categorical model and the Beta-Negative Binomial model, we refer to the hyperparameters λ , in $\boldsymbol{\alpha} = (\lambda + 1)\mathbf{1}_{N_f}$, and $\beta_1 = \beta_2$ as *smoothing parameters*. We consider multiple different settings of smoothing parameters in our experiments, described below.

5.1.2 Information Theoretic Models

The first information theoretic model we consider uses the *pointwise-mutual information* (PMI; Church & Hanks 1990) between a verb and a frame.

$$\text{PMI}(v, f) = \log \frac{\mathbb{P}(v, f)}{\mathbb{P}(v) \cdot \mathbb{P}(f)}$$

To compute this quantity, we assume the Dirichlet-Categorical model described above and obtain MAP estimates for the parameters of the joint distribution $\mathbb{P}(v, f)$. The parameters of the marginal distributions $\mathbb{P}(v)$ and $\mathbb{P}(f)$ can then be obtained from the joint. We estimate these distributions using the same definition of the smoothing parameter λ .

The second information theoretic model we consider uses the terms of the G statistic, as described by Dunning (1993). This amounts to scaling the PMI by the relevant conditional probability—in our case, the probability of the frame given the verb—better controlling for relatively poorer probability estimates for low frequency verbs.

$$G(v, f) = \mathbb{P}(f \mid v) \cdot \text{PMI}(v, f)$$

As for PMI, we estimate these distributions using the same definition of the smoothing parameter λ .

5.2 Experiments

We compute MAP estimates for the parameters of the Dirichlet-Categorical model θ_v and the Beta-Negative Binomial model π_v, r_v with smoothing parameters $\lambda \in L = \{0, 0.1, 0.2, 0.5, 1, 2, 5, 10, 20, 50\}$ and $\beta \in L - \{0\}$.¹² We compute PMI and G using MAP estimates based on the same settings of λ .

¹² The smoothing parameter β cannot be set to zero, since the Beta distribution requires strictly positive parameters.




Figure 6: Mean variance explained in normalized acceptability judgments in 10-fold/10-fold nested cross-validation for each model and smoothing parameter.

We regress the normalized acceptability judgments for each verb in each frame on each of these representations in a multivariate ridge regression—i.e. a linear regression with L2 regularization—implemented in sklearn (Pedregosa et al. 2011). To set the regularization parameter $\alpha \in \{0.01, 0.1, 0.2, 0.5, 1, 2, 5, 10\}$, we use a 10-fold cross-validation. To compute the generalizability of this model, we nest this cross-validation within another 10-fold cross-validation and compute the mean R^2 (variance explained) on the held out datasets in this outer cross-validation.

5.3 Results

Figure 6 shows the mean R^2 across the 10 cross-validation folds for each model and smoothing parameter. We see that the Beta-Negative Binomial model ($\beta = 1$) is the best-performing model, with the PMI model ($\lambda = 5$) a close second. The Dirichlet-Categorical model ($\lambda = 0$) is a close third, performing only slightly (but reliably) worse. The G model consistently does more poorly than the other models—possibly because it too strongly downweights the scores for low frequency verbs (unlike the other measures, which attempt to filter out frequency to some extent). But even for the best-performing models, the scores are quite low. This suggests that, while the joint frequency of a verb and a frame carries some information about the

acceptability of that verb in that frame, it is far from enough to determine that acceptability.

One question that arises here is whether this poor performance is due to verb-frame pairs that received highly variable judgments. In this case, we should expect a positive correlation between judgment variability (as defined in Section 4) and models' absolute error. We test this hypothesis using our best-performing model's absolute error on the held-out data for each fold of the cross-validation. Rather, than finding a positive correlation, we find a weak (but reliable) negative Spearman rank correlation of -0.197 (95% CI = [-0.205, -0.188]).

This suggests that highly variable judgments are actually slightly easier to predict than less variable judgments. One reason this may come about is that more variable judgments tend to have normalized acceptabilities near the center of the acceptability scale, which arises from the fact that high variability is a consequence of extreme responses from participants that average out to the middle of the scale. This means that if a model incorrectly predicts a very high or very low acceptability score it will tend to be less wrong for the highly variable predicates in the middle of the scale than for predicates that participants were more certain about.¹³

A similar question arises with respect to frequency: the poor performance we observe could be due to poor estimates for the distributions of low frequency verbs. In this case, we should expect a negative correlation between verb frequency and models' absolute error. We again test this hypothesis using our best-performing model's absolute error on the held-out data for each fold of the cross-validation. We instead find a reliably positive Spearman rank correlation here, though it is extremely weak 0.018 (95% CI = [0.009, 0.027]). This suggests that poor estimates of verbs' distributions—at least their frequency distributions—is not the cause of our models' poor performance.

5.4 Discussion

What is the source of the models' low performance then? We believe it is likely due to a systematic bias in the kinds of information frequency distributions contain. Specifically, we posit that those aspects of a verbs'

¹³ One should not infer from this discussion that there is a perfect correlation between normalized acceptability and variability. There are also many verb-frame pairs that have normalized acceptabilities in the middle of the scale and low variability because they always receive middling ratings. Verbs near the two extreme do tend to have low variability, however, since otherwise they will not average out to a high value.

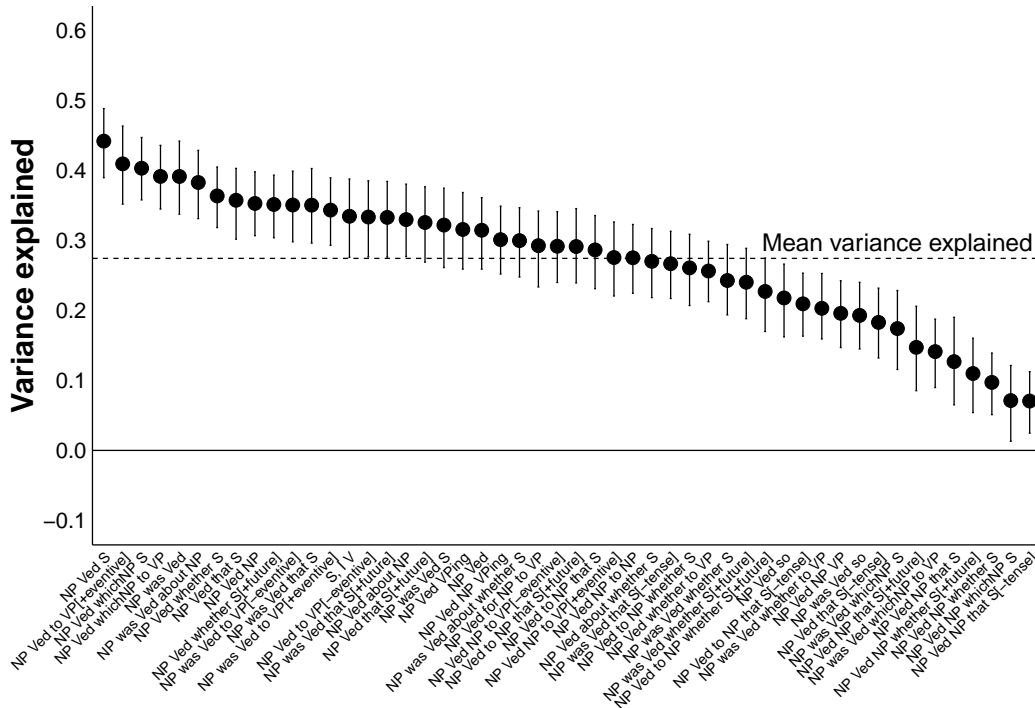


Figure 7: Variance explained by best-performing Beta-Negative Binomial model ($\beta=1$) broken out by frame.

distributions that are predictable either from their abstract syntactic properties or from their meaning will not necessarily be directly encoded in their frequency distributions. That is, one will not necessarily observe all frames a particular verb is acceptable in insofar as the acceptability of that verb in that frame is predictable from its acceptability in another frame (see Grimshaw 1981; Pinker 1984; 1989; Lasnik 1989; Kako 1997; Lidz et al. 2004 for how this might work). Conversely, a verb being acceptable in a frame does not entail observing that verb in that frame, unless that acceptability is not predictable from its meaning.

One piece of evidence for this comes from which frames our models perform worst on. Figure 7 plots the R^2 for the best-performing Beta-Negative Binomial model ($\beta = 1$) broken down by frame. We see that the model systematically does more poorly in predicting verbs’ acceptability in frames involving direct and indirect objects and a tensed embedded clause. This is consonant with our hypothesis insofar as verbs’ acceptability in these frames is predictable from some semantic property—plausibly in this case, whether the verb is communicative or not.

Part of this hypothesis appears to fly in the face of previous work in the syntactic bootstrapping literature demonstrating that distributional cues are useful for inferring a word’s meaning (Landau & Gleitman 1985; Gleitman 1990; Naigles 1990; 1996; Naigles et al. 1993; Fisher et al. 1991; Fisher 1994; Fisher et al. 1994; 2010; Lederer et al. 1995; Gillette et al. 1999; Snedeker & Gleitman 2004; Lidz et al. 2004; Gleitman et al. 2005; Papafragou et al. 2007; White 2015; White et al. 2017a; 2018). But this conflict is only apparent. A key part of our hypothesis is that acceptability is not *directly* encoded in frequency distributions. But certain components of that distribution may be observed for particular verbs, and the observation of that component may imply the acceptability of others. For instance, there is a relatively strong correlation between acceptability in an *NP Ved NP that S* frame and acceptability in an *NP Ved NP whether S* and so it is generally a safe bet that if a verb is acceptable in one it will be acceptable in another.

What this view implies is that, insofar as abstract syntactic and semantic properties reveal themselves in verb’s subcategorization frame frequency distributions, it may be possible to infer those properties from regularities observable across those distributions. So far, the results are consistent only with our hypothesis H2 from Section 2, and not with its direct counterpart: we apparently *cannot* predict acceptability in selectional behavior from frequency. We consider various models of this abstraction in the next section.

6 Abstracting Frequency

Constructing useful abstractions of frequency distributions is a major component of much work in NLP. Abstraction techniques can take myriad forms, both probabilistic and neural. We consider four popular abstraction techniques, selected to be roughly analogous to the probabilistic and information theoretic models presented in the last section.

Our main goal in doing this is to determine the extent to which these abstractions represent acceptability directly, by which we mean that the space of abstractions and acceptability are homomorphic. As noted in the last section, the set of homomorphisms in a vector space are just the linear functions, and so as in the last section, we will attempt to predict acceptability using linear regression on the different representations we consider. We point this out because, for some of the representations we construct, it is common practice to learn nonlinear functions to a quantity of interest; and while this can be useful for understanding whether a particular abstraction

implicitly contains information about a quantity, it does not tell us the extent to which that abstraction is potentially a representation of that quantity in an algebraic sense.

6.1 Models

The first model we consider is Latent Dirichlet Allocation (LDA; Blei et al. 2003), which is analogous to the Dirichlet-Categorical model presented in Section 5. This model is closely related to Alishahi & Stevenson’s (2008) model of verb learning. It assumes that each verb is probabilistically associated with a set of K latent syntactic and semantic properties via a conditional Categorical probability distribution $\mathbb{P}(k | v) = \theta_{vk}$ and that each frame is probabilistically associated with that same set of properties via a conditional Categorical probability distribution $\mathbb{P}(f | k) = \phi_{kf}$. The probability of seeing a verb v in a frame f is then modeled via these two distributions.

$$\mathbb{P}(f | v) = \sum_k \mathbb{P}(f | k) \mathbb{P}(k | v) = \sum_k \theta_{vk} \phi_{kf}$$

As for the Dirichlet-Categorical model from Section 5, the parameters θ_v and ϕ_k are assumed to be distributed Dirichlet.

The second model we consider is logistic factor analysis (LFA) with a negative binomial likelihood (see Zhou 2018). This model is closely related to the Poisson Factor Analysis (Zhou et al. 2012; Zhou & Carin 2015) model proposed as a model of syntactic bootstrapping by White (2015) and further developed in White et al. (2017a). This model is analogous to the Beta-Negative Binomial model presented in Section 5, using the same likelihood function but modeling Π via two matrices $\mathbf{U} \in \mathbb{R}^{N_V \times K}$ and $\mathbf{A} \in \mathbb{R}^{K \times N_F}$, where N_V is the number of verbs and N_F is the number of frames.

$$\pi_{vf} = \text{logit}^{-1} \left(\sum_k u_{vk} a_{kf} \right)$$

Similar to Θ in LDA, one way to view \mathbf{U} is as encoding verbs’ abstract syntactic and/or semantic properties; and similar to Φ in LDA, one way to view \mathbf{A} is as encoding the syntactic properties of frame along with whatever aspect of verbal semantics project onto that frame (White & Rawlins 2016). As for the Beta-Negative Binomial model presented in Section 5, we infer \mathbf{U} , \mathbf{V} , and the rate parameters \mathbf{r} of the Negative Binomial likelihood using gradient descent.

The third model we consider uses Global Vectors (GloVe; Pennington et al. 2014), which is a popular *word embedding* method in NLP. GloVe itself is not directly analogous to the PMI method from Section 5, but it is closely related (Levy & Goldberg 2014; Suzuki & Nagata 2015). In essence, it is a factor analysis of the log cooccurrence counts c_{ij} for words i and j .¹⁴

$$\mathbb{P}(c_{ij} \mid \mathbf{W}, \mathbf{W}', \mathbf{b}, \mathbf{b}') = \mathcal{N}(\log c_{ij}; \mathbf{w}_i \cdot \mathbf{w}'_j + b_i + b'_j)$$

As for LDA and LFA, \mathbf{W} represents a relation between verbs and latent syntactic and/or semantic properties and \mathbf{W}' represents the relation between these properties and frames.

We consider two versions of this GloVe-based model. The first uses pre-trained GloVe to compute a neural bag of words (NBoW) representation for each sentence (Iyyer et al. 2015).¹⁵ In NBoW, the point-wise mean of the vector for each word in a multiset is computed. In our case, this multiset is the multiset of words in each sentence of MegaAcceptability. We then predict the acceptability for that sentence from its NBoW representation.

In addition to using pretrained GloVe, we train our own GloVe embeddings on the basis of the VALEX verb-frame counts c_{vf} . This yields an embedding for each verb and each frame on a K -dimensional latent space. We consider the same settings of K for this space as for the LDA and LFA models.

The final model we consider uses contextual word embeddings produced using pretrained Bidirectional Encoder Representations from Transformers (BERT; Devlin et al. 2018).¹⁶ A full technical explication of BERT is not possible here, but in essence, BERT consists of multiple layers of interacting neural network modules known as a transformers (Vaswani et al. 2017) that are trained to predict the probability of a word given the surrounding words in the sentence as well as sentence ordering in a document. This means that BERT’s representation of each word in a sentence contains some amount of information about other words in the sentence.

We use BERT to encode each sentence in MegaAcceptability and then extract the embedding of the sentence start token—i.e. the *classifier* token

¹⁴ It diverges slightly from a standard factor analysis in downweighting the contributions of low frequency words by a factor of $f(c_{ij}) = \min\left(1, \frac{c_{ij}}{c_{\text{cutoff}}}\right)^\alpha$. In the pretrained models, c_{cutoff} is set 100 and α is set to $\frac{3}{4}$. In the models we train ourselves, we set c_{cutoff} to 10 and retain the same α , since VALEX contains an order of magnitude fewer observations than the cooccurrence matrices pretrained GloVe is trained on.

¹⁵ Pretrained GloVe is available at <https://nlp.stanford.edu/projects/glove/>. We specifically use the uncased, 300 dimensional vectors trained on 42 billion words of Common Crawl.

¹⁶ Pretrained BERT models are available at <https://github.com/google-research/bert>. We specifically use the BERT-base-uncased models.

([CLS])—following standard practice (Devlin et al. 2018).¹⁷ This is analogous to using NBoW with GloVe embeddings in that the sentence start token contains information about all words in the sentence (along with their positions) due to the way the model is trained.

6.2 Experiments

We compute MAP estimates for the parameters of LDA using the default hyperparameters in the sklearn package, and we compute Maximum Likelihood Estimates (MLEs) for the parameters of the logistic factor analysis model and the GloVe representations we train on VALEX. For all three models, we consider numbers of latent components $K \in \{2, 5, 10, 15, 20, 25, 30, 35, 40, 45, 50\}$. For LDA and LFA, we additionally concatenate the predicted distributions over the subcategorization frames and the best-performing normalized distributions from the Dirichlet-Categorical ($\lambda = 0$) and Beta-Negative Binomial ($\beta = 1$) models, respectively.

As in Section 5, we regress the normalized acceptability judgments for each verb in each frame on each of these representations in a multivariate ridge regression—i.e. a linear regression with L2 regularization. To set the regularization parameter $\alpha \in \{0.01, 0.1, 0.2, 0.5, 1, 2, 5, 10\}$, we use a 10-fold cross-validation. To compute the generalizability of this model, we nest this cross-validation within another 10-fold cross-validation and compute the mean R^2 (variance explained) on the held out datasets in this outer cross-validation.

6.3 Results

Figure 8 plots the mean R^2 across the 10 cross-validation folds for each model and number of latent components. The black dashed line shows the best performing model from Section 5: Negative Binomial ($\beta = 1$).

We see that BERT is far and away the best performing model with LFA a distant second. In turn, LFA ($K = 10$) reliably outperforms the Beta-Negative Binomial model ($\beta = 1$) by approximately two points as well as the best-performing LDA ($K = 40$) and GloVe ($K = 20$) models by six points.

Figure 9 plots the mean R^2 across the 10 cross-validation folds for the BERT model, broken out by frame. We see that, in contrast to the analogous Figure 7, frames involving NP direct objects do not show systematically poor

¹⁷ We also experimented with extracting the embedding of the clause-embedding verb in the sentence. The results were the same.




Figure 8: Mean variance explained in normalized acceptability judgments in 10-fold/10-fold nested cross-validation for each model and number of latent components.

performance, suggesting that BERT may be able to capture some regularity about selection of a direct object that the models based on frame frequencies were not able to.

6.4 Discussion

At a high level, these results suggest two things. First, there is some amount of information in verbs' subcategorization frame frequency distributions that is not accessible directly from those distributions themselves—even after various forms of clever normalization. Accessing that information requires some amount of abstraction of the frequencies, confirming both hypothesis H1 and H2 from Section 2.

Second, the amount of extra information that can be gleaned from the subcategorization frame frequency distributions alone is relatively small—especially compared to the gains obtained in using models, such as BERT, that additionally have access to the particular lexical items that cooccur with a verb (see Grimshaw 1994; Pinker 1994; Resnik 1996; White et al. 2017c for reasons this might be true). But not just any model of lexical cooccurrence will do, since pretrained GloVe, which does have access to such cooccurrence statistics, is one of the worst performing models of all.




Figure 9: Variance explained by BERT model broken out by frame.

These results, as of yet, do not make strong commitments as to the nature of the abstraction that confirms H2. Part of this poor performance of the shallow (non-BERT) models may be a product of the heavy constraints that we place on the space of functions that we considered from abstractions to acceptabilities. But this makes the good performance of BERT even more surprising because, though state-of-the-art performance has been demonstrated on multiple NLP tasks using its embeddings, those models invariably learn nonlinear mappings from the embeddings to the quantity of interest—as is common practice for neural methods in NLP (see [Goldberg 2017](#) and references therein). While the results are consistent with ideas from linguistic theory and acquisition about what these abstractions might be like—BERT is typically thought to be rich enough that substantial syntactic/semantic information is present in it—it will require substantially more investigation to understand exactly how this model is predicting acceptability so well. We leave this as an open question for future work.

7 Conclusion

This paper addresses the question of how direct the relationship between well-formedness and linguistic experience is, focusing in particular on lexical knowledge. To do this, we developed the *bleaching method* for scaling standard acceptability judgment experiments to very large sets of verbs. After validating this method against more standard methods, we deployed it on 1,000 clause-embedding verbs in 50 syntactic frames to create the MegaAcceptability dataset, which is publicly available at megaattitude.io under the auspices of the MegaAttitude Project. Using this dataset, which we take to exhaust the set of clause-embedding verbs in English, we found that the relationship between acceptability and subcategorization frame frequency is surprisingly weak and that shallow abstractions of the data yield miniscule improvements in the prediction of acceptability. Deeper abstractions, however, can do surprisingly well at predicting acceptability, though we still see quite a bit of variation.

We take our results to imply that accounts of how knowledge of c-selection is acquired must posit something beyond simple smoothing or shallow factorization, as previous computational accounts have done (Alshahi & Stevenson 2008; Barak et al. 2012). One form this might take is to rely exclusively on deep, domain-general abstraction mechanisms, like BERT. Another is to enrich shallow domain-general factorization models with tunable domain-specific biases (White et al. 2017a). There is an inherent trade-off between these approaches: (i) Occam’s razor implores us to posit as few inherent biases as possible; but (ii) the data hungriness of deep abstraction mechanisms makes a strictly domain-general model suspect, unless it comes with biases specific to language learning. A hybrid approach is almost certainly necessary, and we believe the MegaAcceptability dataset will prove useful in evaluating the success of such approaches.

Supplementary files

The MegaAcceptability dataset is available at megaattitude.io.

Funding information

This research was funded by the following National Science Foundation grants: BCS-1748969/BCS-1749025 (*The MegaAttitude Project: Investigating selection and polysemy at the scale of the lexicon*), DDRIG BCS-1456013

(*Learning attitude verb meanings*), INSPIRE BCS-1344269 (*Gradient symbolic computation*) as well as the JHU Science of Learning Institute.

Acknowledgements

The authors wish to thank audiences at SALT 26, DGfS 2017, Johns Hopkins University, University of Rochester, and Stanford University for helpful discussion of this work.

Competing interests

The authors have no competing interests to declare.

Authors' contributions

White and Rawlins collaborated in the design of the materials for the MegaAcceptability dataset. White developed and implemented all models and analyses.

References

- Adger, David. 2003. *Core Syntax: A minimalist approach*. Oxford University Press.
- Agresti, Alan. 2014. *Categorical Data Analysis*. John Wiley & Sons.
- Alishahi, Afra & Suzanne Stevenson. 2008. A computational model of early argument structure acquisition. *Cognitive Science* 32(5). 789–834.
- Altmann, Gerry & Yuki Kamide. 1999. Incremental interpretation at verbs: Restricting the domain of subsequent reference. *Cognition* 73(3). 247–264.
- Anand, Pranav & Valentine Hacquard. 2013. Epistemics and attitudes. *Semantics and Pragmatics* 6(8). 1–59.
- Anand, Pranav & Valentine Hacquard. 2014. Factivity, belief and discourse. In Luka Crnič & Uli Sauerland (eds.), *The Art and Craft of Semantics: A Festschrift for Irene Heim*, vol. 1, 69–90. Cambridge, MA: MIT Working Papers in Linguistics.
- Aslin, Richard, Jenny Saffran & Elissa Newport. 1998. Computation of conditional probability statistics by 8-month-old infants. *Psychological Science* 9(4). 321–324.

- Barak, Libby, Afsaneh Fazly & Suzanne Stevenson. 2012. Modeling the acquisition of mental state verbs. In *Proceedings of the 3rd Workshop on Cognitive Modeling and Computational Linguistics*. 1–10. Montreal, Canada: Association for Computational Linguistics.
- Bard, Ellen Gurman, Dan Robertson & Antonella Sorace. 1996. Magnitude estimation of linguistic acceptability. *Language* 32–68.
- Blei, David M., Andrew Y. Ng & Michael I. Jordan. 2003. Latent dirichlet allocation. *The Journal of Machine Learning Research* 3. 993–1022.
- Bresnan, Joan. 2007. Is syntactic knowledge probabilistic? Experiments with the English dative alternation. In Sam Featherston & Wolfgang Sternefeld (eds.), *Roots: Linguistics in search of its evidential base*, vol. 96, 77–96. Walter de Gruyter.
- Bresnan, Joan, Anna Cueni, Tatiana Nikitina & R Harald Baayen. 2007. Predicting the dative alternation. In Gerlof Bouma, Irene Kramer & Joost Zwarts (eds.), *Cognitive Foundations of Interpretation*, 69–94. Chicago: University of Chicago Press.
- Briscoe, Ted & John Carroll. 1997. Automatic Extraction of Subcategorization from Corpora. In *Proceedings of the Fifth Conference on Applied Natural Language Processing (ANLC '97)*. 356–363. Stroudsburg, PA, USA: Association for Computational Linguistics. <https://doi.org/10.3115/974557.974609>. <https://doi.org/10.3115/974557.974609>. Event-place: Washington, DC.
- Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- Chomsky, Noam. 1973. Conditions on transformations. In S. Anderson & P. Kiparsky (eds.), *A Festschrift for Morris Halle*, 232–286. New York: Holt, Rinehart, & Winston.
- Church, Kenneth W. & William A. Gale. 1995. Poisson mixtures. *Natural Language Engineering* 1(02). 163–190.
- Church, Kenneth Ward & Patrick Hanks. 1990. Word association norms, mutual information, and lexicography. *Computational Linguistics* 16(1). 22–29.
- Clark, Alexander, Gianluca Giorgolo & Shalom Lappin. 2013a. Statistical representation of grammaticality judgements: The limits of n-gram models. In *Proceedings of the Fourth Annual Workshop on Cognitive Modeling and Computational Linguistics*. 28–36.
- Clark, Alexander, Gianluca Giorgolo & Shalom Lappin. 2013b. Towards a statistical model of grammaticality. In *Proceedings of the 35th annual conference of the cognitive science society*. 2064–2069.

- Clark, Alexander & Shalom Lappin. 2011. *Linguistic Nativism and the Poverty of the Stimulus*. Wiley-Blackwell.
- Conneau, Alexis, Germán Kruszewski, Guillaume Lample, Loïc Barrault & Marco Baroni. 2018. What you can cram into a single vector: Probing sentence embeddings for linguistic properties. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, vol. 1. 2126–2136. Association for Computational Linguistics. <https://aclanthology.info/papers/P18-1198/p18-1198>.
- Devlin, Jacob, Ming-Wei Chang, Kenton Lee & Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv:1810.04805 [cs]* <http://arxiv.org/abs/1810.04805>. ArXiv: 1810.04805.
- Dunning, Ted. 1993. Accurate methods for the statistics of surprise and coincidence. *Computational Linguistics* 19(1). 61–74.
- Featherston, Sam. 2005. Magnitude estimation and what it can do for your syntax: Some wh-constraints in German. *Lingua* 115(11). 1525–1550.
- Featherston, Sam. 2007. Data in generative grammar: The stick and the carrot. *Theoretical Linguistics* 33(3). 269–318.
- Fillmore, Charles John. 1970. The grammar of hitting and breaking. In R.A. Jacobs & P.S. Rosenbaum (eds.), *Readings in English Transformational Grammar*, 120–133. Waltham, MA: Ginn.
- Fine, Alex B. & T. Florian Jaeger. 2013. Evidence for implicit learning in syntactic comprehension. *Cognitive Science* 37(3). 578–591.
- Fisher, Cynthia. 1994. Structure and meaning in the verb lexicon: Input for a syntax-aided verb learning procedure. *Language and Cognitive Processes* 9(4). 473–517.
- Fisher, Cynthia, Yael Gertner, Rose M. Scott & Sylvia Yuan. 2010. Syntactic bootstrapping. *Wiley Interdisciplinary Reviews: Cognitive Science* 1(2). 143–149.
- Fisher, Cynthia, Henry Gleitman & Lila R. Gleitman. 1991. On the semantic content of subcategorization frames. *Cognitive Psychology* 23(3). 331–392.
- Fisher, Cynthia, D. Geoffrey Hall, Susan Rakowitz & Lila Gleitman. 1994. When it is better to receive than to give: Syntactic and conceptual constraints on vocabulary growth. *Lingua* 92. 333–375.
- Garnsey, Susan M., Neal J. Pearlmutter, Elizabeth Myers & Melanie A. Lockety. 1997. The contributions of verb bias and plausibility to the comprehension of temporarily ambiguous sentences. *Journal of Memory and Language* 37(1). 58–93.

- Gibson, Edward & Evelina Fedorenko. 2010. Weak quantitative standards in linguistics research. *Trends in Cognitive Science* 14(6). 233–234.
- Gibson, Edward & Evelina Fedorenko. 2013. The need for quantitative methods in syntax and semantics research. *Language and Cognitive Processes* 28(1-2). 88–124.
- Gillette, Jane, Henry Gleitman, Lila Gleitman & Anne Lederer. 1999. Human simulations of vocabulary learning. *Cognition* 73(2). 135–176.
- Gleitman, Lila. 1990. The structural sources of verb meanings. *Language Acquisition* 1(1). 3–55.
- Gleitman, Lila R., Kimberly Cassidy, Rebecca Nappa, Anna Papafragou & John C. Trueswell. 2005. Hard words. *Language Learning and Development* 1(1). 23–64.
- Goldberg, Yoav. 2017. *Neural Network Methods for Natural Language Processing*, vol. 37 (Synthesis Lectures on Human Language Technologies). Morgan & Claypool Publishers.
- Grimshaw, Jane. 1979. Complement selection and the lexicon. *Linguistic Inquiry* 10(2). 279–326.
- Grimshaw, Jane. 1981. Form, function and the language acquisition device. In C.L. Baker & John J. McCarthy (eds.), *The Logical Problem of Language Acquisition*, 165–182. Cambridge, MA: MIT Press.
- Grimshaw, Jane. 1990. *Argument Structure*. Cambridge, MA: MIT Press.
- Grimshaw, Jane. 1994. Lexical reconciliation. *Lingua* 92. 411–430.
- Gruber, Jeffrey Steven. 1965. *Studies in Lexical Relations*. Cambridge, MA: Massachusetts Institute of Technology dissertation.
- Hacquard, Valentine & Alexis Wellwood. 2012. Embedding epistemic modals in English: A corpus-based study. *Semantics and Pragmatics* 5(4). 1–29.
- Hale, John. 2001. A Probabilistic Earley Parser As a Psycholinguistic Model. In *Proceedings of the Second Meeting of the North American Chapter of the Association for Computational Linguistics on Language Technologies* (NAACL '01). 1–8. Stroudsburg, PA, USA: Association for Computational Linguistics.
- Hankamer, Jorge & Ivan Sag. 1976. Deep and surface anaphora. *Linguistic Inquiry* 391–428.
- Hofmeister, Philip & Ivan A. Sag. 2010. Cognitive constraints and island effects. *Language* 86(2). 366–415.
- Iyyer, Mohit, Varun Manjunatha, Jordan Boyd-Graber & Hal Daumé III. 2015. Deep unordered composition rivals syntactic methods for text classification. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on*

- Natural Language Processing (Volume 1: Long Papers)*, vol. 1. 1681–1691.
- Jackendoff, Ray. 1972. *Semantic Interpretation in Generative Grammar*. Cambridge, MA: MIT Press.
- Kako, Edward. 1997. Subcategorization Semantics and the Naturalness of Verb-Frame Pairings. *University of Pennsylvania Working Papers in Linguistics* 4(2). 11.
- Kann, Katharina, Alex Warstadt, Adina Williams & Samuel R. Bowman. 2019. Verb Argument Structure Alternations in Word and Sentence Embeddings. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2019*. 287–297. <https://doi.org/10.7275/q5js-4y86>. <https://aclweb.org/anthology/papers/W/W19/W19-0129/>.
- Keller, Frank. 2000. *Gradience in grammar: Experimental and computational aspects of degrees of grammaticality*: University of Edinburgh dissertation.
- Kipper-Schuler, Karin. 2005. *VerbNet: A broad-coverage, comprehensive verb lexicon*: University of Pennsylvania dissertation.
- Korhonen, Anna, Yuval Krymolowski & Ted Briscoe. 2006. A large subcategorization lexicon for natural language processing applications. In *Proceedings of LREC*, vol. 6.
- Kush, Dave, Terje Lohndal & Jon Sprouse. 2018. Investigating variation in island effects. *Natural Language & Linguistic Theory* 36(3). 743–779.
- Landau, Barbara & Lila R. Gleitman. 1985. *Language and Experience: Evidence from the Blind Child*, vol. 8. Cambridge, MA: Harvard University Press.
- Lasnik, Howard. 1989. On certain substitutes for negative data. In *Learnability and Linguistic Theory*, 89–105. Springer.
- Lau, Jey Han, Alexander Clark & Shalom Lappin. 2017. Grammaticality, Acceptability, and Probability: A Probabilistic View of Linguistic Knowledge. *Cognitive Science* 41(5). 1202–1241. <https://doi.org/10.1111/cogs.12414>. <https://onlinelibrary.wiley.com/doi/abs/10.1111/cogs.12414>.
- Lederer, Anne, Henry Gleitman & Lila Gleitman. 1995. Verbs of a feather flock together: semantic information in the structure of maternal speech. In M. Tomasello & W.E. Merriman (eds.), *Beyond Names for Things: Young Children's Acquisition of Verbs*, 277–297. Hillsdale, NJ: Lawrence Erlbaum.
- Levin, Beth. 1993. *English Verb Classes and Alternations: A preliminary investigation*. Chicago: University of Chicago Press.
- Levin, Beth & Malka Rappaport Hovav. 2005. *Argument Realization*. Cambridge: Cambridge University Press.

- Levy, Omer & Yoav Goldberg. 2014. Neural word embedding as implicit matrix factorization. In *Advances in Neural Information Processing Systems*. 2177–2185.
- Levy, Roger. 2008. Expectation-based syntactic comprehension. *Cognition* 106(3). 1126–1177.
- Lidz, Jeffrey, Henry Gleitman & Lila Gleitman. 2004. Kidz in the 'hood: Syntactic bootstrapping and the mental lexicon. In D.G. Hall & S.R. Waxman (eds.), *Weaving a Lexicon*, 603–636. Cambridge, MA: MIT Press.
- Linzen, Tal & T. Florian Jaeger. 2016. Uncertainty and expectation in sentence processing: evidence from subcategorization distributions. *Cognitive science* 40(6). 1382–1411.
- Maye, Jessica, Janet F Werker & LouAnn Gerken. 2002. Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition* 82(3). B101–B111.
- McRae, Ken, Michael J. Spivey-Knowlton & Michael K. Tanenhaus. 1998. Modeling the influence of thematic fit (and other constraints) in on-line sentence comprehension. *Journal of Memory and Language* 38(3). 283–312.
- Naigles, L., Henry Gleitman & Lila Gleitman. 1993. Syntactic bootstrapping and verb acquisition. In Esther Dromi (ed.), *Language and Cognition: A Developmental Perspective*. (Human Development Series), Norwood, NJ: Ablex.
- Naigles, Letitia. 1990. Children use syntax to learn verb meanings. *Journal of Child Language* 17(2). 357–374.
- Naigles, Letitia. 1996. The use of multiple frames in verb learning via syntactic bootstrapping. *Cognition* 58(2). 221–251.
- Papafragou, Anna, Kimberly Cassidy & Lila Gleitman. 2007. When we think about thinking: The acquisition of belief verbs. *Cognition* 105(1). 125–165.
- Pearl, Lisa & Jon Sprouse. 2013. Syntactic islands and learning biases: Combining experimental syntax and computational modeling to investigate the language acquisition problem. *Language Acquisition* 20(1). 23–68.
- Pedregosa, Fabian, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot & Edouard Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12. 2825–2830.
- Pennington, Jeffrey, Richard Socher & Christopher D. Manning. 2014. Glove: Global Vectors for Word Representation. In *Proceedings of the*

- 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP). 1532–1543. Doha, Qatar: Association for Computational Linguistics.
- Pesetsky, David. 1982. *Paths and Categories*: Massachusetts Institute of Technology dissertation.
- Pesetsky, David. 1991. Zero syntax: vol. 2: Infinitives.
- Pinker, Steven. 1984. *Language Learnability and Language Development*. Harvard University Press.
- Pinker, Steven. 1989. *Learnability and Cognition: The Acquisition of Argument Structure*. Cambridge, MA: MIT Press.
- Pinker, Steven. 1994. How could a child use verb syntax to learn verb semantics? *Lingua* 92. 377–410.
- Poliak, Adam, Aparajita Haldar, Rachel Rudinger, J. Edward Hu, Ellie Pavlick, Aaron Steven White & Benjamin Van Durme. 2018. Collecting Diverse Natural Language Inference Problems for Sentence Representation Evaluation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Brussels, Belgium: Association for Computational Linguistics.
- Rawlins, Kyle. 2013. About ‘about’. *Semantics and Linguistic Theory* 23. 336–357.
- Resnik, Philip. 1996. Selectional constraints: An information-theoretic model and its computational realization. *Cognition* 61(1). 127–159.
- Ross, John Robert. 1972. Act. In Donald Davidson & Gilbert Harman (eds.), *Semantics of Natural Language*, 70–126. Springer.
- Ross, John Robert. 1973. Slifting. In Maurice Gross, Morris Halle & Marcel-Paul Schützenberger (eds.), *The Formal Analysis of Natural Languages*, 133–169. The Hague: Mouton de Gruyter.
- Saffran, Jenny, Richard Aslin & Elissa Newport. 1996a. Statistical learning by 8-month-old infants. *Science* 274(5294). 1926–1928.
- Saffran, Jenny, Elissa Newport & Richard Aslin. 1996b. Word segmentation: The role of distributional cues. *Journal of Memory and Language* 35(4). 606–621.
- Schütze, Carson T. & Jon Sprouse. 2014. Judgment data. In Robert J. Podesva & Devyani Sharma (eds.), *Research Methods in Linguistics*, 27–50. Cambridge University Press.
- Snedeker, Jesse & Lila Gleitman. 2004. Why it is hard to label our concepts. In D. Geoffrey Hall & Sandra R. Waxman (eds.), *Weaving a Lexicon*, 257–294. Cambridge, MA: MIT Press.
- Sorace, Antonella & Frank Keller. 2005. Gradience in linguistic data. *Lingua* 115(11). 1497–1524.

- Spivey-Knowlton, Michael & Julie C Sedivy. 1995. Resolving attachment ambiguities with multiple constraints. *Cognition* 55(3). 227–267.
- Sprouse, Jon. 2007. Continuous acceptability, categorical grammaticality, and experimental syntax. *Biolinguistics* 1. 123–134.
- Sprouse, Jon. 2011. A validation of Amazon Mechanical Turk for the collection of acceptability judgments in linguistic theory. *Behavioral Research* 43. 155–167.
- Sprouse, Jon & Diogo Almeida. 2013. The empirical status of data in syntax: A reply to Gibson and Fedorenko. *Language and Cognitive Processes* 28(3). 222–228.
- Sprouse, Jon, Carson T. Schütze & Diogo Almeida. 2013. A comparison of informal and formal acceptability judgments using a random sample from Linguistic Inquiry 2001–2010. *Lingua* 134. 219–248.
- Sprouse, Jon, Matt Wagers & Colin Phillips. 2012. A test of the relation between working-memory capacity and syntactic island effects. *Language* 88(1). 82–123.
- Sprouse, Jon, Beracah Yankama, Sagar Indurkha, Sandiway Fong & Robert C. Berwick. 2018. Colorless green ideas do sleep furiously: gradient acceptability and the nature of the grammar. *The Linguistic Review* 35(3). 575–599.
- Suzuki, Jun & Masaaki Nagata. 2015. A unified learning framework of skip-grams and global vectors. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, vol. 2. 186–191.
- Trueswell, John C., Michael K. Tanenhaus & Christopher Kello. 1993. Verb-specific constraints in sentence processing: separating effects of lexical preference from garden-paths. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 19(3). 528.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser & Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*. 5998–6008.
- Wang, Alex, Amapreet Singh, Julian Michael, Felix Hill, Omer Levy & Samuel R. Bowman. 2018. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. *arXiv:1804.07461* <http://arxiv.org/abs/1804.07461>.
- Warstadt, Alex, Amanpreet Singh & Samuel R Bowman. 2018. Neural Network Acceptability Judgments. *arXiv preprint arXiv:1805.12471* .

- Wells, Justine B., Morten H. Christiansen, David S. Race, Daniel J. Acheson & Maryellen C. MacDonald. 2009. Experience and sentence processing: Statistical learning and relative clause comprehension. *Cognitive Psychology* 58(2). 250–271.
- White, Aaron Steven. 2015. *Information and Incrementality in Syntactic Bootstrapping*. College Park, MD: University of Maryland dissertation.
- White, Aaron Steven. 2019. Nothing's wrong with believing (or hoping) whether.
- White, Aaron Steven, Rachel Dudley, Valentine Hacquard & Jeffrey Lidz. 2014. Discovering classes of attitude verbs using subcategorization frame distributions. In Hsin-Lun Huang, Ethan Poole & Amanda Rysling (eds.), *Proceedings of the 43rd annual meeting of the North East Linguistic Society*. 249–260.
- White, Aaron Steven, Valentine Hacquard & Jeffrey Lidz. 2017a. The labeling problem in syntactic bootstrapping: Main clause syntax in the acquisition of propositional attitude verbs. In Kristen Syrett & Sudha Arunachalam (eds.), *Semantics in Acquisition* (Trends in Language Acquisition Research (TiLAR)), in press. John Benjamins Publishing Company.
- White, Aaron Steven, Valentine Hacquard & Jeffrey Lidz. 2018. Semantic Information and the Syntax of Propositional Attitude Verbs. *Cognitive Science* 42(2). 416–456.
- White, Aaron Steven, Pushpendre Rastogi, Kevin Duh & Benjamin Van Durme. 2017b. Inference is Everything: Recasting Semantic Resources into a Unified Evaluation Framework. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, vol. 1. 996–1005.
- White, Aaron Steven & Kyle Rawlins. 2016. A computational model of S-selection. *Semantics and Linguistic Theory* 26. 641–663.
- White, Aaron Steven, Philip Resnik, Valentine Hacquard & Jeffrey Lidz. 2017c. The contextual modulation of semantic information.
- Zhou, M. & L. Carin. 2015. Negative Binomial Process Count and Mixture Modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 37(2). 307–320. <https://doi.org/10.1109/TPAMI.2013.211>.
- Zhou, Mingyuan. 2018. Nonparametric Bayesian Negative Binomial Factor Analysis. *Bayesian Analysis* 13(4). 1065–1093. <https://doi.org/10.1214/17-BA1070>. <https://projecteuclid.org/euclid.ba/1510801993>.
- Zhou, Mingyuan, Lauren Hannah, David Dunson & Lawrence Carin. 2012. Beta-Negative Binomial Process and Poisson Factor Analysis. In *Artificial Intelligence and Statistics*. 1462–1471. <http://proceedings.mlr.press/v22/zhoul2c.html>.

A Validation Normalization

To normalize the acceptability judgments collected in the replication experiment (Section 3), we fit an ordinal (linked logit) mixed effects model to the ratings from both datasets, with fixed effects for VERB, FRAME, and their interaction and random unconstrained cutpoints for each participant (Agresti 2014).

We parameterize this model by assuming that each pairing of a verb v and a frame f is associated with some true real-valued acceptability a_{vf} and that each participant p is associated with a way of binning these real-valued acceptability judgments, where each bin corresponds to a particular scale rating. These bins are defined by cutpoints c_p for each participant p , where the bin corresponding to the worst rating—in our case, 1—is to the interval $(-\infty, c_{p1}]$ and the bin corresponding to the best ratings—in our case, 7—is the interval (c_{p6}, ∞) . For all other ratings i , the corresponding bin for participant p is $(c_{p(i-1)}, c_{pi}]$. Alternatively, we say that $c_{p0} = -\infty$ and $c_{p7} = \infty$ for all participants p .

Similar to a binary logistic regression, which we can think of as having just two bins defined by a single cut point, we define the probability of giving a particular participant p giving a response r_{pvf} to verb v and frame f (assuming true acceptability a_{vf}) based on these cutpoints. First, we define the cumulative density function.

$$\mathbb{P}(r_{pvf} \leq i) = \text{logit}^{-1}(c_{pi} - a_{vf})$$

Then, from the cumulative density function, we can reconstruct the probability for each response i .

$$\mathbb{P}(r_{pvf} = i) = \mathbb{P}(r_{pvf} \leq i) - \mathbb{P}(r \leq (i - 1))$$

From this, the (log-)likelihood of the data immediately follows. In addition, we assume an Exponential prior on the distances between adjacent cutpoints for each participant. We estimate the parameter λ of this prior over the course of fitting the model, thus making this a mixed effects model.

We estimate the true acceptabilities A for all verb-frame pairs and the cutpoints for all participants C by using gradient descent to maximize the sum of the likelihood of the data, the prior on cutpoints, and a small smoothing term, under the constraint that the mean of the third cutpoint is locked to zero, thus making the parameters identifiable. All analyses use the resulting acceptabilities A .

B MegaAcceptability Normalization

As for our replication of [White et al.](#)'s dataset, to measure interannotator agreement, we compute the Spearman rank correlation between the responses for each pair of participants that did the same list. This yields a mean correlation of 0.415 (95% CI: [0.412, 0.418]), which is more than 10 points lower than the agreement obtained in the replication.

Part of the reason for this is likely that [White et al.](#)'s—and consequently, our replication—contained mostly high frequency verbs, whereas the MegaAttitude dataset contains many low frequency verbs that participants are likely less certain about. Another source of this low agreement is likely a higher rate of poor participants in these data. This is evidenced by the fact that the agreement scores have nontrivial left skew, with a median correlation of 0.455 (95% CI: [0.450, 0.460]).

To mitigate the effect of poor participants, we downweight the influence of those participants' responses in constructing the normalized acceptability for each verb-frame pair. Our approach amounts to using the ordinal model-based normalization described in Section 3, but weighting the likelihood of the ordinal model by participant quality scores on [0, 1] derived from pairwise agreement between participants.¹⁸

One simple way of deriving such a score would be to take the mean interannotator agreement for all pairs an participant occurs in and then normalize those means to lie on [0, 1]. This simple approach is problematic, however, since most participants only rate one list and so, if a good participant rates a list rated by mostly bad participants, that participant will be assigned a low quality score.

To address this issue, we derive an participant quality score by first fitting a linear mixed effects model with random intercepts for participant and list to the Spearman rank correlations, then extracting the Best Linear Unbiased Predictors for the participant intercepts. We then z -score these scores and squash them to [0, 1] using the normal cumulative distribution function. This participant quality score is thus high when an participant tends to show high agreement with other participants, adjusting for the effect of the particular list.

¹⁸ This procedure differs from the procedure used by [White & Rawlins \(2016\)](#) for the same dataset in that they filter participants with agreement under a particular threshold. Our approach can be seen as a soft version of their thresholding approach, wherein the influence of participants' responses drops off smoothly as a function of their overall agreement with other participants.

We combine these log-likelihoods into single variability score by computing their mean, weighted by the participant quality score of the participant who provided the rating.

C Method for Adding Verbs

Seven verbs—*manage*, *fail*, *neglect*, *refuse*, *help*, *opt*, *deserve*—were unintentionally excluded from our large-scale experiment due to a coding error. We do not include these verbs in the analyses presented in the body of the paper because it is nontrivial, within the method described above, to build lists that include them without reconducting a large portion of the study.

Because we would like to have data about these verbs for future work, we instead evaluate an alternative method for adding missing verbs to our dataset. In this method, we test a single verb in all of the frames of interest within the same list.

To evaluate how this method compares to a method wherein verbs are intermixed, we constructed a list for each of the 30 pilot verbs from Section 3 paired with each of the 50 frames from the MegaAcceptability data (Section 4). We find that the average pairwise agreement by list is actually higher in this experiment than in our original replication, with a mean Spearman rank correlation of 0.65 (95% CI = [0.63, 0.67]). This higher agreement is due to a few annotators who did many lists showing high agreement with each other, since when we fit the linear mixed effects model described in Appendix B to these correlations, we find an expected correlation of 0.54, which is very close to the correlation found in our validation experiments (Section 3).

To compare the agreement between the normalized ratings from the MegaAcceptability dataset to those from this one-verb-per-list dataset, we applied the normalization used for the MegaAcceptability dataset (Appendix B) to these data and then computed the correlation by verb. Figure 10 shows this agreement which is extremely high across all verbs.

We take this as an indicator that testing one verb per list—at least in this set of frames—produces results that are just as valid as intermixing verbs. We thus tested the seven verbs above using this method and include them with the full dataset.




Figure 10: Correlation by verb between mean normalized verb-frame acceptability in MegaAcceptability and one-verb-per-list dataset.