

## Author



# Frequency of What: How Simple is the Story of Syntax Acquisition?

**Amanda N. Nili**

*English, Psychology & Social Behavior*

Amanda Nili met her mentor at a Psi Chi induction ceremony, and Professor Pearl offered her the opportunity to participate in the Computation of Language Laboratory as a means of completing her Linguistics minor. Her project is unique in that it is an empirical analysis of the validity of a fairly ubiquitous assumption about syntax acquisition in the field of linguistics. The work is also characterized by a discussion of the implications of this assumption for the methodology of linguistics researchers. After graduation, Amanda plans to apply to graduate programs in Clinical Psychology and pursue a career in academia and clinical work.

## Abstract

This study investigates an implicit assumption held by some language acquisition researchers that the frequency of linguistic structures correlates with how acceptable native speakers find these structures. The basic foundation of this assumption is that speakers have learned which structures are acceptable by how often they occur in the input. The authors compare Sprouse and Almeida's acceptability judgment data of a wide range of linguistic structures with the frequencies of those same structures in child-directed speech. The utterances studied are from six American English corpora in the CHILDES database, which contains 170,263 child-directed utterances. Notably, the frequencies of very surface structural representations were not found to correlate with acceptability data. Future work can investigate whether acceptability judgments correlate with the frequencies of more abstract linguistic categories (e.g., using semantic abstraction or using additional syntactic abstraction), which would suggest that speakers use input frequency as a foundation for their linguistic knowledge, but do so in a more sophisticated way.

## Faculty Mentor



Amanda Nili's research project investigates one of those assumptions that is often implicit in the field of language science, but rarely explicitly mentioned. As such, this assumption has subtly shaped the direction of research for theories of how we represent linguistic knowledge in our minds, how we deploy that linguistic knowledge in real time, and how we learn that linguistic knowledge in the first place. Amanda's research not only highlights that this assumption exists, but pursues one of the testable predictions of it. Her findings are illuminating and suggest that this assumption is not obviously a good one when it comes to how we learn our linguistic knowledge.

## Key Terms

- ◆ Frequency-Acceptability Gap
- ◆ Linguistic Behavior
- ◆ Syntax Acquisition

**Lisa S. Pearl**

*School of Social Sciences*

## Introduction

One explanation for why certain structures in a language sound better (*i.e.*, are more “grammatical”) than others to native speakers is that the better-sounding structures have been encountered more frequently. This simple relationship between frequency and acceptability—namely, that the former is the only thing that really determines the latter—has been considered viable by the linguistics community for quite some time, although it has only recently been investigated empirically. More importantly, from the standpoint of language acquisition, since the acceptability of linguistic structures varies somewhat from language to language, these grammaticality preferences must be learned. Notably, the linguistic input that children encounter, sometimes called “motherese” or “child-directed speech,” is linguistically different from the linguistic input that adults encounter (typically called “adult-directed speech”). Thus, to assess this simple explanation of how structure acceptability preferences are learned, this study builds on previous work with adult-directed speech and examines the relationship between structure acceptability and the frequency of structures in child-directed speech.

## Grammaticality and Acceptability

With the aim of assessing grammaticality of structures, we have used acceptability judgment data collected by Sprouse and Almeida (2012). These acceptability scores are intended to represent an individual’s grammar of the language (*i.e.* North American English), although Pearl and Sprouse (2013) note that such acceptability scores include additional factors, and are not limited to the theoretical grammar of that language.<sup>1</sup> Rather, these scores are a reflection of linguistic behavior—naïve native speakers are revealing preferences about actual linguistic knowledge, which should have a basis in the grammar of the language studied as the grammar is meant to be a description of native speaker linguistic knowledge. Notably, there can be differences between what is grammatical and what is acceptable. Consider, for example, the phenomenon of center embedding of relative clauses in North American English. The grammar of this language permits center embedding of relative clauses (*e.g.* “The penguin *that Elise saw* was cute.”), and theoretically permits this multiple times in a single sentence (*e.g.* “The penguin *that Elise who Olaf who Rachel loves knows saw* was cute.”). The difference between these two examples is one of acceptability—while both are technically grammatical

and demonstrate center embedding of relative clauses, the latter example (which has three center-embedded clauses), is less acceptable and typically impossible for a native speaker to follow. That is, we have trouble parsing the multiple subjects in “The penguin *that Elise who Olaf who Rachel loves knows saw* was cute.”, although this structure does not violate any of the theoretical grammar rules of our language. Because it is actual linguistic behavior that matters for a theory of language acquisition, we use acceptability scores as a stand-in for grammaticality, recognizing that additional factors matter for acceptability. While factors beyond structure do matter for acceptability, we might expect that if a simple relationship exists between structural frequency and grammaticality, we should see this relationship reflected in acceptability scores as well.

## The Simple and Not-So-Simple Stories

The simple story asserts that our brains do little more than simple statistical analyses with the base frequencies of the utterances we encounter. However, this theory is not empirically supported; in fact, the simple story of syntax acquisition is more of a hypothesis, one with very important implications for the methods by which linguistic research is conducted.

Alternatively, the not-so-simple story is our collective term for several different theories about why the base frequencies of utterance occurrences in a corpus do not match how acceptable speakers find these utterances. These theories look at more complicated ways of studying frequency, and explain the appearance of what Kempen and Harbusch (2005) call a “frequency-grammaticality gap” by offering different—typically more abstract—accounts for the ways that our brains process the structures we hear.

Below, we expand on the core aspects of the simple and not-so-simple stories for the relationship between frequency and acceptability of utterances, and review samples of literature indicating which theory different researchers (often implicitly) believe.

### *The Simple Story*

The simple story is that our brains take minimally abstracted information (*e.g.*, parsing sentences as parts of speech, rather than structural syntactic categories) from all linguistic input. Then, the story goes, we calculate a frequency score from that information to determine how “good” we think these perceived utterances are. The simple theory is that base frequency of structures correlates perfectly

1. Pearl and Sprouse (2013) note that these “other factors...include semantic plausibility, lexical properties, and parsing difficulty.” The measures designed to prevent semantic interference with the data’s reflection of grammaticality are discussed later in the paper.

with acceptability of those same structures. If this theory were true, then we could create plots like Figures 1 and 2, where frequency is plotted against acceptability, and expect a perfect positive linear correlation, as shown in Figure 1, so that data would only appear in Quadrants I and III, as in Figure 2. In particular, low frequency utterances (-Frequency) should have low acceptability (-Acceptability) while high frequency utterances (+Frequency) should have high acceptability (+Acceptability).




Figure 1  
The Simple Story's Prediction




Figure 2  
Where Data Should Fall (Simple Story)

There is more to the simple theory than this theorized relationship (where the Pearson correlation coefficient,  $r$ , which measures how related two variables are, is equal to 1). While a perfect correlation (or even a very high correlation, scoring above .70) would be a great support for the simple theory, the most important thing to understand is not the prediction that  $r = 1$ ; that value is the goal of every theory discussed in this paper. Instead, the most important

component of both the simple and not-so-simple stories of grammar acquisition is how each define the two variables, or axes on the above figures. There are multiple ways of considering the frequency of structures in corpora, and there are many ways to search for frequency in these datasets. The simple theory looks at surface forms of structures (*i.e.* those that we can collect without having to abstract too much information from the utterance). The most superficial approach might be that our brains collect individual words and tabulate the frequency with which one word follows another in exactly the order the words appear in a given utterance. Were this version true, we would find utterances acceptable only because we had heard that exact series of words before. To illustrate: “The dog barks” would be acceptable only because we have heard “The dog barks” before, and many more times than “The dog speaks.” To be sure, this is one example in which the semantics matter too (*i.e.* barking is an action exclusive to dogs and seals when not speaking figuratively, while speaking is an action less frequently attributed to dogs).

Still, this simple story disregards the function of syntax, and the fact that we produce novel combinations of words often based on the grammatical categories into which those words fall (*e.g.*, noun, verb). For this reason, and because no authors reviewed in this paper assert that frequency of individual words is responsible for utterance acceptability, we investigate a slightly more sophisticated version of the simple story, one which abstracts words into their grammatical categories. For example, “The dog barks” becomes a simple subject composed of a definite article and singular noun, followed by an intransitive verb (*i.e.*,  $\text{det}_{\text{def}} + \text{noun}_{\text{sg}} + \text{vb}_{\text{intrans}}$ ). Such a surface form is a way of studying syntax that incorporates fairly little abstraction.

#### *Who Believes it and Who Doesn't*

The simple story of syntax acquisition—that we analyze the most surface forms of utterances we hear (making minimal abstraction about them) and then translate the calculated frequency of those surface forms into an acceptability score—is not supported by empirical studies on the subject. The frequency of utterance appearance in corpora (when we abstract as little as possible from each utterance) does not correlate perfectly or even well with the acceptability scores that naïve native speakers report. What is important to note here is that all studies comparing frequency with acceptability find some form of discrepancy between those data, regardless of the level of abstraction at which utterances are studied. Those who reject a simple story of syntax acquisition account for the discrepancy by providing theories of different language learning and production

models to account for the divergent acceptability scores. In contrast, those who accept the simple story do so by criticizing the methods employed by those researchers that find a discrepancy. This is how Kempen and Harbusch (2005) assert an implicit version of the simple theory: they object to findings of a discrepancy between acceptability and frequency in a study by Keller (2001), in which some structures have average acceptability scores yet do not appear even once in the corpus. Specifically, Kempen and Harbusch accept that the structures are not in any of Keller’s corpora, but they believe that the naïve native speakers who provided acceptability scores did so erroneously, scoring utterances as ungrammatical that were in reality acceptable, and vice versa (Kempen & Harbusch 2005). This bad-method accusation is an important problem for the field of linguistics, as naïve native speakers are fairly popular sources for acceptability data (Sprouse & Almeida 2012). Furthermore, as discussed earlier, the difference between grammaticality and acceptability makes it more useful to study the acceptability data that speakers can provide—who better to ask about how naïve native speakers judge linguistic data than those very speakers?

Ultimately, scientists in every field are tasked to account for the data they find—for Kempen and Harbusch, this account is not a modification of their (simple) theory, but rather an attack on methodology. Below, we explore a few alternative explanations for the gap between frequency and acceptability. Unlike the Kempen and Harbusch (2005) account, these studies offer explanations of syntax acquisition which account for the empirical data, rather than arguing with its method of collection.

*The Not-So-Simple Example Story: Linear Optimality Theory.* Keller (2001)—in a study of various grammatical constraints and their violations in English, German, and Greek—finds a gap between the frequency values of produced structures and the acceptability data collected. Keller’s data are not limited to a single syntactic phenomenon, but instead cover a broad range of cross-linguistic grammar constraints. These constraints are tested in the “correct” (*i.e.* unviolated) form, in addition to testing the acceptability and frequency of their violated counterparts. For example, a test of Subject-Verb Agreement would look something like this:

Trish has painted a picture of Arthur. (1)

Trish have painted a picture of Arthur. (2)

In this example, Utterance 1 represents a correct version, while Utterance 2 is the violated form. Keller’s participants make acceptability judgments about both of these utterances,

and Keller sorts different constraints according to how unacceptable the violation of each constraint is.

Keller accounts for the “gap” between unviolated (highly acceptable) structures and low corresponding frequency as evidence for Linear Optimality Theory, a theory that groups constraints into types and assumes a gradient grammaticality dependent on something more than the base frequencies at which a structure is produced.<sup>2</sup> Linear Optimality Theory (Keller 2001) is the theoretical accounting for discrepancies between acceptability and frequency data by sorting grammatical constraint into two types, of varying significance when violated: “hard” and “soft” constraints. To account for the apparent lack of relationship between acceptability and frequency, Keller posits a “soft constraint” that can be violated and receive a lower acceptability score, but will still be produced somewhat frequently. One example of this is Definite Article Use (where the use of indefinite articles when a definite article would be more appropriate is a violation of the soft constraint).

#### Definite Article Use

Which friend has Trish painted a picture of? (3)

Which friend has Trish painted the picture of? (4)

Although this is not a violation of a hard constraint, according to Keller, the asterisk notation is used to clarify which example violates (with mild acceptability) the constraint discussed here. Clear notation is necessary because soft constraint violations are not as offensive as those of hard constraints [as in (2) above]. However, Keller finds that the utterance in (4) is less acceptable to naïve native speakers.

In contrast, a hard constraint accounts for low frequency utterances that are unacceptable: these are the structures that are rarely, if ever, produced. Subject-Verb Agreement (featured on the previous page) is a hard constraint. The constraint violation of (2) in this example should be much more obvious to the reader, as it is to participants in Keller’s study. This constraint was found to be “hard”—*i.e.*, resulting in great unacceptability when violated—cross-linguistically. In fact, Keller’s study finds minimal cross-linguistic variation of hard constraints.

2. The production frequency of a structure is relevant to this study as a means of examining the linguistic input a child receives. For our purposes, “production” implies frequency of linguistic perception of those same structures, and so frequency of production will be used to discuss linguistic input in the later sections of this paper. However, the age of audience for which speech is produced (*i.e.*, adult-directed versus child-directed) is further explored later in this paper.



*The Not-So-Simple Story: Response to Methodological Critiques.* As alluded to earlier in this section, Kempen and Harbusch's (2005) claim of bad methodology creating a gap between frequency and acceptability has important implications for the field of linguistics research. Claims against methodology are not new, nor are they yet resolved; debate over rigor in experimentation is an ongoing topic of interest in this field (Hickok 2010). Specifically, the debate over how reliable methods of collecting acceptability (*i.e.*, "judgment") data is more than half a century old (Sprouse & Almeida 2012). In defense of linguistic methodology, Sprouse and Almeida (2012) investigate both formal and informal methods of collecting acceptability data; the two methods share results at a 98% overlap rate. This finding, that the method is in fact good, contradicts the Kempen and Harbusch (2005) account for a gap (*i.e.*, that the method is bad). That explanation for discrepancies between frequency and acceptability scores, based on an assumed simple relationship between these two variables, does not hold up to empirical support for the methods in question.

*Another Not-So-Simple Story: Jurafsky (2002) & Featherston (2004).* An alternative to questioning the methods is accounting for the gap by re-examining the theory. This is the approach that Jurafsky (2002) takes when accounting for what he calls a "mismatch" between the variables. Jurafsky's account is a probabilistic model of production, in which frequency is treated as something much more complicated than the base frequencies of parts of speech appearing in a certain order (*i.e.* the simple account of syntax learning). By this account, the gap found in empirical study of structure frequency and its acceptability score is the result of acceptability data being abstract and frequency scores being concrete in comparison. To provide an explanation for this gap, linguists would need to find a way to translate between the acceptability and the frequency data. This would then allow for a fair comparison between the two very different kinds of data (Jurafsky 2002).

A probabilistic model of syntax that offers an additional explanation of the low correlation between frequency and acceptability is Featherston's (2004) model of structure selection. In this account, each potential syntactic form competes with other potential forms (each an alternative way of expressing the same thought, but all of varying grammaticality), for selection by the speaker. According to this model, the best structure is not always chosen, so frequency of a structure's appearance in a corpus cannot be expected to correlate perfectly with the acceptability score of that same structure—the two measures of the structure are not related, because some additional factor influences

whether a given structure is produced (*i.e.*, acceptability does not perfectly determine what structures are produced by speakers, or, the frequency with which structures are produced does not exclusively determine how acceptable that structure is).

### *Investigating the Simple Story for wh-Dependencies*

Pearl and Sprouse (2013) compare frequency and acceptability data using English to study a specific syntactic phenomenon: *wh*-dependencies. Studying these kinds of structures (and violations of the rules by which they abide), requires a fairly abstract representation of structure. For example, the *wh*-dependency "Who did she like?" is abstracted to representations at the phrasal level [*e.g.*, complementizer phrases (CP), inflectional phrases (IP), and verb phrases (VP), as in (5)]. For Pearl & Sprouse (2013), the *wh*-dependency is described by the sequence of these phrases between the *wh*-word and the gap (the place where the *wh*-word is interpreted).

[<sub>CP</sub> Who did [<sub>IP</sub> she [<sub>VP</sub> like \_]]]? (5)

*wh*-dependency structure = IP-VP

Thus, to study *wh*-dependencies, a search operates at the phrasal level of syntax. This higher level of abstraction is important to consider, because the story might still be "simple" in the sense that frequency of a structure determines how acceptable it is, if we define structure as something more abstract than parts of speech. This would be the conceptualization of structure according to generative grammar. Nonetheless, it still does not offer the kind of perfect correlation that a simple theoretical account predicts: Pearl & Sprouse (2013) also find a gap between adult acceptability data and frequencies collected from adult-directed speech for the *wh*-dependencies they investigate.

### The Frequency-Acceptability Gap in Adult-Directed Speech

In previous investigations of native speaker intuitions about language structure, researchers have used both formal (magnitude estimation) and informal (yes-no) acceptability judgment data because of how easily this kind of information can be collected and used in an assessment of grammar (Sprouse & Almeida 2012). While the methods of collecting acceptability judgment data have been attacked for a lack of proven reliability, a methodological defense of these collected judgments is possible. Sprouse and Almeida (2012) find that both methods are very nearly equal in terms of

reliability;<sup>3</sup> however, not all methods are equally appropriate for every study. All acceptability scores used in the present study were collected by the magnitude estimation method, in which participants are asked to assign relative ratings of acceptability to each utterance. This is a preferable method for this study because the theory tested here is one of gradient grammaticality, where some grammatical structures are more or less grammatical than other grammatical structures. Although the yes-no method can give reliable data, and acceptability judgments do not differ significantly from those scores acquired by magnitude estimation, the logic of yes-no methodology is binary and ignores the gradience that we intend to study (Sprouse & Almeida 2012; Featherston 2004). To assess gradient grammaticality, we must use data that were collected with the possibility of relative scoring. All acceptability scores used in this study are z-scores, meaning that they indicate the better-or-worse-than quality of every structure assessed, and are either above or below a mean (*i.e.*, average) value of acceptability. To control for semantic influence and assess the structure's grammaticality specifically, structures are assessed multiple times, in multiple instantiations; each z-score is calculated from the average score of those multiple assessments. This allows us to eliminate, to some extent, semantic interference (although the influence of semantics becomes relevant in our later discussion of the *kind* of frequency that we study).

#### *Prior Work: A Pilot Study by Pearl & Sprouse (in prep.) on Adult-Directed Input*

The present study follows up on previous work by Pearl and Sprouse (in prep.) that investigated adult-directed speech and studied the relationship between frequency and acceptability for a larger range of structures, rather than a particular subset of structures—*e.g.*, *wh*-questions (Pearl & Sprouse 2013). Pearl & Sprouse (in prep.) use the same adult judgment data that we will use here, which is taken from Sprouse and Almeida's (2012) work with Adger's *Core Syntax*, a linguistics textbook. This range of structures allows a broader assessment of the veracity of the simple story: if syntactic intuitions about acceptability are acquired by how frequently we perceive minimally abstracted versions of linguistic structure, then this should occur for all syntactic phenomena (although the earlier discussion of previous work demonstrates that the simple story is not supported in studies of other languages, or smaller subsets of structure in North American English). While Pearl and Sprouse (in prep.) have not completed a formal assessment

of their data, preliminary analysis suggests that the simple story of syntax acquisition does not hold up when structure frequencies in adult-directed speech are compared with the adult judgment data (*i.e.*, acceptability scores) of those same structures. Instead, the authors find that the frequency-acceptability gap occurs in adult-directed speech, even when studying a large range of syntactic structures.

### A Need for Frequency Data from Child-Directed Speech Corpora

Previous work by Pearl and Sprouse (in prep.) demonstrates the lack of empirical support for a simple account of syntax acquisition when studying adult-directed speech. However, it is important to analyze the actual input for children during the period of language acquisition if we want to understand how we come to have our grammaticality intuitions. This is particularly important given that there are several known differences between adult-directed and child-directed speech.

We examine these differences—and their relation to the simple story of syntax acquisition—in the following sections. We also describe the corpora of child-directed speech that we use to study frequency in child-directed speech.

#### *Differences between Child-Directed and Adult-Directed Speech*

One common form of child-directed speech, “motherese,” has many apparent differences from adult-directed speech. These differences occur at phonetic (*e.g.*, emphasized vowel sounds) [Fernald 1985], lexical (*e.g.*, increased use of monosyllabic words) [Yang 2004], and shallow syntactic levels (*e.g.*, use of Determiner+Noun constructions, rather than pronouns) [Furrow 1978]. Given these known differences, one potential explanation for the frequency-grammaticality gap that would still support the simple story is that child-directed speech also consists of different frequencies of structure occurrences which correlate better with acceptability data (Kempen & Harbusch 2005). In order to gain a better understanding of how useful the simple story is as a theory of syntax acquisition, it is imperative to determine if the frequency-acceptability gap, a notable problem for this simple story, exists when frequencies are collected from the actual input that children receive, rather than the input adults receive and produce. Pearl and Sprouse (2013) compare child-directed speech and adult-directed speech at a more abstract syntactic level of representation for *wh*-questions (involving phrasal structure) and interestingly find little difference at this level of abstraction.

3. Discrepancies between yes-no and magnitude estimation methods occur at a rate of approximately 2%. For Sprouse and Almeida (2012), a discrepancy occurred any time that yes-no methods reported a different score from that of the magnitude estimation (the assumption of the study being that magnitude estimation methods are more valid).

Additional support for the marked difference between child-directed speech and adult-directed speech comes from the Vygotsky's theory of the *zone of proximal development* (Vygotsky 1978). This developmental psychology account is one of scaffolding; child-directed speech responds to the child's level of acquisition by adjusting the level of difficulty to accommodate what the child has learned and still needs to learn. Thus, it is considerably different from adult-directed speech. The application of Vygotsky's theory to syntax acquisition could mean that frequencies of most structures studied should be very different from those that Pearl and Sprouse (in prep.) collect in adult-directed speech, at least when comparing the data for more complex syntactic structures.

Interestingly, Yang (2004, 2011) provides evidence that there does *not* seem to be a frequency-acceptability gap when the learner tracks the frequency of very abstract structural representations (called linguistic parameters). The frequency of unambiguous "linguistic signatures" for these structural representations correlates very well with how early children learn the structural representations. If we assume that the earlier a structure is learned, the earlier a child finds that structure acceptable, then this indicates that frequency correlates quite well with acceptability. The key is the frequency of *what*, a question we return to at the end of the paper.

Our approach investigates structural representations that are less abstract than Yang's linguistic parameters. In particular, the current study examines the large range of structures presented in Adger's *Core Syntax* (Sprouse & Almeida 2012), and thus seeks to account for a relationship (or lack thereof) between the frequencies of different types of syntactic phenomena at a very shallow level, and their associated acceptability scores.

### Corpora from CHILDES

The child-directed speech corpora consist of portions of the CHILDES datasets that have speech directed at children of the ages indicated in Table 1.

All ages are within an appropriate period for studying language acquisition. While this study makes no comment on the debate between a critical and a sensitive period of language acquisition, the age range does not approach (or surpass) any of the barrier ages implicated in either theory (Newport 2002).<sup>4</sup> A total of 170,263 utterances from these

4. All age range data are written in terms of years and months. For example, 2;3-4;10 denotes that the youngest age in that corpus's population is two years and three months, with the oldest age at four years and ten months.

Table 1  
CHILDES Corpora

Corpus Name	Age Range	Number of Utterances
Brown-Adam	2;3-4;10	26,280
Brown-Eve	1;6-2;3	14,245
Brown-Sarah	2;3-5;1	46,948
Soderstrom	0;6-1;0	21,334
Suppes	1;11-3;11	35,906
Valian	1;9.20-2;8.24	25,550

corpora are mined for frequency data in order to provide a large enough sample size for our work to have meaningful implications in the field of language acquisition research.

## Methods of Assessing Frequency

To assess the frequency of each structure in the corpora, it is first necessary to determine how to annotate and search for each structure. For the purposes of an investigation into the simple story of syntax acquisition, it is important that our queries involve minimal abstraction made at the structural level; meanwhile, each of these searches must capture what is most important about the structure (*e.g.*, a subject-verb agreement violation or the use of an intransitive verb). To identify the relevant structures in child-directed speech, the structures can be translated into queries written for Tregex,<sup>5</sup> a linguistic pattern-matching utility. We designed these to search for those salient lexical and syntactic features of the sentences previously scored for acceptability.

In the sections below, we discuss how queries were formulated (and the implications of our annotations for the kind of frequency assessed), as well as how the raw frequency values are normalized and smoothed to allow for accurate visual comparison between frequency and acceptability scores.

### Queries

All queries are created to search a dataset of utterances annotated independently of acceptability studies (*e.g.*, the annotation used for the CHILDES database does not mark a violation of Subject-Verb Agreement as such; instead, we must search for all plural nouns followed by a singular verb form).

To illustrate the ways in which we annotate utterances while making minimal abstraction about their structures, we provide the following examples:

5. <http://nlp.stanford.edu/software/tregex.shtml>

A sentence that received a high acceptability score:

“The scissors are lost.”

(6)

[<sub>S</sub> [<sub>NP</sub> [<sub>DT</sub> The] [<sub>NNS</sub> scissors]] [<sub>VP</sub> [<sub>AUX</sub> are] [<sub>VBN</sub> lost]]]

The salient features of this sentence (S) are the noun phrase (NP), which is headed by a definite article “the” (DT) and contains a plural noun “scissors” (NNS), and the verb phrase (VP) which consists of an auxiliary “are” (AUX) and past-participle “lost” (VBN).

The query design for this structure is: (/S/ <1 (/NP/ < /NNS/) <2 (/VP/ < (/AUX/ . /VBN/))), where we look for a sentence that has an NP with a plural noun as its first child and a VP with an auxiliary verb immediately followed by a participle. This would identify sentences such as “The scissors are lost” and “The scissors is lost.” We then hand-check the identified sentences to see if they follow the salient aspects of the structure we want. In this case, “The scissors are lost” would count for our search while “The scissors is lost” would not.

A sentence that received a low acceptability score:

“The pigs grunts.”

(7)

[<sub>S</sub> [<sub>NP</sub> [<sub>DT</sub> The] [<sub>NNS</sub> pigs]] [<sub>VP</sub> [<sub>VBZ</sub> grunts]]]

The salient feature of this sentence (S) is its subject-verb agreement error. In this particular structure, the subject-verb agreement error consists of a noun phrase (NP) that includes a plural noun (NNS), followed by a verb phrase (VP) that includes the third-person singular verb form (VBZ).

The query design for this structure is: /S/ < ((/NP/ < NNS) \$+ (/VP/ < VBZ)), where we look for a sentence that has an NP with a plural noun as one of its children, and a VP with the verb form appropriate for a singular noun (in the third person). This would identify sentences such as “The pigs grunts.” and exclude sentences like “The pig grunts.” or “The pigs grunt.” To ensure that all identified forms were instantiations of subject-verb agreement errors (and not erroneously annotated), we hand-check the results here, as well.

To ensure that queries are correctly written, we use a test file containing correctly parsed versions of the sample sentences, against which every query design is checked (*i.e.*, if a query is written correctly, it should always find at least one result in the test file).

### Calculating Frequency Values

Once frequency values were collected by the methods described previously, we normalized these frequencies (*e.g.*, a structure of raw frequency 20, divided by the total number of utterances in the corpora 170,263, becomes 0.000117465). However, these normalized frequencies were very small numbers (*e.g.*, 0.000117465), and many of the structures were not found at all in the corpora (*i.e.*, they had frequencies of zero). We smoothed the data by adding +0.5 to all normalized frequency scores (*e.g.*, 20 became 20.5, and 0.000117465 became 0.00012042), so that, when taking  $\log_{10}$  of the normalized frequency scores, we would take the  $\log_{10}$  of 0.5 instead of zero (which is undefined). It was necessary to take  $\log_{10}$  of all the values in order to graph these very small numbers against their much larger acceptability score counterparts. This results in larger absolute values (although they all become negative numbers as a result) [*e.g.*,  $\log_{10}(0.000117465) = -3.930090286$ ].

## Results

After collecting and adjusting the frequency scores, we entered these data into an Excel file, to allow for visual comparison of acceptability and frequency, as well as calculation of an *r* score to assess the correlation between these two kinds of data. In the sections below, we explore the different kinds of mismatches between frequency and acceptability that appear in our data.

### Comparison of structure acceptability and frequency in CHILDES corpora

In Figure 3, we see an absence of the positive linear relationship predicted by the simple story (and pictorially represented in Figure 1).




Figure3

The x-axis represents ( $\log_{10}$ ) frequency scores while the y-axis represents adult acceptability scores. Each point on this graph represents a single structure.



Although there are some structures for which the simple story seems to be true (*i.e.* their acceptability scores align with frequency values, in compliance with a positive linear relationship) there are many that do not follow this pattern. Recall Figure 2: in a graph where the x-axis is frequency (from very low to very high) and the y-axis is acceptability (also from very low to very high), the simple story predicts that data points that fall only in Quadrants I (high frequency/high acceptability) and III (low frequency/low acceptability). However, as Figure 3 demonstrates, there are many exceptions to this assumed relationship.

The following sections discuss the different types of data not predicted (or explained) by the simple story.

*Quadrant II Violations: High Acceptability and Low Frequency.* One high acceptability structure that has a low frequency value is *Subject (nominative pronoun) + “have”-auxiliary + transitive verb past participle + object (accusative pronoun)*, or “She has kissed her.” The acceptability score of this structure is fairly high for the data (0.80), while its raw frequency value is very low (*i.e.*, it only appears twice in the CHILDES corpora). This structure, like other Quadrant II violations, demonstrates that minimally abstracted structures can be infrequent but still highly acceptable.

*Quadrant IV Violations: Low Acceptability and Some Frequency.* If our knowledge of acceptability is based solely on how often we hear structures, then all unacceptable structures should have low frequencies. However, the utterance “Letter is on the table.” and its structure (*NP = [singular count noun with no determiner] + verb + PP, nothing after*) have an acceptability score of -0.9 and a raw frequency of 7. Notably, this unacceptable structure has the same raw frequency value as the highly acceptable utterance “Joss’s idea is brilliant.” (which has an acceptability score of 1.07). Thus, this data point is unexpected under the simple story.

*That Line in Quadrants II and III: Acceptability Variation in No-Frequency Structures.* The data present with a very different kind of linear relationship than that predicted by the simple story: 168 of 219 structures searched for in the CHILDES corpora have frequencies of zero. The following examples show the considerable variation in acceptability between these structures:

- (1) “Peter is pigs.” is an instantiation of the structure: *Subject (name) + be + object (plural noun)*. This structure has a very low acceptability score of -1.20.

- (2) “His analysis of her was flawed.” is an instantiation of the structure: *Noun phrase (possessive + noun + PP-of-simple-noun) + verb + participle*. This structure has a low acceptability score of -0.39.
- (3) “Kim should leave for work on time.” is an instantiation of the structure: *Subject= (singular noun + name) + modal + intrans-verb + PP=[preposition+noun] + PP=[preposition+noun]*. This structure has a very high acceptability score of 1.13.

The 168 structures that fall into this vertical line represent an  $r$  of 0; these structures show a lack of relationship between the frequency of minimally abstracted structures and acceptability data.

### Discussion of $r$ and the Structures that Support the Simple Story

The simple story predicts a perfect correlation, or an  $r$  of 1. The correlation between frequency values of utterances in the CHILDES corpora and corresponding acceptability data in this study is not perfect; the  $r$  value is 0.509997592. This is a positive relationship, but it is not a linear one.

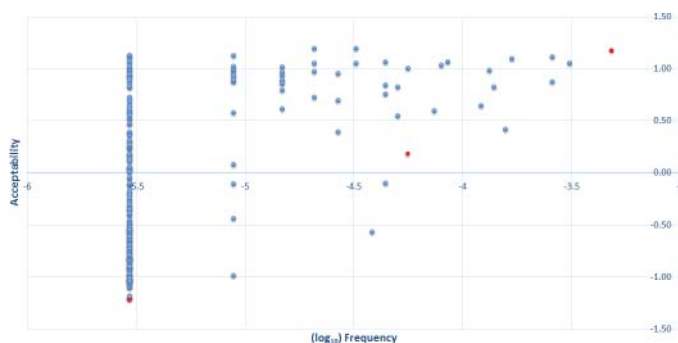


Figure 4  
When the Simple Story Holds

However, the  $r$  value is considered very strong, and some of the structures do fall into the simple story’s predicted line. Figure 4 highlights three example structures that represent the predicted frequency-acceptability relationship of the simple story.

In the following section, we discuss the data points that the simple story correctly predicts.

*Correlated Scores.* The unacceptable, no-frequency utterance, “It’s arrived first that Julie and Jenny” represents an expletive-*it* located at the head of an intransitive verb phrase, with the subjects of that phrase located at the end of sentence, after complementizer-*that*. This is a structure not permitted

by the English grammar, or (according to its acceptability score), by naïve native speakers of that language. Because it has the lowest acceptability score, this structure should also have the lowest possible frequency (zero), and it does.

The mildly acceptable (0.18), moderately frequent utterance (appearing 9 times), “I worry if the lawyer forgets his briefcase at the office” represents a simple subject, a main verb that takes an embedded-if clause with a transitive verb, a direct object, and a locational prepositional phrase. There is no obvious violation of the known rules of English grammar, and naïve native speakers score it as more acceptable than average, so this structure supports the simple story’s claims.

The most acceptable (1.18), most frequent utterance (appearing 81 times), “Genie bought the mirror” represents a subject that consists of a name, a transitive verb, and its direct object (which consists of a determiner and a noun). The fact that the most acceptable structure also appears the most frequently in the corpora offers some support to the simple story of syntax acquisition (as does the fact that the least acceptable structure has the lowest possible frequency). These sample structures show that the simple story is not merely wrong; however, the numerous exceptions to this story do suggest that something more complicated is happening in the process of syntax acquisition. We posit that these well-correlated structures are evidence for the simple story because surface analysis of these structures gives a learner all of the information he or she needs, without abstraction or a learning bias to support the raw frequency. In short, the simple story works for some structures, but apparently there are many structures that we understand in terms beyond the surface ones laid out by this simple story.

## Discussion: Frequency of What?

### *Levels of Abstraction*

There are multiple levels of abstraction at which we can study structures in order to calculate their frequencies. Notably, certain violations can only be studied at a lower or higher level of abstraction than the one we generally adopted for the purpose of this study. For example, Semantic Category Violations such as “The book ran” do not seem unacceptable when we only account for the syntactic structure. However, on the semantic level, they sound terrible; books, we know, do not run. We cannot adequately calculate the frequency of this structure by simply searching for a subject composed of a determiner and a noun, followed by an intransitive verb. There is more to understanding this utterance (and its ungrammaticality) than shallow, minimal-

ly abstracted structure. In our study, we account for these errors by requiring a more abstract kind of data—any calculation for this structure only included inanimate nouns combined with animate verbs. This account includes a very abstract understanding of syntax, one that incorporates some semantic information, rather than the more shallow one with which we approached other structures.

Because we do not know how abstract the analyses of children acquiring language are, it is impossible to know which level of abstraction is the appropriate one. Studies like the present one can evaluate a theory by collecting frequency data using the terms (*i.e.*, level of abstraction) set forth by that theory. What we have done here is assess the simple story of syntax acquisition—one that claims very minimal abstraction is made about the structures children hear. With the aforementioned exception of Semantic Category Violations, we annotated and analyzed structure at a very shallow level of abstraction. As discussed earlier, there are some structures for which the simple story is an adequate account. However, for those kinds of violations (*i.e.*, the frequency-acceptability gap), the simple story is not saying enough. Very little of the data are accounted for by the simple story—it is more exception than rule—and for that reason, it is not an adequate account of how we acquire our intuitions about which structures are acceptable and which are not.

*Links to the Language Acquisition Process.* Our previous discussion of Yang (2004, 2011) shows that there are alternate, and more abstract, ways of understanding syntax and the language acquisition process. Yang’s (2004) findings suggest that more abstract structure frequencies correlate better with acceptability, especially when age of acquisition is considered (*i.e.*, the sooner children learn more abstract structures, the sooner they view those structures as acceptable).

### *Future Directions for Research*

We find that the simple story of syntax acquisition is an inadequate explanation of all the data. When these data are collected by sound research methods, there is not enough support for a theory of acquisition that claims children analyze only the base frequencies of shallow structure. In the following sections, we discuss ways that future studies might better understand the nature of first-language syntax acquisition.

*Grammaticality According to Children.* Children do, at certain levels, receive different input. Prior research finds that the differences between child- and adult-directed speech are not significant at a more abstract syntactic level (Pearl &

Sprouse 2013). Preliminary investigation also shows that there are negligible differences between child- and adult-directed speech when working with much more shallow structures (Pearl & Sprouse in prep.). However, the corpora used provide a fairly wide age range (from six months to five years and one month). This fact may account for the lack of evidence supporting Vygotskian expectations of the data, which would assume different structures being used in speech directed at younger children as opposed to older children.<sup>6</sup>

Another consideration is that, while the input is similar,<sup>7</sup> the interpretation of that input may be very different. Adults do not speak very differently at a structural level whether the audience is a child or another adult, but their audiences are potentially very different. As Yang (2004) notes, child learning biases can help account for an otherwise apparent poverty of stimulus; it is plausible that children have different (less mature) knowledge than adults do, which affects what structural input they attend to. Child acceptability judgments could offer a more fair comparison between acceptability and (shallow) frequency scores, providing support for the simple story. More importantly, these child acceptability data could give researchers insight into how children acquire the syntax of their first languages.

*A More Sophisticated Theory.* The present study does not support a simple story of syntax acquisition. Shallow structure frequencies do not correlate with many adult acceptability judgments. Our literature review shows that alternative accounts for acquisition—*i.e.*, those that study more abstract syntactic information—are able to close the frequency-acceptability gap. For this reason, we believe that a more sophisticated theory would better account for those data unexplained by the simple story, and that the phrase with which we have titled this larger section (“Frequency of what?”) is essential to understanding the role of frequency in syntax acquisition. It is not enough to say that frequency of a structure very basically determines the acceptability of that structure; instead, future research should investigate the level of abstraction (and whether that varies between different syntactic phenomena) necessary to account for a larger portion of the data.

6. An age range has the potential to average (and thereby diminish the visibility of) any differences that occur at developmental stages, although the necessary input for a child of six months is likely different, at least at the level of syntactic structure complexity, from that provided to a child of five years.

7. A finding that supports the validity of prior studies that used adult-directed speech only.

## Acknowledgements

I would like to thank my lab-mates in the Computation of Language Laboratory for their contributions to this project in its early stages. I would also like to thank my research advisor, Lisa Pearl, for being so supportive and giving me such expert advice every step of the way.

## Works Cited

- Fernald, Anne. “Four-Month-Old Infants Prefer to Listen to Motherese.” 1985. Web. 30 June 2014. Retrieved from: <http://waldron.stanford.edu/~babylab/pdfs/Fernald%201985.pdf>
- Furrow, David, Katherine Nelson, and Helen Benedict. “Mothers’ speech to children and syntactic development: Some simple relationships.” *Journal of child language* 6.03 (1979): 423–442.
- Hickok, Greg. “Weak quantitative standards in linguistics research? The Debate between Gibson/Fedorenko & Sprouse/Almeida [Msg 6].” June 18, 2010. Web. 18 June 2014. Message posted to <http://www.talkingbrains.org/2010/06/weak-quantitative-standards-in.html>
- Keller, Frank. “Gradience in grammar: Experimental and computational aspects of degrees of grammaticality.” 2001: 1–388.
- Kempen, Gerard, and Karin Harbusch. “The relationship between grammaticality ratings and corpus frequencies: A case study into word order variability in the midfield of German clauses.” *Linguistic evidence: Empirical, theoretical, and computational perspectives* (2005): 329–49.
- Newport, Elissa L. “Critical periods in language development.” (2002). Retrieved from <http://www.bcs.rochester.edu/people/newport/Newport-ECS-A0506.pdf>
- Pearl, Lisa, and Jon Sprouse. “Syntactic islands and learning biases: Combining experimental syntax and computational modeling to investigate the language acquisition problem.” *Language Acquisition* 20.1 (2013): 23–68.
- Sprouse, Jon, and Diogo Almeida. “Assessing the reliability of textbook data in syntax: Adger’s Core Syntax.” *Journal of Linguistics* 48.03 (2012): 609–652.
- Vygotsky, Lev S. “Mind in society: The development of higher psychological processes.” Cambridge, MA: Harvard University Press, 1987.

Yang, Charles D. "Universal Grammar, statistics or both?." Trends in cognitive sciences 8.10 (2004): 451–456.

Yang, Charles D. "Computational Models of Syntactic Acquisition." Wiley Interdisciplinary Reviews: Cognitive Science, 2011. Web. 24 May 2014. doi: 10.1002/wcs.1154.