Computational Linguistics in the Netherlands Journal 14 (2025) 191-214 Submitted 12/2024; Published 07/2025

From ‘nijntje’ to ‘konijn’: New methods for analyzing stress pattern acquisition

Nienke Wessel[*] NIENKE.WESSEL@RU.NL
Marc van Oostendorp[*] MARC.VANOOSTENDORP@RU.NL
[*]*Radboud University, Nijmegen, The Netherlands*

Abstract

In order to gain more insight into stress pattern acquisition by children, we are in need of large datasets and analysis tools that can analyze those data sets. This paper suggests a new way to analyze PhonBank corpora by automatically loading in the annotated data and plotting it in different manners, focusing on stress pattern acquisition. This way, we can more easily track changes in children’s development through time. This method is applied to three longitudinal and two cross-sectional corpora of Dutch, German, and English. The results show that there are some aspects of previously proposed theories that are in need of further investigation. Most notably, past theories have assumed that stages of stress pattern development were the same across different words, while we found some evidence for stages of development per word. This suggests that further research is needed into how exactly these stages differ for different words, and what determines whether a child learns the correct stress pattern for a word sooner or later. Furthermore, comparing our findings to those of previous investigations which used manual analysis methods provides us with the opportunity to discuss the drawbacks and benefits of our analysis tool.

1. Introduction

Prosodic stress plays several important roles in languages. For starters, it can help us determine where words start or end in continuous speech (Curtin et al. 2005). Sometimes, a stress pattern difference is the only difference between two otherwise equal words. For example, in English, there are a number of noun-verb minimal pairs in which the stress determines the word’s meaning. Compare ‘a rebel’ to ‘to rebel’.

English famously has complex rules for stress assignment in words. While most native speakers are not consciously aware of these rules, they do seem to be able to apply them to both new words and nonsense words. This suggests that when learning the language as children, they acquire these rules on a subconscious level. The question is how children do this.

Studies have shown that children tend to learn stress patterns in phases, as evidenced by their mistakes (Fikkert 1994, Kehoe and Stoel-Gammon 1997). For example, Dutch children initially produce utterances such as /ˈnɛinə/ (forming the basis for the famous ‘Nijntje’, or ‘Miffy’ in English) before producing the correct /koˈnɛin/. Past research has found evidence for these phases by manually sifting through production data produced by children of different ages. This has proven to be time-intensive and also does not allow us to easily compare different large datasets in the same way, while this is necessary to form and test robust hypotheses. Some investigations have therefore looked into ways of automating this process (Rose and MacWhinney 2014, Rose and Stoel-Gammon 2015, McAuliffe et al. 2019, Hall et al. 2019). The current investigation aims to further our insights into the development of stress acquisition by revisiting past datasets through new methods. Instead of manually identifying phases, the current investigation aims to use computer programs to create visualizations of the development of individual children and groups of children in order to provide insight into how they have progressed through time in a more complete manner than was possible with manual analysis.

The contribution of this investigation is threefold. First of all, the investigation proposes new ways of analyzing child production data. This has several benefits, most notably that it allows us to look at the data in a different way and thus highlights new aspects of the data that have previously been unexplored. Another benefit of the proposed method is that it is easy to apply it to all datasets that have been sufficiently annotated. This allows us to compare different datasets through the lens of the same analysis with little extra effort from the researcher's side. This relates to the second contribution of this paper: it provides different analyses on multiple different datasets that have not been compared in this manner before. The comparison is also multilingual, comparing Dutch, German, and English speaking children. Future analyses are not limited to these languages, but can be done on any language. Finally, this research also contributes by providing more insight into stress acquisition. The results here suggest that global phases that hold for all words at the same time are hard to locate and that it therefore perhaps might make more sense to look at phases at a word level. As the results from this investigation provide more insight into how children learn stress patterns, there is also societal relevance to them, as this can be used to improve learning and help us notice earlier when children are struggling with learning these types of patterns.

2. Literature review

There is quite some literature on stress and its acquisition. This section first explains the necessary theory to understand the phonological analyses performed in this paper, and then continues to discuss related work, with a focus on work with similar methodologies.

2.1 Stress patterns

Many, but not all, languages have *word stress.* This is also the case for Dutch, German, and English, the languages on which we concentrate here. This means that some syllables stand out phonetically from others (Rietveld and Heuven 2016). In our notation, we use ˈ to indicate that the following syllable carries stress. All lexical words in stress languages have a syllable with primary stress. If a word has three or more syllables, it can also have secondary stress.

Generally, a stressed syllable is paired up with one or more unstressed syllables into what we call a *foot.* Languages can either make pairs in which the first syllable in a foot carries stress (e.g. *ˈwinter*), or in which the second syllable carries stress (e.g. *misˈtake*) . We call the first type of feet *trochaic* and the second type *iambic.*

In English, Dutch, and German, the standard foot is the trochee. Most words, like *ˈwinter*, are trochees. Nonetheless, words like *misˈtake* with surface iambic patterns also exist. The exact rules governing stress allocation are rather complicated in all three languages.

Before delving into this, it is good to be clear on some terminology pertaining to the 'heaviness' of a syllable. Syllables can be 'light', i.e. open without a long vowel (CV); 'heavy', i.e. containing either a long vowel or being closed (CVV or CVC); or 'superheavy', i.e. containing a long vowel and a closing consonant (CVVC) or two closing consonants (CVCC). In general, the heavier a syllable, the more likely it is to be stressed.

We will now briefly discuss stress assignment rules in our three languages as so far they are relevant for the following discussion. The focus here is on two-syllable, morphologically simple words. Trommelen and Zonneveld (1999) provide a more elaborate overview of all types of words for those interested.

Stress in the West-Germanic languages Dutch, German and English is generally trochaic. The main stress is on one of the last three syllables of a word (Kager 1989, Féry 1998). A schwa can never be stressed. Words with superheavy final syllables or final diphthongs have final stress (Kager 1989, 227). Some other generalizations exist, but those are not relevant for the remainder of this paper. See e.g. Booij (1999), Kager (1989) or Trommelen and Zonneveld (1999).

For words consisting of only two syllables (the focus of this paper), these general rules mean that the stress is generally on the first syllable in all three languages. If the final syllable of a word is superheavy or a diphthong, the word starts to show a iambic pattern. A syllable is considered superheavy if it contains a diphthong and is closed (i.e. there is a coda). Some Dutch examples include *ko'nijn* ('rabbit'), *gi'taar* ('guitar'), and *ba'naan* ('banana').

Furthermore, stress in West-Germanic is not always governed by rules. For example, the Dutch *ka'non* ('cannon') and *'canon* ('literary canon') are both pronounced /kanɔn/, and only differ in stress placement. This shows that stress is also partially lexical.

One difference between English on the one hand and German and Dutch on the other is that long vowels in the English also attract stress while German or Dutch do not truly have long vowels (Trommelen and Zonneveld 1999). Furthermore, English phonological rules make use of word classes. Nouns are more likely to have stress on the first syllable while verbs are more likely to have stress on the second syllable. Compare word pairs like: 'a rebel' vs. 'to rebel', or 'the content' vs. 'to content' (Guion et al. 2003).

2.2 Stress acquisition

In learning the stress pattern of a language, children have to derive the rules described in the previous section from the input they receive. As children make overgeneralization errors (see below), we know that they actually learn rules as opposed to simply learning how to assign stress lexically. They are furthermore able to apply the rules to nonsense words (Oh et al. 2011), further suggesting that they learn rules rather than all words by heart.

Still, one of the central questions in child language acquisition literature and stress acquisition literature specifically is to what extent children come with a born sense of how language is structured (e.g. that languages can be stressed, and that feet contain two syllables, see e.g. Archibald (1995) for a coverage of some rules), or that they derive everything from input. The different theories that have been proposed for how children learn stress make different assumptions about this.

Here, we focus on two main theories. The first theory has been proposed by Fikkert, which in its core is still a leading theory on stress development, although some amendments have been proposed. We will first discuss her theory and then discuss some proposed changes.

Fikkert (1994) analyzed the speech production from twelve Dutch children through time. She notes several things of importance on the basis of her data. First of all, she identifies truncation as an important phenomenon present in the productions of the children. This is not simply the deletion of the unstressed syllable; if that were the case, we would expect it to take place in the same way in trochaic disyllables (initial stress) and iambic disyllables (final stress). Fikkert shows that this is not the case in her data; iambic disyllables are more commonly truncated than trochaic disyllables (Fikkert 1994, 25). Similarly, she finds that children make more mistakes in stress realizations in iambic productions (Fikkert 1994, 25).

On the basis of these observations and other observations, Fikkert developed a theory (1994, 26) in which children go through four developmental stages in learning how to stress iambic disyllables:

1. In the first stage, children produce only the stressed syllable of the iambic word. *Konijn* ('rabbit') is pronounced as e.g. /ˈnɛin/ or /ˈtɛin/ instead of the adult production /koˈnɛin/.[1] We can analyze this as the target being mapped on a trochaic template. This is further supported by that children sometimes add an extra vowel (usually the schwa) to the left-over syllable to match the trochaic template, such as /ˈnɛinə/.

2. In the second stage, children produce all adult syllables but place stress on the initial syllable instead of the final. E.g. *gitaar* is pronounced as /ˈsitaː/ instead of /χiˈtar/.

3. In the third stage, children produce both syllables with equal stress level.

4. In the fourth and final stage, children reach an adult like stress production.

1. All examples come from Fikkert (1994).

Fikkert notes that the stages are discernible from their relative frequencies. Several studies after Fikkert have found evidence to confirm these stages in general, but often suggest small amendments (see also Joanisse and Curtin (1999) and Gerken (1994)).

Fikkert works from a universal grammar perspective where children already have default metrical settings, e.g. that the trochee is the basic metrical template (see also Adam and Bat-El (2009)). Some authors have proposed competing theories. For example, al Huneety et al. (2023) argue that, based on their investigation into children learning Ammani Arabic, children begin without any bias for a certain type of stress. See also Tzakosta (2004) for Greek. Archibald (1995) argues that children do not come with a preconceived idea about trochaic feet, but rather that they simply parse from right to left and cut off syllables at the stress. Related to this, Özçelik (2011) finds further arguments against incorporating feet as a preset variable, as certain languages do not have feet.[2]

One thing left unexplained by Fikkert's theory is that children seem to be in multiple stages at the same time. Joanisse and Curtin (1999) and Grimm (2008) give possible ways to understand this within Optimality Theory.

There are also people arguing that it is possible to learn stress patterns without a need for (preset) parameters at all. For example, Daelemans et al. (1994) trained a system which only had word representations and had to learn stress assignment from those without being able to infer information about structures like metric feet, iambs or trochees. Their system managed to obtain 90% accuracy.

2.3 Methodologically related work

From the previous discussion, it follows that there are different theories about stress acquisition. The question now is how we can test these theories. A recent trend in research into this topic has been to use larger datasets and analyze them automatically. This section discusses some work that has been put forward towards this goal.

The work can be divided into two related subsets. The first is the creation of large(r) corpora and annotation standards for corpora. The second is the creation of analysis software. As this paper proposes analysis software, we focus on this second set.

Several types of analysis software have been proposed. For example, *ISCAN* (McAuliffe et al. 2019) is a system that allows the user to (semi-)automatically analyze different corpora in the same way on a phonetic level. It focuses on analyzing sound waves, and thus has a different use case compared the system we propose.

Another similar analysis tool is *Phonological CorpusTools* (Hall et al. 2019). This system is also designed to analyze larger corpora and provides phonological and phonetic information corpus entries. It was built for specific types of analyses, such as calculating probability distributions of sounds dependent on different contexts. Compared to the system proposed in the current investigation, it does not allow for the analysis of stress level patterns.

The most similar project to the current investigation is *Phon* (Rose and MacWhinney 2014, Rose and Stoel-Gammon 2015). *Phon* is a computer program designed to work on the so-called PhonBank corpora. It aligns child productions and models to provide opportunities for analysis. It can, for example, be used to see how certain consonants have been produced over time in a child's development. Compared to the system in the current investigation, it is not as suited for studying stress development. While it does allow for syllabification, the system proposed in the current investigation is more suited for studying development as it builds a vocabulary dictionary for the child at every measurement point, allowing the researcher to study more carefully how the complete picture changes. Furthermore, the software proposed in the current investigation allows for more

2. Also Kehoe argues for multiple amendments and caveats of Fikkert's theory for trochaic languages (Kehoe and Stoel-Gammon 1997, Kehoe 1998, Kehoe 2001)

detailed plotting, allowing us to analyze many different stress patterns and their change throughout time.

It should be mentioned that most research performed with these types of systems has focused on phonological development and not prosodic development. See, for example, the different chapters in *The Oxford Handbook of Corpus Phonology* (Durand et al. 2014). The current investigation is, to the best of our knowledge, the first to focus on building an automatic way to analyze stress pattern development.

3. Methodology

In this part of the paper, we describe the method used for the analysis. We first discuss the corpora analyzed and then move on to a discussion on how these corpora have been processed.

3.1 Description of corpora

This investigation made use of several parts of the PhonBank corpus (TalkBank n.d.). This corpus is part of the larger CHILDES/TalkBank corpus and contains annotated files of spoken child language. Files in this corpus are often annotated for phonological and prosodic information.

From this corpus, we used five corpora. We have chosen to only look at corpora for the Dutch, German, and English language. Dutch was chosen because it is the same corpus as used by Fikkert (1994). English and German are languages from the same language family with relatively similar stress patterns to Dutch. Fikkert's theory is also relevant for them, as shown by previous research (Grimm 2008, Kehoe and Stoel-Gammon 1997).

The CLPF corpus is a corpus by Paula Fikkert (1994) and Clara Levelt (1994) which follows twelve Dutch children, approximately from age 1 to age 3. Recordings of the children were made at home and are of spontaneous speech. We used all available data.

The Grimm corpus has been collected by Grimm (2008) and follows four children from Germany in their second life year. The speech is generally spontaneous, although the researchers did occasionally use plastic animals with long German names to elicit tri- and quadrisyllabic words. We used all available data from this corpus, except for the files where no date of recording was available.

The Bracci (2020) and Fox-Boyer (2019) corpora are both cross-sectional corpora of picture description tasks from children in Kindergarten. The Bracci corpus had fifteen participants from Luxembourg (most of which were multilingual), while the Fox-Boyer corpus contains 32 participants from Germany who all speak standard German.

The Providence corpus by Demuth (2004) contains data from six children in the United States from years 1 to 3 of their lives. All children were monolingual and one of them was later diagnosed with autism.

We looked at all other corpora in PhonBank in our three languages, but they were not suitable for our purpose, mostly because they did not have sufficient stress annotations.

3.2 Corpus processing

The corpora were read and analyzed by a computer program written in the Python programming language. All code is available on the first author's Github page[3].

The first analysis step was to parse the xml files that are available on the PhonBank website. From this, we created a dataset which contained the following data points: name of the child, time of measurement, the word uttered (in Latin script), the model pronunciation of the word in IPA (i.e. how a native adult speaker would pronounce the word), and the pronunciation realized by the child in IPA. Table 1 shows the number of data points for the corpora used.

3. `https://github.com/NienkeWessel/PhonBank-stress-analysis`

	CLPF	Grimm	Providence	Bracci	Fox-Boyer
From	Fikkert (1994)	Grimm (2008)	Demuth (2004)	Bracci (2020)	Fox-Boyer (2019)
Data type	Longitudinal	Longitudinal	Longitudinal	Cross-sectional	Cross-sectional
Language	Dutch	German	US English	German (Luxem.)	German (Germany)
Nr of data points	18636	7770	164213	1305	3092
Nr of iamb entries	599	363	1213	60	95

Table 1: Data of the corpora employed in this paper.

A second script then used the IPA transcriptions to determine the stress pattern of the word. The program codes iambs and trochees as [False, True] and [True, False] respectively, where True indicates a syllable is stressed and False indicates a syllable is not stressed.

In order to properly distinguish between the different phases from Fikkert (1994), it was also annotated whether a realization ends in a schwa. For example, from the examples in the previous chapter, /ˈnɛinə/ (stage 1 in Fikker (1994)) and /ˈsitaː/ (stage 2) both map to [True, False], where we can only determine the stage by looking at whether the production ends in a schwa or not.

Having these representations of both the productions and models, we looked at how a certain model pronunciation pattern (e.g. [False, True], so a iambic disyllable) is actually produced (e.g. [True] (truncation), [True, False] (overgeneralization of trochaic structure) or [False, True] (correct pronunciation)). We were especially interested in how this changes over time, as this indicates at which of the stages described by Fikkert (1994) the child is. We often plotted moving averages instead of the raw data. This is to allow underlying patterns to emerge without being distracted by flukes and outliers. It is always indicated whether the raw data or moving average is plotted.

4. Results

This section aims to apply the method described in the previous chapter to the CLPF corpus and the other corpora in order to see to what extent we manage to obtain the same findings as Fikkert (1994) with the new method. The goal is twofold. First of all, this shows us to what extent we can replicate the findings by Fikkert. Second, this allows us to investigate the strengths and limitations of the new method. Furthermore, by applying the same method to other corpora, we can compare the different languages.

The results consist of three subsecctions. The first subsection is concerned with truncation, where an attempt is made to replicate a part from Fikkert (1994). The second part looks at the developmental phases as proposed by Fikkert (1994) and sees to what extent we can find them in the corpora. The third part looks at the phases per word, rather than a total development of a child.

4.1 Truncation

As an initial testing ground, we try to replicate Table 1 from page 201 from Fikkert (1994). In this table, Fikkert showed that iambs are more likely to be truncated than trochees. We first discuss the results on the CLPF corpus (which is roughly the same corpus that Fikkert used) and then on the other corpora.

4.1.1 CLPF corpus

The results from the truncation study can be found in Table 2.

A few things are noteworthy. First of all, the script finds significantly more iambs and trochees in the dataset (right two columns) than reported in the table from Fikkert (left two columns). If we filter out all words with secondary stress (which are likely compounds), and words ending in 'en' or

	Trochees						Iambs					
	Manual		Script	(filtered)	Script	(original)	Manual		Script	(filtered)	Script	(original)
Eva	9	(7/78)	9.2	(32/347)	13.2	(20/152)	100	(27/27)	77.8	(21/27)	72.7	(8/11)
Elke	4	(6/152)	11.3	(37/328)	11.8	(28/238)	94	(59/63)	89.1	(49/55)	91.7	(22/24)
Jarmo	3	(7/219)	11.8	(67/566)	14.5	(54/372)	92	(46/50)	68.6	(24/35)	65.0	(13/20)
Tirza	7	(7/106)	6.3	(33/509)	7.0	(16/228)	62	(37/60)	52.5	(21/40)	55.0	(11/20)
Noortje	5	(9/188)	8.8	(33/373)	6.6	(19/286)	57	(41/72)	54.5	(24/44)	47.1	(8/17)
Robin	1	(3/283)	6.3	(40/636)	7.4	(24/323)	41	(26/63)	33.3	(18/54)	26.9	(7/26)
Tom	5	(6/120)	10.6	(54/511)	8.9	(30/338)	48	(42/87)	43.4	(33/76)	48.6	(18/37)
Leon	1	(4/274)	2.7	(16/582)	4.4	(16/360)	29	(35/120)	22.9	(22/96)	22.0	(13/59)
Leonie	0	(0/41)	7.9	(17/216)	12.0	(10/83)	14	(4/29)	8.0	(2/25)	10.5	(2/19)
Enzo	4	(7/194)	8.1	(30/369)	9.6	(18/188)	13	(14/105)	28.0	(14/50)	31.0	(9/29)
Catootje	2	(3/130)	5.4	(32/597)	7.0	(23/327)	11	(16/144)	8.8	(8/91)	2.9	(1/35)
David	N/A		5.4	(11/205)	7.9	(10/127)	N/A		18.5	(5/27)	22.2	(4/18)
Totals	N/A		8.9	(268/3022)	7.7	(402/5239)	N/A		36.8	(116/315)	38.9	(241/620)

Table 2: Frequency of truncation per child. Manual refers to the manual analysis as performed by Fikkert (1994); Script (filtered) are the numbers found by the script after filtering out some cases with complex morphology; Script (original) contains all instances found by the script. The first column reproduces the percentages of truncations from Fikkert (1994, p. 201), with the second column reproducing the absolute numbers. David was not included in this table, so we have no data for him to compare to. The third to sixth columns show the output of the script utilized in this paper. Column three and four are the results after filtering out most words that are not morphologically simplex. Column five and six are the original numbers, found by simply counting all occurrences.

'je' (both morphological endings in Dutch), we are still left with huge differences. These numbers are shown in the middle two columns.

There are some possible reasons for these differences. First of all, Fikkert's analysis only looked at cases in which the weak syllable is truncated and not where any syllable is truncated, which is what the script does. Second, there might be differences in defining what counts as truncation and what not.

In order to investigate the difference a bit more closely, we manually looked at some parts of the data. We chose to investigate the data for Leonie more closely as she had zero cases of trochaic truncation in the paper, while the script did find some cases. Looking at these cases manually, we see there are morphologically simplex words that have been truncated. For example, we find two instances of truncated 'borstel' (/ˈbɔɹstəl/, brush): /ˈhɔʃ/ and /ˈbɔːʃ/. After discussing this with Paula Fikkert herself, we believe this is likely the result of that Fikkert used only a subset of the corpus in her original investigation for Table 1 in her paper.

Despite the count differences, the general pattern described by Fikkert *does* emerge here. Indeed, we see that iambs are truncated significantly more often than trochees.

We have run the same analysis on the other corpora. The tables for those corpora are in Appendix A. Grimm has very similar results to Fikkert for the four children in the corpus. Fox-Boyer and Bracci have little truncation at all in their limited word set, but do suggest a similar pattern. It therefore seems that German children are also more likely to truncate iambs. The Providence corpus also produces similar results for English. We again find that iambs are more prone to truncation than trochees.

These findings were expected; all three languages are trochaic languages with similar stress assignment systems.

4.2 Developmental phases

This subsection is concerned with the developmental phases for iambic words as described by Fikkert. If we plot the production of a child through time, we expect this child to go through the different phases. In the notation used by the script, where T denotes stress and F denotes no stress, we expect the following:

- Phase 1: truncation, sometimes supplemented by addition of an extra schwa to fit the trochaic template. Truncation is denoted as 'T+F', meaning only one syllable is left[4]. The addition of a schwa is denoted as 'T@'.
- Phase 2: all adult syllables are produced, but stress is on the initial syllable rather than the final syllable. This is denoted as 'TF'.
- Phase 3: both syllables are produced with equal stress. This is denoted as 'TT'.
- Phase 4: adult-like production is reached, with stress on the final syllable, denoted as 'FT'.

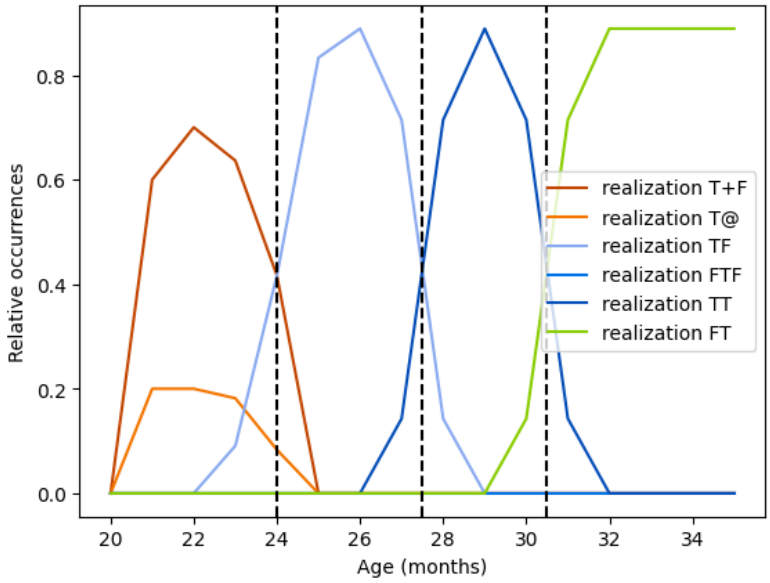

Figure 1: Hypothesis based on Fikkert (1994). The dotted lines indicate the end and start of a phase. The x axis denotes the age in months; the y axis the proportion of productions.

If we plot the production over time, we expect the proportion of productions of each kind to change as the child goes through the different phases. If a child perfectly follows the theory, we expect to see a plot like the one depicted in Figure 1. The child moves from phase to phase (with some overlap, as found by Fikkert and others), until it reaches adult production. We do not expect any other forms to be produced, although (as we will see later), other forms like 'FTF' do occur occasionally[5]. The duration of the phases could be longer or shorter than depicted here, but the phases should be discernible.

The following subsection discusses to what extent we encountered this pattern in our datasets.

4. Due to annotation inconsistencies in the corpora, single syllable words do not always have an annotation for stress. That is why we combine both stressed and unstressed single syllable words.
5. The graph includes 'FTF', as it turned out that children *do* produce this form. We would, based on Fikkert's theory, not expect them to produce this, which is why the production is hypothesized to be 0.

4.2.1 CLPF corpus

In figure 2, the results for the CLPF corpus are shown. Several things are to be noted here. First of all, we see that the relative graphs do not add up to 100 exactly. Looking at the data, it turns out that children occasionally produce forms that are not part of any phase, most often 'FTF', but occasionally also other forms. None of the present theories can explain why this happens. These occurrences are relatively rare, so they will be disregarded in the rest of the analyses.

Second, we note that a comparison of different children is somewhat difficult because not all children were measured over the same time period.

Third, we note that the phases are difficult to make out. Some children show a clearly demarcated stage 1, where the [False, True] pattern is realized in truncated form [True]. For example, Leon initially shows a large use of these truncated forms, which quickly declines after that. While Noortje initially seems to truncate 100% of the iambs, this percentage decreases around the age of 30 months. Yet, it is rather difficult to identify an end to stage 1 for most other children. We do see that most children do show an increase in their correct productions as they get older (the light green line), as was expected; they are learning how to correctly stress iambs.

Comparing these graphs to Figure 1, the phases described by Fikkert are not easily discernible, and especially phases 2 and 3 are particularly difficult to discern. Based on the underlying theoretical framework of UG, we would expect these phases to be clearer. After all, if certain parameters are set, we do not expect children to produce forms that do not fit with those parameters. Some overlap between neighboring phases might be expected, but many of these children produce utterances from three or four phases at the same time.

These conclusions can be explained in several ways. First of all, data is limited. The results are therefore sensitive to all kinds of flukes in the data, such as annotation mistakes or chance occurrences. Clearer patterns could emerge if more data were available per child. Grouping all children together could be one way of enabling that. As a small intermezzo, in order to see to what extent this might be the case, we also plotted all children together in Figure 3. Note that there is less data at the lower and higher ends of the x-axis, as not all children were measured at the full scale of the x-axis. On average, truncation is by far the most common way to produce iambs. Furthermore, as children grow older, they produce more adult-like iambs. The other two phases as described by Fikkert (1994) are hard to distinguish.

Before delving deeper into possible explanations, we first analyze to what extent the above findings also hold for the other corpora under investigation.

4.2.2 Other corpora

The development of individual children through time can only be plotted in longitudinal corpora, namely CLPF, Grimm and Providence. Figures 4 and 5 show the results of the analysis from Figure 2 on the Grimm and Providence corpora respectively.

The Grimm corpus does show some phases, but they are in different orders per child. Sandra starts out with mostly truncations and then moves on to two-syllable productions with stress sometimes on the initial and sometimes on the final syllable. Wiglaf seems to first produce trochaic productions before moving to a stage of truncation. Looking at the absolute data, we see that this concerns four productions. Eleanora shows the same pattern even more clearly. She starts out with more trochaic productions before moving on to truncations. After that, a mix of different productions emerges. At the end, it seems like the correct iambic production is emerging. Nele shows the same pattern as Wiglaf and Eleonora.

Moving on to the Providence corpus, we first look at the individual children in figure 5. For each of these children, we see that truncations seem to be the most common at the start (except for Alex and Lily), after which most productions are correct. However, during this stage, many children also produce trochees as well as producing both syllables stressed. What is perhaps most surprising is that here, there seems to be some U-shaped development occasionally, where a child

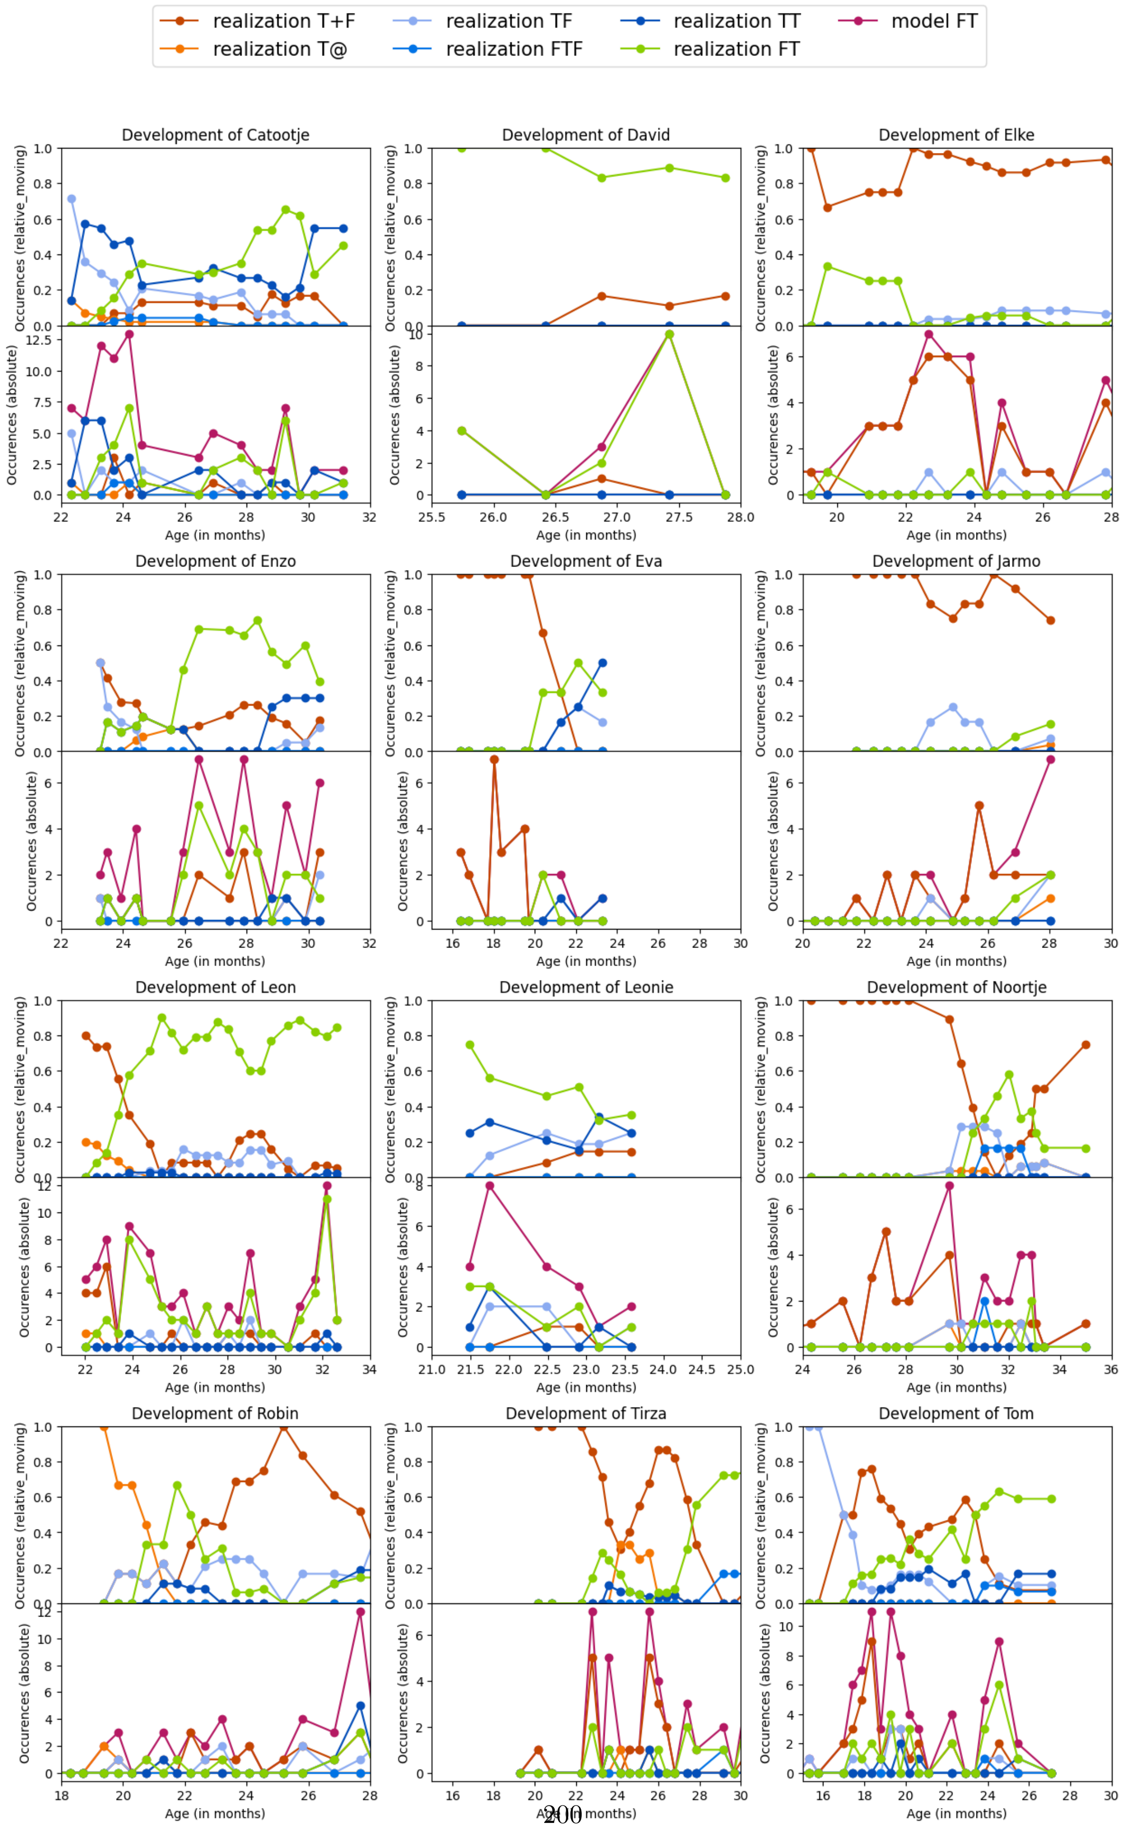

Figure 2: Approximation of child development for the children in the CLPF corpus of iambic words. The even rows show the relative proportion of each production type at the different ages (in months). The odd rows show the absolute numbers of each occurrence. NB: the children were tested on different age ranges (as shown in the x-axis).

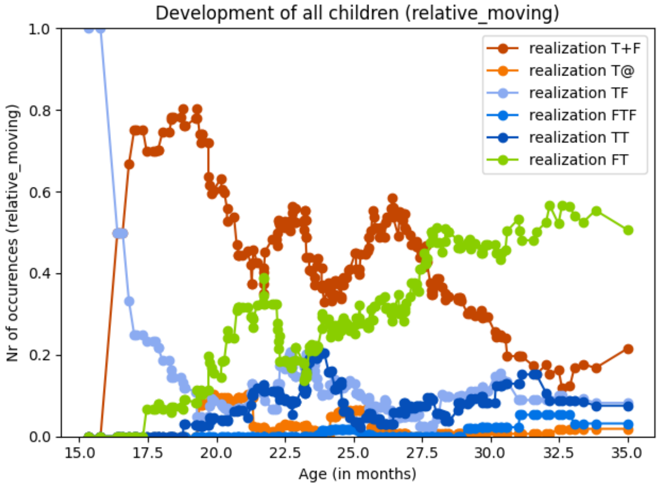

Figure 3: Approximation of child development for all children combined in the CLPF corpus of iambic words. This is the relative proportion of productions throughout time.

moves from correct productions to overgeneralizations, back to correct productions. This is most visible in Ethan's graph, between month 30 and 35, but also partly visible for Violet and, to some extent, Naima.

Combining all children in the Grimm corpus and all children in the Providence corpus (Figure 6), we can distinguish the stages more clearly. Again, the results for the Grimm corpus suggest that we should rearrange the order of the phases as compared to Fikkert's proposal for German. Initially, children apply a trochaic template to almost all words, before applying truncation. Only at the end of the data collection period (which ends at a little after two years of age), we see the start of correct productions.

For the Providence corpus, we see the same things as for individual children. An initial peak in truncations is followed by a slow increase in correct predictions. We also see a small U-shaped development around month 28, where more trochees are produced.

4.2.3 Preliminary discussion

There are several noticeable findings. First of all, we notice that we get quite different results for all three of the corpora. There might be several reasons for this. For starters, the datasets are limited, which makes the analysis more sensitive to flukes in the data. Second, there might be differences in the languages that cause these problems. As of now, we are unsure what types of difference might cause these different developments, especially since the languages are so similar in stress assignment (in particular Dutch and German), but this could be the topic of further investigations.

The second and perhaps most striking observation is that the phases are difficult to discern. That the phases are not clearly visible here does not automatically mean they do not exist. A downside of the combined visualization is that all children are aggregated and aligned based on their age, while we know that children go through developmental phases at different speeds. One should first align the children properly and shift their developments to plot them in a truly aggregated way. On the other hand, as discussed above, it was not possible to find the phases in the individual children either.

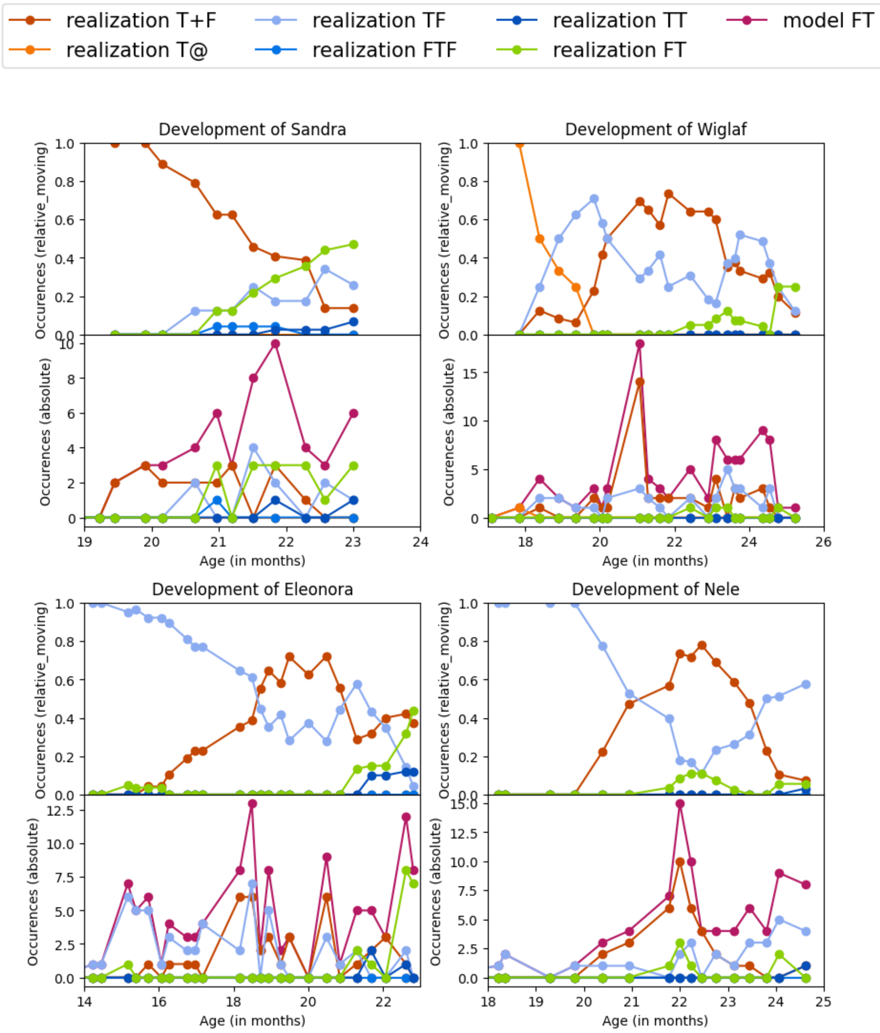

Figure 4: The relative and absolute development per child in the Grimm corpus. See figure 2 for a more detailed description.

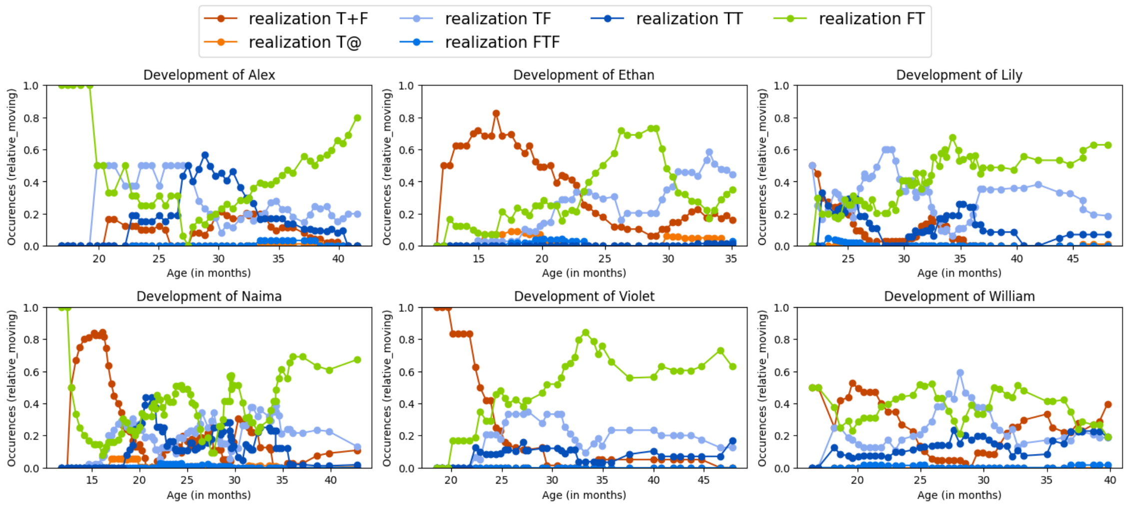

Figure 5: The relative development per child in the Providence corpus. See figure 2 for a more detailed description.

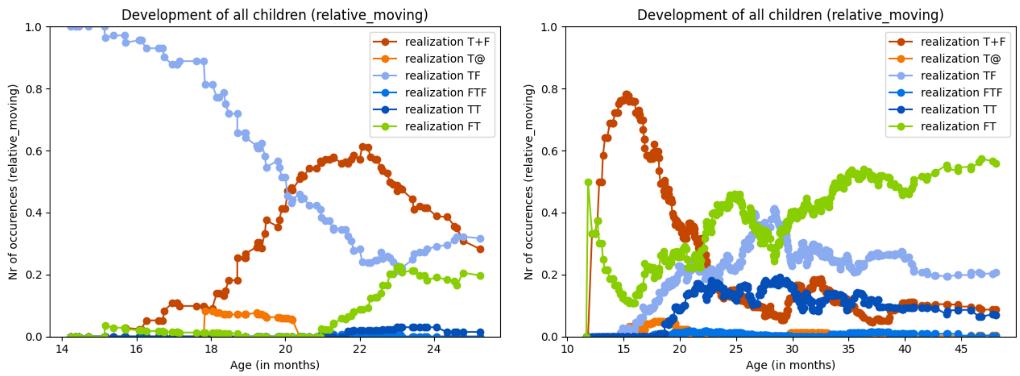

Figure 6: Approximation of child development for all children combined in the Grimm corpus (left) and Providence corpus (right) for iambic words. This is the relative proportion of productions throughout time.

Returning to the possible explanations for the difficulty discerning phases in each child, it is possible that there are data points that the script includes but should have excluded. Manual exploration of the dataset does indeed find some of these cases, but these are negligible in number.

It is also possible that the theory we are trying to support is not accurate in its entirety. Based on Fikkert's paper, there indeed seems to be support for the theory at word level, but perhaps not at an aggregated level. It therefore seems interesting to investigate to what extent phases are more demarcated when plotted *per word* per child instead of all together. This is explored in section 4.3.

4.3 Word-level development

As discussed above, it is difficult to distinguish phases throughout the whole development of a child. Given that previous research, among which Fikkert (1994), has found compelling evidence for the existence of phases and gives examples of this, we should carefully think about what might be happening. A possible explanation could be that the phases do exist at word-level, but that a child goes through the stages at different moments for different words. This hypothesis is investigated in this subsection. The method employed in this paper is especially well-suited to investigate a hypothesis such as this one, because it allows for easy alignment in plotting. An attempt is made to visualize word-level development to the extent possible with the limited data. After all, a drawback of this method is that one needs words that are produced sufficiently often throughout time to allow us to map out how the production of those words develops. There are sadly not many available in the datasets. Nonetheless, some words exist.[6] Like in the previous sections, the results for the CLPF corpus will be discussed first, after which the other corpora are covered.

4.3.1 CLPF corpus

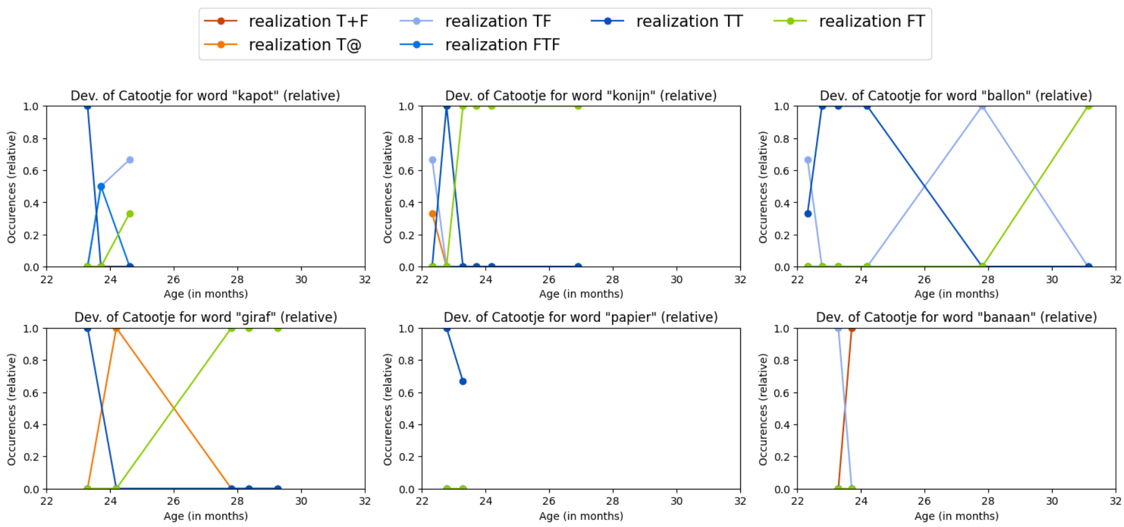

Figure 7: Catootje's development for six different words. The x-axis have been aligned to be the same for all words.

Figure 7 shows the development of Catootje for a few of those words. We notice that for some words (most clearly visible for *konijn*), Catootje reaches adult-like production at an age of 23-25

6. The selection criterium to include a word for this analysis was: the word was used by a child at least twice at two different times. This is the bare minimum needed to be able to compare the production of a word at the two different times.

months. The word *giraf* gets the first correct productions a little before month 28. For the word *ballon*, however, she still produces incorrect forms at least up to 28 months, with the first correct production appearing between month 28 and 32. This suggests that learning how to assign stress might not be a process that is only governed by universal grammar settings or by one general learning process, but might be related to the individual words. This is particularly interesting when taking into account that *giraf* and *ballon* both have heavy but not superheavy final syllables, and are in that sense quite similar.

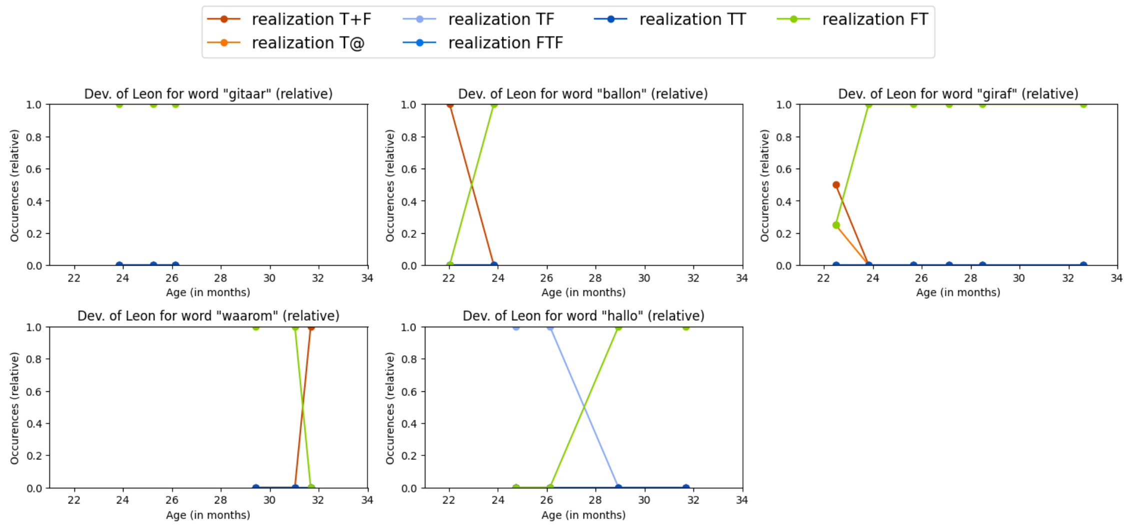

Figure 8: Leon's development for five different words. The x-axis have been aligned to be the same for all words.

Catootje is not the only child for which we encounter words like this. Looking at Leon, for whom some words are depicted in figure 8, we see that *gitaar*, *ballon* and *giraf* are properly produced from age 24 onward, while *hallo* is improperly produced until Leon is about 29 months of age. Similarly, for Tom (depicted in figure 9), we see that *trompet* is produced properly in all instances while *kameel* is not and the production of *marmot* changes from 20 to 24 months.

These findings therefore suggest that more research is necessary into determining to what extent stress pattern acquisition is a holistic process or a process influenced by individual words. On the one hand, the findings argue that there is something to say for the latter: the production of different words seems to develop at different paces. On the other hand, the data is scarce and easily influenced by outliers. Therefore, in order to properly investigate this question further, more data is needed.

4.3.2 Other corpora

The same word level analysis has been performed on the Grimm and Providence corpora. Similar patterns emerge. For example, in the Grimm corpus, Sandra (figure 10) shows different developments per word. *Ballon* is produced correctly earlier than *Kamel* or *Gela*.

The Providence corpus (American English) also provides us with similar findings. In this corpus, we see more volatile productions: some words are correctly produced, and then incorrectly again. Besides that, we again find large differences per word in the same child. These persist if we take word-type (noun, verb, or something else, which matters in English stress assignment) into account.

Besides the Grimm and Providence corpora, we can also take a look at two cross-sectional corpora. While cross-sectional corpora are not ideal for the type of research that is being performed

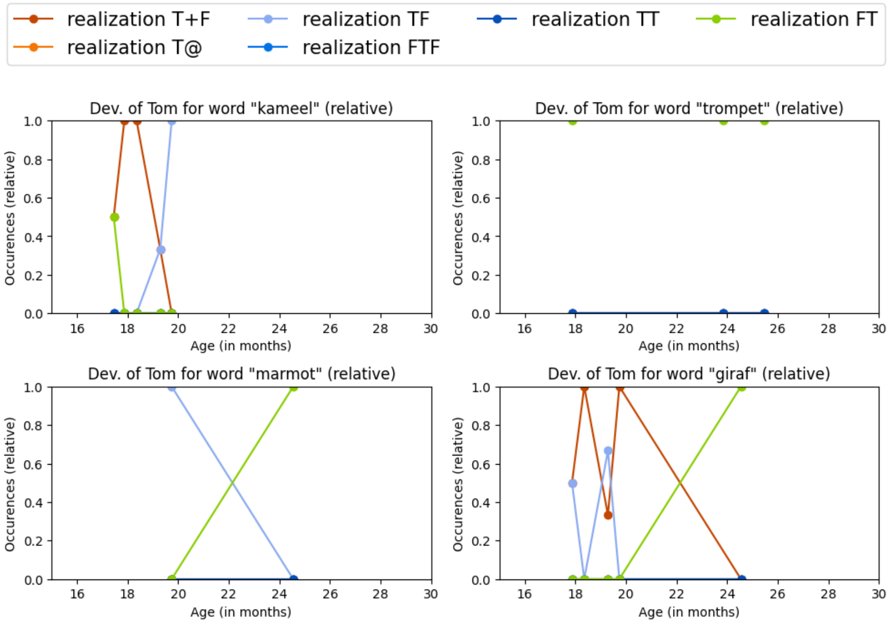

Figure 9: Tom's development for four different words. The x-axis have been aligned to be the same for all words.

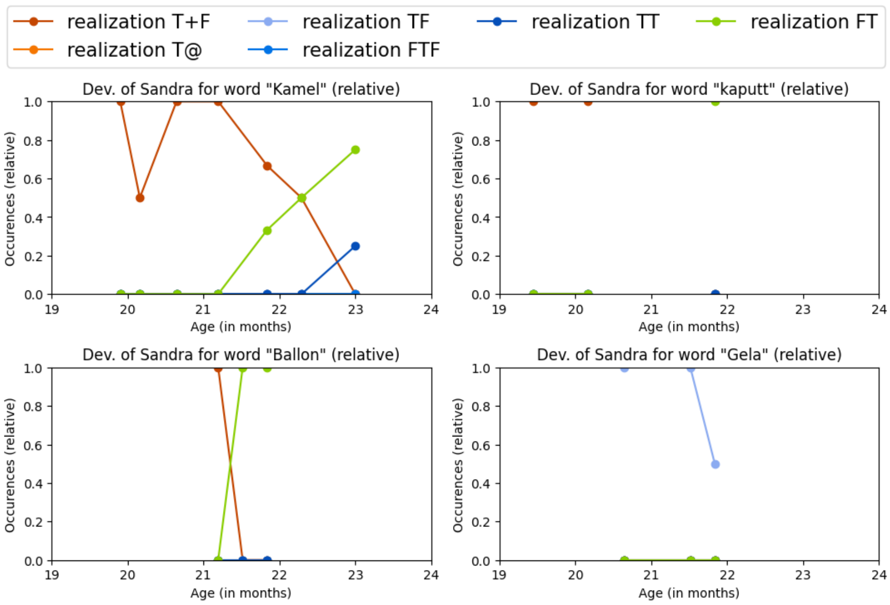

Figure 10: Sandra's development for four different words. The x-axis have been aligned to be the same for all words.

in this paper in general, they can provide supporting evidence. In our case, the Bracci and Fox-Boyer corpora of German speaking children indeed do so. The benefit of these corpora is that children were asked to produce specific words, meaning that we can track developments of these words throughout time. A drawback of the corpora is their cross-sectional nature. We can no longer interpret the graphs as the development of a single child through time, as every data point comes from an individual child. Nonetheless, we can gain some information from these graphs anyways as long as we interpret it as a sort of average development.

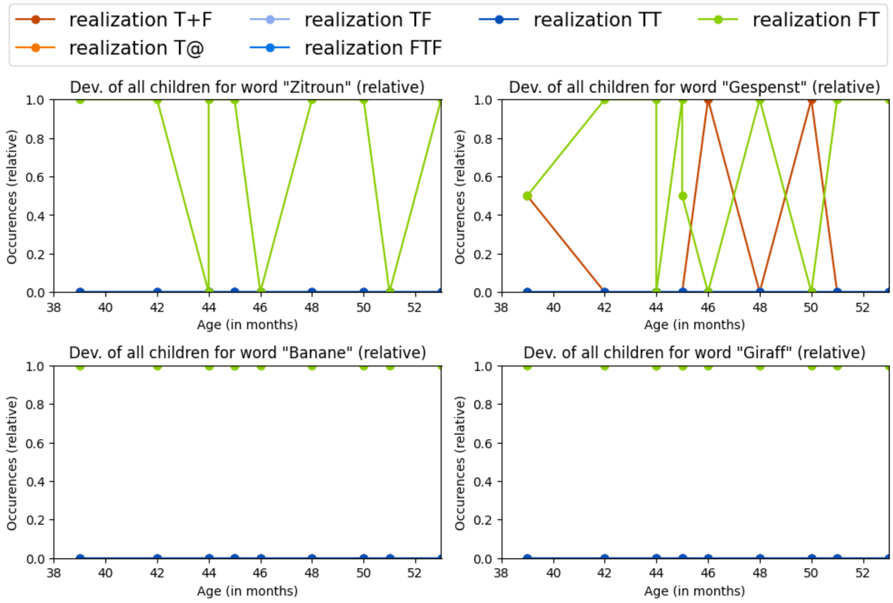

Figure 11: Relative proportion of production for the four bisyllabic iambic words in the Bracci corpus.

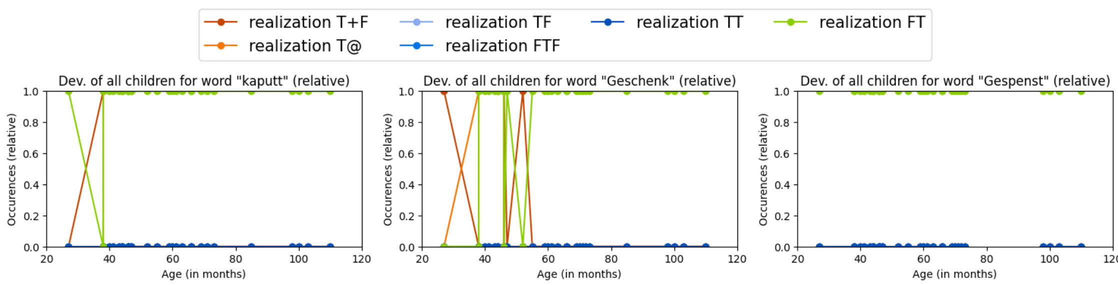

Figure 12: Relative proportion of production for the three bisyllabic iambic words in the Fox-Boyer corpus.

As for the Bracci corpus, we indeed find differences between words. These results are shown in figure 11. The words 'Giraff' and 'Banane' are always stressed correctly by all children, while the words 'Zitroun' and 'Gespenst' are not. Similarly, for the Fox-Boyer corpus, shown in figure 12, the word ‘Gespenst’[7] is correctly produced by all children and the word ‘kaputt’ is correctly produced in all except for one child, while the word ‘Geschenk’ is often produced incorrectly up to when the children are about 60 months of age.

Overall, these findings together suggest that there might be differences between words within the same child.

5. Discussion

The current investigation set out to investigate stress pattern acquisition in children by employing new analysis methods. The previous sections laid out this new method which uses phonetic transcriptions to automatically determine what stress pattern was produced by a child and then collect different types of productions. These can then be graphed in different ways to investigate the development of the child. Several conclusions follow from the investigation.

First of all, our investigation replicated some past findings. We found that iambs are more difficult for children learning Germanic languages than trochees. In all of our corpora, children truncate iambs more frequently than they truncate trochees, and they are more likely to assign stress in a non-adultlike way. It is important to note that while this held for all of the languages investigated, it is not necessarily universal. al Huneety et al. (2023) found that children can also start with a iambic template when learning Ammani Arabic and Tzakosta (2004) found that children learning Greek do not seem to show any template. Our findings are therefore likely the result of the abundance[8] of trochees in the Germanic languages.

Second, while past theories of acquisition focus on universal stages of development, we found evidence for stages of development per word. Universal stages consist of stages in which children develop their linguistic knowledge to include different information such as knowledge of prosodic feet in their production of words and do so for all words at the same time. After having developed such a notion, the theory goes, children should be able to use them in their production of all words to which the notion applies. It has been difficult to find evidence for this in our results. This might be due to the nature of the analysis, which has some drawbacks (see below). However, not everything can be explained by these drawbacks. A possible way to understand the data better might be to consider prosodic development at word level. Our data analysis shows that this fits better with the data at hand for all three of our different languages. It therefore seems wise for future investigations to delve deeper into this possibility. First of all, it is necessary to replicate the findings here with more fine-grained data (see also below on limitations of the data sets). Second, if the findings are replicated, we can wonder why some words are learned earlier than others. Aspects such as frequency effects come to mind. We also noted that a word like *hello* in English was more difficult to learn, perhaps because (depending on what theory of stress you believe in for English) it can be considered an actual exception as the last syllable is not superheavy. It is also grammatically difficult to categorize, which could also lead to phonologically special behavior. Kehoe and Stoel-Gammon (1997) furthermore already suggested that the onset of a syllable can also be of influence in how difficult a stress assignment is to learn. This could also be examined further. Another aspect in need of investigation is how the development of individual words correlates to the development of the consciousness of linguistic constructs such as prosodic feet.

A third conclusion to be drawn from this investigation is that the method proposed has benefits and drawbacks. The benefit of the current method is that it is particularly easy to extract all kinds of data or cases from the dataset. For example, it becomes easy to find cases that do not adhere to a proposed theory. This method of working therefore allows us to test stress acquisition theories in

7. Note that the Fox-Boyer corpus has mostly older children than the Bracci corpus. This might have caused the different results for the two corpora; while the word ‘Gespenst’ might have been difficult for the children in Bracci, it no longer is for the children in Fox-Boyer.
8. Note that in the paper by Tzakosta (2004) it was noted that iambs were more common in the earliest words that a child learns, even though the language is trochaic.

new and easier ways. On the other hand, we noted that real life data is often messy and it becomes harder to find certain patterns in the data set that might be easier to find manually. A benefit of the method is that it allows for quick inclusion of more corpora as long as they have been annotated sufficiently. As soon as the chosen analysis method has been programmed out, it is simply a matter of reading in another corpus and pulling it through the pipeline to analyze that corpus.

The method is limited by the data available. The different data sets used each come with their own limitations and do not cover the whole receptive or productive capability of a child. For example, Fikkert (2020) herself has noted that the one hour twice a month in her dataset is limiting. Because the dataset does not cover all capabilities of children, it turned out to be too difficult to test hypotheses like the one of Yang (2016), which was an original goal of this investigation. The data set also limited us by not allowing all types of analyses. While it would have been interesting to test the hypotheses put forward by Kehoe and Stoel-Gammon (1997) on threesyllabic words and secondary stress, this was not possible due to limitations in the data. First of all, the difference between primary and secondary stress was often not annotated consistently in the corpora. Second, the corpora also did not contain many three syllable words of the forms that Kehoe makes predictions about. This made it impossible to perform such analyses. This is left for future research, for researchers with more suitable data sets.

In general, a drawback of the method comes from its automatic nature. The benefit of manual analysis is that you see every data point and can carefully consider what to take into account. While this automatic method is supported by manual inspection before, during, and after the development and application of the algorithm, not all data points have been reviewed by the researcher. Mistakes are therefore more likely to happen. For example, it is impossible in our script to know what syllable has been truncated, as it does not align sounds. Automatically aligning sounds has proven to be complicated in other systems such as Phon (Rose and MacWhinney 2014, Rose and Stoel-Gammon 2015), especially pertaining to child language[9]. Furthermore, the current system is better suited for testing theories than for developing new theories; you need to know what you are looking for when using this method. On the other hand, manual analysis is also prone to mistakes, albeit of a different type. The method proposed here is therefore mostly suited for large-scale analyses where some small mistakes can be forgiven.

Besides the method in question, there are also limitations to the research in general. Most notably, this investigation, like many before, only focuses on Indo-European, Germanic languages. The conclusions drawn from this might not apply to other languages in other language families. Future research should therefore investigate to what extent conclusions drawn here and in related research hold up for different languages (Kidd and Garcia 2022). It would be interesting to include the other languages that are available in PhonBank, which should be relatively easy given that they are annotated in the same way as the corpora used here.

References

Adam, Galit and Outi Bat-El (2009), When Do Universal Preferences Emerge in Language Development? The Acquisition of Hebrew Stress, *Brill's Annual of Afroasiatic Languages and Linguistics* **1** (1), pp. 255–282.

al Huneety, Anas, Bassil Mashaqba, Mais Abu Alhala, Mohammed Nour Abu Guba, and Khalid G. Al-Shdifat (2023), Acquisition of stress in the speech of Ammani Arabic-speaking children, *Ampersand* **10**, pp. 100094. https://linkinghub.elsevier.com/retrieve/pii/S2215039022000133.

Archibald, John (1995), The Acquisition of Stress, *in* Archibald, John, editor, *Phonological Acquisition and Phonological Theory*, Psychology Press, United States of America. Google-Books-ID: ryIBAwAAQBAJ.

9. But see advances in the field of sequence alignments, such as Prokić et al. (2009).

Booij, Geert (1999), *The Phonology of Dutch*, Clarendon Press.

Bracci, Anne (2020), PhonBank Luxembourgish Bracci Corpus. https://phon.talkbank.org/access/German/Bracci.html.

Curtin, Suzanne, Toben H. Mintz, and Morten H. Christiansen (2005), Stress changes the representational landscape: evidence from word segmentation, *Cognition* **96** (3), pp. 233–262. https://www.sciencedirect.com/science/article/pii/S0010027704001726.

Daelemans, Walter, Gert Durieux, and Steven Gillis (1994), The Acquisition of Stress: A Data-Oriented Approach, *Computational Linguistics* **20** (3), pp. 31.

Demuth, Katherine (2004), CHILDES English Providence Corpus. https://phon.talkbank.org/access/Eng-NA/Providence.html.

Durand, Jacques, Ulrike Gut, and Gjert Kristoffersen, editors (2014), *The Oxford Handbook of Corpus Phonology*, Oxford University Press. https://academic.oup.com/edited-volume/38637.

Fikkert, Paula (1994), *On the Acquisition of Prosodic Structure*, PhD thesis, Rijksuniversiteit Leiden, The Hague.

Fikkert, Paula (2020), De verwerving van woordklemtoon in perspectief, *Nederlandse Taalkunde* **25** (2-3), pp. 295–304. Publisher: Amsterdam University Press. https://www.aup-online.com/content/journals/10.5117/NEDTAA2020.2-3.012.FIKK.

Fox-Boyer, Annette (2019), PhonBank German Fox-Boyer Corpus. https://phon.talkbank.org/access/German/FoxBoyer.html.

Féry, Caroline (1998), German Word Stress in Optimality Theory, *Journal of Comparative Germanic Linguistics* **2** (2), pp. 101–142. Publisher: Springer. https://www.jstor.org/stable/23738532.

Gerken, Louann (1994), A metrical template account of children's weak syllable omissions from multisyllabic words, *Journal of Child Language* **21** (3), pp. 565–584.

Grimm, Angela (2008), *The Development of Early Prosodic Word Structure in Child German*, PhD thesis, Universität Potsdam.

Guion, Susan G., J.J. Clark, Tetsuo Harada, and Ratree P. Wayland (2003), Factors Affecting Stress Placement for English Nonwords include Syllabic Structure, Lexical Class, and Stress Patterns of Phonologically Similar Words, *Language and Speech* **46** (4), pp. 403–426. Publisher: SAGE Publications Ltd. https://doi.org/10.1177/00238309030460040301.

Hall, Kathleen Currie, J. Scott Mackie, and Roger Yu-Hsiang Lo (2019), Phonological CorpusTools: Software for doing phonological analysis on transcribed corpora, *International Journal of Corpus Linguistics* **24** (4), pp. 522–535. Publisher: John Benjamins. https://www.jbe-platform.com/content/journals/10.1075/ijcl.18009.hal.

Joanisse, Marc and Suzanne Curtin (1999), Dutch stress acquisition: OT and connectionist approaches, *Proceedings of SWOT*.

Kager, René Willibrord Joseph (1989), *A Metrical Theory of Stress and Destressing in English and Dutch*, PhD Thesis, Utrecht University, Utrecht. https://dspace.library.uu.nl/handle/1874/8680.

Kehoe, Margaret (1998), Support for metrical stress theory in stress acquisition, *Clinical Linguistics & Phonetics* **12** (1), pp. 1–23. http://www.tandfonline.com/doi/full/10.3109/02699209808985210.

Kehoe, Margaret and Carol Stoel-Gammon (1997), The Acquisition of Prosodic Structure: An Investigation of Current Accounts of Children's Prosodic Development, *Language* **73** (1), pp. 113. https://www.jstor.org/stable/416597?origin=crossref.

Kehoe, Margaret M. (2001), Prosodic Patterns in Children's Multisyllabic Word Productions, *Language, Speech, and Hearing Services in Schools* **32** (4), pp. 284–294. http://pubs.asha.org/doi/10.1044/0161-1461

Kidd, Evan and Rowena Garcia (2022), How diverse is child language acquisition research?, *First Language* **42** (6), pp. 703–735. http://journals.sagepub.com/doi/10.1177/01427237211066405.

Levelt, Clara (1994), *On the Acquisition of Place*, PhD thesis, Rijksuniversiteit Leiden, The Hague.

McAuliffe, Michael, Arlie Coles, Michael Goodale, Sarah Mihuc, Michael Wagner, Jane Stuart-Smith, and Morgan Sonderegger (2019), ISCAN: a System for Integrated Phonetic Analyses Across Speech Corpora. Conference Name: 19th International Congress of Phonetic Sciences (ICPhS 2019) ISBN: 9780646800691 ISSN: 0241-0669 Meeting Name: 19th International Congress of Phonetic Sciences (ICPhS 2019) Pages: 1322-1326 Place: Canberra, Australia Publisher: Australasian Speech Science and Technology Association Inc. https://eprints.gla.ac.uk/183719/.

Oh, Grace E., Susan Guion-Anderson, and Melissa A. Redford (2011), Developmental Change in Factors Affecting Stress Placement in Native English-Speaking Children, *Proceedings of the 17th International Congress of Phonetic Sciences* pp. 1522–1525.

Prokić, Jelena, Martijn Wieling, and John Nerbonne (2009), Multiple sequence alignments in linguistics, *Proceedings of the EACL 2009 Workshop on Language Technology and Resources for Cultural Heritage, Social Sciences, Humanities, and Education - LaTeCH-SHELT&R '09*, Association for Computational Linguistics, Athens, Greece, pp. 18–25. http://portal.acm.org/citation.cfm?doid=1642049.1642052.

Rietveld, A. C. M and Vincent van Heuven (2016), *Algemene fonetiek*, vierde, herziene druk ed., Uitgeverij Coutinho, Bussum. OCLC: 1003284573. https://search.ebscohost.com/login.aspx?direct=truescope=sitedb=nlebkdb=nlabkAN=1562638.

Rose, Yvan and Brian MacWhinney (2014), The PhonBank Project: Data and Software-Assisted Methods for the Study of Phonology and Phonological Development, *in* Durand, Jacques, Ulrike Gut, and Gjert Kristoffersen, editors, *The Oxford Handbook of Corpus Phonology*, Oxford University Press, p. 0. https://doi.org/10.1093/oxfordhb/9780199571932.013.023.

Rose, Yvan and Carol Stoel-Gammon (2015), Using PhonBank and Phon in Studies of Phonological Development and Disorders, *Clinical linguistics & phonetics* **29** (0), pp. 686–700. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4774542/.

TalkBank (n.d.), PhonBank. https://phonbank.talkbank.org/.

Trommelen, Mieke and Wim Zonneveld (1999), Word Stress in West-Germanic Languages, *in* van der Hulst, Harry, editor, *Word Prosodic Systems in the Languages of Europe*, De Gruyter Mouton, Berlin, pp. 478–515.

Tzakosta, Marina (2004), *Multiple Parallel Grammars in the Acquisition of Stress in Greek L1*, PhD Thesis, Universiteit Leiden, Leiden.

Yang, Charles (2016), *The Price of Linguistic Productivity: How Children Learn to Break the Rules of Language*, The MIT Press, Cambridge, Massachusetts.

Özçelik, Öner (2011), *Representation and Acquisition of Stress: the Case of Turkish*, PhD Thesis, McGill University, Montréal, Québec, Canada.

	TF		FT	
Sandra	9.2	(67/727)	41.7	(25/60)
Wiglaf	5.5	(38/689)	44.9	(44/98)
Eleonora	13.7	(109/798)	32.2	(38/118)
Nele	7.1	(50/702)	46.0	(40/87)
Totals	9.1	(264/2916)	40.5	(147/363)

Table 3: Frequency of truncation per child for the Grimm corpus according to the script employed in this study. See description of table 2 for more details.

	TF		FT	
Alex	16.8	(390/2315)	14.2	(20/141)
Ethan	17.9	(303/1691)	38.1	(61/160)
Lily	7.6	(260/3405)	7.7	(20/261)
Naima	9.2	(397/4319)	17.6	(70/398)
Violet	14.7	(156/1062)	9.6	(11/115)
William	25.9	(419/1617)	28.6	(48/168)
Totals	13.4	(1925/14409)	18.5	(230/1243)

Table 4: Frequency of truncation per child for the Providence corpus according to the script employed in this study. See description of table 2 for more details.

Appendix A. Truncations for other corpora

The truncation tables for the other corpora are shown in the tables below. No attempt was made to filter out morphologically complex instances, because the rules are different in Dutch on the one hand, and German and English on the other. It would have required a more thorough examination of German and English morphology. In the following, TF = 'True False', i.e. stress on the first syllable (a trochaic pattern), FT = 'False True', i.e. stress on the last syllable (an iambic pattern).

	TF		FT	
3028	0.0	(0/18)	0.0	(0/4)
3025	0.0	(0/18)	25.0	(1/4)
1005	0.0	(0/18)	0.0	(0/4)
3022	0.0	(0/18)	0.0	(0/4)
3023	0.0	(0/18)	0.0	(0/4)
1006	0.0	(0/18)	0.0	(0/4)
3024	0.0	(0/18)	0.0	(0/4)
2017	0.0	(0/18)	0.0	(0/4)
2011	0.0	(0/18)	0.0	(0/4)
2027	0.0	(0/18)	0.0	(0/4)
1002	0.0	(0/18)	0.0	(0/4)
2008	0.0	(0/18)	0.0	(0/4)
2015	0.0	(0/18)	25.0	(1/4)
1001	0.0	(0/18)	25.0	(1/4)
1003	0.0	(0/18)	0.0	(0/4)
Totals	0.0	(0/270)	5.0	(3/60)

Table 5: Frequency of truncation per child for the Bracci corpus according to the script employed in this study. See description of table 2 for more details.

	TF		FT	
TD PL-II C01	0.0	(0/41)	0.0	(0/3)
TD PL-II C02	0.0	(0/42)	0.0	(0/3)
TD PL-II C03	0.0	(0/44)	0.0	(0/3)
TD PL-II C04	0.0	(0/44)	0.0	(0/3)
TD PL-II C05	0.0	(0/42)	33.3	(1/3)
TD PL-II C06	0.0	(0/41)	0.0	(0/3)
TD PL-II C07	0.0	(0/41)	33.3	(1/3)
TD PL-II C08	0.0	(0/41)	0.0	(0/3)
TD PL-II C09	0.0	(0/41)	0.0	(0/3)
TD PL-II C10	3.8	(2/52)	25.0	(1/4)
TD PL-II C11	0.0	(0/42)	0.0	(0/3)
TD PL-II C12	0.0	(0/41)	0.0	(0/3)
TD PL-II C13	0.0	(0/45)	0.0	(0/4)
TD PL-II C14	0.0	(0/41)	0.0	(0/3)
TD PL-II C15	2.3	(1/43)	0.0	(0/3)
TD PL-II C16	0.0	(0/45)	33.3	(1/3)
TD PL-II C17	0.0	(0/41)	0.0	(0/3)
TD PL-II C18	2.4	(1/41)	0.0	(0/3)
TD PL-II C19	0.0	(0/41)	0.0	(0/3)
TD PL-II C20	0.0	(0/36)	0.0	(0/3)
TD PL-II C21	0.0	(0/40)	0.0	(0/3)
TD PL-II C22	0.0	(0/41)	0.0	(0/3)
TD PL-II C23	0.0	(0/37)	0.0	(0/2)
TD PL-II C24	0.0	(0/40)	0.0	(0/3)
TD PL-II C25	0.0	(0/38)	0.0	(0/3)
TD PL-II C26	2.5	(1/40)	0.0	(0/3)
TD PL-II C27	0.0	(0/38)	0.0	(0/2)
TD PL-II C28	0.0	(0/41)	0.0	(0/3)
TD PL-II C29	0.0	(0/41)	0.0	(0/2)
TD PL-II C30	0.0	(0/41)	0.0	(0/3)
TD PL-II C31	0.0	(0/41)	0.0	(0/3)
TD PL-II C32	0.0	(0/41)	0.0	(0/3)
Totals	0.4	(5/1324)	4.2	(4/95)

Table 6: Frequency of truncation per child for the Fox-Boyer corpus according to the script employed in this study. See description of table 2 for more details.