Title

Frequency vs. Salience in First Language Acquisition: The Acquisition of Aspect Marking in Chintang

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

Authors

Mazara, Jekaterina
Stoll, Sabine

Publication Date

2021

Frequency vs. Salience in First Language Acquisition: The Acquisition of Aspect Marking in Chintang

Jekaterina Mažara (jekaterina.mazara@uzh.ch)
Department of Comparative Language Science & Center for the Interdisciplinary Study of Language Evolution,
University of Zurich, Switzerland

Sabine Stoll (sabine.stoll@uzh.ch)
Department of Comparative Language Science & Center for the Interdisciplinary Study of Language Evolution,
University of Zurich, Switzerland

Abstract

Frequency of occurrence in the input is a main factor determining the ease of acquisition in first language learners. However, little is known about the factors relevant for the acquisition of low-frequency items. We examine the use of aspectual markers in a longitudinal corpus of Chintang (Sino-Tibetan, Nepal) children (ages 2;1-4;5). Only 7.7% of all Chintang verbs are overtly marked for aspect. Chintang has three aspect markers, one of which is substantially more frequent than the others. One of the low-frequency markers is positionally and prosodically more salient, appearing at the word-boundary. Using a Bayesian beta-binomial model, we assess the distribution and flexibility of use of aspectual markers in the input and children's production. Our analysis shows that the most frequent marker is acquired earliest, as predicted. For the low-frequency markers, position, segmentability and uniformity are better predictors of ease of acquisition.

Keywords: language acquisition; grammatical productivity; verb morphology; frequency; salience; modeling; corpus study

Introduction

Language learning is a gradual process which takes several years of exposure to surrounding and child-directed speech and constant interaction (Tomasello, 2003; Lieven, Pine, & Baldwin, 1997). For languages that use grammatical morphemes, this process includes identifying morphemes in the input and learning how to combine them into word forms. Questions concerning the factors determining order of acquisition have been central to language acquisition research from the earliest corpus studies (cf. R. Brown, 1973), and has since been studied for a host of linguistic phenomena in L1 and L2 acquisition. Multiple factors can play a role in the trajectory and speed of this development. Here, we focus on frequency distributions (cf. Naigles & Hoff-Ginsberg, 1998; Theakston, Lieven, Pine, & Rowland, 2004; Ambridge, Kidd, Rowland, & Theakston, 2015 for an overview), including presence of an element in a wide variety of contexts (Naigles & Hoff-Ginsberg, 1998; Küntay & Slobin, 2002), and salience of cues, while acknowledging that both concepts can themselves be broken down into many individual variables. Salience is a concept that is frequently invoked but notoriously fuzzy and varied in its definition (see e.g. Goldschneider & DeKeyser, 2001). It includes ease of segmentation from the speech stream (Monaghan & Christiansen, 2010), salience due to rare occurrence (Stoll & Bickel, 2013), prosodic structure (Christophe & Dupoux, 1996), and position of the word or marker in the phrase or word unit (Freudenthal, Pine, Aguado-Orea, & Gobet, 2007; Freudenthal, Pine, & Gobet, 2010; Longobardi, Rossi-Arnaud, Spataro, Putnick, & Bornstein, 2015).

Frequency of use is known to be one of the main determinants for the acquisition of elements (Ambridge et al., 2015); the more frequent an element or category the easier and faster it is expected to be acquired. So far, most attention has been given to forms that are acquired earliest. In this paper, we focus on a relatively low-frequency grammatical feature and the order of acquisition of the least-used markers of the category. We examine the input distributions and production patterns of aspectual markers in Chintang, a polysynthetic Sino-Tibetan language spoken in Nepal. Grammatical aspect is used to express the inner temporal structure of individual events. While this category is expressed obligatorily in some languages (e.g. Slavic languages), thus presenting the children with many exemplars of structure and use, other languages express aspectual meaning infrequently. Chintang overtly marks aspect infrequently, using three markers that differ along two dimensions: frequency and positional and prosodic salience. Here, we assess the role of frequency and salience on the development of use of these three markers. First, we compare the frequency distributions of the markers in the ambient language to the target children's production to examine whether the children's acquisition trajectory follows the input. For the less frequent two markers, we then fit a Bayesian beta-binomial model to characterize the flexibility of form use in children as well as their preference for one of the markers.

Chintang and its verbal morphology

Chintang[1] Verb structure is fairly complex: verb forms are inflected for a large number of categories, expressed by affixes. Verbs inflect for agreement with up to two arguments (in three persons, three numbers, and clusivity), tense, aspect, mood, and polarity. These markers are encoded in up to 4 prefix and 11 suffix positions. The number of morphemes in verb forms can also vary greatly. In the corpus, the number

[1] Chintang is both the name of the village and the language spoken by its inhabitants. It is located in the foothills of the Himalayas in the district of Dhankutā in Eastern Nepal and belongs to the Kiranti group of Sino-Tibetan. There are around 6,000 Chintang speakers, who are bilingual in Nepali (Indo-European). At the time of recording the language was still transmitted to children and used in many everyday situations. However, children learn Nepali, the lingua franca of Nepal, early in life.

of morphemes in child-surrounding speech varies between as few as 2 and as many as 10. A further complication for learners is added by the fact that the relative position of affixes within a word-unit can vary freely without a change in meaning (Bickel et al., 2007). These factors add up to a very large number of morphemes and morpheme combinations [2], which leads to sparser distributions of individual categories and a lower cue availability and reliability (cf. MacWhinney, 1978; MacWhinney, Pléh, & Bates, 1985) for most markers save for a few highly frequent ones.

Aspect markers in Chintang Aspect is morphologically marked on a mere 7.7% of verbs in the input Chintang children hear; the marking of aspect is not obligatory. Chintang has three aspectual markers, which have several allomorphs[3]. Two markers appear within the verb unit (*-ŋs* and *-akt*), while a third one is an enclitic (*=ta*), and thus, consistently appears after the verb unit, at the word boundary.

The first word-internal morpheme, *-ŋs*, expresses a meaning similar to the English present perfect (see examples 1 and 2).[4] It marks events that happened in the past but resulted in something that can still be perceived at speech time. The marker *-ŋs* appears predominantly with past tense.

(1) *them wad-o-**ŋs**-e?*
what put.on-3P-**PRF**-IND.PST
What is (he) wearing? [CLLDCh1R12S02]

(2) *miʔmuŋ na*
a.little TOP
*dam-u-mai-phak-u-okt-a-**ŋs**-e?*
scratch-3nsS/A-NEG-scratch-3P-NEG-PST-**PRF**-IND.PST
naŋ copt-a-kha-nu-m-c-u-m!
BUT see-IMP-NMLZ-2/3p-1/2nsA-d-3P-1/2nsA
A bit has not been dug yet, look! [CLLDCh1R02S03b]

The second word-internal aspectual marker, *-akt*, marks imperfective (i.e. focusing on internal structure of an event) durative or habitual events and actions. According to Schikowski, 2012, it is compatible with all verb forms, but in practice, most *-akt* forms appear with past marking (see example (3)).

(3) *ba balti bhayu yuw-a-**k**-e.*
DEM.PROX bucket DEM.PROX LOC.ACROSS
be.there-PST-**IPFV**-IND.PST
This bucket was here. [CLLDCh2R01S02a]

These examples illustrate the variety in length of the word units as well as the variation in number of other markers and length of word-units. This variation might complicate the learning task, since children cannot rely on a marker's position, nor the similarity of the verb forms it appears in. Additionally, *-akt* has three allomorphs and can appear as *-yakt*, *-akt*, or *-kt*.

[2]Stoll, Mazara, & Bickel, 2017 counted 4,745 unique combinations of grammatical markers.

[3]The description of Chintang aspect is based on Schikowski, 2012 as well as talks by Balthasar Bickel and corpus exploration.

[4]All examples are taken from children's input.

The third aspectual marker is an enclitic *=ta*. It appears reliably in the same position, at the end of the verb form, and is phonologically less integrated with the verb form than word-internal morphemes. *=ta* expresses non-past ongoing actions. While it also denotes imperfective meaning, in contrast to *-akt*, *=ta* cannot express habitual or generic situations, nor can it be used to express ongoing events in the past or future (Bickel et al., 2005), for example (4).

(4) *kanchi a-ca-no**=ta** elo?*
youngest.daughter 2-eat-NPST**=IPFV** EMPH.Q
Kanchi, are you eating? [CLLDCh3R01S03]

Since *-ŋs* is relatively frequent in input, we expect that children will use this marker more productively earlier on. For the lower frequency morphemes, however, we expect frequency to be a bad predictor and salience features to play a role. Since salience can encompass a whole host of subdomains, we focus on those that constitute a contrast between the markers *=ta* and *-akt*, which could otherwise appear with the same selection of lexemes. These features are

1. Positional salience: *=ta* appears at the word-boundary in 100% of cases, while *-akt* varies in its position within the verb
2. Prosodic salience: *=ta* is prosodically less bound to the word-unit and is therefore easier to perceive
3. Transparency (MacWhinney, 1978): *=ta* does not vary in its surface form, while *-akt* has three allomorphs and is in part similar to the *past* morpheme, which can appear in an adjacent slot in the same verb form

Based on these criteria, we expect *=ta* to be acquired earlier.

Data

Table 1: Age spans of the target children and number of words and verbs produced by the children and surrounding adults

Focal Child	Age span	Number of hours	N(tokens)			
			Child		Ambient	
			words	verbs	words	verbs
3	2;1 - 3;6	42	24,222	4,254	118,210	35,654
4	2;1 - 3;5	45	23,408	4,309	128,986	37,072
5	3;0 - 4;5	48	37,446	8,011	105,305	28,029
6	2;11 - 4;3	48	32,897	6,225	157,635	42,235

Our data stem from a longitudinal audio-visual corpus (approx. 900k words) of language acquisition by 6 Chintang children, who were recorded for 4 hours in monthly intervals (Stoll et al., 2019). Here, we use the data of two children between the ages of 2;1 and 3;6 and two children between 3;0 and 4;5[5]. Since the children are mostly outside, surrounded by many different people, the corpus includes a varied sample of ambient language. For the analyses in this paper, each

[5]For more information, see http://www.clrp.uzh.ch

child and their ambient language were examined separately. Table 1 provides an overview of the transcribed and analyzed recording hours and number of tokens for each child and the surrounding speakers in their recordings.

Frequency Distributions

When frequency of occurrence of aspectual marking is calculated as a simple percentage of all verb forms, it only occurs in 7.7% of all verbs in the ambient language, and 5.2% of children's verb forms. The amount increases with age, see Figure 1.

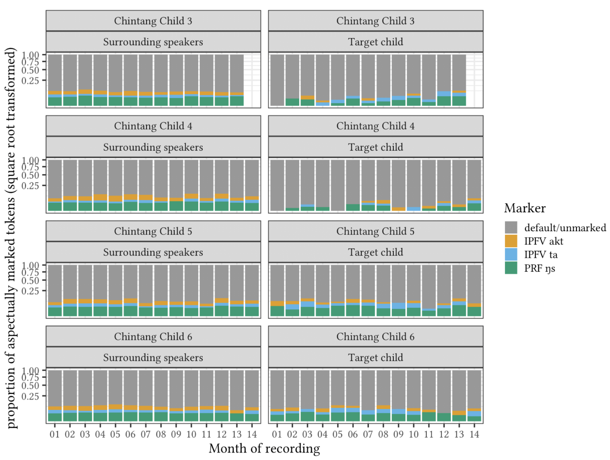

Figure 1: Proportion of aspectually marked verb forms out of all verb forms used by target children and in the ambient language. (Children 3 & 4 are between 2;0 and 3;5, children 5 & 6 between 3;0 and 4;6.)

While these input proportions appear low, *-ŋs* occurs upwards of 40 times per hour of recording. This means that, on average, every 1.5 minutes the child is confronted with this marker. The markers *-akt* and *=ta*, on the other hand, are used less frequently and occur only 9.6 and 3.5 times per hour, respectively. Still, if a marker occurs several times per hour, this suggest that the child is constantly reminded about it, which supports memory consolidation. The frequent repetition of words or morphemes in small time units are known to facilitate the acquisition of these elements (Schwartz & Terrell, 1983; Childers & Tomasello, 2002). However, if we assume that frequency is the main relevant factor for acquisition, we would predict that *-ŋs* will be the easiest to acquire. For the other two markers, *-akt* and *=ta*, predictions will depend less on frequency, since they are both similarly infrequent. Even though *-akt* is more frequently encountered in the input and might, therefore, presumably, be easier to learn, the numbers are so low that other factors might outweigh the small advantage of slightly higher frequency.

Figure 1 shows that the proportion of aspectually marked verb forms increases in the older target children. As predicted, *-ŋs* is the marker used most frequently by all focal children. However, it also illustrates two more tendencies. i) All markers appear, albeit infrequently, in children's production from around the age of 2;4-2;6; ii) for the infrequent markers, both target children age groups exhibit a pattern that does not align with the input. In adult language *-akt* is used more frequently. The focal children, however, show a preference for the use of *=ta*.

Frequency in typical contexts

Another way of measuring frequency of occurrence is to compute the proportion of marker use in the typical contexts of use (Stoll & Bickel, 2012). For the three aspectual markers, the typical contexts are i) *-ŋs* with past marking, ii) *-akt* with past marking, and iii) *=ta* with non-past marking. Other contexts are attested, but much less frequently. Table 2 shows the raw numbers of verbs with past and non-past marking used in the ambient language, as well as in the two age groups of focal children, and Table 3 shows the percentage of typical contexts in which one of the markers appears. Examining the use of markers in the typical contexts allows us to further compare the behaviour of the learners with that of the speakers in their input.

Table 2: Number of past and non-past tokens in the ambient language and the two age groups.

speaker group	past	non-past
Ambient language	31,756	34,582
Children 3 & 4	2,508	1,801
Children 5 & 6	3,767	3,942

Table 3: Percentage of the typical contexts of occurrence that carry the respective aspectual marker.

	Ambient language	Children 3 & 4	Children 5 & 6
-ŋs + PST	23.5%	5.8%	18.8%
-akt + PST	4.5%	0.3%	1.5%
=ta + NPST	1.6%	1.8%	3.0%

Table 3 shows that overall the 3;0 to 4;6 year-olds use aspectual marking more frequently in the relevant contexts than the younger children. However, they are not quite at the level of adults yet. For *-ŋs*, the proportion of marked forms in children's production is below that in the input, yet it matches the pattern in that it is the most frequently used form by far. For the other two markers, however, children have the opposite tendency and use *=ta* more frequently (especially the older group).

Lexical diversity used with aspectual markers

Token frequencies in themselves are not informative enough, since a child could be using the same form over and over again without displaying flexible use of morphology. Therefore, we also assess the distribution of the number of lexical types used with the three markers by the focal children as well as surrounding speakers.

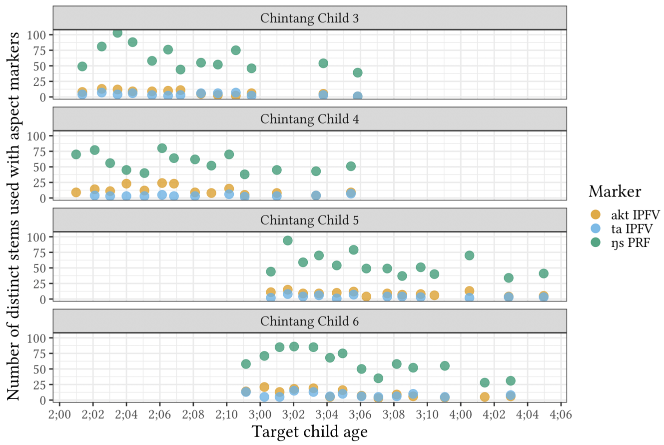

(a) Ambient language

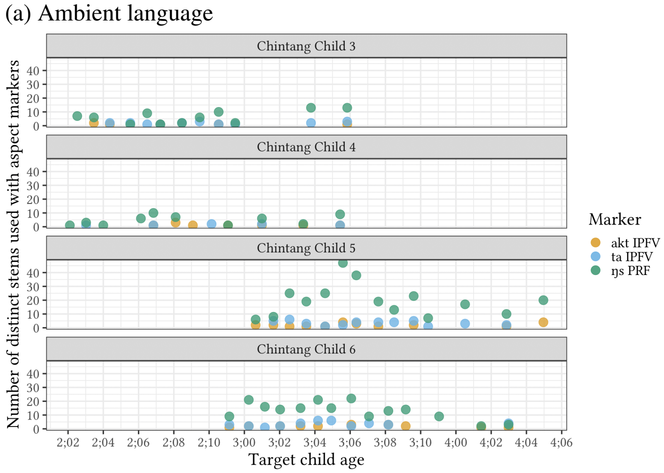

(b) Children's production

Figure 2: Number of distinct stems used with aspect markers.

Figure 2a shows that *-ŋs* occurs with the greatest variety of stems in the input. A greater variety of contexts in which a marker appears can also increase learnability by facilitating the segmentation of the element from the word-unit (see for example Küntay & Slobin, 1996; P. Brown, 1998b; Cameron-Faulkner, Lieven, & Tomasello, 2003; Stoll, Abbot-Smith, & Lieven, 2009; Moran et al., 2019).

Figure 2b illustrates the development of flexibility. The younger group uses all markers less and with fewer distinct stems. In the older group, thee use of *-ŋs* increases substantially. Additionally, the children seem to show a slight preference for the use of *=ta*. However, due to the low frequencies of *-akt* and *=ta*, it is hard to observe their use in individual recording sessions. Therefore, to assess the flexibility of stem use and preference for one of the two low-frequency markers, we run a Bayesian beta-binomial model.

Beta-Binomial Model

We established that *-ŋs* is the most abundantly and flexibly used marker in the input, and this is also mirrored in the production of the focal children. In the following model, we therefore disregard this marker and instead focus on the low-frequency markers *-akt* and *=ta*. While there is a slight difference in their frequency distributions, it might not be large enough to predict the speed of acquisition. However, as mentioned above, they differ in several salience feature, which should make *=ta* easier to acquire.

Since Chintang verbs are not obligatorily marked for aspect, we cannot do an analysis of errors of omission. Therefore, we use a a Bayesian beta-binomial model, which allows us to analyses the flexibility of use of the two markers across verb stems as well as users' preference for one or the other marker (cf. Meylan, Frank, Roy, & Levy, 2017).

Beta-binomial models assume that collections of binary discrete data are generated by shape parameters of an underlying beta distribution (a family of distributions generating probabilities between 0 and 1), which determine the bias toward a particular variant.

We model the use of every stem in the sample with either *=ta* or *-akt* as a draw from a beta-binomial distribution, where 0 represents the use of *=ta* and 1 the use of *-akt*. Each stem can display a certain preference for one of the two markers, or be equally likely to be used with both. We use Rstan (Stan Development Team, 2018) to infer the posterior distributions of the shape parameters of the beta distribution assumed to have generated the observed data in our sample. If the inferred shape parameters correspond to a U-shaped beta distribution, it means that each stem tends to be used with either one of the two markers, but not with both. A bell-shaped distribution centered around 0.5 indicates that stems are used with both markers, and the narrower the peak of the bell, the more uniform or even the use of the stems is. When the distribution is skewed towards the left or the right side, the speaker shows a preference for the use of the value coded as 0 (here: *=ta*) or 1 (here: *-akt*) and shows little flexibility.

We ran an initial model with uninformative priors, drawing the two shape parameters from a $Gamma(1,1)$ distribution, which generates positive numbers concentrated around one with relatively high dispersion. We inferred the shape parameters for each child and their surrounding speakers separately. In a second model, we chose priors based on the distributions empirically attested in the input to simulate the assumption that children's development should follow the distributions encountered in their ambient language.

Finally, since the aspectual markers typically occur with specific tense marking: (*-ŋs* and *-akt* with past, and *=ta* with non-past), we ran a third model to exclude the possibility that children's preference for *=ta* is simply an effect of their preferential use of non-past forms. This possibility is based on the assumption that children might prefer referring to "here and now" situations over past events. This model used an uninformative prior; non-past is coded as 0 and past as 1.

Results

Uninformative prior Table 4 shows the means and credible intervals of the fit for each speaker/speaker group, while Figure 3 shows the density curves (0 represents *=ta*, 1 *-akt*) based on the inferred posterior shape parameters of the model with an uniformative prior[6]. The modes and interpretations of the results are given in Table 5.

Table 4: Posterior parameter means and 95% credible intervals for each speaker/speaker group for the model with an uninformative prior.

Speaker(s)	Mean	Credible interval
Child 3	$\alpha = 0.47; \beta = 1.99$	[0.08, 1.39]; [0.44,4.87]
Child 4	$\alpha = 1.02; \beta = 1.12$	[0.15, 2.81]; [0.18,3.09]
Child 5	$\alpha = 0.31; \beta = 0.63$	[0.09, 0.78]; [0.17,1.62]
Child 6	$\alpha = 0.57; \beta = 1.13$	[0.13, 1.48]; [0.26,2.98]
Input 3	$\alpha = 1.93; \beta = 1.32$	[0.69, 3.97]; [0.53,2.57]
Input 4	$\alpha = 2.32; \beta = 0.78$	[0.92,4.69]; [0.34,1.52]
Input 5	$\alpha = 2.81; \beta = 1.65$	[1.30,5.28]; [0.81,2.93]
Input 6	$\alpha = 2.14; \beta = 1.19$	[1.03,3.93]; [0.61,2.04]

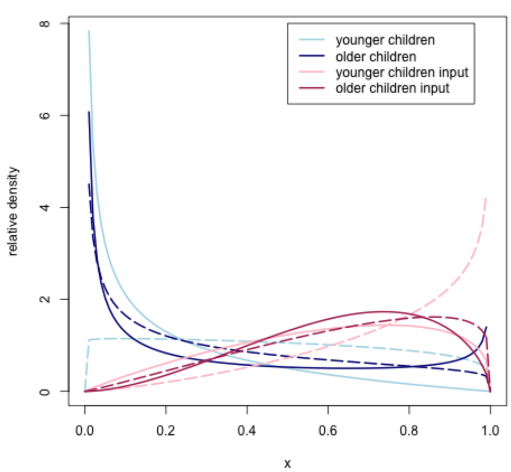

Figure 3: Distribution based on the inferred shape parameters which characterizes aspect marker preference and flexibility for each child's production as well as the input. The line type differences indicate individual children/their input within the same age group.

Table 5 shows that children do in fact prefer the use of *=ta*. Even Child 5, who displays a bimodal distribution, shows a higher density around 0, i.e. more use of *=ta*. Adults show a preference for *-akt* across the board. They, unlike the focal

Table 5: Modes of the distributions based on the inferred shape parameters and interpretation of the results.

Speaker(s)	Mode	Interpretation
Child 3	0	preference for *=ta*
Child 4	0.09	preference for *=ta*
Child 5	Bimodal (0,1)	uses both markers, stem use not flexible
Child 6	0	preference for *=ta*
Input 3	0.74	preference for *-akt*, some flexibility
Input 4	1	preference for *-akt*
Input 5	0.74	preference for *-akt*, some flexibility
Input 6	0.86	preference for *-akt*, some flexibility

children, exhibit some flexibility in their use of stems with the aspectual marking.

Prior skewed towards *-akt* The posterior distributions inferred in the second model, with a skewed prior, are given in Figure 4.

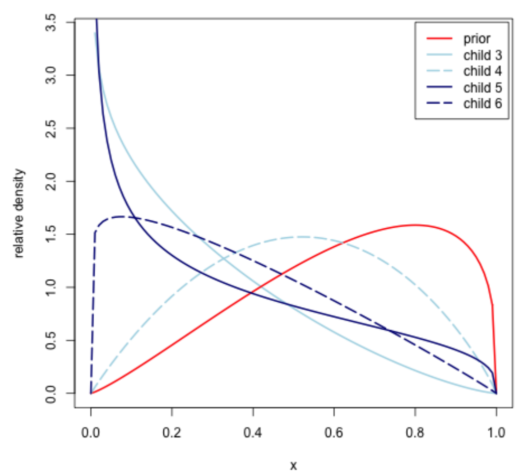

Figure 4: Distribution based on the inferred shape parameters which characterizes aspect marker preference and flexibility for each child's production using an empirical prior based on input distributions. The line type differences indicate individual children within the same age group.

We use this model to test how robust the data are against a prior that represents the assumption that frequency distributions in the input are the best predictor of ease of acquisition. The only child for whom the mode changed is Child 4. However, this was an expected result, since Child 4 uses fewer than 20 tokens of *-akt* and *=ta* markers over the entire recording period, which makes the data not powerful enough to counteract the prior.

Model of past and non-past use Figure 5 shows the posterior distributions of the model that models the choice between non-past and past verb forms. If the preference for *=ta* is due to a preference for non-past, the posterior distribu-

[6]We run 4 chains with 4000 iterations each, of which 2000 each are the burn-in period, sampling at every step. All models converged with Rhat between 1 and 1.01.

tion should be skewed towards 0. However, all but one line are bell-shaped and centered close to 0.5, which means that most of the children do not show any preference for either past or non-past and use tense-marking flexibly across stems. Counter-intuitively, perhaps, the younger two children skew towards past forms, albeit Child 4 to a lesser degree.

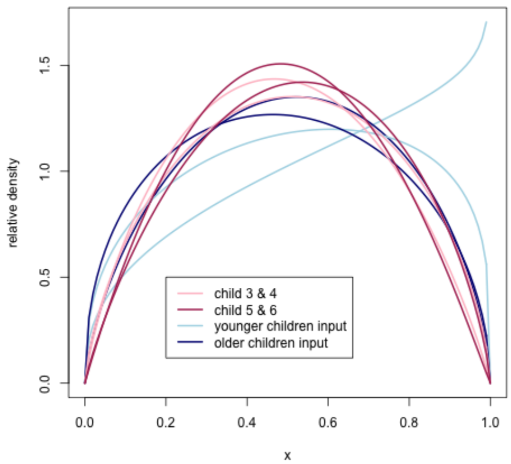

Figure 5: Distributions based on inferred shape parameters characterizing the use of past and non-past marking.

Discussion and outlook

We compared the acquisition of three aspectual markers in Chintang along two dimensions: frequency, and salience. Both *-ŋs* and *-akt* appear within the word unit, making them less positionally and prosodically salient. *-ŋs*, however, occurs relatively frequently, while the frequency of *-akt* is very low. On the other hand, *-akt* and *=ta* are both far less frequent than *-ŋs*. They differ in their position and prosodic properties, since *=ta* appears at the word-boundary, which is a highly salient position. This approach allows a direct comparison between the effects of different predictors of ease of acquisition. As predicted, the most frequent marker in the input is the one used most frequently by children early on. For the low-frequency markers, frequency in the input was not a reliable predictor. Instead, *=ta*, the marker that is positionally and prosodically more salient and represents a more reliable cue due to its fixed position and surface form, is preferred by children, while adults show a preference for the more opaque and less salient marker *-akt*.

Disentangling frequency from other factors is not always straightforward and many studies focus on the effect of frequency or find frequency to be the strongest predictor (cf. Rowland, Pine, Lieven, & Theakston, 2003; Ambridge et al., 2015). In many cases, the question revolves around which items are acquired first (cf. Naigles & Hoff-Ginsberg, 1998), which might obscure effects of other variables due to the predictive strength of frequency. If frequency is, in fact, the strongest predictor of ease of acquisition, then we might miss the relative effects of other factors, which might play a greater role for lower-frequency items.

When it comes to the relative effect of various variables of salience, many accounts of first language acquisition focus on word or construction learning (Naigles & Hoff-Ginsberg, 1998; Longobardi et al., 2015), while the effect of salience on morphology acquisition is discussed more frequently in the domain of L2 acquisition (cf. Goldschneider & DeKeyser, 2001; Ellis, 2017). This is most likely due to the nature of early language acquisition and the studied languages, where words are easily segmented. Because of this, polysynthetic languages are especially interesting to look at because their rich morphology provides an interesting test-case for the relative effects of various predictors such as frequency but also positional or prosodic salience. Effects of perceptual salience have been found in the early linguistic stage of children learning a number of polysynthetic languages, e.g. Mithun, 1989 for Mohawk, Pye, 1992 for Quiché Mayan, P. Brown, 1998a for Tzeltal. On the other hand, Courtney & Saville-Troike, 2002 find no effect of salience on early verb form production in Quechua and Navajo, stating that children often omit salient affixes.

To be able to make statements about the relative effect of each of these factors, we must consider language-specific properties and distributions in the input as well as the properties and distributions of each studied morpheme. Chintang presents many opportunities to further examine the multifactorial nature of learning in detail due to its multitude of affixes and complex verb forms. Future work will include other carefully matched morphemes that might allow us to disentangle this puzzle further.

Acknowledgments

This work was supported by the European Research Council (ERC Consolidator Grant, ACQDIV 615988, to S. Stoll). We are grateful to Chundra Cathcart for methodological help and Balthasar Bickel for sharing his analyses of Chintang morphology.

References

Ambridge, B., Kidd, E., Rowland, C. F., & Theakston, A. L. (2015). The ubiquity of frequency effects in first language acquisition. *Journal of Child Language*, *42*(02), 239–273.

Bickel, B., Banjade, G., Gaenszle, M., Lieven, E., Paudyal, N., Poppitz, S., ... Stoll, S. (2005). *Morphological quirks in kiranti: new findings in chintang and puma.* (Unpublished handout from a talk at DOBES 2005)

Bickel, B., Banjade, G., Gaenszle, M., Lieven, E., Paudyal, N., Rai, I. P., ... Stoll, S. (2007). Free prefix ordering in chintang. *Language*, *83*, 44-73.

Brown, P. (1998a). Children's first verbs in tzeltal: evidence for an early verb category. *Linguistics*, *36*, 713-753.

Brown, P. (1998b). Conversational structure and language acquisition: the role of repetition in tzeltal adult and child speech. *Journal of Linguistic Anthropology*, *8*, 197-221.

Brown, R. (1973). *A first language: the early stages*. Cambridge, MA: Harvard University Press.

Cameron-Faulkner, T., Lieven, E., & Tomasello, M. (2003). A construction based analysis of child directed speech. *Cognitive Science*, *27*, 843-873.

Childers, J. B., & Tomasello, M. (2002). Two-year-olds learn novel nouns, verbs, and conventional actions from massed or distributed exposures. *Developmental Psychology*, *38*, 967-978.

Christophe, A., & Dupoux, E. (1996). Bootstrapping lexical acquisition: the role of prosodic structure. *The Linguistic Review*, *13*(3-4), 383–412.

Courtney, E. H., & Saville-Troike, M. (2002). Learning to construct verbs in navajo and quechua. *Journal of Child Language*, 623-654.

Ellis, N. (2017). Salience in usage-based sla. In *Salience in second language acquisition* (pp. 21–40). Routledge.

Freudenthal, D., Pine, J. M., Aguado-Orea, J., & Gobet, F. (2007). Modeling the developmental patterning of finiteness marking in english, dutch, german, and spanish using mosaic. *Cognitive Science*, *31*(2), 311–341.

Freudenthal, D., Pine, J. M., & Gobet, F. (2010). Explaining quantitative variation in the rate of optional infinitive errors across languages: a comparison of mosaic and the variational learning model. *Journal of Child Language*, *37*, 643–669.

Goldschneider, J. M., & DeKeyser, R. M. (2001). Explaining the "natural order of l2 morpheme acquisition" in english: A meta-analysis of multiple determinants. *Language learning*, *51*(1), 1–50.

Küntay, A. C., & Slobin, D. I. (1996). Listening to a turkish mother: some puzzles for acquisition. In D. I. Slobin, J. Guo, & A. Kyratzis (Eds.), *Social interaction, social context, and language: essays in honor of susan ervin-tripp* (p. 265-286). Hillsdale, NJ: Erlbaum.

Küntay, A. C., & Slobin, D. I. (2002). Putting interaction back into child language: examples from turkish. *Psychology of Language and Communication*, *6*, 5–14.

Lieven, E. V., Pine, J. M., & Baldwin, G. (1997). Lexically-based learning and early grammatical development. *Journal of Child Language*, *24*, 187–219.

Longobardi, E., Rossi-Arnaud, C., Spataro, P., Putnick, D. L., & Bornstein, M. H. (2015). Children's acquisition of nouns and verbs in italian: Contrasting the roles of frequency and positional salience in maternal language. *Journal of child language*, *42*(1), 95.

MacWhinney, B. (1978). The acquisition of morphophonology. *Monographs of the society for research in child development*, 1–123.

MacWhinney, B., Pléh, C., & Bates, E. (1985). The development of sentence interpretation in hungarian. *Cognitive Psychology*, *17*(2), 178-209.

Meylan, S. C., Frank, M. C., Roy, B. C., & Levy, R. (2017). The emergence of an abstract grammatical category in children's early speech. *Psychological science*, *28*(2), 181–192.

Mithun, M. (1989). The acquisition of polysynthesis. *Journal of Child Language*, *16*(2), 285-312.

Monaghan, P., & Christiansen, M. H. (2010). Words in puddles of sound: Modelling psycholinguistic effects in speech segmentation. *Journal of child language*, *37*(3), 545–564.

Moran, S., Lester, N. A., Gordon, H., Küntay, A., Pfeiler, B., Allen, S., & Stoll, S. (2019). Variation sets in maximally diverse languages.

Naigles, L. R., & Hoff-Ginsberg, E. (1998). Why are some verbs learned before other verbs? effects of input frequency and structure on children's early verb use. *Journal of child language*, *25*(1), 95–120.

Pye, C. (1992). The acquisition of k'iche'maya. *The crosslinguistic study of language acquisition*, *3*, 221–308.

Rowland, C. F., Pine, J. M., Lieven, E. V., & Theakston, A. L. (2003). Determinants of acquisition order in wh-questions: Re-evaluating the role of caregiver speech. *Journal of Child Language*, *30*(3), 609–635.

Schikowski, R. (2012). *Chintang morphology*. Unpublished manuscript.

Schwartz, R. G., & Terrell, B. Y. (1983). The role of input frequency in lexical acquisition. *Journal of Child Language*, *10*, 57-64.

Stan Development Team. (2018). *Rstan: the r interface to stan. r package version 2.17. 3.*

Stoll, S., Abbot-Smith, K., & Lieven, E. (2009). Lexically restricted utterances in russian, german and english child-directed speech. *Cognitive Science*, *33*, 75-103.

Stoll, S., & Bickel, B. (2012). How to measure frequency? different ways of counting ergatives in chintang (tibeto-burman, nepal) and their implications. In F. Seifart et al. (Eds.), *Potentials of language documentation: methods, analyses, utilization* (p. 84-90). Manoa: University of Hawaii Press.

Stoll, S., & Bickel, B. (2013). The acquisition of ergative case in chintang. In E. L. Bavin & S. Stoll (Eds.), *The acquisition of ergativity* (p. 183-208). Amsterdam: John Benjamins Publishing Company.

Stoll, S., Lieven, E., Banjade, G., Bhatta, T. N., Gaenszle, M., Paudyal, N. P., ... Bickel, B. (2019). *Audiovisual corpus on the acquisition of chintang by six children.*

Stoll, S., Mazara, J., & Bickel, B. (2017). The acquisition of polysythetic verb forms in chintang. In M. Fortescue, M. Mithun, & N. Evans (Eds.), *Handbook of polysynthesis* (p. 495-517). Oxford University Press.

Theakston, A. L., Lieven, E. V. M., Pine, J. M., & Rowland, C. F. (2004). Semantic generality, input frequency and the acquisition of syntax. *Journal of Child Language*, *31*, 61-99.

Tomasello, M. (2003). *Constructing a language: a usage-based theory of language acquisition*. Harvard, MA: Harvard University Press.