Frequency and informativity of phonological input directed to children in the first four years of life

Arjun Pawar (arjun7@ucla.edu)
Department of Linguistics, UCLA, 335 Portola Plaza
Los Angeles, CA 90095 USA

Margaret Cychosz (mcychosz4@ucla.edu)
Department of Linguistics, UCLA, 335 Portola Plaza
Los Angeles, CA 90095 USA

Abstract

Information theory characterizes how signals are optimized for transmission from source to receiver across noisy channels, yet little is known about how these principles manifest when the receiver's capabilities change over time. Using child-directed speech as a natural experiment, we analyzed >7.5 million phones in North American English caregiver speech to children aged 3-44 months (N=218) from the CHILDES database. We found that while the relative frequency of individual phones remained stable over this developmental time period, phonological informativity increased from early infancy (3-8 months) through toddlerhood (27-32 months), before plateauing in the preschool years. This result suggests that speech directed to children sounds less redundant, with a phonological structure that is harder to predict in context, as children progress through early childhood. Our findings demonstrate how linguistic signals may be optimized to accommodate receiver (child) characteristics, with implications for both general principles of information transmission and theories of how children carve out linguistic representations and patterns from limited, noisy input.

Keywords: child-directed speech; information theory; phonology; informativity, corpus linguistics

Introduction

Effective communication balances competing demands: redundancy in signals ensures robust transmission and can overcome noise, but too much redundancy risks inefficiency in the signal for both source and receiver. In human communication, this framework, known as INFORMATION THEORY, manifests as modifications to speech and language that optimize the trade-off of effort in production and comprehension (Lindblom, 1990; Piantadosi et al., 2011). While information theory has frequently been studied for communication between adult speakers (Gibson et al., 2019; Cohen Priva, 2015), little is known about how these principles unfold when the receiver's processing capacity develops over time, such as over the course of early language acquisition (c.f. Tal et al. 2024).

Information theory and human communication

Information theory quantifies the abstract notion of 'information' which is transmitted as data or signals across a channel. Under this theory, any transmitted signal or message carries some information and some uncertainty which can be probabilistically quantified.

Producing and comprehending language can be seen as a practical application of information theory where the speaker is the transmitter of the message and the listener is the receiver, in the presence of noise. Noise in this context could mean external (ambient) noise or internal constraints which pose speech processing difficulties.

Information-theoretic measures such as relative frequency, informativity, and predictability (defined in more detail below) have been employed in cognitive science for multiple levels of linguistic analysis to understand how natural language is structured and used to facilitate communication (Mahowald et al., 2013). Broadly, the *predictability* of a linguistic unit (word or phone) correlates with how probable it is to observe that unit given its context i.e. $-\log_2 \mathbb{P}(\text{phone}|\text{context})$. The *informativity* of a linguistic unit averages the predictability across all contexts, weighted by context occurrence. So a more informative unit could be interpreted as one carrying more 'information' in the signal (see Methods for formulae). Work in this area has concluded that across multiple languages, not only does word length correlate with frequency i.e. shorter words are more frequent (Zipf, 1949), but length also positively correlates with informativity (Piantadosi et al., 2011). And within the human speech stream, the speech signal is optimized for redundancy (via duration and prosodic structure) to improve communicative robustness between speakers (Aylett & Turk, 2004).

These results likewise extrapolate to the phonological level (Bell et al., 2009). Informativity can affect phonological reduction (i.e. phonemes are articulated less distinctively or with reduced stress) as less informative segments have a greater chance of getting reduced or in the extreme case, deleted (Cohen Priva, 2015). Work at the phonological level has also considered how 'functional load'—another probabilistic measure that estimates how many minimal pairs a given phoneme is responsible for maintaining—is correlated with historical tendencies in sound change e.g. a high functional load reduces likelihood of phoneme mergers (Wedel et al., 2013).

While various information-theoretic measures have been used until now, Cohen Priva & Jaeger (2018) performed comprehensive evaluations of different information-theoretic measures in an attempt to disassociate the inter-related measures of frequency, predictability, and informativity for psycholinguistic experimentation. The authors found that spurious effects of frequency might arise when informativity is not controlled for, at least in analyses of adult communication.

Altogether, such studies show that the impact of language use on linguistic structure at different levels (syntactic, lexical, phonological) can be explained through the lens of com-

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

municative optimization in the production and comprehension of language. We innovate upon previous work by applying an information-theoretic framework to child language acquisition and treat phrases, words, and individual sounds as 'data packets' that are transmitted from a speaker (caregiver) to a listener (child). It remains unclear if and how speech directed to children, or child-directed speech (CDS), is optimized in a similar manner as previous work has demonstrated for adult-directed speech. On the one hand, CDS could be optimized to facilitate efficient communication, such as maintaining phonemic contrasts, as is the case for adult-directed interactions (Wedel et al., 2013). However, CDS differs from adult-directed communication because of its implicit intent (supplying learning data to a child), with clear changes in acoustics, phonology, word use, and grammar over developmental time (Cychosz et al., 2021; Kitamura & Lam, 2009; Kalashnikova & Burnham, 2018). Overall, it is unclear how well findings concerning the efficiency of inter-adult speech and language communication extrapolate to speech directed to children.

CDS—often characterized by slower utterances, produced with exaggerated speech contrasts, and more frequent words from denser phonological neighborhoods—has long been considered an important component for children's early language learning (Soderstrom, 2007; Snow, 1972; Jones et al., 2023). Increased exposure to CDS has been shown to facilitate vocabulary growth and word segmentation (Thiessen et al., 2005) as well as lexical processing skills (Weisleder & Fernald, 2013), at least in North American English samples. From an information-theoretic perspective, little is known about how CDS changes from infancy through early childhood and helps language overcome noise during formative periods of language acquisition. Modeling communication between a caregiver and a child using information-theoretic measures could contribute to our understanding of one of the most fundamental questions in language acquisition: how do infants and children acquire robust linguistic systems from limited data in relatively short periods of time? Work in this area is limited, though a recent study measuring lexical entropy rates found that CDS becomes less lexically redundant between 7 - 24 months (Tal et al., 2024) suggesting that CDS could be lexically tailored to the child's developmental stage.

Current Study

Here, we characterize the phonological structure of CDS in North American English by examining its changes over the course of early language acquisition and we extend upon Tal et al. (2024) in three ways: (1) modeling changes in CDS at the *phonological* level, instead of lexical, to understand how the sound structure of CDS changes over developmental time, (2) sampling from more corpora, and (3) comparing how two different information-theoretic measures manifest in CDS over developmental time. This study thus investigates whether CDS is similarly redundant across the early years *at the phonological level*, with implications for how very young children are able to process the sounds of their ambient language and eventually carve out robust phonological representations. Our main research questions are: (1) how does the phonological structure of CDS change across the first four years of childhood, (2a) what is the relationship between informativity and frequency, and (2b) how does the relationship between informativity and frequency change over this time?

Methods

Data sources

We computed the frequency and informativity of phonological input to children using CDS transcripts from `CHILDES`, an openly-available database of child language corpora (MacWhinney, 2000). Candidate corpora for our analysis needed to contain (semi-) naturalistic North American English caregiver speech directed to typically-developing children aged 3-45 months (i.e. not book reading). Corpora additionally had to be transcribed to the word level, and we limited selection to corpora from the last 25 years. This resulted in the 5 corpora listed in Table 1 (Newman et al., 2016; Demuth et al., 2006; Rollins, 2003; Brent & Siskind, 2001; Van Kleeck, 2004).

Corpus	Age Range (mos.)	N unique children
Brent Siskind	6-15	17
Newman Ratner	7-24	121
Providence	12-36	6
Rollins	3-30	54
VanKleeck	36-48	20

Table 1: Corpora chosen for this study

To evaluate the effect of child age on the information-theoretic measures, we classified the children into 6-month bins (e.g. labeled as 3-8 months, inclusive). Six-month bins ensure there are a sufficient number of unique children within each bin (minimum 7 children) while allowing evaluation of developmental changes in CDS.

Data preparation

Corpora transcripts were scraped using the `childespy` package that leverages the `childes-db` utility (Sanchez et al., 2019). To map word-level corpus transcripts to their phonemic forms, we used the Montreal Forced Aligner English Grapheme-to-Phoneme (G2P) models (McAuliffe et al., 2017) which convert word-level transcripts to phonemic transcriptions in IPA (International Phonetic Alphabet). To ensure that the phonemic transcription was uniform across all corpora (i.e. level of phonetic transcription and treatment of allophones), we converted even those corpora that were already phonemicized (e.g. Providence) using this pipeline.

Following phonemicization, an extensive data cleaning pipeline was developed: each corpus was filtered so that only caregiver (e.g. not child, not researcher) utterances remained. Next, standardized CHILDES transcription codes indicating non-speech sounds were removed. At this point, we used

Corpus	Age Bin (mos)	Word Tokens	Candidate Phones	N children*
Brent/Siskind	3-8	5283	13951	2
	9-14	368361	1024941	16
	15-20	20697	57018	8
Newman/Ratner	3-8	150113	409668	123
	9-14	241078	663382	120
	15-20	77917	212451	53
	21-26	213830	580853	120
Providence	9-14	73638	212516	4
	15-20	268433	764414	6
	21-26	311600	881519	6
	27-32	338439	957946	6
	33-38	216250	610817	6
	39-44	71742	201342	5
Rollins	3-8	11466	29271	13
	9-14	46687	121173	25
VanKleeck	33-38	1066	2905	1
	39-44	12623	32705	6

Table 2: Availability of data. Candidate phones refers to the total no. of available phones, from which we drew samples. *Children can be counted in multiple age bins

the dictionary-based and rule-based G2P models– `English (US) MFA dictionary v3.1.0` and `English (US) MFA G2P model v3.0.0` respectively. The dictionary-based approach involves a large dictionary of common words and their attested pronunciations in American English. The rule-based `G2P` is a finite state machine algorithm which uses phonological rules to produce an IPA transcription for any input word. First, an initial pass was made by the MFA pronunciation dictionary (93% of words recognized and phonemicized). The remaining 7% (i.e. unrecognized words) were manually inspected: proper nouns and non-verbal expressions (*muah, huh*) were removed. The rule-based `g2p` model was then run on the remaining words. Finally, we corrected allophonic variants in the `g2p` output that are not standard in American English: for example, [ɫ] → [l]; [a] → [ɑ]; [k^h] → [k] (select allophones such as [ɾ] were retained). The full pipeline can be found in the Github repository available at this link. See Table 2 for descriptive information about the number of children, word- and phone-level counts for each age bin.

Determining sample size

To ensure that any potential developmental was due to *age*, and not other potential confounds such as the amount of data within a given 6-month bin, the number of unique children, or the number of corpora, we followed Tal et al. (2024) in employing a quantitative criterion to determine the ideal number of phones to analyze for each 6-month bin. For each age bin, we sampled via the following process: we ensured that no utterance from the transcript was broken up while sampling i.e. complete lines from the conversation were taken; then, in increments of 3k, we took different sample sizes from 3k to 300k phones, computed our measures (relative frequency and informativity) to get sample means, and calculated the standard deviation of these mean values in a sliding window-size of 10 samples. The sample size was recorded when the standard deviation of both mean informativity and mean frequency fell below SD=0.01 (see Figure 1 for illustration). The largest of the minimum sample size required across all age bins was then taken as the standard sample size for all bins (Tal et al., 2024): this was 81000 phones in our dataset.

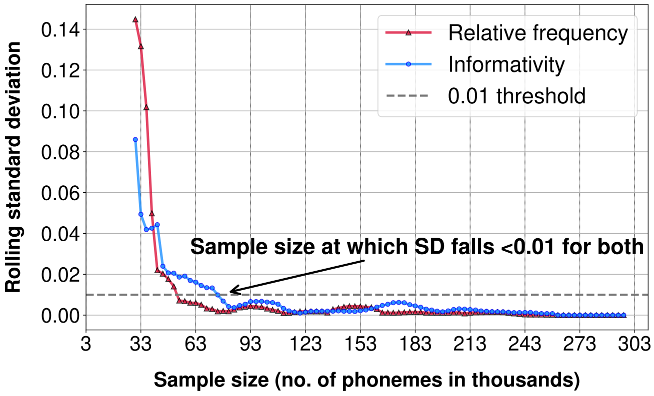

Figure 1: Example of how the ideal sample size of phone samples to be drawn was determined for each age bin. Each point represents the rolling SD taken over a window size of 10 samples where the sample sizes increased in increments of 3000. SD for every 10 such samples was plotted. The sample size at which both curves were < 0.01 was recorded.

In this way, we could be confident in the stability of our estimates of frequency and informativity within each bin and that our estimates of the measure were not due to potential sampling biases. We could also more confidently attribute effects of child age to changes in the information-theoretic measures rather than other attributes of each bin (i.e. sample size). In Results, we take additional measures to evaluate potential confounds in the construction of the bins. Finally, we randomly drew 100 samples at each age bin, and treated outlier samples appropriately by re-sampling multiple times and computing measures again.

Computing measures at the phonological level

To compute the phone-level information theoretic measures, we applied formulae following the techniques outlined in Cohen Priva & Jaeger (2018). We used the definition of 'context' to be all the preceding phones *in the same word* (e.g. for 'father' [fɑðɚ] the context of [ɑ] is [f] and the context of [ɚ] is [fɑð]). We computed negative log probability (relative frequency) and informativity at the phone level over the processed transcripts (standard number of phones within each age bin). These measures bear the unit 'bits' which is standard in information theory (Shannon, 1948).

[1] Relative frequency (–log probability) of a phone X: We used the negative log probability of finding phone X:

$$-\log_2\left(\frac{\text{total number of occurrences of phone } X}{\text{total number of phones in sample}}\right)$$

Because of the negative log, lower values are assigned to phones which are *more probable* in a corpus. This can be considered non-contextual phonemic surprisal (Brodbeck et al., 2018).

[2] Informativity of a phone X: Informativity depends on predictability, so we first calculated predictability as follows: The **predictability** of a phone X in a context C:

$$-\log_2\left(\frac{\text{total occurrences of phone } X \text{ in context } C}{\text{total occurrences of context } C}\right)$$

For example, if [kæt] appeared 99 times but [kæp] appeared once, and no other word began with [kæ], then the predictability of [p] under context [kæ] is $-\log_2(0.01) = 6.64$ bits while the predictability of [t] under context [kæ] is $-\log_2(0.99) = 0.01$ bits. This means that [p] is more informative in that context. For our analysis, we used informativity instead of predictability (which is context-specific) as informativity averages predictability across multiple contexts.

We computed **informativity** of phone X as:

$$\sum_{C\in\text{contexts}} \mathbb{P}(C|X)\cdot \text{predictability}(X,C) \quad \text{where}$$

$$\mathbb{P}(C|X) = \frac{\text{total occurrences of phone } X \text{ in context } C}{\text{total occurrences of phone } X}$$

Informativity averages out the predictability scores across contexts; more informative phones have higher values.

Results and analysis

How do the frequency and informativity of the phonological structure of CDS change across the first four years of childhood?

We computed the frequency and informativity for N=100 samples/age bin (81k phones/sample) via bootstrapping (randomly sampling with replacement). One-way ANOVAs evaluating the effect of child age (treated categorically as it was binned in 6-month intervals) on each measure (frequency and informativity, separately) revealed significant effects of child age on phonological informativity ($F(6) = 929$, $p < .001$, $\eta_p^2 = 0.89$) and relative frequency ($F(6) = 25$, $p < .001$, $\eta_p^2 = 0.18$); however, as seen in Figures 2 and 3, this effect was sustained for informativity, but not frequency (which we further evaluated via post-hoc comparisons). This result indicates that while the informativity of phonological structure in CDS changes systematically over development (as phones become less predictable in context), the relative frequencies of individual phones in speech to children remain stable.

To identify which age-related changes drove this effect, we conducted Bonferroni-corrected pairwise comparisons between consecutive age bins. We found significant differences in the informativity of phonological structure in CDS between each pair of consecutive age bins from 3-8 months through 33-38 months (Table 3). Specifically, informativity showed a steady increase from early infancy (3-8 months, M=3.27, SD=0.02) through toddlerhood (27-32 months, M=3.43, SD=0.03), before leveling off in the preschool period, suggesting that the sounds of CDS are becoming less redundant as children age.

Evaluating confounds with age We considered that it was possible that our results were not due to developmental change in CDS characteristics, but instead due to random chance of comparing samples across different ages/corpora. To address this, we conducted two control "scrambling" analyses following Tal et al. (2024). In the first, we preserved the grouping structure of our samples (N=100/age bin), but randomly reassigned the age labels to these groups (Figure 4): e.g. all samples originally from the 3-8 month bin might be relabeled as 33-38 months, while maintaining their grouping. We then repeated our statistical analyses on these shuffled data. In the second control analysis, we randomly reassigned age labels to each individual sample, independent of its original agebin. Unlike the first analysis, this approach allowed samples from the same original bin to be assigned to different age categories– e.g., two samples originally from the 15-20 month bin could be reassigned to different age bins: one to 3-8 months and another to 33-38 months.

Both control analyses supported our findings. Under the first shuffling analysis, the developmental trajectory (increasing informativity with age) disappeared. Although the effect of age was, unsurprisingly, still significant, the actual direction of the effect was inconsistent (alternating between positive and negative changes in informativity), suggesting that the trajectory of change in informativity relied on the increasing age of the child. In the second analysis, the effect of age on relative frequency ($F(6) = 0.41, p = 0.87, \eta_p^2 = 0.004$) and informativity ($F(6) = 0.87, p = 0.52, \eta_p^2 = 0.008$) was not significant and the effect sizes (η_p^2) were negligible. These results suggest that our observed developmental changes in informativity reflect genuine age-related differences rather than random variation in the data due to other sources.

How does the relationship between relative frequency and informativity change across the first four years of childhood?

Frequency and informativity are positively correlated, so some effects of informativity that we have measured in CDS might be attributed to frequency (Cohen Priva & Jaeger, 2018). To address this, we computed both the relationship between these measures as well as how this relationship changed over development. For each age bin, we measured the correlation between a given sample's average phone frequency and informativity (N=100 samples/bin). Here a positive correlation indicates that more frequent phones (i.e. lower on the negative log frequency scale) were less informative. We found a positive correlation between the measures for all age bins (mean Pearson's $r = 0.23$, SD=0.13), but this relationship was only reliably or marginally significant for the

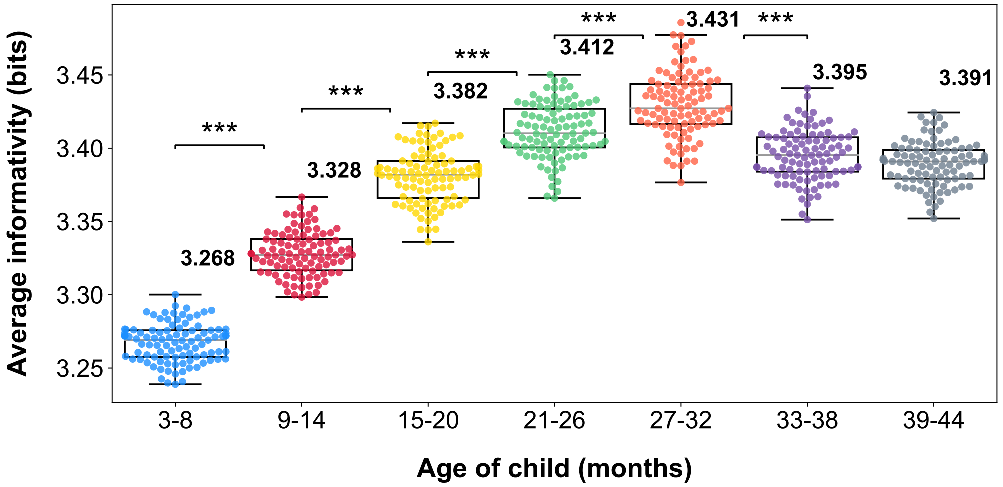

Figure 2: Phonological informativity by child age. Each point represents the average informativity of phones in a randomly-drawn sample of 81k phones. 100 such samples were drawn per age bin with replacement. Mean informativity is printed next to each box (placement systematically shifted for visibility). The boxes represent interquartile range with the median marked with a gray line. Asterisks mark statistical significance between consecutive age bins.

first 4 consecutive age bins (3-8, 9-14, 15-20, 21-26 mos.; p=0.01-0.07). For older age bins (27-32 and 39-44 mos.), the relationship between the measures was no longer significant.

Finally, we were interested in how the relationship between frequency and informativity in CDS compared to the relationship between these measures in typical adult-to-adult speech (estimated at $r = 0.5$, based on sample size = 0-200k word tokens drawn from transcribed telephone conversations and sociolinguistic interviews (Cohen Priva & Jaeger, 2018)). The relationship between frequency and informativity appears to be somewhat higher in speech between adults than in CDS, a fact that we return to in the discussion.

Discussion and Conclusion

This study used $>$ 7.5M phones of CDS as a natural experiment to understand how source signals (speech from caregivers) are optimized when the receiver capabilities (the child learner's cognitive capabilities) change over time. Results showed that the informativity of phonological content in CDS—the sounds of speech directed to children—became more informative, and less redundant, as children aged (between 3-32 months). Changes then leveled off in the preschool years (after 33 months), *above and beyond changes in frequency*: there was no reliable effect of age on the relative frequency of phones.

The developmental trajectory observed here—from redundant, more predictable sounds in early infancy to increasing informativity into the preschool years—suggests that caregivers (implicitly) optimize the sounds of CDS to correspond to children's developing phonological processing skills. This developmental pattern could have a bootstrapping effect as, for example, redundancy earlier in life may help infants establish foundational phonological representations and phonotactic constraints in their native language(s). These early foundations may then serve a critical purpose as children age and are required to accommodate more information-dense speech input both in the form of more diverse word types/tokens and phonotactic structures (see next paragraph), but also parse the signal through noisier channels (e.g. speech from a novel interlocutor, processing in a new listening environment).

The findings have implications for children's developing phonological skills as well. As children age and the number of unique word types in their input increases (Rowe, 2012), each individual speech sound in the words that children hear will carry more information. The flip side of this is that individual phones are also less predictive in context (i.e. higher

Reference age bin	Comparison age bin	Informativity		Relative frequency	
		t statistic	Difference in means	t statistic	Difference in means
3-8	9-14	-28.07	-0.06***	0.99	0.0
9-14	15-20	-21.03	-0.05***	-1.69	0.0
15-20	21-26	-11.16	-0.03***	-4.39	-0.01***
21-26	27-32	-6.42	-0.02***	0.19	0.0
27-32	33-38	11.96	0.04***	-1.67	0.0
33-38	39-44	1.45	0.0	8.10	0.02***

Table 3: Bonferroni adjustment applied after ANOVA. Significant differences in means indicated with *** for $p < .001$

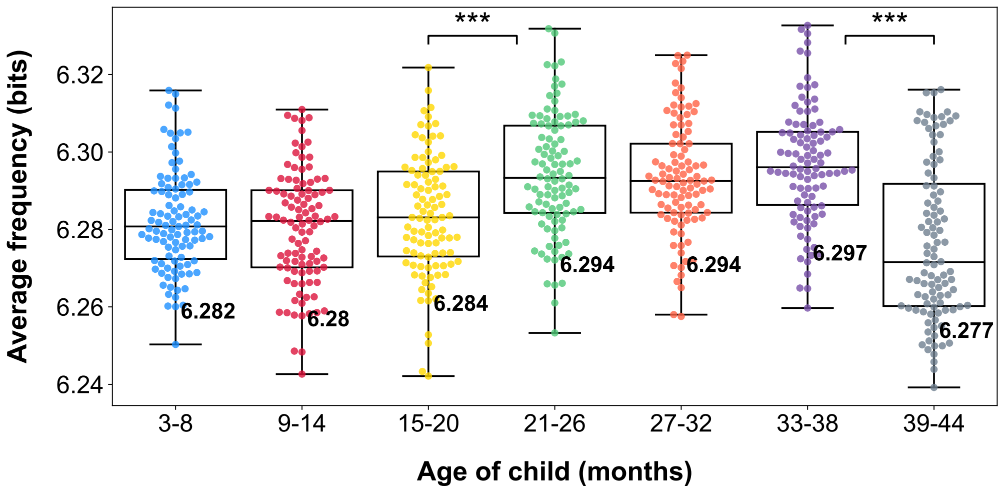

Figure 3: Relative frequency of phones by child age. Each point represents the average relative frequency in a randomly-drawn sample (81k phones). 100 such samples were drawn per age bin with replacement. Mean frequency is printed next to each box in black (placement systematically shifted for visibility). The boxes represent interquartile range with the median marked with a black line. Asterisks mark statistical significance between consecutive age bins.

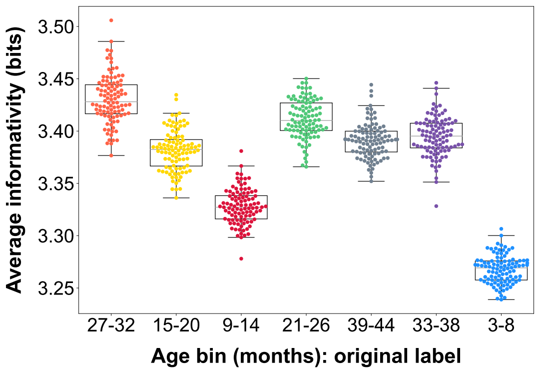

Figure 4: Example of bin scrambling—reordered samples with grouping intact (first scrambling analysis)

surprisal), and require greater resources to process as CDS unfolds around and to the child in real time (as evidenced from in-lab studies of nonword repetition and real-time spoken word processing among children of this age (Edwards et al., 2004; McMurray et al., 2022).

Finding that the sounds of CDS become more informative as children age begs the question of what may drive this change. One explanation could be the documented age-related changes in lexical input over this time period as caregivers gradually introduce more morphologically-complex and infrequent words (Soderstrom, 2007). These changes diversify the structure of phonological content that children hear as these words may contain more diverse and less predictable phonological sequences, both due to their morphological structure (e.g. affixes introduce new phone combinations) and their general phonological properties (e.g. longer words with more varied syllable structures) (Piantadosi et al., 2011; Hay, 2003). Thus, caregivers might implicitly accommodate the child's phonological processing capabilities via lexical choice. However, we acknowledge that due to the limitations posed by the data, our analysis relied on a more idealized phonemic state of CDS input. A finer-grained analysis could include effects of speaking rate and reduction.

Our results corroborate the results of Tal et al. (2024) who likewise found a reduction in redundancy with age. They attributed the observed changes not simply to increases in lexical diversity, but to decreases in multi-word sequences. Finding recurring evidence for reduced redundancy over developmental time, at multiple levels of language, lends greater support not simply to theories of communicative efficiency (Pate & Goldwater, 2015; Piantadosi et al., 2011), but specifically to how efficiency is modulated based on receiver characteristics over relatively long periods of time. Our results evaluating the relationship between the measures diverged slightly from Cohen Priva & Jaeger (2018) who found a moderate relationship in adult-directed speech ($r = .5$), compared to the average $r = .23$ (varying by age) in our data. These differences could be due to the size of the corpora used and/or the level of annotation within the corpora—for example, at least one of the corpora used in Cohen Priva & Jaeger is phonetically transcribed to exclude under- or-unrealized segments (Pitt et al., 2007).

In sum, we used CDS as a natural experiment to show that transmission of linguistic signals across noisy channels (here due to a child's processing limitations) is optimized, hence balancing complexity and efficiency in the sounds of language directed to children. This works suggests that the dynamic tuning of speech input complexity guides the transmission of speech signals for phonological learning.

References

Aylett, M., & Turk, A. (2004). The smooth signal redundancy hypothesis: A functional explanation for relationships between redundancy, prosodic prominence, and duration in spontaneous speech. *Language and Speech*, *47*(1), 31–56. doi: 10.1177/00238309040470010201

Bell, A., Brenier, J. M., Gregory, M., Girand, C., & Jurafsky, D. (2009, January). Predictability effects on durations of content and function words in conversational English. *Journal of Memory and Language*, *60*(1), 92–111. doi: 10.1016/j.jml.2008.06.003

Brent, M. R., & Siskind, J. M. (2001). The role of exposure to isolated words in early vocabulary development. *Cognition*, *81*(2), B33-B44. doi: 10.1016/S0010-0277(01)00122-6

Brodbeck, C., Hong, L. E., & Simon, J. Z. (2018, December). Rapid Transformation from Auditory to Linguistic Representations of Continuous Speech. *Current Biology*, *28*(24), 3976–3983.e5. Retrieved 2025-01-15, from `https://linkinghub.elsevier.com/retrieve/pii/S096098221831409X` doi: 10.1016/j.cub.2018.10.042

Cohen Priva, U. (2015). Informativity affects consonant duration and deletion rates. *Laboratory Phonology*, *6*(2), 243–278.

Cohen Priva, U., & Jaeger, T. F. (2018). The interdependence of frequency, predictability, and informativity in the segmental domain. *Linguistics Vanguard*, *4*(s2). doi: 10.1515/lingvan-2017-0028

Cychosz, M., Edwards, J. R., Bernstein Ratner, N., Torrington Eaton, C., & Newman, R. S. (2021). Acoustic-lexical characteristics of child-directed speech between 7 and 24 months and their impact on toddlers' phonological processing. *Frontiers in Psychology*, *12*, 712647. doi: 10.3389/fpsyg.2021.712647

Demuth, K., Culbertson, J., & Alter, J. (2006, June). Word-minimality, Epenthesis and Coda Licensing in the Early Acquisition of English. *Language and Speech*, *49*(2), 137–173. doi: 10.1177/00238309060490020201

Edwards, J., Beckman, M. E., & Munson, B. (2004). The interaction between vocabulary size and phonotactic probability effects on children's production accuracy and fluency in novel word repetition. *Journal of Speech Language and Hearing Research*, *57*, 421–436.

Gibson, E., Futrell, R., Piantadosi, S. P., Dautriche, I., Mahowald, K., Bergen, L., & Levy, R. (2019, May). How Efficiency Shapes Human Language. *Trends in Cognitive Sciences*, *23*(5), 389–407. Retrieved 2025-01-13, from `https://linkinghub.elsevier.com/retrieve/pii/S1364661319300580` doi: 10.1016/j.tics.2019.02.003

Hay, J. (2003). *Causes and Consequences of Word Structure*. New York and London: Routledge.

Jones, G., Cabiddu, F., Barrett, D. J. K., Castro, A., & Lee, B. (2023, June). How the characteristics of words in child-directed speech differ from adult-directed speech to influence children's productive vocabularies. *First Language*, *43*(3), 253–282. Retrieved 2025-01-15, from `https://journals.sagepub.com/doi/10.1177/01427237221150070` doi: 10.1177/01427237221150070

Kalashnikova, M., & Burnham, D. (2018, September). Infant-directed speech from seven to nineteen months has similar acoustic properties but different functions. *Journal of Child Language*, *45*(5), 1035–1053. doi: 10.1017/S0305000917000629

Kitamura, C., & Lam, C. (2009). Age-Specific Preferences for Infant-Directed Affective Intent. *Infancy*, *14*(1), 77–100. doi: 10.1080/15250000802569777

Lindblom, B. (1990). Explaining Phonetic Variation: A Sketch of the H&H Theory. In W. J. Hardcastle & A. Marchal (Eds.), *Speech Production and Speech Modelling* (pp. 403–439). Dordrecht: Springer Netherlands. doi: 10.1007/978-94-009-2037-8_16

MacWhinney, B. (2000). *The CHILDES Project: Tools for Analyzing Talk* (3rd Edition ed.). Mahwah, NJ: Lawrence Erlbaum Associates.

Mahowald, K., Fedorenko, E., Piantadosi, S. T., & Gibson, E. (2013, February). Info/Information theory: Speakers choose shorter words in predictive contexts. *Cognition*, *126*(2), 313–318. Retrieved 2025-01-23, from `https://linkinghub.elsevier.com/retrieve/pii/S0010027712002107` doi: 10.1016/j.cognition.2012.09.010

McAuliffe, M., Socolof, M., Mihuc, A., Wagner, M., & Sonderegger, M. (2017). *Montreal Forced Aligner.*

McMurray, B., Apfelbaum, K. S., & Tomblin, J. B. (2022). The slow development of real-time processing: Spoken-word recognition as a crucible for new thinking about language acquisition and language disorders. *Current Directions in Psychological Science*, *31*(4), 305–315. doi: 10.1177/09637214221078325

Newman, R. S., Rowe, M. L., & Bernstein Ratner, N. (2016, September). Input and uptake at 7 months predicts toddler vocabulary: The role of child-directed speech and infant processing skills in language development. *Journal of Child Language*, *43*(5), 1158–1173. doi: 10.1017/S0305000915000446

Pate, J., & Goldwater, S. (2015). Talkers account for listener and channel characteristics to communicate efficiently. *Journal of Memory and Language*, *78*, 1–17.

Piantadosi, S. T., Tily, H., & Gibson, E. (2011). Word lengths are optimized for efficient communication. *Proceedings of the National Academy of Sciences*, *108*(9).

Pitt, M., Dilley, L., Johnson, K., Kiesling, S., Raymond, W., Hume, E., & Fosler-Lussier, E. (2007). *Buckeye Corpus of Conversational Speech* (2nd release ed.). Columbus, OH: The Ohio State University: Department of Psychology.

Rollins, P. (2003). Caregiver contingent comments and subsequent vocabulary Comprehension. *Applied Psycholinguistics*, *24*, 221–234.

Rowe, M. L. (2012). A Longitudinal Investigation of the Role of Quantity and Quality of Child-Directed Speech in Vocabulary Development: Child-Directed Speech and Vocabulary. *Child Development*, *83*(5), 1762–1774. doi: 10.1111/j.1467-8624.2012.01805.x

Sanchez, A., Meylan, S. C., Braginsky, M., MacDonald, K. E., Yurovsky, D., & Frank, M. C. (2019, August). childes-db: A flexible and reproducible interface to the child language data exchange system. *Behavior Research Methods*, *51*(4), 1928–1941. Retrieved 2025-01-13, from `http://link.springer.com/10.3758/s13428-018-1176-7` doi: 10.3758/s13428-018-1176-7

Shannon, C. E. (1948, July). A mathematical theory of communication. *Bell System Technical Journal*, *27*(3), 379–423. Retrieved 2024-11-20, from `https://ieeexplore.ieee.org/document/6773024` doi: 10.1002/j.1538-7305.1948.tb01338.x

Snow, C. E. (1972). Mothers' Speech to Children Learning Language. *Child Development*, *43*(2), 549–565. doi: 10.2307/1127555

Soderstrom, M. (2007, December). Beyond babytalk: Re-evaluating the nature and content of speech input to preverbal infants. *Developmental Review*, *27*(4), 501–532. doi: 10.1016/j.dr.2007.06.002

Tal, S., Grossman, E., & Arnon, I. (2024, August). Infant-directed speech becomes less redundant as infants grow: Implications for language learning. *Cognition*, *249*, 105817. doi: 10.1016/j.cognition.2024.105817

Thiessen, E. D., Hill, E. A., & Saffran, J. R. (2005). Infant-Directed Speech Facilitates Word Segmentation. *Infancy*, *7*(1), 53–71.

Van Kleeck, A. (2004). *CHILDES English Van Kleeck Corpus.* TalkBank. Retrieved 2025-01-31, from `https://childes.talkbank.org/access/Eng-NA/VanKleeck.html` doi: 10.21415/T54883

Wedel, A., Kaplan, A., & Jackson, S. (2013, August). High functional load inhibits phonological contrast loss: A corpus study. *Cognition*, *128*(2), 179–186. doi: 10.1016/j.cognition.2013.03.002

Weisleder, A., & Fernald, A. (2013, November). Talking to children matters: Early language experience strengthens processing and builds vocabulary. *Psychological Science*, *24*(11), 2143–2152. doi: 10.1177/0956797613488145

Zipf, G. K. (1949). *Human behavior and the principle of least effort.* Addison-Wesley Press.