Naturalistic speech supports distributional learning across contexts

Kasia Hitczenko[a,1] **and Naomi H. Feldman**[b]

[a]Laboratoire de Sciences Cognitives et Psycholinguistique, Département d'Études Cognitives, ENS, EHESS, CNRS, PSL Research University, Paris, France; [b]Department of Linguistics and UMIACS, University of Maryland, College Park, USA

This manuscript was compiled on March 26, 2022

At birth, infants discriminate most of the sounds of the world's languages, but by age one, infants become language-specific listeners. This has generally been taken as evidence that infants have learned which acoustic dimensions are contrastive, or useful for distinguishing among the sounds of their language, and have begun focusing primarily on those dimensions when perceiving speech. However, speech is highly variable, with different sounds overlapping substantially in their acoustics, and after decades of research, we still do not know what aspects of the speech signal allow infants to differentiate contrastive from non-contrastive dimensions. Here, we show that infants could learn which acoustic dimensions of their language are contrastive, despite the high acoustic variability. Our account is based on the cross-linguistic fact that even sounds that overlap in their acoustics differ in the contexts they occur in. We predict that this should leave a signal that infants can pick up on, and show that acoustic distributions indeed vary more by context along contrastive dimensions compared with non-contrastive dimensions. By establishing this difference, we provide a potential answer to how infants learn about sound contrasts, a question whose answer in natural learning environments has remained elusive.

phonetic learning | language acquisition | distributional learning

Languages differ in the speech sound inventories they use to reliably convey meaning. For example, Hindi has a distinction between unaspirated retroflex [ʈ] vs. dental [t̪] that is used to distinguish meanings (e.g., [ʈal] means 'postpone', while [t̪al] means 'beat'), but English does not. Adult speakers are generally tuned to the properties of the language(s) they speak. For example, while most adult Hindi speakers can hear the difference between [ʈ] and [t̪], most English-speaking adults cannot (1). Because speech sound inventories differ across languages, listeners must learn about the speech sounds of their language from the input they hear.

The first signs of this phonetic learning appear within the first year of life. During their first couple of months, infants can discriminate most sounds of the world's languages, showing similar perceptual abilities regardless of their language experience. For example, both newborn English-learning and Hindi-learning infants can hear the distinction between retroflex [ʈa] vs. dental [t̪a], a speech contrast that exists in Hindi, but not English. However, over the course of the first year of life, this changes. Infants become language-specific listeners, starting to more closely resemble adults in their discrimination abilities. Their ability to discriminate non-native contrasts (e.g. between retroflex [ʈa] and dental [t̪a]) for English-learning infants) declines (2–4), whereas their ability to discriminate contrasts in their own language improves (5).

These perceptual changes have generally been taken as evidence that infants are learning which acoustic dimensions are contrastive in their language: that is, which acoustic dimensions have multiple categories along them (6). Speech sounds differ in how they are acoustically produced and one or more acoustic dimensions will be used to signal differences between sound contrasts. The idea, then, is that infants become aware of which acoustic dimensions are used to contrast the meaningful sounds in their language, and begin primarily focusing on those dimensions when perceiving speech.

Decades of research into how infants learn about contrastiveness in their first year of life has built a wealth of knowledge in this area; however, we still do not know what aspects of the speech signal allow infants to make these inferences from the acoustically variable speech they hear in their daily lives. One of the most well-studied current proposals for how infants learn which dimensions of their language are contrastive is known as distributional learning (6). It proposes that infants learn the contrastive dimensions of their language by tracking the frequency distribution of sounds along acoustic cue dimensions. If an infant observes a bimodal (two-peaked) distribution along a dimension, then they learn that the dimension is contrastive, whereas if an infant observes a unimodal (one-peaked) distribution, then they learn that the dimension is not contrastive. This account has experimental support: distributions of sounds affect infants' discrimination in the lab (6–9). In addition, when bimodality is present in the input, computational models successfully learn correct contrasts (10, 11). However, a key assumption underlying this proposal is that contrastive dimensions do indeed exhibit bimodality, and while this is the case for some contrasts, recent work looking at naturalistic speech corpora has shown that this is

Significance Statement

Languages differ in the speech sounds they use and humans need to learn which sounds their language uses. This learning starts early. By one, infants have already tuned into their language(s): their ability to hear sound distinctions from their language improves, while they often lose the ability to hear other sound distinctions. Understanding how this early learning proceeds is important as it serves as a foundation for later development; however, it has proven difficult to identify a learning mechanism that works on the true input infants hear. Here, we present an account for how infants learn the speech contrasts of their language and show that the necessary signal is present in naturalistic speech, advancing our understanding of early language learning.

K.H. and N.H.F designed the research; K.H. implemented the analyses; K.H. and N.H.F wrote the manuscript

The authors declare they have no conflict of interest.

[1]To whom correspondence should be addressed. E-mail: kasia.hitczenko@ens.psl.eu

not a universal property of child-directed speech (12, 13).

For example, in Japanese, vowel length is contrastive (14), meaning that two different words like /toko/ (*bed*) and /toko:/ (*travel*) can be distinguished solely by how long a vowel is. However, analyses of a spontaneous corpus of Japanese child-directed speech reveal that the distribution along the duration dimension is unimodal despite being contrastive (12) (Fig. 1a; note that infants do not have access to the individual color-coded short and long vowel distributions shown in this figure, only the combined overall distribution). A similar finding has been reported for Dutch vowel length (13), as well as many other contrasts (5, 15, 16). That is, although infants are able to use distributional information for learning when available, it is not available for all of the contrasts they learn about, so distributional learning is not sufficient.

Many follow-up theories have been proposed to explain how infants learn in cases where bimodality is not present. This has included theories arguing that bimodality might be present when considering only the most prominent sounds (e.g., stressed vowels) (17), when normalizing for effects of neighboring sounds or other factors (18), or when using word-level, visual, or referential information (13, 19–23). While many of them have experimental support and work on controlled lab speech, over the past 40 years, it has proven difficult to identify a learning mechanism that works on the true speech infants hear. (13) takes an important step in that direction by showing that, in Dutch, average vowel durations by word type are often longer in word types with long vowels than word types with short vowels (and, thus, that short and long vowels may be separable). However, we still do not have a measure that consistently separates vowels with a contrast from vowels without a contrast across corpora, languages, and vowel qualities. This problem is so extreme that recent work has suggested that infants might not actually be learning how many phonetic categories there are along a dimension at all, because this signal is not present in their input in a way that they have access to (24).

In this paper, we show that the necessary signal to learn which acoustic dimensions are contrastive may be present in naturalistic input and accessible to infants. Our proposal takes advantage of the contextual information of a sound, which infants are sensitive to (20–22, 25–29). In this work, we take the context of a sound to include factors like its neighboring sounds, its prosodic position in a word/utterance (i.e., if it immediately borders a word or utterance boundary), and its word frame; however, we think of context more broadly as any information that listeners track about where a sound occurs or who spoke it. When an acoustic dimension is contrastive, there are multiple categories along it and the relative proportion of those categories may differ across contexts (e.g., if two categories are present, one context may be 50% category 1 and 50% category 2, whereas another context may be 90% category 1 and only 10% category 2). We show that such differences in category frequency—which are extremely common across languages (30–32)—can help infants distinguish contrastive from non-contrastive dimensions.

We test our proposal on two test cases, Japanese and Dutch, which have been most problematic for both distributional learning and additional previous theories, and show that our proposal explains how infants could nonetheless learn the contrast from information available to them within their first year of life. Complemented by previous findings that (1) infants are sensitive to distribution shapes and contextual information, and (2) changes in the relative proportion of sounds across contexts is a cross-linguistically widespread property of sound categories, these results are promising and suggest that infants may be able to learn about contrastiveness from naturalistic speech input, thus pointing towards a possible answer to a long-standing question in the field.

Distributional Learning Across Contexts

The inspiration for our proposal comes from a finding showing that the context a sound occurs in (neighboring sounds, prosodic position, speaker, etc.) is predictive of its identity: just knowing what context a Japanese vowel appears in can predict its length with around 95% accuracy (33). This means that short and long vowels appear in different proportions in different contexts. Most contexts have almost all short vowels (e.g. Context 1 in Fig. 1b), whereas some contexts have almost all long vowels (e.g. Context 2 in Fig. 1b), and some are in between (e.g. Contexts 3-4 in Fig. 1b). Figure 1b reveals that these changes in the relative proportion of short and long vowels can change the overall shape of the frequency distribution in the context. All of the distributions in Figure 1b are unimodal, despite the fact that there are two categories. Thus, they would not be conducive to the distributional learning theory proposed by (6). However, this is only one aspect of a distribution's shape, and across contexts, the distributions differ in how wide or peaky they are, where they peak, and so forth. This arises because of two facts: (i) when a dimension is contrastive, the overall frequency distribution in each context is the sum of the short vowel distribution and the long vowel distribution, (ii) short and long vowels have different distributions, as can be seen in Figure 1. Taken together, this means that in a language like Japanese, where there is a contrast, we would expect different relative proportions of short vowels and long vowels across different contexts, and since short vowels and long vowels have different acoustic distributions, we would expect the overall distribution to change across different contexts. On the other hand, in a language like French (where there is no length contrast), shape changes cannot arise from different relative proportions of short and long vowels because there is no short vs. long vowel distinction.

In light of this, we propose that infants might learn that a dimension is contrastive by tracking the acoustic distribution along that dimension across different contexts. They could compare the shapes of the distributions across those contexts, and infer that a dimension is contrastive if the shape varies substantially across contexts, but infer that it is not contrastive if the shape is largely the same across contexts. We operationalize a sound's context as (i) its (immediately) neighboring sounds, its prosodic position (whether it falls at a word or utterance boundary), and its quality (learned before length), or (ii) its word frame, due to evidence that infants are sensitive to this information in their input (20–22, 25–29, 34). However, we are not tied to these particular factors. Any contextual factors that infants track, and that change the relative proportion of sound category membership, could work.

It is important to note that the learning outcome of this proposal is the same as in (6), but differs from the learning outcomes of some phonetic learning theories that have arisen since then (19). In particular, the learning outcome here is whether

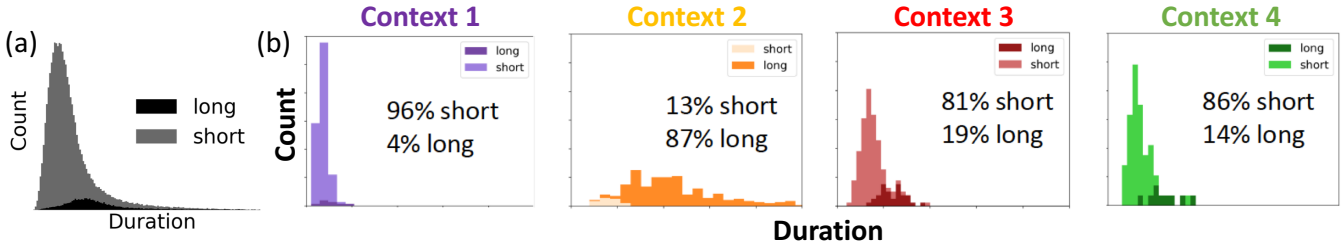

Fig. 1. (a) The frequency distribution of sounds along the duration dimension in Japanese is unimodal, despite vowel length being contrastive. (b) Vowel frequency distributions along duration, for four Japanese contexts (defined by prosodic position, neighboring sounds, and vowel quality). The relative proportion of phonemically short and long vowels changes substantially across contexts, which results in differently shaped distributions. The short vs. long categories are color-coded for the reader's benefit. Infants (and our analyses) do not have access to this color information when learning, only the overall distributions.

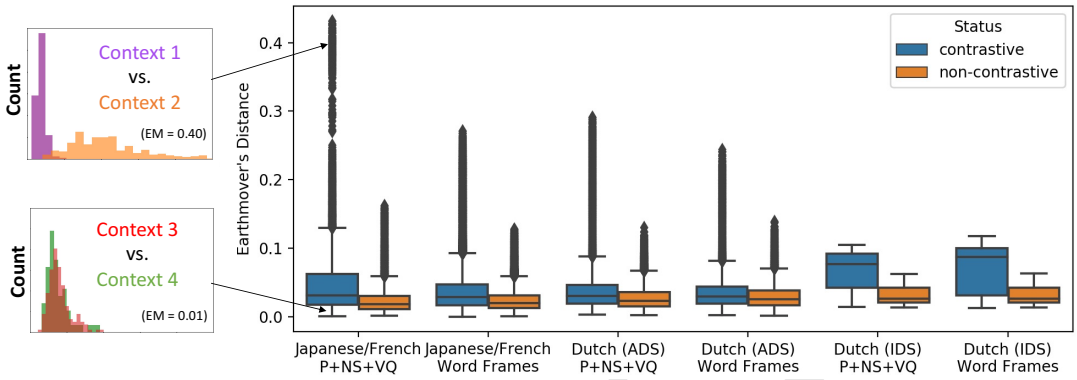

Fig. 2. Distribution of Earthmover's distances by test case. Each datapoint represents the pairwise Earthmover's distance (EM) between distributions from two different contexts (e.g., we show the comparison for Fig. 1b's Context 1 vs. Context 2, which has a high Earthmover's distance, and Fig. 1b's Context 3 vs. Context 4, which has a small Earthmover's distance). Across all test cases, the tail of the contrastive boxplot (left) is longer than that of the non-contrastive boxplot (right), suggesting that there are more extreme distribution shape changes across contexts when the acoustic dimension is contrastive. P+NS+VQ = prosodic position + neighboring sounds + vowel quality.

or not an acoustic dimension is contrastive - i.e. whether it is used to distinguish multiple categories. The learning outcome of some other theories included this knowledge implicitly, but often also included knowledge about what those categories were. Certainly listeners eventually learn about the categories, and a full learning account would need to eventually explain how that happens. However, the discrimination behavior infants exhibit in (2) does not require them to have learned categories (2, 24, 35), so we follow the original literature and focus on how infants learn which dimensions are contrastive in their language.

In what follows, we ask whether the necessary signal for this learning account is present in naturalistic speech; that is, whether there are larger distribution shape changes across contrastive dimensions than non-contrastive dimensions. We focus on three test cases, which each involve some data in which vowel length is contrastive, and some data in which vowel length is not contrastive. We look at vowel length for two reasons. First, it has a largely agreed upon primary cue (duration) that can be easily extracted from any annotated corpus. Second, it is possibly the best known case of extreme overlapping categories that cannot be explained by previous theories like distributional learning (12). We focus on the test cases that have been problematic for past phonetic learning theories, but argue in the General Discussion that this same approach to phonetic learning is likely to be effective across a wide range of languages and contrasts.

Results

Our results confirm that in spontaneous speech there are more extreme distribution shape changes across contexts when a dimension is contrastive than when it is not.

For all of the corpora we study, we extract the acoustic distributions across a number of contexts, and compare them pairwise, using Earthmover's distance (36), a commonly used metric of distribution shape difference (see Supplementary Materials for discussion of results using an alternative metric, KL divergence, instead). We operationalize 'context' in two different ways, both of which rely on information that infants at the relevant age are sensitive to: (i) a combination of prosodic position, neighboring sounds, and vowel quality (P+NS+VQ) and (ii) word frames (WF).

We first compare a spontaneous speech corpus of Japanese (which has a vowel length contrast) against a spontaneous speech corpus of French (which does not). We then test two spontaneous Dutch corpora. Dutch has the property that a subset of its vowels has a length contrast, whereas a different subset does not. Comparing the subset that has a contrast against the subset that does not allows us to control for any effects that may arise due to differences in how the French and Japanese corpora were collected and annotated. Two of our tests examine adult-directed speech (ADS) corpora because they allow us to test this proposal on large-scale, spontaneous speech corpora which do not exist for infant-directed speech (IDS), but we include results from a small corpus of infant-directed Dutch as well.

A. Japanese vs. French ADS. We first compared Japanese and French, defining context as a combination of prosodic position, neighboring sounds, and vowel quality (Fig. 2). Each datapoint contributing to the boxplot represents the Earthmover's distance between a pair of contextual acoustic distributions. For example, the comparison between Context 1 and Context

Context	Percent Long	Count	Frequency Rank
Phrase-initial, word-final /e/	64.7	1357	18
Phrase-initial, phrase-final /a/	56.7	255	95
Phrase-initial, phrase-final /e/	87.9	244	100

Table 1. Information about the Japanese contexts that drive the tail in the case of the P+NS+VQ analysis, including what percentage of vowels in that context are long, how many times that context occurred (Count), as well as its frequency rank out of all contexts that occurred.

2 in Fig. 1b has a high Earthmover's distance, whereas the distance for Contexts 3 and 4 is much small because they are very similar.

The boxplot corresponding to Japanese (where vowel length is contrastive) has a much larger tail, extending upwards towards large Earthmover's distances, than the boxplot corresponding to French (where vowel length is not contrastive). This means that, as predicted, there are many more pairs of contexts that have substantially different shapes (like Context 1 vs. Context 2 in Fig. 1b) when there is a contrast than when there is not. The maximum distance, the mean distance, and the distance variance are all larger for Japanese than French (max = 0.43 vs. 0.16; mean = 0.05 vs. 0.02; variance = 0.003 vs. 0.0004). Analyzing the contents of the tail in Japanese reveals that the tail is driven by contexts that have a much higher percentage of long vowels than observed overall and that occur frequently in the input (see Table 1 for frequency counts and ranks of the contexts that drive the signal).

These same patterns hold when we continue looking at French vs. Japanese, but instead use word frames as contexts. As before, there are more contexts with more extreme distribution shape changes in Japanese than French (i.e. along contrastive than non-contrastive dimensions), as seen by the longer tail in the second pair of boxplots in Fig. 2. As before, the maximum distance, the mean distance, and the distance variance are all larger for Japanese than French (max = 0.27 vs. 0.12; mean = 0.04 vs. 0.02; variance = 0.001 vs. 0.0002).

B. Dutch ADS and IDS. To test our proposal using a within-language comparison, we compare the subset of Dutch vowels that do contrast in length and the subset of Dutch vowels that do not. We find that the predicted pattern still holds - and it holds for both ways of defining context and both the ADS and IDS corpora (Fig. 2). This confirms that the results are not merely an artifact of using different corpora, as in the French vs. Japanese case, but seem to reflect something inherent to the existence or nonexistence of categories along an acoustic dimension. It is worth noting that Dutch-learning infants would not be able to perform this exact analysis to learn whether there is a length contrast, because they would not yet know enough to separate the vowels into contrastive and non-contrastive subsets. We return to the issue of what a learning account might look like in the Discussion. Meanwhile, we conclude from this analysis that the signal our account predicts exists in Dutch: contrastive dimensions differ from non-contrastive dimensions.

Despite the qualitative similarity in results across all test cases, the scale of the difference in tail length varies. For example, in the French vs. Japanese P+NS+VQ case, the maximum Earthmover's distance in the contrastive Japanese case is 0.43, whereas for the other ADS cases, the maximum is less than 0.3. In the Dutch IDS corpus, which only has 284

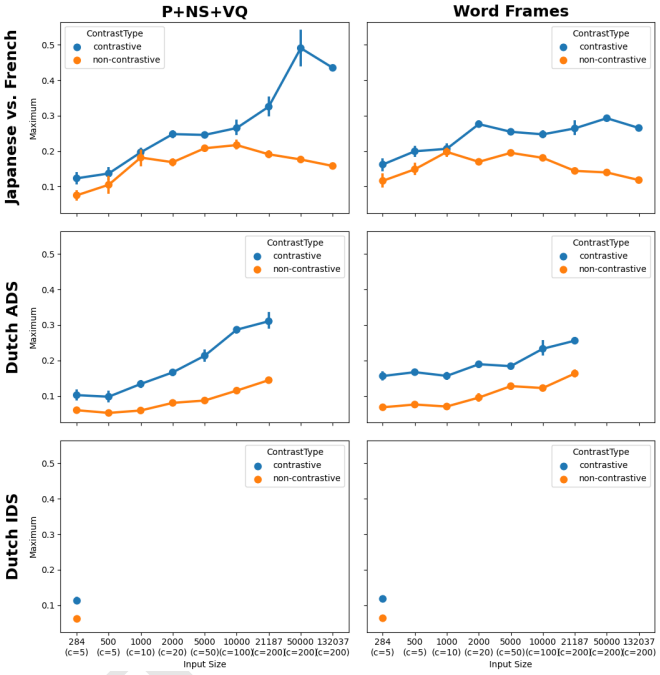

Fig. 3. Bootstrapped analyses reveal that observed differences between contrastive (top/blue line) vs. non-contrastive (bottom/orange line) dimensions are meaningful, but that input size does matter. "C" refers to the number of contexts included in the analysis. The maximum input size for which data is shown depends on the corpus size: 284 for Dutch IDS, 21187 for Dutch ADS, and 132037 for Japanese vs. French.

vowel tokens, the maximum is only around 0.1. One possibility is that these differences arise because of the large differences in corpus size. The Japanese vs. French corpora considered 132,037 tokens and the Dutch ADS corpus considered 21,187 tokens, but the Dutch IDS corpus only considered 284 tokens.

C. Corpus size analyses. To test how corpus size impacts results, we used bootstrap samples to run each analysis 50 times for 10 different corpus sizes ranging from the size of smallest corpus (284) to the size of the largest corpus (132,037). This also allowed us to test how much the size of the tail varied, and whether differences observed between contrastive vs. non-contrastive cases were meaningful. Fig. 3 shows these results when calculating the maximum Earthmover's distance across all 50 runs; analogous plots for mean are provided in the Supplementary Materials. First, this analysis reveals that the differences observed are meaningful: across many runs, at large enough corpus sizes, the contrastive line is higher than the non-contrastive line. That being said, in the Japanese vs. French case, the difference does not emerge until around 2000 vowel tokens have been observed, so input size does matter. Second, this analysis reveals that differences in scale may be partially, but are not entirely, due to corpus size. When subsetted to the size of the Dutch ADS corpus, the Japanese vs. French word frame maximum matches the remaining ADS results. However, the results are less clear for Dutch IDS: subsetting the Dutch ADS corpus to the size of the Dutch IDS corpus yields results more in line with each other for the P+NS+VQ analysis, but less so for the word frames analysis.

From a learning perspective, this means that an ideal learner would need to observe around 2000 vowel tokens and track the acoustic distribution within the 20 most frequent contexts in order to observe the difference (though we discuss potential ways to reduce the memory demands of the proposal next as

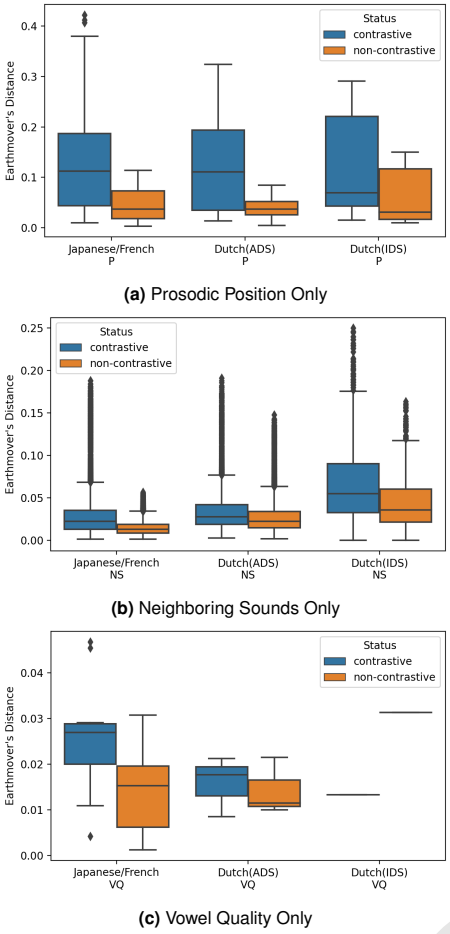

Fig. 4. Results are similar when we relax our assumptions about infants' knowledge. Here, instead of studying a combination of factors, we study (a) prosodic position (P), (b) neighboring sounds (NS), and (c) vowel quality (VQ) individually (left: contrastive dimension; right: non-contrastive dimension). While less clear for vowel quality, we see the same pattern of results in the prosodic position and neighboring sounds only cases.

well as in the General Discussion).

D. Relaxing our assumptions about infants' knowledge. Our analyses so far have been conducted assuming that (a) infants can track acoustic distributions across combinations of three contexts (prosodic position, neighboring sounds, and vowel quality), (b) infants can perfectly identify neighboring consonants, and (c) infants can perfectly segment words from speech. Although some of these assumptions have yet to be tested (e.g., we don't know whether infants can track distributions along multiple contextual dimensions), we know that others are likely overestimating infants' prior knowledge when learning about contrastive dimensions. Here, we show that the same qualitative results still emerge even when we weaken these three assumptions.

The first assumption we revisit is whether infants can track distributions across combinations of contexts (i.e., prosodic position, neighboring sound, and vowel quality). We test what happens when we study each of these three factors individually. Fig. 4 shows that, especially for prosodic position, but also for the other factors, the same patterns, for the most part, emerge. While it is still quite conceivable that infants track

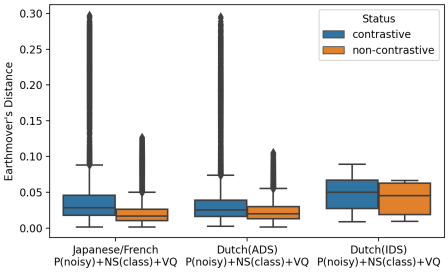

Fig. 5. Results are similar when we relax our assumptions about infants' knowledge. Here, we define neighboring sounds by their broad class (stop, fricative, etc.) and add noise to the prosodic position with a noise rate of 20% (left: contrastive dimension; right: non-contrastive dimension).

distributions across combinations of contexts, these results suggest that this need not be the case in order for our account to be successful. It suffices to track distributions across individual contexts (e.g., word frames or prosodic position). From a learning perspective, this means that infants would only need to be tracking the acoustic distribution across ∼4-10 contexts (rather than ∼20 as observed before), and these contexts include extremely prominent contexts (e.g., utterance-final, utterance-initial tokens).

The next two assumptions we revisit are whether infants can perfectly encode the identity of neighboring sounds, and whether they have a solid enough grasp on word segmentation to have access to the prosodic position information we use. On the one hand, prosodic boundaries are one of the first signals that infants are sensitive to (26, 27); however, there is concurrent evidence that infants make missegmentation errors (37–39). To address these two assumptions, we test what happens when we re-run the P+NS+VQ simulations, with neighboring sounds defined by their broad class (i.e., stop, fricative, vowel, approximant, etc.) rather than their particular identity (e.g., /k/, /g/, /b/) and with noise added to the prosodic information (we simulate a 20% error rate here, such that 20% of the time, the infant misrepresents the prosodic position of the vowel, but the results generalize across error rates). Fig. 5 shows that these differences do not qualitatively change the results, suggesting that even with a more rudimentary grasp on contextual factors, infants could still use this method to learn the sound contrasts.

Overall, we show that the necessary linguistic knowledge and capabilities can be considerably reduced and yet the correct finding still emerges. This suggests that this finding is a robust one that immature learners could learn from even in noisy learning environments.

E. Analyses with long vowels removed. Finally, to test whether these results arose because of the contrast, we removed all vowels labeled as long from the corpora and reran the same analyses. We predicted that removing the long vowels would cause the tail for the contrastive dimension to disappear, such that the results for the contrastive dimension with long vowels removed (i.e. with the contrast artificially removed) would resemble those of the non-contrastive dimension. As can be seen in Fig. 6, in Japanese and in Dutch IDS (the two hand-annotated corpora we use), the tail disappears or is reduced once the long vowels are removed, suggesting that it is at least partially the presence of the long vowels that causes these large changes in distribution shape. However, this is not case in Dutch ADS: the contrastive dimension still

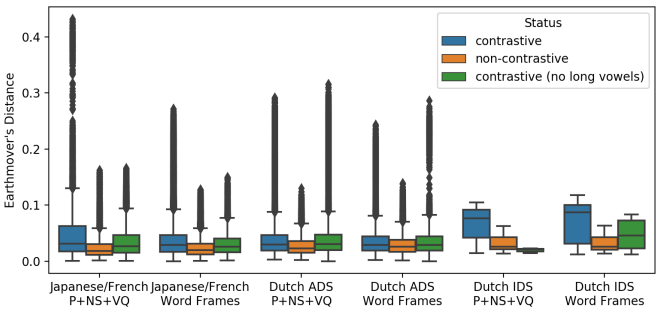

Fig. 6. Results from artificially removing long vowels. Within each case, the left boxplot corresponds to the contrastive dimension, the middle boxplot corresponds to the non-contrastive dimension, and the right boxplot corresponds to the contrastive dimension with long vowels removed. We observe that the tail length is reduced for the Japanese/French and Dutch IDS analyses; however, against predictions, not for the Dutch ADS analysis.

has a longer tail even when long vowels are removed. While this is not predicted by our account, there are a number of reasons why we may observe this result here. One possibility is that the annotations are imperfect. The remaining datasets studied here include hand-corrected segmental annotations, while the Dutch ADS data included force-aligned annotations which were not validated for their duration. Another possibility is that it has to do with the fact that these data were phonetically annotated, rather than phonemically annotated. That is, some phonemically long vowels were marked as being short. In fact, (40) reports that as many as 20% of word tokens that had long vowels underwent a shortening process. Especially combined with the fact that these data were automatically annotated, this could mean that we are unable to actually remove all long vowels and that some long vowels remain which are driving the differences observed. Finally, another possibility is that the presence of long vowels in a context changes the distribution of short vowels. For example, if a context is 50% short vowels and 50% long vowels, then the short vowels may be pronounced with shorter durations than in a context with 90% short vowels and only 10% long vowels (i.e. to better differentiate the vowel types). If this is the case, then even if we remove the long vowels, we should expect to see differences in the short vowel distributions across contexts. However, this finding is nonetheless different from what our account predicts and more work should be done with hand-annotated corpora to see whether this finding remains and, if so, what is driving it.

Nonetheless, across the board, we do observe that the contrastive dimensions have a longer tail than the non-contrastive case, suggesting that contrastive dimensions exhibit more extreme distribution shape changes across contexts (even though all distributions are unimodal) than non-contrastive dimensions. These results show that contrastive dimensions look different than non-contrastive dimensions, and that infants would learn the correct generalization about their language by using this signal. As such, this is one of the first pieces of signal that has been shown to successfully differentiate contrastive and non-contrastive dimensions using spontaneous speech.

Discussion

In this paper, we proposed a new account for how infants could learn which acoustic dimensions of their language are contrastive. The idea is that infants track the frequency distributions of sounds in different contexts, and that they learn that an acoustic dimension is contrastive if the distribution shape along that dimension varies substantially across different contexts. For this account to work, it needs to be the case that the distribution shape varies more across contexts when a dimension is contrastive than when it is not. We tested this prediction in three test cases, with two different definitions of context. Across the board, our results show that the distribution shape along an acoustic dimension changes more across different contexts when that dimension is contrastive than when it is non-contrastive. This is a signal that differentiates contrastive and non-contrastive dimensions, and it is the kind of signal that listeners are likely to be sensitive to. As such, this is one of the first phonetic learning accounts that has been shown to work on spontaneous data and suggests that infants could be learning which acoustic dimensions are contrastive after all. In the remainder of the paper, we discuss the promise and open questions of this proposal, including its generalizability, evidence on whether infants have the necessary sensitivities, and how infants could use this signal to learn.

Generalizability. We considered the test cases of Japanese and Dutch vowel length because they are famous problem cases for many of the phonetic learning theories that already exist (12, 41). However, they are unique contrasts in a number of ways. They have *low functional load* in that they are not frequently used to distinguish different meanings and they have particularly overlapping acoustic distributions. In addition, the Japanese contrast is primarily signaled by one acoustic cue (i.e. duration), while most contrasts are signaled by more (42, 43) and, as around 90% of Japanese vowels are short, it is less balanced than many other contrasts. Future work will need to test the generalizability of this proposal. Nonetheless, we think it is likely that this signal will generalize to other contrasts because the signal we illustrate in Fig. 1b does not stem from these idiosyncrasies. Rather it arises because of a handful of contexts that have particularly different frequency distributions and we think these exist because of properties of language have phonologists have argued are universal (30).

All languages are thought to have phonotactic or co-occurrence constraints. For example, for the English contrast [n]-[ŋ] (e.g. *sin* vs. *sing*), [n] can occur at the beginning or end of a syllable, while [ŋ] can only occur at the end of syllables. This means that the relative proportion of [n]'s and [ŋ]'s will change across those contexts, and could show patterns like in Fig. 1. Similarly, sometimes sounds will be pronounced differently based on their context. For example, if we consider the contrast between [n]-[m], the [n] sound never occurs before the sounds [p,b], but is instead pronounced as [m] in those contexts (as in *impossible*). These types of phonological alternations also create differences in which sounds occur in which contexts and are cross-linguistically widespread (30–32). Finally, there are systematic regularities based on the words in the language. For example, for the [p]-[b] contrast in English, [b] is more likely in the word frames _aby and _ar (*baby* is a word, but *paby* is not and *bar* is more frequent than *par*), whereas [p] is more likely in the word frames _lay and _in (*play* is a word, but *blay* is not and *pin* is more frequent than *bin*). Taken together, phonotactics, phonological alternations, and word regularities all create systematic regularities in the contexts that sounds occur in, such that different sounds occur in different contexts and different contexts are made up of

different ratios of sounds. This is currently thought to be true of all contrasts - even those that are equally balanced, have higher functional load, or are multidimensional - and will be the case in adult- as well as infant-directed speech. Indeed, (33) showed the context is similarly predictive of which category a sound belongs to in both adult-directed and infant-directed speech.

Nonetheless, while our results are promising, they ultimately come from one acoustic dimension, one contrast type, and only three languages. Furthermore, there are additional complexities that we have not considered here, like how extralinguistic factors (e.g., speaking style/register, emotional content of speech, speech rate) could affect the signal and how this proposal interfaces with other learning strategies that infants have been shown to use (e.g., using visual and referential information).

To test generalizability, future work should replicate these findings on other data sets as they become available (we have made our code publicly available to facilitate this effort). Particular focus should be placed on replicating these results on a large ecological corpus of infant-directed speech, replicating these results on other contrasts whose acquisition is difficult to explain (e.g., Filipino nasals (44)), studying Dutch to understand why removing long vowels does not always change the result, as well as replicating these results in languages where the identity of a particular sound cannot be predicted as well from its context, due to having fewer phonotactic and other systematic restrictions (45). It will also be important to test this proposal on contrasts that are signaled by multiple acoustic dimensions (as the contrasts we study here are unique in primarily being signaled by one: duration). In order for this account to work in those cases, we would need to observe the signal we report here along at least one of the acoustic dimensions that signal the contrast (e.g., along VOT or F0 for stop contrasts). Alternatively, rather than operating over individual acoustic dimensions, infants could search for this signal along composite acoustic dimensions that are discovered from the input by combining dimensions that are highly correlated in the input (e.g., as discussed in (35, 46)). Finally, we note that Earthmover's distance can be calculated over multidimensional distributions, so it should be possible to scale this approach up to multidimensional contexts if so desired; however, for that to work, the infant would have to first identify the relevant combinations of acoustic dimensions they should focus on. Overall, it will be important to replicate these findings across highly variable corpora that adequately represent the full range of speech types that infants could encounter.

Finally, it will also be important to test that these results do not *overgeneralize* and wrongly label non-contrastive dimensions as contrastive. Here, it will be particularly important to test behavior on allophonic variation, where a particular sound is realized differently depending on the context it occurs in. While this pattern is similar to that of different phonemes occurring in different contexts, our analyses provide preliminary evidence that allophonic variation is not labeled as contrastive. In particular, in French, vowel length varies allophonically (vowels are lengthened depending on the following consonant) (47), yet our analyses reliably treated French (allophonic variation in vowel length) differently from Japanese (phonemic variation in vowel length). This suggests that this method may correctly differentiate contrastive and allophonic variation, though it will be important to study this further and understand how it does so (if it does). We offer two speculative reasons why allophonic variation may not be detected to the same degree as phonemic variation, though they will need to be tested. First, this proposal relies on contrasting sounds having sufficiently different acoustic distributions (so that changing the relative proportion of the sounds changes the shape of the distribution). It is possible that allophonic variation changes how a sound is produced to a lesser degree than phonemic variation, though this may be difficult to assess. A second possibility is that allophony often, though not always, affects all of the sounds produced in a particular context (e.g., the duration of all French vowels is affected by neighboring consonants). This may lead to a shift in distribution between contexts, without a change in shape, which may lead to smaller distribution shape changes as measured by Earthmover's distance. Certainly, more research studying how this proposal handles allophony will be crucial.

Could infants do this?. This proposal places higher computational and memory demands on infant learners than many past theories have (e.g., distributional learning). Infants would need to be able to track distribution shapes across many contexts and then compare their shapes pairwise.

(1) *Could infants track distributions across different contexts?* Though this has not been tested, we know that listeners track the shape of frequency distributions overall, and use this for phonetic learning and processing. For example, infants make different phonetic inferences depending on whether the distribution they hear is bimodally- or unimodally- shaped (6). Another study showed that adult listeners are sensitive to the variance of the sound distributions they are exposed to, another property of a distribution's shape (48). In that experiment, adults categorized sounds differently depending on whether they heard a distribution with high or low variance, though this has not been tested in infants.

In addition, we know that listeners are sensitive to the context of a sound and use it for phonetic learning and processing. Both toddlers 12 months and older and adults have been argued to track acoustic distributions across speakers (which can be thought of as a context), can adapt to speakers who have different accents (i.e. different distributions of sounds) (49–55), and mirror the speech of their interlocutors. In addition, infants are sensitive to phonotactics (56), as well as phonological alternations - the fact that sounds tend to be pronounced differently in different contexts (57, 58). Additionally, multiple studies have shown that infants use the word frame of a sound in phonetic learning (20, 22). That is, infants seem to assign acoustically similar sounds to different categories if they occur in different word frames, suggesting that infants can track the context that a sound occurs in and use it for phonetic learning. Adult speech perception is affected by contextual factors, like neighboring sounds. For example, (59) showed that Japanese perception of whether the final vowel in CoC'V was phonemically short or long depended on the identity of both C and C'. Furthermore, (60) showed that adults continually track how informative particular acoustic cues are and will selectively reweight these cues in some contexts but not others (e.g., when categorizing /b-p/ in beer-pier, but not for /d-t/ in deer-tier and vice versa). Indeed, listeners may even track information across contexts defined by multiple factors.

For example, in English, to learn whether a stop consonant will be aspirated or not, the speaker must track whether the sound is voiceless or voiced (akin to Vowel Quality), whether it occurs at the start of a stressed syllable or not (akin to Prosodic Position), and what the neighboring sounds are and then notice the change in pronunciation/distribution across those different contexts.

Taken together, while these results do not provide direct evidence that infants track distribution shapes across contexts, they suggest that listeners can track complex statistical regularities across complex contexts and use them in real-time phonetic learning and processing.

(2) *Could infants compare distribution shapes?* Finally, the last skill necessary for the proposal is for infants to be able to compare distributions. While this has again not been directly tested, one possible clue is that listeners seem to reweight acoustic cues depending on how variable/informative they are, with cues that have narrower distributions being more informative than cues with wider distributions (48, 61, 62). Another possible clue is that toddlers and adults are able to identify when they need to adapt their representations to speech they hear. Being able to identify an accent implies that listeners can identify when the speech they are hearing differs from the speech they usually hear, a computation that is likely to involve tracking at least some properties of the distributions (63).

Overall, given infants' demonstrated sensitivity to distribution shape and to changes across context, there is good reason to believe that infants could be sensitive to the type of distributional information that our account assumes, but future work should test whether infants/listeners can track distributions across different contexts (defined by one or more contextual factors) and compare distributions' shapes. In addition to testing whether infants can, in theory, perform the computations this account requires, future work should also test whether infants actually use them to learn about contrastiveness in the way we propose here. One approach would be to test this experimentally, by exposing infants to acoustic distributions that differ or remain the same across contexts and seeing if this affects their learning/behavior. Another approach would be to use cross-linguistic corpora to identify contrasts that should be easier/harder to learn according to our proposal and compare that against age of acquisition and speech perception/production data. For example, controlling for degree of acoustic overlap, this theory would predict that contrasts that have stronger phonotactic restrictions or that are more predictable (i.e., it is easier to predict which member of the contrast occurred based on the context it occurred in) should be easier to learn through this method. These approaches will allow us to overcome the next big hurdle for this account, which is determining whether infants use this signal to learn.

Reducing the computational complexity of the proposal. It is also possible that the memory and computational restrictions of the proposal could be reduced. On the one hand, we saw that considering individual contexts (e.g., just the most frequent word frames or just prosodic position) was still effective, as was introducing some parsing errors and considering broader segment classes rather than individual neighboring sounds. On the other hand, online approximations or metrics that do not require the whole distribution to be tracked in order to get a measure of distribution shape distance could also reduce the computational and memory complexity of the proposal. For example, rather than exactly representing the distribution, this proposal could operate over a compressed representation of the distribution that keeps track of how many points fall within larger bins/bands (similar to reducing the number of bins in a histogram). It is likely that we would still observe the critical pattern even with this less detailed representation of distribution shape, and, as this only requires keeping track of one number per bin (its count), it could reduce the size of the representation of a distribution to just 5-10 numbers. Even considering all 200 of the most frequent contexts, this could involve storing as few as 1000 numbers. In addition, there may be a way to zero in on the contexts that yield the necessary signal without doing all of the pairwise comparisons represented in the boxplots. It is possible that the "key" contexts that drive the signal are overrepresented in the outliers of the overall distribution (i.e., particularly short or particularly long vowels could be more likely to occur in a context that drives the tail). If this were true, infants could arrive at the same signal we observe here, by focusing in on the contexts of outlier sounds, rather than tracking the distribution across all contexts. Finally, another possibility is that infants could compare contextual distributions using higher-order measures of distribution shape (e.g., variance) rather than tracking the entire distribution. The fact that the distribution shape changes across contexts could also mean that the variance of the distributions changes across contexts. If so, infants could pick up on this difference without encoding the entire distribution across contexts.

Moving from signal to learning account. The data we report plots Japanese and French side-by-side, but most Japanese infants do not get French input to compare against. Assuming that this pattern generalizes to other contrasts and that infants have the necessary sensitivities to detect this signal, how could infants actually use it to learn?

One possibility is that infants use a built-in threshold to determine whether a dimension is contrastive: if the metric (this could be something like the average, range, variance, or maximum Earthmover's distance) exceeds the threshold, they learn the dimension is contrastive; otherwise, they learn that it is not. Another possibility is that infants compare against other acoustic dimensions of their own language (instead of against other languages, as we did). If these metrics turn out to be larger for all contrastive dimensions than all non-contrastive dimensions, infants could easily separate contrastive vs. non-contrastive dimensions.

One complication for this possibility, however, is that the metrics we report are sensitive to the scale of a dimension, making it difficult to compare across dimensions with different scales (e.g., formants vs. duration). To overcome this problem, we tried z-scoring the dimensions, but found that the key effect partially disappeared: the effect was retained for French vs. Japanese and when considering only the most frequent contexts, but when more contexts were considered, the pattern reversed in Dutch ADS. This happened because z-scoring is sensitive to variance, and as there was more variability along the contrastive than non-contrastive dimensions, z-scoring led to artificially lowered Earthmovers' distances along the contrastive (more variable) dimension. Nonetheless, a method that standardizes the scales with less sensitivity to the overall variance could allow for comparison across dimensions.

Finally, given the overall distribution along an acoustic dimension, infants could have a probabilistic model of how different they should expect distributions to be across contexts, depending on whether the overall distribution is made up of one vs. two categories. With this, they could compare how likely a one-category vs. two-category solution is to have generated the observed Earthmover's distances (controlling for phonological processes, it would be unlikely for one category to produce extremely different distribution shapes as in Fig. 1b).

One issue that should be considered in future work is whether learners consider aggregate distributions across all of the vowel qualities when comparing distribution shapes across contexts. In our analyses of Japanese and French, we have assumed that they do, and in fact, vowel quality is one of the contextual factors we analyze. However, this creates a potential problem for Dutch, because only a subset of Dutch vowels contrast in length. If Dutch infants were to initially rely on the aggregated distribution shape comparison approach we put forth for Japanese, this might lead them to conclude that vowel duration is a contrastive dimension that they should tune into (as the context pairs that showed high Earthmover's distance would still be in this analysis), but they would not realize that only some of the vowels contrast. It is possible that they could later learn which specific vowel qualities contrast in length. It is also possible that Dutch infants use a different strategy entirely for discovering the vowel length contrast: those Dutch vowels that contrast in length also contrast in vowel quality (e.g., [ɑ] vs. [aː] contrast in Dutch, but [a] vs. [aː] contrast in Japanese). Having already separated [ɑ] from [aː] using their vowel qualities, Dutch infants could simply notice that these vowels differ systematically in their durations, without doing any distribution shape comparisons of the type we propose here. Ultimately, we remain agnostic as to how exactly infants learn. Given the complexity of the task infants are faced with as well as past experimental findings showing infants use many types of information in phonetic learning, it likely involves a combination of strategies (e.g., using word-level, visual, referential, and other distributional information in addition to the types of analyses we report here). Having established that a signal exists in naturalistic data, we hope future research will investigate how this signal might best be used for learning and how it integrates with other promising accounts.

Conclusion

Infants need to learn which acoustic dimensions of their language are contrastive in order to learn the sound system of their language. However, we still do not know what aspects of naturalistic input provide the necessary signal for them to do so. In this paper, we propose a potential account for how infants learn this and show that there is a signal about whether a dimension is contrastive in noisy, spontaneously produced input. This account is particularly promising for two reasons. First, the signal that we pick up on is a direct consequence of multiple categories exhibiting properties that hold true across most languages, so we think it is likely that this result will generalize to other contrasts. In addition, the signal is something that even infants may be sensitive to. Past work has shown that infants track the shapes of overall frequency distributions, and know about how sounds are likely to sound in different contexts (6, 57, 58). Adults have been shown to track distributions across situations (i.e. across different talkers) (49). In conclusion, we show that even when two sounds overlap acoustically, the fact that they occur in different contexts leaves signal to their contrastiveness. These results provide initial support for a phonetic learning account that works on highly acoustically variable spontaneous speech.

Materials and Methods

Methods. For each test case, one of the datasets (contrastive or non-contrastive) was larger than the other. To correct for this, we only considered the first N tokens of the larger dataset, where N was the size of the smaller dataset. We extracted the duration, the primary acoustic cue to length, of each vowel token in seconds, rounding to the same degree of precision. In addition, we extracted all contextual information that was available across all of the corpora we study and that infants of the relevant age are sensitive to:

- **Vowel quality**: For Japanese, this was: /a/, /e/, /i/, /o/, or /u/. For French, this was: /a/, /e/, /i/, /o/, /u/, /y/, /ø/, /ã/, /õ/, or /œ̃/. For Dutch, this was: /ɑ-a/, /ɔ-o/, /œ-ø/, /ɛ-e/, /ɪ/, /u/, /y/, or /i/. The first four listed pairs are differentiated by quality and length, but we do not incorporate these vowel quality differences into this paper. Vowel quality is thought to be learned before vowel length (25).
- **Prosodic position**: We represented prosodic position (a vowel's position relative to prosodic boundaries) with four indicator values: (1) whether the vowel was word-initial or not, (2) whether the vowel was word-final or not, (3) whether the vowel was phrase-initial or not, and (4) whether the vowel was phrase-final or not. Infants have been shown to be sensitive to prosodic boundaries quite early (26, 27).
- **Neighboring sounds**: We extracted the identity of the immediately previous sound and the immediately following sound, as labelled by the phonetic transcription, ignoring length information. Again, vowel length contrast is thought to be learned later than other types of contrasts (25).
- **Word frame**: We extracted the word frame that the vowel occurred in, excluding all length information. For example, one word frame could have been [b_i_ru], which would include both [biːru] and [biru]. We chose to include word frames, as infants know and can segment words early (28, 29, 34, 64), and use word frames in phonetic learning (20–22).

We looked at two main ways of defining context, though we do not have any commitments about which contexts infants would compute over. In the first way of defining context, we used a combination of vowel quality, prosodic position, and neighboring sounds (e.g. /o/ vowels that follow a /t/ and precede a /k/ that are word- and phrase-internal) - though we also consider each of these three contextual factors individually. This combined set of factors corresponds to the subset of factors considered in (33) that were available for the corpora we study and that infants are most sensitive to. In the second way of defining context, we used word frames, as has been done in (19) and (13) among others.

Because most contexts occur very infrequently, we looked at a subset of all possible contexts. We subsetted the contexts in two qualitatively different ways: either by taking the top X most frequent contexts, or by taking all contexts that had at least N tokens, varying X and N. Results were qualitatively similar in all cases, so we present results from including the 200 most frequent contexts for the French vs. Japanese and Dutch ADS analyses, and only the 5 most frequent contexts for the Dutch IDS analysis due to its much smaller size. Once we had the contexts, we extracted the vowel duration frequency distributions in each context (examples shown in Fig. 1).

We compared the shape of each pair of contextual frequency distributions, using a metric known as Earthmover's distance or Wasserstein distance (65, 66), which is commonly used to measure the difference in shape between two distributions (see Supplementary Materials for methods and results using KL divergence instead). Earthmover's distance is often talked about in terms of two piles of dirt, which represent the two distributions being compared. In this

context, Earthmover's distance can be thought of as the minimum cost of turning one earth pile into the other, where cost corresponds to a combination of the amount of earth being moved as well as the distance it has to be moved. In other words, the distance is the minimum average distance a piece of dirt will have to be moved in order to turn one pile into the other. A higher distance means there was a greater shape mismatch. We plot the distribution of Earthmover's distances, and report its mean, variance, and maximum.

Having found that the distribution of Earthmover's distances for contrastive dimensions had a longer tail than for non-contrastive dimensions, we qualitatively analyzed the contents of this long tail to determine which individual contexts led to the pattern observed. We identified the contexts that showed up most frequently in the tail and analyzed how frequent they were (both in terms of absolute count and their frequency ranking relative to all contexts), as well as what the relative frequency of short and long vowels was in each of these key contexts.

To assess the reliability of these differences between contrastive and non-contrastive dimensions were reliable, we used bootstrap statistics. We sampled (with replacement) particular vowel tokens with their contexts, creating a new contrastive and a new non-contrastive dataset. We then recalculated Earthmover's distances across the contexts using these bootstrap samples, repeating the process 50 times. We plot the maximum Earthmover's distance with the standard deviation, which allows us to observe the reliability of these differences. To study the effect of the input corpus size, we varied the number of vowel tokens sampled from 284 vowel tokens (the size of the Dutch IDS corpus) to 132037 (the size of the French vs. Japanese corpora).

In our final simulation, we relaxed our assumptions about infants' prior knowledge. First, while previous analyses used neighboring segment identity directly (e.g., /k/, /t/, /s/), this simulation only used the segment's broad class (e.g., stop, fricative, vowel, etc.), which infants are more sensitive to. Second, to simulate imperfect segmentation, we added noise to the prosodic position factor. Prosodic position is represented with four indicator values (depending on whether the vowel in question is word/utterance-initial and word/utterance-final). To add noise, we changed 20% of these values (making sure that the resulting prosodic position was real - e.g., sounds considered to be utterance-final were necessarily also considered to be word-final). We then used the same procedure from above with these updated factors.

Corpora. The French vs. Japanese analysis compared the Corpus of Spontaneous Japanese against the Nijmegen Corpus of Casual French. The Dutch analyses looked at the Ernestus Corpus of Spontaneous Dutch (ADS) and the Levelt/Fikkert corpus (IDS).

Corpus of Spontaneous Japanese (CSJ). The CSJ is a large corpus of spontaneously produced adult-directed speech (67). Around 90% of the speech consists of spontaneously produced monologues about academic field, their favorite memory, and so forth. The remaining 10% consists of spontaneous dialogues either in free conversation with the experimenter or engaged in a task. Our analysis focuses on the core portion of the corpus, which was force-aligned and hand-corrected with the segmental information required for our analyses (see (67) for more details). The core portion consists of 811,731 total vowel tokens of which 89.1% are phonemically short and 10.9% are phonemically long, but only the first 132,307 tokens were used to match the size of the French corpus.

Nijmegen Corpus of Casual French (NCCFr). The NCCFr is a corpus of spontaneously produced adult-directed speech (68). Unlike the CSJ, however, the NCCFr consists exclusively of conversational speech between close friends. Topics included upcoming exams, travel plans, an ongoing strike, and so forth. The corpus consists of speech by 46 French speakers, and includes 132,307 vowel tokens. The corpus was orthographically transcribed by two professional transcribers. The corpus was transcribed at the segmental level by Martine Adda Decker (p.c. with M. Ernestus, January 14, 2019).

Ernestus Corpus of Spontaneous Dutch (ECSD). The ECSD consists of adult-directed, conversational speech, with speakers talking with a friend, at first freely, and then engaged in a task-oriented discussion (40). The corpus has speech by 20 different speakers, and includes 60,955 tokens with a length contrast and 21,187 tokens without. Professional transcribers created a orthographic transcription of the interactions, which was manually aligned to the speech. The corpus was also phonetically transcribed using a forced alignment model (details can be found in (69)). Validations revealed a 14% discrepancy between manual annotations and forced-aligned annotations, which is in the range of human disagreement. However, these analyses did not directly validate durational information, so it is unclear how accurate annotations of the start and end points of the phones are. This could introduce some noise into our analyses, as it could affect how accurate the vowel durations are, and how accurately we can determine which word a vowel belonged to.

Fikkert/Levelt/Swingley IDS corpus. We also tested our account on a corpus of Dutch IDS collected by Fikkert and Levelt (70, 71). The annotated portion of this corpus is small: it contains a total of 300 utterances, with a total of only 1296 vowel tokens, but each of the contrastive and non-contrastive datasets had to be subsetted to 284 to make equally sized subsets. The corpus consists of naturalistic longitudinal speech interactions with one child (Catootje) aged 1;10. The corpus was transcribed at the word level. Time-aligned phonetic annotations were created by Dan Swingley (DS) (13). Given the transcriptions, the speech toolkit HTK (72) was used to estimate the boundaries of the phones using the HVITE forced-alignment tool. The output of the forced-alignment tool was manually corrected by DS, a speaker of Dutch. KH time-aligned the word-level transcription to the time-aligned phonetic transcriptions based on the location of the phones in Praat (73).

Code for all analyses is available on Github: http://github.com/khitczenko/contextual-dl. Data are available by request from the researchers who control their distribution.

ACKNOWLEDGMENTS. This work was supported by NSF grants IIS-1421695, BCS-1734245, and NRT #1449815. We thank Thomas Schatz, Mirjam Ernestus, Dan Swingley, and Francisco Torreira for help acquiring corpora, as well as Micha Elsner, Bill Idsardi, Jeff Lidz, Reiko Mazuka, Rochelle Newman, Thomas Schatz, Dan Swingley and the reviewers for helpful feedback and discussion.

1. Werker JF, Gilbert JH, Humphrey K, Tees RC (1981) Developmental aspects of cross-language speech perception. *Child development* pp. 349–355.
2. Werker JF, Tees RC (1984) Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant Behavior and Development* 7(1):49–63.
3. Kuhl PK, Williams KA, Lacerda F, Stevens KN, Lindblom B (1992) Linguistic experience alters phonetic perception in infants by 6 months of age. *Science* 255:606–608.
4. Tsuji S, Cristia A (2014) Perceptual attunement in vowels: A meta-analysis. *Developmental psychobiology* 56(2):179–191.
5. Narayan C, Werker JF, Beddor PS (2010) The interaction between acoustic salience and language experience in developmental speech perception: Evidence from nasal place discrimination. *Developmental Science* 13(3):407–420.
6. Maye J, Werker JF, Gerken L (2002) Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition* 82(3):B101–B111.
7. Maye J, Weiss DJ, Aslin RN (2008) Statistical phonetic learning in infants: Facilitation and feature generalization. *Developmental science* 11(1):122–134.
8. Yoshida KA, Pons F, Maye J, Werker JF (2010) Distributional phonetic learning at 10 months of age. *Infancy* 15(4):420–433.
9. Cristia A (2018) Can infants learn phonology in the lab? a meta-analytic answer. *Cognition* 170:312–327.
10. Vallabha GK, McClelland JL, Pons F, Werker JF, Amano S (2007) Unsupervised learning of vowel categories from infant-directed speech. *Proceedings of the National Academy of Sciences* 104(33):13273–13278.
11. McMurray B, Aslin RN, Toscano JC (2009) Statistical learning of phonetic categories: insights from a computational approach. *Developmental science* 12(3):369–378.
12. Bion RA, Miyazawa K, Kikuchi H, Mazuka R (2013) Learning phonemic vowel length from naturalistic recordings of Japanese infant-directed speech. *PLOS ONE* 8(2):e51594.
13. Swingley D (2019) Learning phonology from surface distributions, considering Dutch and English vowel duration. *Language Learning and Development* 15(3):199–216.
14. Vance TJ (1987) *An introduction to Japanese phonology*. (SUNY Press).
15. Narayan C (2008) The acoustic–perceptual salience of nasal place contrasts. *Journal of Phonetics* 36(1):191–217.
16. Swingley D, Alarcon C (2018) Lexical learning may contribute to phonetic learning in infants: A corpus analysis of maternal spanish. *Cognitive science*.
17. Adriaans F, Swingley D (2012) Distributional learning of vowel categories is supported by prosody in infant-directed speech in *Proceedings of the Annual Meeting of the Cognitive Science Society*. Vol. 34.
18. Dillon B, Dunbar E, Idsardi W (2013) A single-stage approach to learning phonological categories: Insights from inuktitut. *Cognitive Science* 37(2):344–377.
19. Feldman NH, Griffiths TL, Goldwater S, Morgan JL (2013) A role for the developing lexicon in phonetic category acquisition. *Psychological Review* 120(4):751.

20. Feldman NH, Myers EB, White KS, Griffiths TL, Morgan JL (2013) Word-level information influences phonetic learning in adults and infants. *Cognition* 127(3):427–438.
21. Swingley D (2009) Contributions of infant word learning to language development. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 364(1536):3617–3632.
22. Thiessen ED (2007) The effect of distributional information on children's use of phonemic contrasts. *Journal of Memory and Language* 56(1):16–34.
23. Yeung HH, Chen LM, Werker JF (2014) Referential labeling can facilitate phonetic learning in infancy. *Child Development* 85(3):1036–1049.
24. Schatz T, Feldman NH, Goldwater S, Cao XN, Dupoux E (2021) Early phonetic learning without phonetic categories: Insights from large-scale simulations on realistic input. *Proceedings of the National Academy of Sciences* 118(7).
25. Sato Y, Sogabe Y, Mazuka R (2010) Discrimination of phonemic vowel length by Japanese infants. *Developmental Psychology* 46(1):106.
26. Christophe A, Dupoux E, Bertoncini J, Mehler J (1994) Do infants perceive word boundaries? An empirical study of the bootstrapping of lexical acquisition. *The Journal of the Acoustical Society of America* 95(3):1570–1580.
27. Christophe A, Mehler J, Sebastián-Gallés N (2001) Perception of prosodic boundary correlates by newborn infants. *Infancy* 2(3):385–394.
28. Jusczyk PW, Aslin RN (1995) Infants' detection of the sound patterns of words in fluent speech. *Cognitive Psychology* 29(1):1–23.
29. Jusczyk PW, Houston DM, Newsome M (1999) The beginnings of word segmentation in English-learning infants. *Cognitive Psychology* 39(3-4):159–207.
30. Greenberg JH, Ferguson CA, Moravcsik EA (1978) *Universals of human language: phonology*. (Stanford University Press) Vol. 2.
31. Moreton E (2002) Structural constraints in the perception of english stop-sonorant clusters. *Cognition* 84(1):55–71.
32. Gómez DM, et al. (2014) Language universals at birth. *Proceedings of the National Academy of Sciences* 111(16):5837–5841.
33. Hitczenko K, Mazuka R, Elsner M, Feldman NH (2020) When context is and isn't helpful: A corpus study of naturalistic speech. *Psychonomic Bulletin & Review* pp. 1–37.
34. Bergmann C, Cristia A (2016) Development of infants' segmentation of words from native speech: A meta-analytic approach. *Developmental science* 19(6):901–917.
35. Feldman NH, Goldwater S, Dupoux E, Schatz T (2021) Do infants really learn phonetic categories? *Open Mind* pp. 1–19.
36. Rubner Y, Tomasi C, Guibas LJ (1998) A metric for distributions with applications to image databases in *Sixth International Conference on Computer Vision*. (IEEE), pp. 59–66.
37. Babineau M, Shi R (2011) Processing of french liaisons in toddlers. *BUCLD 35 Proceedings. Cascadilla Press, Somerville, MA* pp. 25–37.
38. Mattys SL, Jusczyk PW (2001) Phonotactic cues for segmentation of fluent speech by infants. *Cognition* 78(2):91–121.
39. Seidl A, Johnson EK (2006) Infant word segmentation revisited: Edge alignment facilitates target extraction. *Developmental science* 9(6):565–573.
40. Ernestus MTC (2000) *Voice assimilation and segment reduction in casual Dutch: A corpus-based study of the phonology-phonetics interface*. (Holland Institute of Generative Linguistics, Utrecht).
41. Antetomaso S, et al. (2017) Modeling phonetic category learning from natural acoustic data in *BUCLD 41: Proceedings of the 41st Annual Boston University Conference on Language Development*.
42. Lisker L (1986) "Voicing" in English: A catalogue of acoustic features signaling /b/ versus /p/ in trochees. *Language and Speech* 29(1):3–11.
43. Narayan C (2013) Developmental perspectives on phonological typology and sound change. *Origins of sound change: Approaches to phonologization* pp. 128–146.
44. Narayan C, Peters A, Woldenga-Racine V (2017) Fragile phonetic contrasts in longitudinal infant-directed speech: Implications for infant speech perception in *BUCLD 42: Proceedings of the 41st Annual Boston University Conference on Language Development*.
45. Pimentel T, Roark B, Cotterell R (2020) Phonotactic complexity and its trade-offs. *Transactions of the Association for Computational Linguistics* 8:1–18.
46. Roark CL, Plaut DC, Holt LL (2022) A neural network model of the effect of prior experience with regularities on subsequent category learning. *Cognition* 222:104997.
47. Tranel B (1987) *The Sounds of French: An Introduction*. (Cambridge University Press).
48. Clayards M, Tanenhaus MK, Aslin RN, Jacobs RA (2008) Perception of speech reflects optimal use of probabilistic speech cues. *Cognition* 108(3):804–809.
49. Maye J, Aslin RN, Tanenhaus MK (2008) The weckud wetch of the wast: Lexical adaptation to a novel accent. *Cognitive Science* 32(3):543–562.
50. Cristia A, et al. (2012) Linguistic processing of accented speech across the lifespan. *Frontiers in Psychology* 3:479.
51. van Heugten M, Johnson EK (2012) Infants exposed to fluent natural speech succeed at cross-gender word recognition. *Journal of Speech, Language, and Hearing Research*.
52. White KS, Aslin RN (2011) Adaptation to novel accents by toddlers. *Developmental Science* 14(2):372–384.
53. Weatherhead D, White KS (2016) He says potato, she says potahto: Young infants track talker-specific accents. *Language Learning and Development* 12(1):92–103.
54. Kuhl PK (1979) Speech perception in early infancy: Perceptual constancy for spectrally dissimilar vowel categories. *The Journal of the Acoustical Society of America* 66(6):1668–1679.
55. Kuhl PK (1983) Perception of auditory equivalence classes for speech in early infancy. *Infant behavior and development* 6(2-3):263–285.
56. Jusczyk PW, Luce PA, Charles-Luce J (1994) Infants' sensitivity to phonotactic patterns in the native language. *Journal of Memory and Language* 33(5):630.
57. Jusczyk PW, Hohne EA, Bauman A (1999) Infants' sensitivity to allophonic cues for word segmentation. *Perception & Psychophysics* 61(8):1465–1476.
58. White KS, Peperkamp S, Kirk C, Morgan JL (2008) Rapid acquisition of phonological alternations by infants. *Cognition* 107(1):238–265.
59. Moreton E, Amano S (1999) Phonotactics in the perception of Japanese vowel length: Evidence for long-distance dependencies in *EUROSPEECH*.
60. Idemaru K, Holt LL (2011) Word recognition reflects dimension-based statistical learning. *Journal of Experimental Psychology: Human Perception and Performance* 37(6):1939.
61. Holt LL, Lotto AJ (2006) Cue weighting in auditory categorization: Implications for first and second language acquisition. *The Journal of the Acoustical Society of America* 119(5):3059–3071.
62. Toscano JC, McMurray B (2010) Cue integration with categories: Weighting acoustic cues in speech using unsupervised learning and distributional statistics. *Cognitive science* 34(3):434–464.
63. Liu L, Jaeger TF (2018) Inferring causes during speech perception. *Cognition* 174:55–70.
64. Bergelson E, Swingley D (2012) At 6–9 months, human infants know the meanings of many common nouns. *Proceedings of the National Academy of Sciences* 109(9):3253–3258.
65. Rubner Y, Tomasi C, Guibas LJ (1998) A metric for distributions with applications to image databases in *Sixth International Conference on Computer Vision (IEEE Cat. No. 98CH36271)*. (IEEE), pp. 59–66.
66. Villani C (2008) *Optimal transport: Old and new*. (Springer Science & Business Media) Vol. 338.
67. Maekawa K (2003) Corpus of Spontaneous Japanese: Its Design and Evaluation in *ISCA & IEEE Workshop on Spontaneous Speech Processing and Recognition*.
68. Torreira F, Adda-Decker M, Ernestus M (2010) The Nijmegen Corpus of Casual French. *Speech Communication* 52(3):201–212.
69. Schuppler B, Ernestus M, Scharenborg O, Boves L (2011) Acoustic reduction in conversational Dutch: A quantitative analysis based on automatically generated segmental transcriptions. *Journal of Phonetics* 39(1):96–109.
70. Fikkert P (1994) *On the acquisition of prosodic structure*. (Ph.D. Thesis).
71. Levelt C (1994) *On the acquisition of place*. (Ph.D. Thesis).
72. Young S, et al. (2002) The HTK Book. *Cambridge University Engineering Department* 3(175):12.
73. Boersma P (2001) Praat: A system for doing phonetics by computer. *Glot International* 5(9/10):341–345.

Supplementary Materials

Figure 3 in the main text reports our bootstrapping analysis results, treating the maximum Earthmover's distance as the metric of interest (i.e., showing how maximum Earthmover's distance varies across 50 runs). In Figure S1, we provide analogous plots treating the mean Earthmover's distance as the primary metric instead. Results are similar: across input sizes the mean contrastive Earthmover's distance is greater than the mean non-contrastive Earthmover's distance, suggesting that observed differences between contrastive vs. non-contrastive dimensions are meaningful, but that input size does not matter.

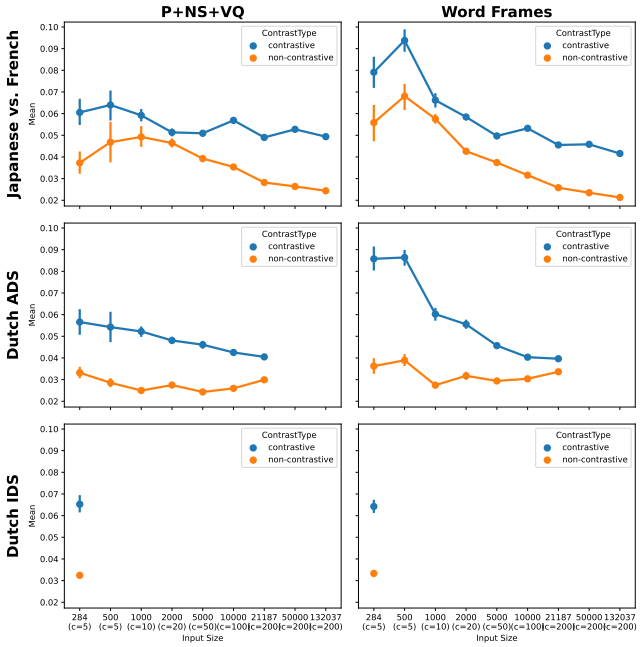

Fig. S1. Results from bootstrapped input size analyses. Across input sizes, the mean contrastive Earthmover's distance is greater than the mean non-contrastive Earthmover's distance. "C" refers to the number of contexts included in the analysis. The maximum input size for which data is shown depends on the corpus size: 284 for Dutch IDS, 21187 for Dutch ADS, and 132037 for Japanese vs. French.

The analyses in the main text use Earthmover's distance to compare distribution shapes across contexts. Figure S2 shows results using an alternative metric, namely KL divergence (or Kullback-Leibler divergence). KL divergence is a measure of how different two probability distributions are. As we do not have access to the closed-form probability distributions within each context (only samples), we estimated KL divergence as follows. We divided the tokens in each context into 10 evenly-sized bins, and smoothed the counts in each bin (adding a count of 1e-5 to each bin). We then used these counts to arrive at a probability distribution over the 10 bins, which were used to calculated KL-divergence. As KL-divergence is not a symmetric measure, for each pair of contexts, we summed the KL-divergence of Context 1 against Context 2 and the KL-divergence of Context 2 against Context 1 to arrive at a final measure, that is plotted in Figure S2. We find that the expected qualitative pattern still emerges across all three test cases even using a different metric. It is, however, interesting that the difference in magnitude across the test cases changes: here, Dutch ADS and IDS show the greatest KL divergences, while Japanese/French showed the greatest Earthmover's distances. This suggests that this effect is largely insensitive to the particular metric used, but which metric is used could make a difference in specific predictions the account makes.

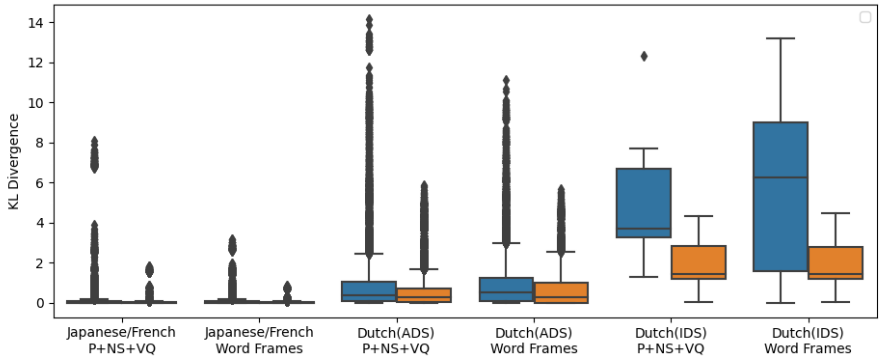

Fig. S2. Distribution of KL divergences by test case. The same qualitative pattern holds, though the difference in magnitude changes across the three test cases. P+NS+VQ = prosodic position + neighboring sounds + vowel quality.