Low-dimensional representation of infant and adult vocalization acoustics

Silvia Pagliarini[1], Sara Schneider[2], Christopher T. Kello[2], Anne S. Warlaumont[1]

[1]Department of Communication, University of California, Los Angeles, USA [2]Cognitive and Information Sciences, University of California, Merced, USA

`pagliarini@ucla.edu, sschneider2@ucmerced.edu, ckello@ucmerced.edu, warlaumont@ucla.edu`

Abstract

During the first years of life, infant vocalizations change considerably, as infants develop the vocalization skills that enable them to produce speech sounds. Characterizations based on specific acoustic features, protophone categories, or phonetic transcription are able to provide a representation of the sounds infants make at different ages and in different contexts but do not fully describe how sounds are perceived by listeners, can be inefficient to obtain at large scales, and are difficult to visualize in two dimensions without additional statistical processing. Machine-learning-based approaches provide the opportunity to complement these characterizations with purely data-driven representations of infant sounds. Here, we use spectral features extraction and unsupervised machine learning, specifically Uniform Manifold Approximation (UMAP), to obtain a novel 2-dimensional spatial representation of infant and caregiver vocalizations extracted from day-long home recordings. UMAP yields a continuous and well-distributed space conducive to certain analyses of infant vocal development. For instance, we found that the dispersion of infant vocalization acoustics within the 2-D space over a day increased from 3 to 9 months, and then decreased from 9 to 18 months. The method also permits analysis of similarity between infant and adult vocalizations, which also shows changes with infant age.

Index Terms: infant vocalization, visualization, UMAP, MFCCs, unsupervised learning

1. Introduction

Human infant vocal repertoires undergo dramatic, multifaceted, and linguistically significant changes over the first two years of postnatal life. At birth, infants produce simple precursors to speech sounds, predominantly quasivowels and other short, quiet sounds [1, 2]. Over subsequent months, the sounds they produce diversify and become more complex, to include a range of loudness, vocal qualities, pitches and pitch contours, durations, and primitive consonant productions. By about 7 months, they begin to produce canonical babbles (which contain speech-like syllables combining consonants and vowels); over the following months and years, additional changes in infant vocal repertoires and vocal motor control continue [1, 3, 4]. This learning provides an essential foundation for linguistic communication, and is believed to be supported both by intrinsically-motivated play and by contingent and often imitative interactions with adult caregivers [5, 6, 7, 8, 9, 10, 11]. Infant language acquisition in general is strongly influenced by auditory inputs from caregivers and other environmental sound sources (e.g., [12, 7, 13, 14]). An infant's auditory environment includes voices from caregivers and other individuals of differing ages and genders plus a wide range of other sound sources, such as animals, physical objects, and electronic devices. The presence of these many different types of sound sources, the high variety of sounds produced by each, and the wide range of sound types produced by infants themselves present a challenge for characterizing infant vocal productions and auditory inputs.

Most characterizations of infant vocalizations are based on acoustic analyses of specific acoustic features (such as fundamental frequency and formant frequencies (e.g., [15, 16]), human listener categorizations into protophone categories (e.g. quasivowel, canonical babble, etc.) [1, 3, 2], or phonetic transcription [4]. A limitation of these methods is that they don't provide a fully comprehensive characterization of a sound. For example, two sounds may be similar in their pitch characteristics but differ in terms of the phonetic features they contain. Or the sounds may belong to the same protophone category but still acoustically differ considerably and be perceived differently by listeners. Data-driven analyses of raw acoustic information could provide a complementary approach.

Sainburg et al. [17] showed the power of using Uniform Manifold Aproximation (UMAP) [18], a machine-learning method for reducing a high-dimensional acoustic dataset into a two-dimensional space, to represent human speech sounds as well as songbird syllables. For human speech, different phonetic categories tended to project to distinct regions of latent space. For birdsong, the repertoires of different songbird species could be represented in the same overall space, enabling comparisons across species.

Here we apply spectral features extraction and UMAP to transform raw acoustic data from daylong home audio recordings of 3- to 18-month-old infants. This generates a two-dimensional spatial representation of infant and caregiver vocalizations. We then use this space to quantify, for each day-long recording (1) the similarity between infant and caregiver vocalizations and (2) the diversity (i.e., the amount of variation) across infant vocal productions on the day of recording. Finally, we assess how these measures vary with infant age.

2. Methods[1]

2.1. Dataset and preprocessing

Our infant and adult vocalization data comes from long-form (10+ daytime hours), child-centered home recordings. 52 infants were recorded longitudinally at 3, 6, 9 and 18 months using the LENA system [19]. A few infants did not complete all four recordings; we focused on the subset of 42 infants who did have complete data. More details about the data collection methods can be found in [16]. A subset of the recordings are available in the Warlaumont corpus [20] within HomeBank [21]. We extracted short audio clips of duration $0.6s$ to $12s$ for sections of the recording labeled by LENA's built-in algorithm [22] as one of the following four types: class *CHNNSP*

---

[1]More details are available at https://github.com/spagliarini/Infant-vocalization-space-Interspeech2022

contains cry/laugh (non-speech-related) vocalization produced by the infant wearing the recorder, class *CHNSP* (child speech-related) contains speech or protophone vocalizations produced by the infant wearing the recorder, class *FAN* (female adult near) contains adult female vocalization, and class *MAN* (male adult near) contains adult male vocalizations. Although LENA also tags other sound sources (noise, television, other adults, other children, overlap), for the present study we included only the infant and their adult caregivers. Table 1 gives the number of instances of each class across the 42 participants.

Table 1: *Total utterances for each LENA sound category*

Age(months)	CHNNSP	CHNSP	FAN	MAN
3	32631	39110	84999	35170
6	29928	43932	70491	29674
9	32007	45246	69802	29629
18	32852	60024	69351	31949

2.2. Low-dimensional vocalization space

2.2.1. Spectral feature extraction

We used openSMILE [23] to extract 13 Mel-frequency cepstral coefficients (MFCCs) from 26 Mel-frequency bands (with $25ms$ frame size and $10ms$ frame rate). Each MFCC's first and second derivatives (i.e., velocity and acceleration) were also computed, generating an additional 26 features. This gave us, for each child or adult vocalization clip, a $39 \times T$ matrix where T was the number of timebins in the clip. We then summed across each matrix's timebins, yielding a 39-dimensional vector for each sound clip.

2.2.2. Dimensionality reduction

From collections of these feature vectors we obtained a 2D representation of the vocalization space using Uniform Manifold Approximation and Projection for Dimension Reduction (UMAP) [18]. Similar to t-distributed stochastic neighbor embedding (t-SNE [24]), UMAP is a dimension reduction technique where the result axes do not represent single, meaningful discriminant features. In particular, UMAP (1) can be used to perform non-linear dimensionality reduction, (2) has a higher computational efficiency than t-SNE, and (3) can be non-local (i.e., it can take into account distances between points that are located far apart from each other). The tuning of UMAP hyperparameters enables more or less local representation of the data. A smaller neighborhood size n_{neigh} means a more locally-focused representation; larger values will push UMAP to look at larger neighborhoods of each point when estimating the manifold structure of the data. The minimum distance min_{dist} parameter determines how far points are allowed to be from each other in the low dimensional representation; this also influences the balance of emphasis on local detail vs. broad structure. We used the UMAP Python package[2] [18].We chose $n_{neigh} = 15$ and $min_{dist} = 0.1$.

2.2.3. Statistical analyses

For each daylong recording, we measured (1) the distance from the centroid of the infant vocalizations to the centroid of the FAN vocalizations (we focused our statistical analyses in this initial study on CHNSP-to-FAN distance because some recordings had extremely few MAN vocalizations), (2) the average

---

[2]https://github.com/lmcinnes/umap

distance from the infant vocalizations to the infant's centroid and (3) the Shannon entropy of *CHNSP* vocalizations.

We then performed (in R, using lme4 [25] and lmerTest [26]) linear mixed effects regressions with infant ID as a random effect and the number of infant vocalizations in the daylong recording as a covariate to test the hypothesis that infant age influences (1) the distance between the centroid of the infant vocalizations and the centroid of the adult female vocalizations, (2) the average distance from individual infant vocalization locations to the infant vocalization centroid, and (3) the Shannon entropy of the infant vocalization locations. We included both linear and quadratic terms for age in each regression.

3. Results

3.1. UMAP-based low-dimensional vocalizations space

Vocalizations clustered together according to the sound type classes, but with a large amount of overlap. This can be observed in the vocalization space obtained from the collection of infant speech-related (*CHNSP*), adult male (*MAN*) and adult female vocalizations (*FAN*) (left panel of Fig. 1a; our analyses below focus on this space) and in a vocalizations space obtained from the collection of both types of infant sounds (*CHNSP* and *CHNNSP*) (right panel of Fig. 1a). Similarly, in a space constructed from a single daylong recording, vocalizations cluster together when comparing the *CHNSP* and *FAN* vocalizations (Fig. 1b). Figure 1b also illustrates the centroids of each class.

3.2. Similarity between infant and adult female vocalizations

We assessed the similarity between an infant's vocalizations on a given day and the corresponding adult female vocalizations that infant was exposed to by calculating the distance from the centroid of the infant vocalizations to the centroid of the adult female vocalizations. There was a statistically significant positive quadratic effect of age on similarity between infant and adult female ($p < .001$), with a decrease in centroid distance from 3 to 9 months followed by an increase at 18 months (Figure 2a). This suggests that infant and adult female vocalizations become increasingly similar from 3 to 9 months of age, and then diverge again by the time infants get to be 18 months old.

3.3. Infant vocalization variability

For each infant recording, we assessed infant vocalization variability by calculating the average distance of the infant's *CHNSP* vocalization locations from the class centroid in the UMAP-generated space. We then asked if this variability changed across infant age (Fig. 2b). Statistical analyses detected both a linear ($p < .001$) and quadratic ($p < .001$) relationship between age and infant vocalization variability. Both effects were negative, indicating that the diversity of sounds infants produced during a day rose to a peak at 9 months and then decreased from then on, with the lowest variability observed at 18 months. Shannon entropy analysis (Fig. 2c) of the infant non cry/laugh vocalizations replicated the inverted-U-shape observed in the centroid analysis (Fig. 2b), with a statistically significant negative quadratic term ($p < .001$).

3.4. tSNE-based vocalizations space

As mentioned in Section 2.2, tSNE is another method that can be used to perform non-linear dimensionality reduction. The space obtained using tSNE in place of UMAP for the *CHNSP*,

(a) Overview of the vocalization space.
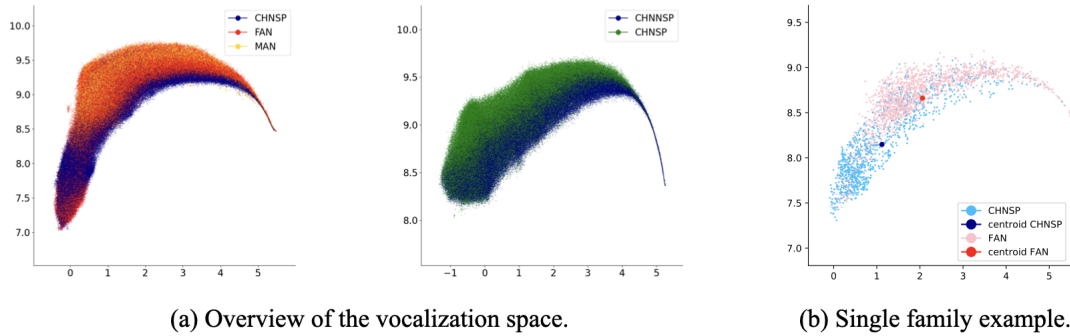
(b) Single family example.

Figure 1: **UMAP-based vocalization space.** *Representation of all vocalization clips in a space generated by applying UMAP to MFCC features summed across the clip. Each point represents a vocal utterance as 39-dimensional vector (13 MFCCs, 13 MFCC velocities, and 13 MFCC accelerations). (a) UMAP space obtained from non cry/laugh infant, adult male and female vocalization clips. Each color represents a class: CHNSP are infant speech and pre-speech non cry/laugh vocalizations (blue dots), MAN and FAN are respectively female (red dots) and male (yellow dots) adult vocalizations, CHNNSP are infant cry/laugh or vegetative. The left panel shows the main space analysed in this paper, constructed based on CHNSP, FAN, and MAN clips. The right panel shows a space constructed based on CHNNSP and CHNSP. (b) Example of a single 18-months old baby recording (extracted from the whole UMAP space). The light blue dots represent the infant vocalizations (CHNSP), the pink dots represent the adult female vocalizations (FAN). The centroid for class CHNSP is shown by the dark blue dot and the centroid for class FAN is the red dot.*
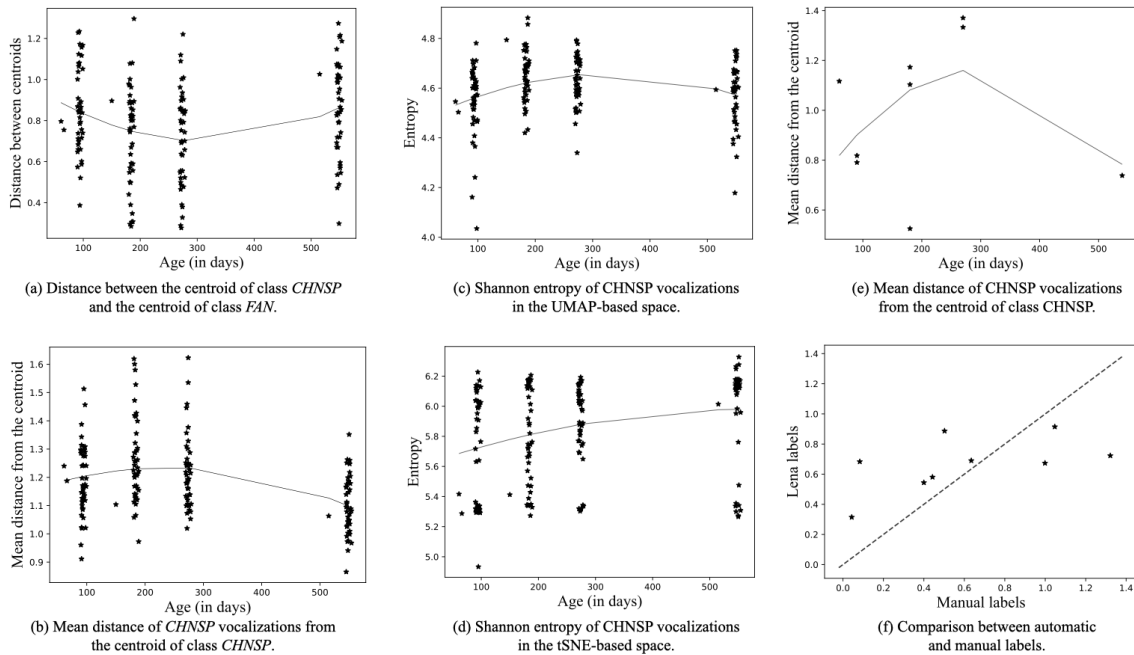

(a) Distance between the centroid of class *CHNSP* and the centroid of class *FAN*.

(c) Shannon entropy of CHNSP vocalizations in the UMAP-based space.

(e) Mean distance of CHNSP vocalizations from the centroid of class CHNSP.

(b) Mean distance of *CHNSP* vocalizations from the centroid of class *CHNSP*.

(d) Shannon entropy of CHNSP vocalizations in the tSNE-based space.

(f) Comparison between automatic and manual labels.

Figure 2: **Age-related changes in UMAP-based measures, comparisons to tSNE, and results for human-validated infant utterances.** *Each star represents data from a particular infant's recording at a given infant age. (a) Similarity between infant speech-related (CHNSP) and adult female (FAN) vocalizations plotted as a function of age. (b) Variability of CHNSP vocalizations in the UMAP-based vocalizations space, measured as distances to the CHNSP centroid, as a function of age. (c) Shannon entropy of CHNSP vocalizations in the UMAP-based vocalizations space versus age. (d) Shannon entropy of CHNSP vocalizations in a tSNE-based vocalizations space versus age. (e) Variability of CHNSP vocalizations in the UMAP-based vocalizations space for a subset of the recordings with child vocalizations verified as such by human listeners. (f) Scatter plot showing the mean distance from centroid measure (for CHNSP) when child vocalizations were identified automatically by LENA (y-axis) versus when those CHNSP clips were manually verified by human listeners as being infant sounds and not mis-classifications by the labeling algorithm (x-axis). The black lines in panels (a-e) represent data fitting obtained using polynomial regression on the represented data. The dashed line in (f) represents hypothetical 1:1 correspondence between the two labeling methods.*

FAN, and *MAN* classes is shown in Figure 3. This figure can be compared with the representation of the same data obtained with UMAP (statistically significant quadratic term: $p < .001$), shown in the left panel of Figure 1a. Interestingly, Shannon entropy analysis (Fig. 2d) of the *CHNSP* vocalizations in the tSNE-generated space ($p = 0.03$) do not appear to have a clear U-shape with respect to age, as was observed in the UMAP case (Figure 2(b-d)). (Given the high discontinuity of the tSNE-based space, we do not compare centroid-based measurements.)
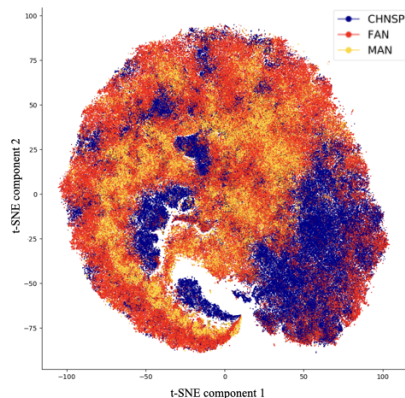

Figure 3: ***tSNE-based vocalizations space.*** *Representation of CHNSP (blue), FAN (red), and MAN (yellow) vocalization clips in a space generated by applying tSNE to MFCC features summed across each clip. Each point represents a vocal utterance as 39-dimensional vector (13 MFCCs, 13 velocities, and 13 acceleration). Vocalizations space obtained from non cry/laugh infant, adult male and female vocalization clips. Each color represents a class: CHNSP are infant speech and pre-speech non cry/laugh vocalizations (blue dots), while MAN and FAN are respectively female (red dots) and male (yellow dots) adult vocalizations.*

3.5. Validation of the dataset

For a subset of our data (9 recordings in total), we worked with a team of ten human listeners to re-label the automatically identified *CHNSP* clips. Listeners assigned a prominence value $P \in \{1, 2, 3, 4, 5\}$ to indicate if they believed the infant was indeed vocalizing in the clip without any other sounds audible ($P = 1$), if the infant was present but so were other sounds and if so, what was the infant sound prominence relative to those other sounds background noise and how much ($1 < P < 5$), or if the infant wearing the recorder did not actually vocalize at all during the clip ($P = 5$). For each clip, we computed the modal value across listeners and defined a threshold for inclusion in the strictest way possible: we validated a clip as infant prominent if its modal value was equal to 1. Otherwise, we considered the clip to be noise and excluded it from the study. We then computed the mean distance of infant vocalizations (*CHNSP*) from the class centroid (Figure 2e), and observed a comparable trend with what we obtained from the whole dataset (Figure 2b). We also observed that the range of this variability measure was larger for the human-validated labels than for the fully automated LENA labels (Fig. 2f). The correlation between the two sets of labels is equal to 0.47.

4. Discussion

We proposed a method based on spectral features extraction (MFCC's + their first and second derivatives, summed across timebins for an utterance) and UMAP [18], an unsupervised machine learning dimensionality reduction method, to represent infants and adult vocalizations in a two-dimensional space. For our dataset, this resulted in a relatively smooth distribution of the data across space (in contrast with tSNE). Statistical analyses revealed significant non-monotonic patterns of change in age. Specifically, (1) the range of sound patterns infants produced (as quantified by distance-to-class centroid and Shannon entropy) increased from 3 to 9 months and then decreased by the time infants were 18-months-old. We also found that similarity between infant and adult vocalizations increased from 3 to 9 months then decreased from 9 to 18 months.

We hypothesize that the increase from 3 to 9 months might relate to the so-called "expansion stage" of infant protophone development as well as to the onset of canonical babbling [1]. This might be tested by establishing if there is a relationship between locations in the UMAP-constructed space and protophone categorizations of the infant sounds. We also note that previous work [15] observed informally (without tests for statistical significance) that as infant age increased from 3 to 9 months, so did the range of infant vocalization durations and the formant frequency ranges. Our findings are consistent with those trends for increasing acoustic variation from 3 to 9 months, but also show that by 18 months, the amount of variation shows a decrease from its 9-month level. It would be informative (1) to compare the relationship between adult female and infant non cry/laugh vocalizations with the relationship between adult female and infant cry/laugh vocalizations, and (2) to connect the UMAP representations to more transparent and commonly used acoustic features, such as duration, formant frequencies, and pitch features.

A limitation is that LENA-generated sound source labels are sometimes incorrect, and even when they are correct, there are sometimes significant other sounds present [27]. To overcome this problem, we are working to obtain a cleaner dataset based on human listener judgments. Preliminary results (Section 3.5) suggest that the overall qualitative patterns may not change drastically, but additional work and data are needed.

In the future, variations on this approach could be explored, including alternative pre-processing and dimensionality reduction methods. For instance, it could be useful to pre-process the audio using end-to-end neural network approaches, perhaps with a training goal of optimizing infant age predictions. This may allow UMAP to be performed on more sophisticated and more functionally-and-practically-relevant acoustic features. This general approach may be useful for characterizing the interactions between infants and caregivers, for studying individual and clinical differences in vocal productions, and for providing additional means for comparison of human infant behavior to that of computational models of vocal learning [28].

6. References

[1] D. K. Oller, *The emergence of the speech capacity.* Lawrence Erlbaum Associates, 2000.

[2] E. H. Buder, A. S. Warlaumont, D. K. Oller, B. Peter, and A. MacLeod, "An acoustic phonetic catalog of prespeech vocalizations from a developmental perspective," *Comprehensive perspectives on child speech development and disorders: Pathways from linguistic theory to clinical practice*, vol. 4, pp. 103–34, 2013.

[3] S. Nathani, D. J. Ertmer, and R. E. Stark, "Assessing vocal development in infants and toddlers," *Clinical Linguistics and Phonetics*, vol. 20, no. 5, pp. 351–369, 2006.

[4] M. M. Vihman, *Phonological development: The origins of language in the child.* Blackwell Publishing, 1996.

[5] K. Bloom, "Quality of adult vocalizations affects the quality of infant vocalizations," *Journal of Child Language*, vol. 15, no. 3, pp. 469–480, 1988.

[6] M. Papoušek and H. Papoušek, "Forms and functions of vocal matching in interactions between mothers and their precanonical infants," *First Language*, vol. 9, no. 6, pp. 137–158, 1989.

[7] P. K. Kuhl, "Early language acquisition: cracking the speech code," *Nature Reviews Neuroscience*, vol. 5, no. 11, pp. 831–843, 2004.

[8] M. H. Goldstein and J. A. Schwade, "Social feedback to infants' babbling facilitates rapid phonological learning," *Psychological Science*, vol. 19, no. 5, pp. 515–523, 2008.

[9] C. Moulin-Frier, S. M. Nguyen, and P.-Y. Oudeyer, "Self-organization of early vocal development in infants and machines: the role of intrinsic motivation," *Frontiers in Psychology*, vol. 4, p. 1006, 2014.

[10] H. L. Long, D. D. Bowman, H. Yoo, M. M. Burkhardt-Reed, E. R. Bene, and D. K. Oller, "Social and endogenous infant vocalizations," *PLOS ONE*, vol. 15, no. 8, p. e0224956, 2020.

[11] A. S. Warlaumont, *Infant vocal learning and speech production.* Cambridge University Press, 2020, pp. 602–631.

[12] M. Papoušek, "Melodies in caregivers' speech: A species-specific guidance towards language," *Early Development and Parenting*, vol. 3, no. 1, pp. 5–17, 1994.

[13] J. Gervain and J. F. Werker, "How infant speech perception contributes to language acquisition," *Language and Linguistics Compass*, vol. 2, no. 6, pp. 1149–1170, 2008.

[14] R. Fusaroli, E. Weed, D. Fein, and L. Naigles, "Hearing me hearing you: Reciprocal effects between child and parent language in autism and typical development," *Cognition*, vol. 183, pp. 1–18, 2019.

[15] R. D. Kent and A. D. Murray, "Acoustic features of infant vocalic utterances at 3, 6, and 9 months," *The Journal of the Acoustical Society of America*, vol. 72, no. 2, pp. 353–365, 1982.

[16] V. P. S. Ritwika, G. M. Pretzer, S. Mendoza, C. Shedd, C. T. Kello, A. Gopinathan, and A. S. Warlaumont, "Exploratory dynamics of vocal foraging during infant-caregiver communication," *Scientific Reports*, vol. 10, no. 1, pp. 1–14, 2020.

[17] T. Sainburg, M. Thielk, and T. Q. Gentner, "Finding, visualizing, and quantifying latent structure across diverse animal vocal repertoires," *PLOS Computational Biology*, vol. 16, no. 10, p. e1008228, 2020.

[18] L. McInnes, J. Healy, and J. Melville, "UMAP: Uniform manifold approximation and projection for dimension reduction," *arXiv preprint arXiv:1802.03426*, 2018.

[19] J. Gilkerson and J. A. Richards, "The lena natural language study," LENA Foundation, Tech. Rep., 2008.

[20] A. S. Warlaumont, G. M. Pretzer, S. Mendoza, and E. A. Walle, "Warlaumont HomeBank corpus," 2016. [Online]. Available: https://doi.org/10.21415/t54s3c

[21] M. VanDam, A. S. Warlaumont, E. Bergelson, A. Cristia, M. Soderstrom, P. De Palma, and B. MacWhinney, "Homebank: An online repository of daylong child-centered audio recordings," *Seminars in Speech and Language*, vol. 37, no. 2, pp. 128–142, 2016.

[22] D. Xu, U. Yapanel, S. Gray, J. Gilkerson, J. Richards, and J. Hansen, "Signal processing for young child speech language development," in *First Workshop on Child, Computer and Interaction*, 2008.

[23] F. Eyben, M. Wöllmer, and B. Schuller, "openSMILE - The Munich Versatile and Fast Open-Source Audio Feature Extractor," in *Proceedings of the 18th ACM international conference on Multimedia*, 2010, pp. 1459–1462.

[24] L. Van der Maaten and G. Hinton, "Visualizing data using t-SNE," *Journal of machine learning research*, vol. 9, no. 11, 2008.

[25] D. Bates, M. Mächler, B. Bolker, and S. Walker, "Fitting linear mixed-effects models using lme4," *Journal of Statistical Software*, vol. 67, no. 1, pp. 1–48, 2015.

[26] A. Kuznetsova, P. B. Brockhoff, and R. H. B. Christensen, "lmerTest package: Tests in linear mixed effects models," *Journal of Statistical Software*, vol. 82, no. 13, pp. 1–26, 2017.

[27] O. Räsänen, S. Seshadri, M. Lavechin, A. Cristia, and M. Casillas, "ALICE: An open-source tool for automatic measurement of phoneme, syllable, and word counts from child-centered daylong recordings," *Behavior Research Methods*, vol. 53, no. 2, pp. 818–835, 2021.

[28] S. Pagliarini, A. Leblois, and X. Hinaut, "Vocal imitation in sensorimotor learning models: a comparative review," *IEEE Transactions on Cognitive and Developmental Systems*, vol. 13, no. 2, pp. 326–342, 2020.