Investigating the relationship between age, gender, and vowel categorization within and across languages:

A comparison of Cantonese, English, and Japanese

Benjamin Munson[1] Andrew R. Plummer[2]

[1] Dept. of Speech-Language-Hearing Sciences, University of Minnesota, Minneapolis, MN, USA
[2] Dept. of Comp. Sci. & Eng., The Ohio State University, Columbus, OH, USA

munso005@umn.edu, plummer@ling.ohio-state.edu

March 07, 2014

Structure of the talk

- Introduction:
  - Why study the perception of vowels produced by a growing vocal tract?
  - What is the bigger-picture question to be answered?
  - Why Cantonese?
- Data Collection:
  - VLAM stimuli,
  - The structure of the experiments.
- Modeling and Analyses:
  - Vowel category response surfaces,
  - Age and Gender Summaries.
- Implications for models of phonological acquisition:
  - The acquisition of vowel normalization,
  - Manifold alignment.

THE DYNAMICS OF PHONOLOGICAL ACQUISITION

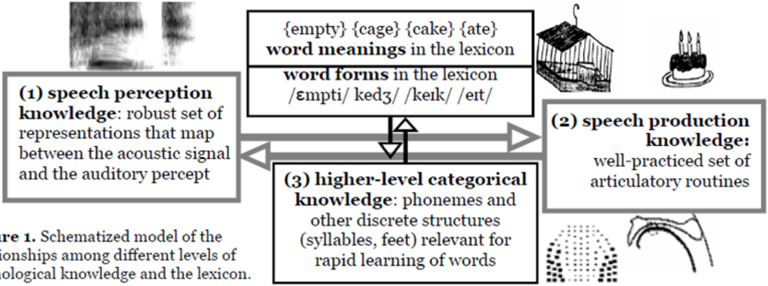

Figure 1. Schematized model of the relationships among different levels of phonological knowledge and the lexicon.

INTRODUCTION

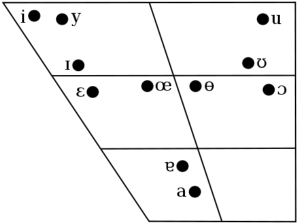

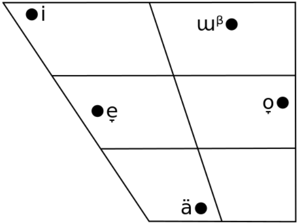

- Languages with a large number of contrasts among vowels, e.g., Cantonese (left) and American English, with eleven monophthongal vowels each, are typologicially rare in relation to languages like Japanese (right), which has only five.
- Juxtaposing languages which exhibit such dramatic differences in vowel inventory size has the potential to illuminate aspects of speech perception that typically fall outside canonical consideration.

INTRODUCTION

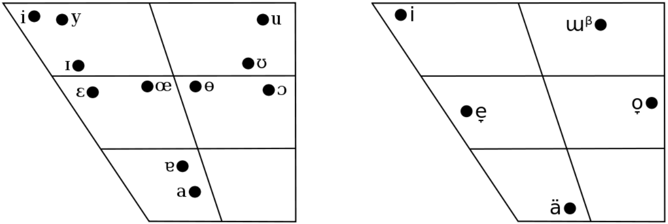

- For example, cross-language comparison of the relative locations and diffuseness of regions in formant space that correspond to vowel categories may reveal cross-language differences in the attention given to finer-grained phonetic detail encoding other categorical information, e.g., age and gender information.
- If so, such comparisons may also reveal within-language differences based on small-group or even individualistic vocal learning experience.

Research Questions

- Is language specificity evident in adults' perceptions/interpretations of young children's vocalizations? How does language specificity in perception change over time? Are there universal tendencies in the perception of children's vocalizations?
- Is language specificity evident in adults' perception of age and gender typicality?
- How do these interact? Do language-specific age- and gender-perception affect how vowels are labeled?

- In this talk, we present an array of data derived from Cantonese, American English, and Japanese listeners' age, gender, and vowel category judgments in response to age-varying sets of synthetic vowels generated from a model of the growing vocal tract.
- The data were collected as part of the Learning to Talk research initiative focused on modeling the development of language-specific speech perception and production:

http://learningtotalk.org/?q=3Dnode/25

Vowel category response modeling

- The main focus of our data analysis is a shape-based statistical modeling methodology which provides the means to parse the formant space for each vocal tract age into listener-specific vowel category regions based on listener-provided category judgments and goodness ratings.
- This listener-specific modeling not only serves as the basis for investigating broader language-specific vowel categorization patterns, but is crucial in modeling the likely caregiver-infant dyad-specific aspects of vowel category acquisition.

Age and Gender Data

- We also examine language-specific differences in the perception of these vowels across vocal tract ages.
- Preliminary analyses focus on a bifurcation in gender rating patterns beginning at age 10, wherein vowel stimuli for this age are rated as either a young male or older female.
- Moreover, the gender rating patterns appear to differ across languages, suggesting a learned, language-specific association between vowel categories and social categories.

Relating vowel quality with age and gender

- We also take steps toward characterizing relationships among the judgments of age and gender and the vowel identifications for the stimulus set where age and gender were the most ambiguous.
- Initial examination of the data suggest that certain relationships between these three variables exhibit greater similarity with respect to the Cantonese and English vowel systems than the Japanese vowel system.
- If so, this suggests that aspects of the vowel systems of Cantonese and English, including their size, may prompt listeners to use different sources of information in disambiguation than does the vowel system of Japanese.

Implications for phonological acquisition

- Finally, we show how shape-based vowel category response modeling can be used in the modeling of phonological acquisition.
- Specifically, we briefly walk through the use of shape-based cognitive models of the acquisition of vowel normalization in capturing the language-specific, and potentially dyad-specific aspects of this acquisition procedure.

Maximal Vowel Spaces and Stimuli

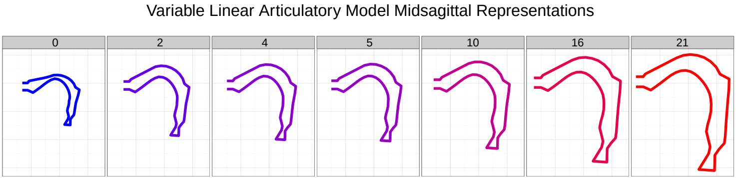

Articulatory Model

- The VLAM (Boë & Maeda, 1998) is a computational model of the articulatory system and its speech production capacities.
- Midsagittal representations are wrought by configuring “articulatory blocks” (Lindblom & Sundberg, 1971; Maeda, 1990, 1991) corresponding to jaw height, tongue body position, tongue dorsum position, tongue apex position, lip protrusion, lip height, and larynx height.
- The VLAM is age-varying and capable of representing vocal tract lengths ranging from those of infants to young adults, calibrated in accordance with age-related “organic variation” (Beck, 1996; Goldstein, 1980).

Maximal Vowel Spaces and Stimuli

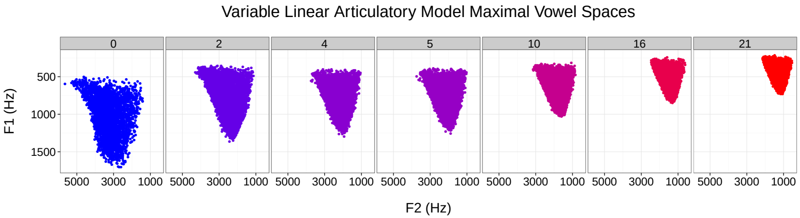

Maximal Vowel Spaces

- Given an age in years, the set of all articulatory configurations of the VLAM at that age that do not result in occlusion of the oral cavity yield a corresponding maximal vowel space (MVS, Boë et al., 1989; Schwartz et al., 2007) for that age.
- We take each MVS to be characterized by a set of formant patterns. Each formant pattern within an MVS is identified with a formant vector whose components are the first three formant frequencies of that formant pattern.
- We fix the following age index $\text{AGES} = \{0.5, 2, 4, 5, 10, 16, 21\}$ for indexing the MVSs discussed below. For each $a \in \text{AGES}$, let $\text{MVS}(a)$ be a dense sampling of the MVS for age a.

Maximal Vowel Spaces and Stimuli

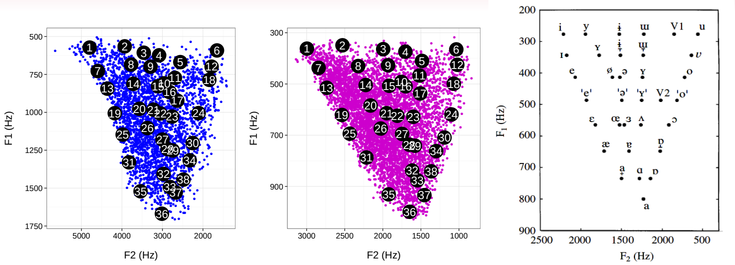

Figure : Prototype sets P(0.5) (left), P(10) (center), and those from Schwartz et al. (1997a).

Stimuli

- The stimuli used in the perceptual experiments (Munson et al., 2010) were *vowel prototypes* (Vallée et al., 1995), or simply *prototypes*, selected from the MVSs for each $a \in$ AGES.
- The selection process (Ménard & Boë, 2000; Ménard et al., 2002) is meant to yield a set of cross-linguistically relevant prototypes in accordance with the dispersion-focalization theory (Schwartz et al., 1997a,b).
- For each $a \in$ AGES, a set of 38 prototypes, denoted by P(a), were selected from MVS(a). Prototypes in P(a) are indexed $\mathbf{p}^i$, $1 \leq i \leq 38$, though the superscript is often dropped, or used in place of $\mathbf{p}$.

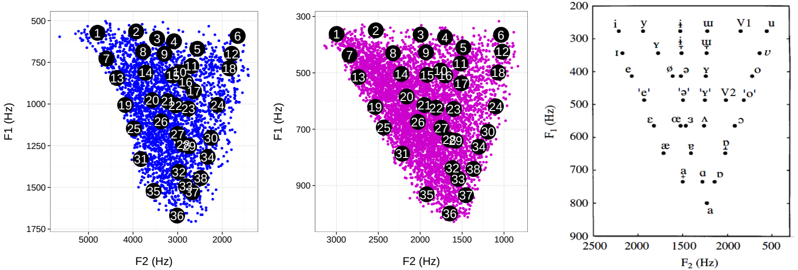

FIGURE : Prototype sets P(0.5) (left), P(10) (center), and those from Schwartz et al. (1997a).

PROTOTYPE 6

Age 0.5 Age 2 Age 4 Age 5 Age 10 Age 16 Age 21

Response Data

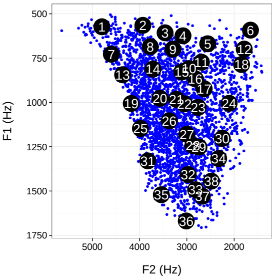

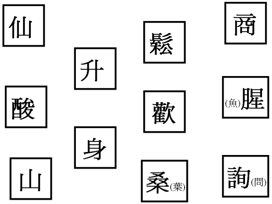

Vowel Category Responses

- Based on the vowel categories within their native languages, Cantonese-, American English- and Japanese-speaking listeners categorized seven sets of synthetic vowels generated from models of the vocal tract ranging in age from six months to 21 years of age.
- The stimuli were words in English, and Cantonese, and kana representing single vowels in Japanese.

Response Data

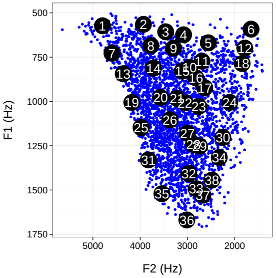

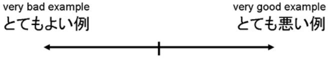

Vowel Category Responses

- After each trial, a visual analog scale (VAS) goodness judgment was collected.
- The numeric goodness rating is meant to capture how well each vowel represented the chosen vowel category.

Response Data

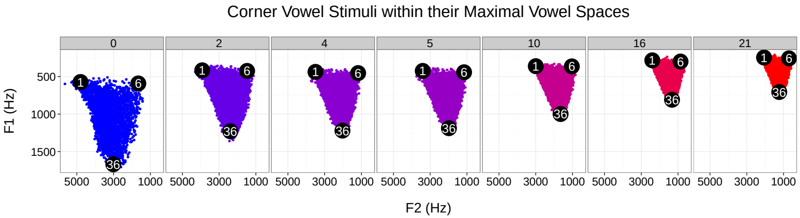

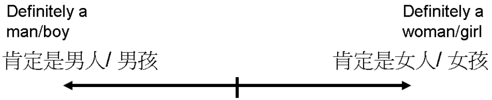

Age and Gender Responses

Listeners also provided numeric age judgments and gender ratings along a visual analog scale ranging from "definitely male" to "definitely female" (or their Cantonese and Japanese equivalents) for subsets of the stimuli corresponding to the corner vowels /i/, /a/, and /u/.

Modeling and Analysis

- Using the rich supply of data, we next present several models and analyses that shed light on within- and cross-language differences in vowel perception.
- We begin with a shape-based modeling approach for reasoning about vowel category perception.
- We then use basic statistics as a starting point for analysis of the age and gender results.
- We conclude with discussion of potential avenues for investigating the interaction between these response data in terms of our models.

Vowel Categorization Data

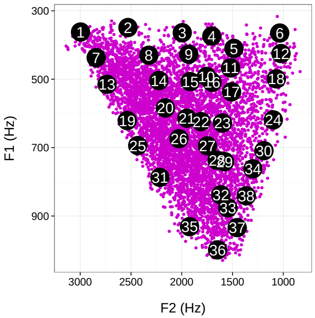

あ い う

え お

Using prototype categorizations and goodness ratings

- For each language, and for each subject, and for each vocal tract age, each prototype has a category associated with it, and a goodness rating.
- We want to use these responses to build models of each subject's knowledge of their language's vowel system.
- To illustrate the approach, we'll focus on the category responses provided by a Japanese subject, named J_{20}, and the 10 year old vocal tract age.

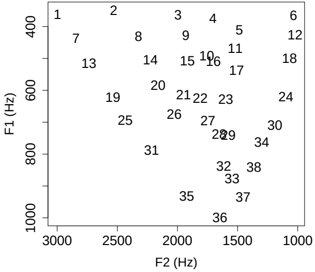

Proto.	Cat.	Good.
1	i	0.88
2	u	0.78
3	u	0.82
4	u	0.85
5	u	0.92
6	u	0.84
7	e	0.16
⋮	⋮	⋮

Organizing Subject Responses

- Vowel category judgements and goodness ratings for subject J_{20} with respect to the age 10 prototypes are shown above.
- Goodness values are normalized to fall between 0 (poorest) and 1 (best) to facilitate interpretation.

Vowel Categorization Data

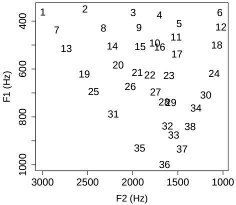

Proto.	Cat.	Good.
1	i	0.88
2	u	0.78
3	u	0.82
4	u	0.85
⋮	⋮	⋮

Can we expand the responses across the vowel space?

- Each response gives us a hint about a subject's knowledge of their language's vowel system, but we need to fill in some gaps.
- For example, according to J_{20} prototype 1 is a really good example of /i/, but what about /u/, /o/, or the other vowels?
- Similarly, the response tells us that prototype 2 is a good example of /u/, but what about /i/, /o/, or the other vowels?
- And what about all of the other prototypes?

EXPANDING THE RESPONSES ACROSS THE VOWEL SPACE

- We can use these responses to model what a subject might know about each prototype with respect to each vowel in the subject's vowel system.
- For each prototype, we create a vector with as many components as there are vowels in the subject's vowel system:

i	u	e	o	a
_	_	_	_	_

  whose components are a function of the subject's response to that prototype.
- To exemplify, for prototype 1 with response (i, 0.88), the vector is

i	u	e	o	a
0.88	0.06	0.06	0.06	0.06

- This function is specified in generality in Plummer et al. (2013).
- Thus, for each prototype, we have a model of a vowel category goodness for each vowel.

Vowel Categorization Data

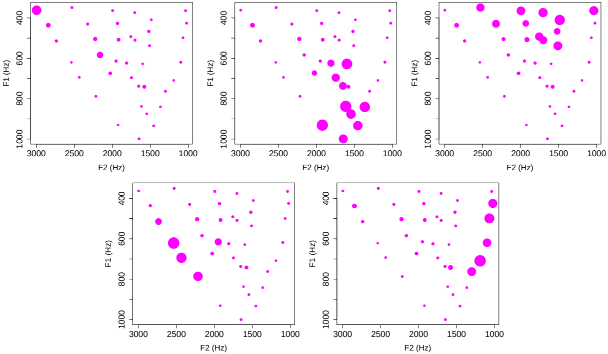

Individual Category Response Functions

- For each subject, and for each age, using these vectors, we can define functions from prototypes to goodness values for each vowel category.
- We can depict these individual category response functions in terms of weighted scatter plots, like the ones shown above for age 10, for subject J_{20}.

Vowel Categorization Data

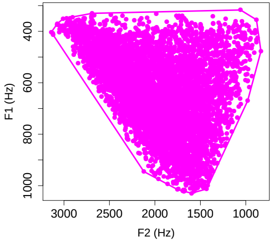

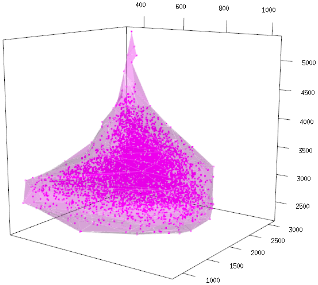

Vowel spaces as shapes, categories as scalar fields

- Now that we have models of vowel category responses for each prototype for each vocal tract age, we can extend these models to the entire maximal vowel space for each age.
- The trick is to think of the maximal vowel spaces as shapes (called manifolds), and the vowel categories as scalar fields over these shapes.
- We can derive these scalar fields from the individual category response functions using statistical models.

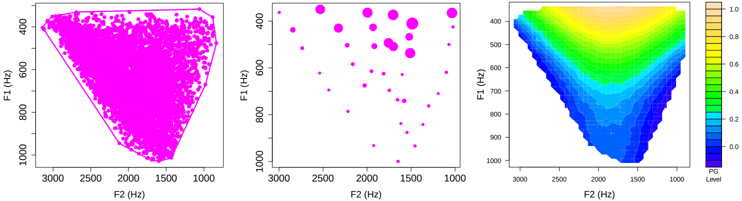

Additive Modeling

- The basic idea is to use additive models (Friedman & Stuetzle, 1981; Buja et al., 1989; Hastie & Tibshirani, 1990; Wood, 2003, 2006) based on smoothing splines (Wahba, 1990; Gu, 2002) to construct a field of responses over a maximal vowel space using an individual category response function.
- For a given vowel category, the response field value for a formant vector in a maximal vowel space is meant to approximate a subject's goodness rating of that vector as an example of that vowel category.

Vowel Categorization Data

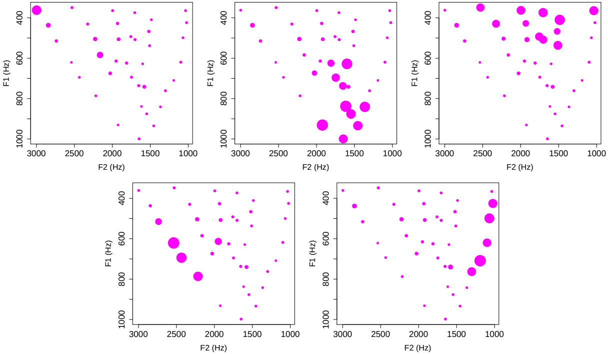

Vowel Category Response Fields

- For each subject, and for each age, using individual category response functions, we can define functions from maximal vowel spaces to goodness values for each vowel category.
- We can depict these vowel category response fields (VCRF) in terms of contour plots, like the ones shown above for age 10, for subject J_{20}.

Vowel Categorization Data

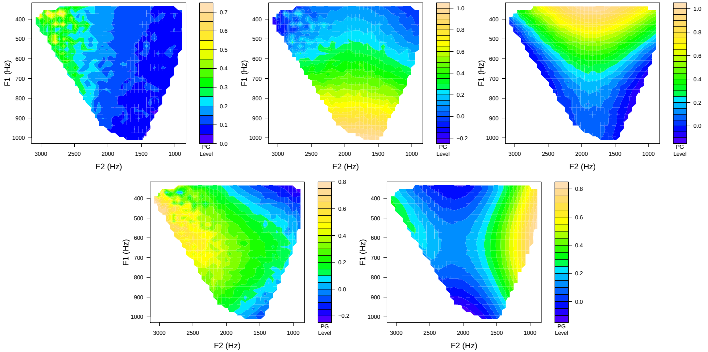

Vowel Category Response Fields

- For each subject, and for each age, using individual category response functions, we can define functions from maximal vowel spaces to goodness values for each vowel category.
- We can depict these vowel category response fields (VCRF) in terms of contour plots, like the ones shown above for age 10, for subject J_{20}.

Vowel Categorization Data

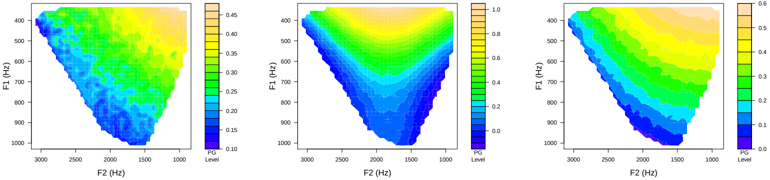

VCRF comparison

- VCRFs allow us to make within and cross-language comparisons concerning each subject's knowledge of the vowel system of their native language.
- Above are VCRFs for a high, back vowel, (left) derived from Cantonese subject C_{10}, (center) from J_{20}, and (right) from J_{04}.
- The differences are visually apparent, but we can also quantify them.
- Quantitative analysis in Plummer et al. (2013) shows that "distances" between VCRFs can capture
  - cross-language differences between the five-vowel systems of Greek and Japanese, and
  - within-language sociolinguistic differences concerning Japanese /u/.

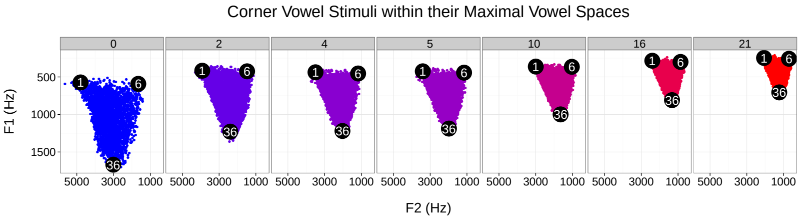

Definitely a man/boy
肯定是男人/ 男孩

Definitely a woman/girl
肯定是女人/ 女孩

ANALYSIS OF AGE AND GENDER RESPONSES

We turn now to analysis of the age and gender responses to corner vowel prototypes.

Age and Gender Data

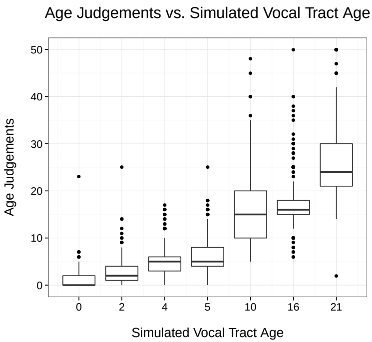

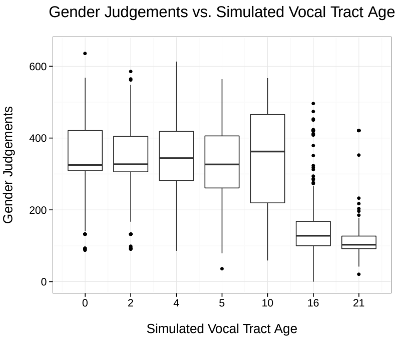

Summary of pooled age and gender responses

- At left, we show a boxplot (median, interquartile range, and full range of values) for the age judgments from all 57 listeners to all 6 trials for each vocal tract age (left).
- At right, we show a boxplot (median, interquartile range, and full range of values) for the gender judgments from all 57 listeners to all 6 trials for each vocal tract age.

Age and Gender Data

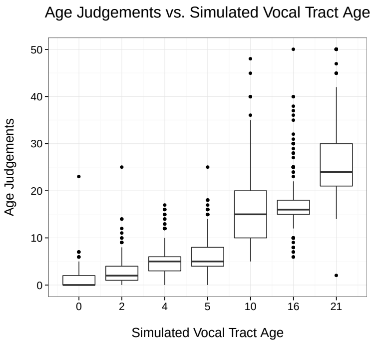

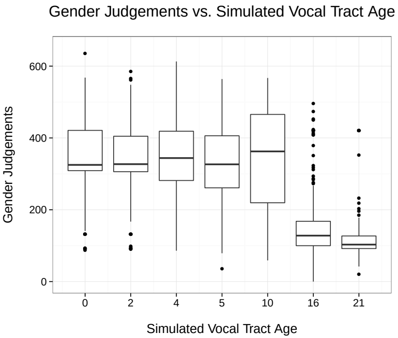

Summary of pooled age and gender responses

- The age judgements generally increase as the vocal tract age of the VLAM increases, with more variable responses for the three oldest vocal tracts, and especially variable responses to vocal tract age 10.
- The gender judgements demonstrate ambiguity in interpretation up to age 10, which is resolved by sexual dimorphism at the simulated age 16.

Age and Gender Data

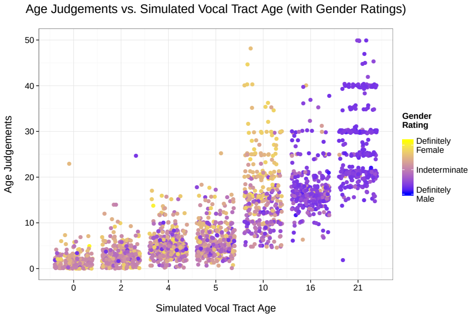

Summary of pooled age and gender responses

- The figure above shows a column scatter plot grouped by simulated vocal tract age for the age judgments from all 57 listeners to all 6 trials for each vocal tract age, with the color indicating the assigned gender rating.
- A bifurcation is evident in the interpretation of age and gender for the age 10 stimuli, which subjects rated as either a younger male or older female, suggesting a nonuniformity in the resolution of variability during processing.

STANDARD LINEAR REGRESSION BY SUBJECT

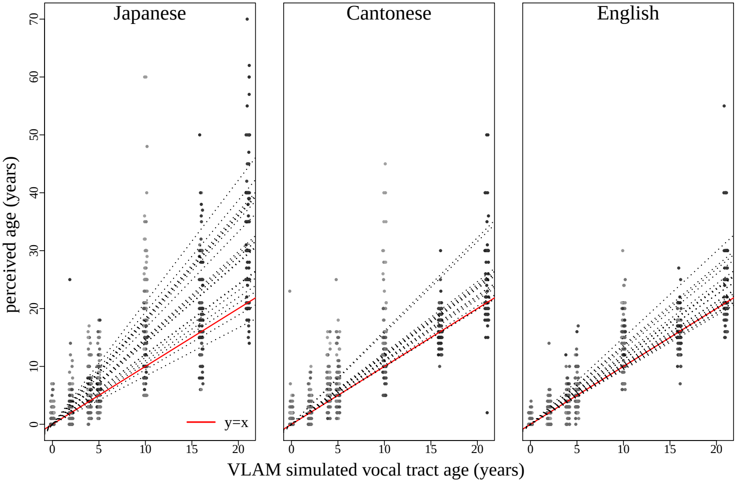

MOTIVATING HIERARCHICAL MODELING

- Standard regression lines for perceived age relative to simulated vocal tract age show a lot of variation both within and across languages, and across subjects.
- The English subject models (left) appear closest to the linear identity model (red), while the Cantonese subject models (center) deviate more, and the Japanese subject models most (right).
- Moreover, the gender responses seem to be contributing to the regression lines being pulled upward, and these also appear to be language-specific.
- We can see this more clearly with a mixed effects model over perceived gender (centered), simulated vocal tract age, and language, and all two way interactions, with random effects for simulated vocal tract age, gender, and their interaction by subject.

FIXED EFFECTS PLOT FACTORING OUT RANDOM EFFECTS FOR THE INTERACTION BETWEEN SIMULATED VOCAL TRACT AGE AND GENDER

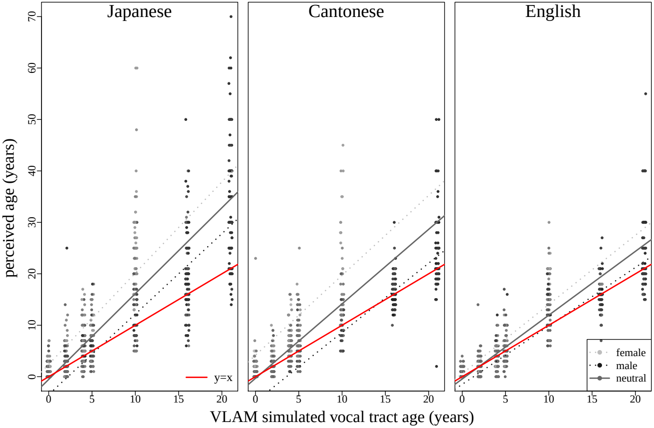

USING THE RANDOM EFFECTS TO MODEL INDIVIDUAL DIFFERENCES

- At present, we're trying to use the space of subject-specific random effects to make comparisons like we did for vowel category response fields.
- The goal is then to relate the distances between individuals in the random effects space with distances in the space of vowel category response fields.

BUT THIS IS TRICKIER THAN IT SEEMS

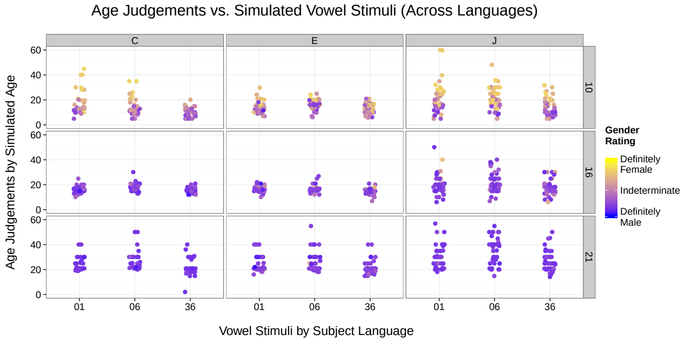

RESPONSES BY VOCAL TRACT AGE, LANGUAGE, AND STIMULI TYPE

- The figure above shows a set of column scatter plots arranged by vocal tract age and language, with data points grouped by vowel stimuli, for the age judgments from all 57 listeners to all 6 trials for vocal tract ages 10, 16, and 21 years, with the color indicating the assigned gender rating.
- The Japanese subjects are assigning an older rating to 21 year old vowels, and this is especially true for stimulus 6. This suggests a complex relation between perceived age and vowel categorization.

Age and Gender Data

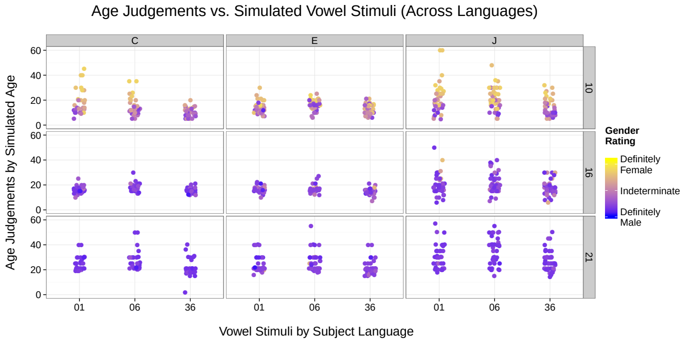

Age and vowel categorization

- The much older responses to stimulus 6 in the Japanese speakers' judgements could be an artifact of the quality of this stimulus, which is [u] rather than the unrounded [ɯ] that is the prototypical value for the high back vowel of Japanese.
- The rounded quality of the [u] might make it be heard as "careful speech" by Japanese listeners (see Okada, 1999), which could suggest an older, more deliberate talker.

Age and Gender Data

		Listener Group	
		Older/Woman	Younger/Man
judgment	ø	52	39
	e	10	11
	ɛ	40	52
	i	13	32
	y	36	30

		Listener Group	
		Older/Woman	Younger/Man
judgment	æ	21	22
	e	8	25
	ɛ	40	28
	i	16	16
	ɪ	21	14

Perceived age, perceived gender, and vowel categorization

- Vowel identification patterns for vocal tract age of 10 differed between individuals who identified that speaker as a younger male and those who identified the speaker as an older woman.
- Exploring these language effects further may reveal the depth of influence of culture on aspects of perception, especially vowel normalization.

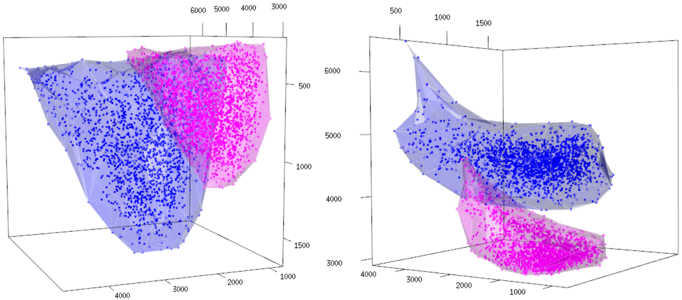

FIGURE : The maximal vowel spaces MVS(0.5) (blue) and MVS(10) (red) in first-three-formant space.

ACQUISITION MODELING

- In order to acquire language, infants must either possess, or come to possess a set of relations between mostly disjoint vowel spaces in order to form their language-specific vowel categories.
- The analyses presented in the previous section seem to suggest that this set is acquired over the course of development.

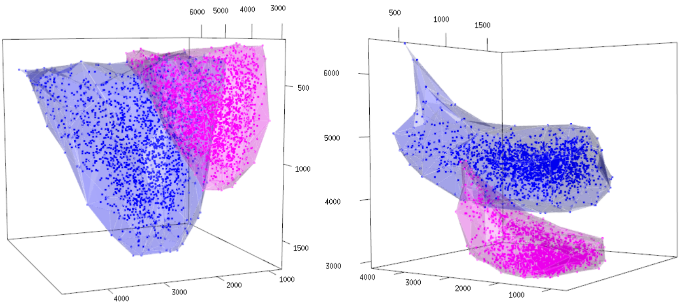

FIGURE : The maximal vowel spaces MVS(0.5) (blue) and MVS(10) (red) in first-three-formant space.

ACQUISITION MODELING

- If so, we can seek ways to model the acquisition of this set of relations that occurs during early infancy.
- Before concluding, we briefly sketch a computational modeling methodology for the acquisition of vowel normalization that makes use of vowel category response fields.

BASIC MODELING COMPONENTS

- We assume a dyad structure with one caretaker, one infant.
- Caretakers are modeled as sets of VCRFs, one set for the infant, and one for the caretaker.
- Infants are modeled as sets of structures, called cognitive manifolds.
- Dyadic vocal interaction is modeled as the weighted pairing of formant vectors from the caretaker's VCRFs with those from the infant's.
- Phonological acquisition is modeled as the development of the infant's cognitive structures in relation to vocal interaction with the caretaker.
- Specifically, the acquisition of vowel normalization is the alignment of cognitive manifolds based on vocal interaction.

JAPANESE CARETAKER J_{20}

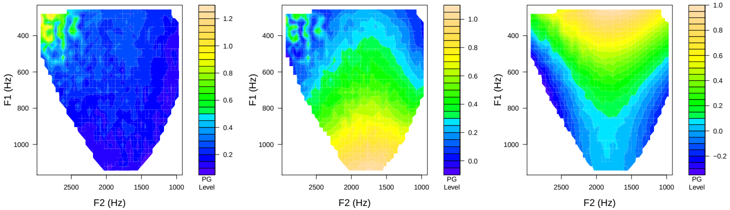

FIGURE : VCRFs over MVS(10).

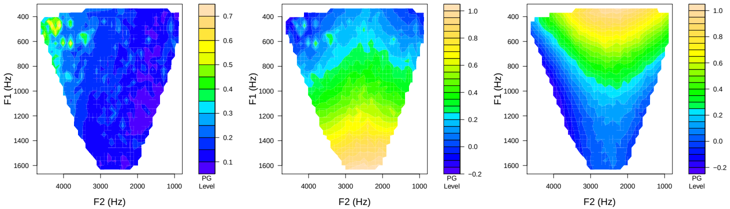

FIGURE : VCRFs over MVS(0.5).

CANTONESE CARETAKER C_{10}

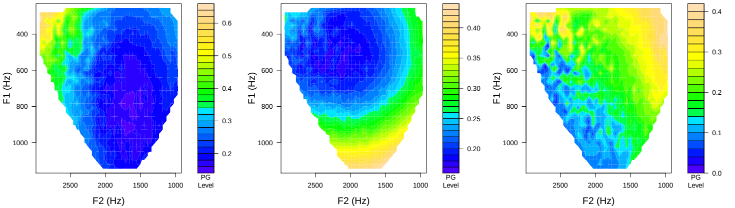

FIGURE : VCRFs over MVS(10).

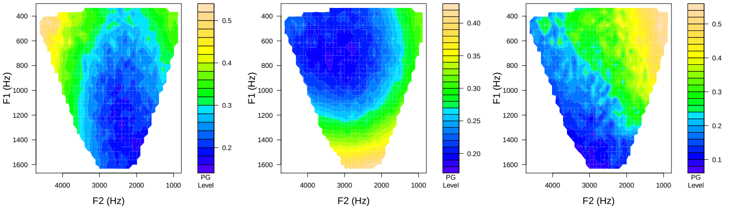

FIGURE : VCRFs over MVS(0.5).

Acquisition of Vowel Normalization

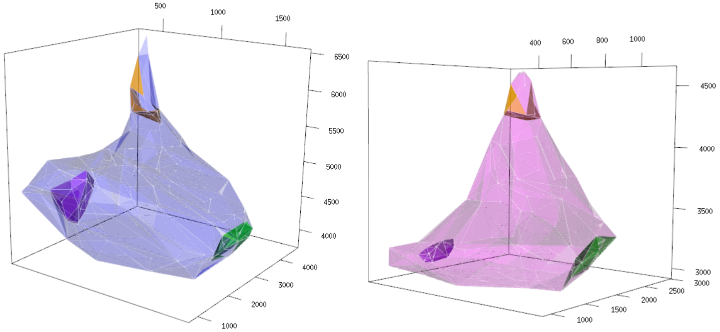

FIGURE : Vocal interaction between J_{20} and infant.

Modeling vocal interaction

- The VCRFs for J_{20} yield weighted pairings of formant vectors which model vocal interaction.
- The weighted pairings are derived from the orange (i), purple (u), and green (a) regions of the maximal vowel spaces.

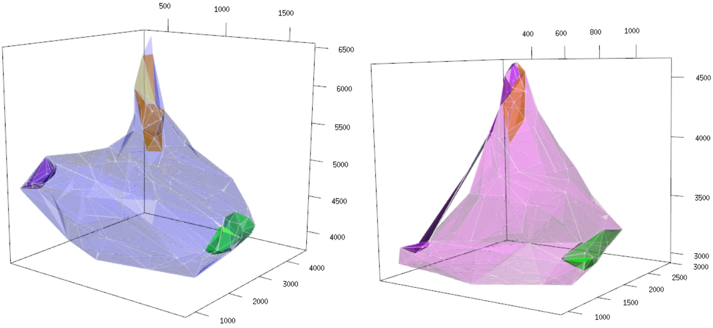

FIGURE : Vocal interaction between C_{10} and infant.

MODELING VOCAL INTERACTION

- Similarly, the VCRFs for C_{10} yield weighted pairings of formant vectors which model vocal interaction.
- The weighted pairings are derived from the orange (i), purple (u), and green (a) regions of the maximal vowel spaces.

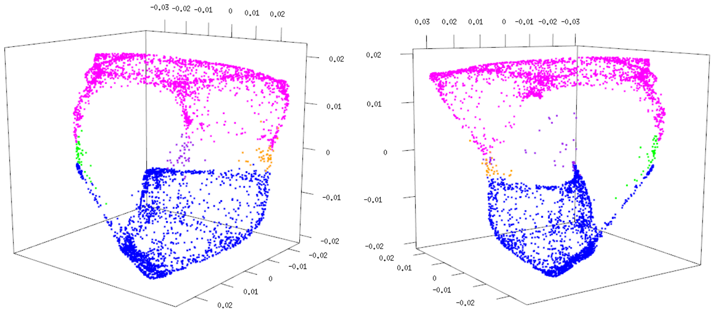

FIGURE : Representations from aligned cognitive manifolds following vocal interaction. The Japanese infant (left) has internalized a different location for [u] than the Cantonese infant (right).

ACQUISITION MODELING

- VCRFs provide the language-specific input to models of the acquisition of these relations.
- The modeling shows that individualistic aspects of the VCRFs also carry through the acquisition computation.
- This suggests that phonological acquisition may be dyad-specific.

RECAP

SUMMARY

- Language specificity seems to be evident in adults' perceptions/interpretations of vocalizations from a wide range of vocal tract ages. We can model this using vowel category response fields derived from responses to VLAM stimuli.
- Language specificity seems to be evident in adults' perception of age and gender typicality. We can model this (more or less) using mixed effects models derived from age and gender responses to the same VLAM stimuli.
- Language-specific age- and gender-perception seem to affect how vowels are labeled, but this is tricky to show, even with the models we have presently developed.

Acknowledgements

- Thanks to Asimina Syrika and Kiyoko Yoneyama for recruiting and testing the Greek- and Japanese-speaking participants.
- Thanks to Hyunju Chung and Chanelle Mays for recruiting and testing the Korean- and the English-speaking participants.
- Thanks to Catherine McBride-Chang for providing the resources to recruit the Cantonese-speaking subjects.
- Work supported by NSF grants BCS-0729306 and BCS-0729277.

Beck, J. M. (1996). Organic variation of the vocal apparatus. In W. J. Hardcastle, & J. Laver (Eds.) *Handbook of Phonetic Sciences*, (pp. 256–297). Cambridge, England: Blackwell.

Boë, L.-J., & Maeda, S. (1998). Modélization de la croissance du conduit vocal. Éspace vocalique des nouveaux-nés et des adultes. Conséquences pour l'ontegenèse et la phylogenèse. In *Journée d'Études Linguistiques: "La Voyelle dans Tous ces États"*, (pp. 98–105). Nantes, France.

Boë, L.-J., Perrier, P., Guérin, B., & Schwartz, J.-L. (1989). Maximal vowel space. In *EUROSPEECH 09*, (pp. 281–284). Paris, France.

Buja, A., Hastie, T., & Tibshirani, R. (1989). Linear smoothers and additive models. *The Annals of Statistics*, *17*(2), 453–510.

Friedman, J. H., & Stuetzle, W. (1981). Projection pursuit regression. *Journal of the American statistical Association*, *76*(376), 817–823.

Goldstein, U. G. (1980). *An articulatory model of the vocal tract of the growing child*. Ph.D. thesis, Massachusetts Institute of Technology.

Gu, C. (2002). *Smoothing spline ANOVA models*. Springer.

Hastie, T. J., & Tibshirani, R. J. (1990). *Generalized Additive Models*. New York: Chapman and Hall.

Lindblom, J., & Sundberg, J. E. F. (1971). Acoustical consequences of lip, tongue, jaw, and larynx movement. *Journal of the Acoustical Society of America*, *50*, 1166–179.

Maeda, S. (1990). Compensatory articulation during speech: Evidence from the analysis and synthesis of vocal-tract shapes using an articulatory model. In W. Hardcastle, & A. Marchal (Eds.) *Speech Production and Speech Modeling*, (pp. 131–149). The Netherlands: Kluwer Academic Publishers.

Maeda, S. (1991). On articulatory and acoustic variabilities. *Journal of Phonetics*, *19*, 321–331.

Ménard, L., & Boë, L.-J. (2000). Exploring vowel production strategies from infant to adult by means of articulatory inversion of formant data. In *Proceedings of the 6th International Conference on Spoken Language Processing (ICSLP 2000)*, (pp. 465–468). Beijing, China.

Ménard, L., Schwartz, J.-L., & Boë, L.-J. (2002). Auditory normalization of French vowels synthesized by an articulatory model simulating growth from birth to adulthood. *Journal of the Acoustical Society of America*, *111*(4), 1892–1905.

Munson, B., Ménard, L., Beckman, M. E., Edwards, J., & Chung, H. (2010). Sensorimotor maps and vowel development in English, Greek, and Korean: A cross-linguistic perceptual categorizaton study (A). *Journal of the Acoustical Society of America*, *127*, 2018.

Okada, H. (1999). Japanese. In *Handbook of the International phonetic Association: a guide to the use of the International Phonetic Alphabet*, (pp. 117–119). Cambridge, U.K.; New York, New York: Cambridge University Press.

Plummer, A. R., Ménard, L., Munson, B., & Beckman, M. E. (2013). Comparing vowel category response surfaces over age-varying maximal vowel spaces within and across language communities. In *Proceedings of INTERSPEECH 2013*.

Schwartz, J.-L., Boë, L.-J., & Abry, C. (2007). Linking dispersion-focalization theory and the maximum utilization of the available distinctive features principle in a perception-for-action-control theory. In M.-J. Sole, P. S. Beddor, & M. Ohala (Eds.) *Experimental Approaches to Phonology*, (pp. 104–124). Oxford University Press.

Schwartz, J.-L., Boë, L.-J., Vallée, N., & Abry, C. (1997a). The dispersion-focalization theory of vowel systems. *Journal of Phonetics*, *25*, 255–286.

Schwartz, J.-L., Boë, L.-J., Vallée, N., & Abry, C. (1997b). Major trends in vowel system inventories. *Journal of Phonetics*, *25*, 233–253.

Vallée, N., Boë, L.-J., & Payan, Y. (1995). Vowel prototypes for UPSID's 33 phonemes. In *Proceedings of ICPhS 2*, (pp. 424–427). Stockholm.

Wahba, G. (1990). *Spline models for observational data*. Society for Industrial and Applied Mathematics.

Wood, S. (2006). *Generalized Additive Models: An Introduction with R*. Chapman & Hall/CRC.

Wood, S. N. (2003). Thin-plate regression splines. *Journal of the Royal Statistical Society (B)*, *65*(1), 95–114.