Vowel Harmony as a Distributional Learning Problem

Spencer Caplan and Jordan Kodner
University of Pennsylvania
spcaplan@sas.upenn.edu, jkodner@sas.upenn.edu

Vowel Harmony in Phonological Theory

- Language-wide vowel alternation pattern
- Present in a large portion of the world's languages
  - Patterns systemic across roots and affixes
  - Languages may have more than one harmony process
- Vowels partitioned into sets
  - Words contain vowels from one set only
  - "Neutral" vowels do not change

- Caused by speading of phonological featues
  - Frontness (**Turkish, Finnish**)
  - ATR (**Mongolian, Javanese, Fula**)
  - Roundness (**Turkish, Walpiri**)
- Spreading may be either:
  - **Left-to-Right** (Turkish, Finnish)
  - **Right-to-Left** (Fula, Pulu)

Finnish Vowels

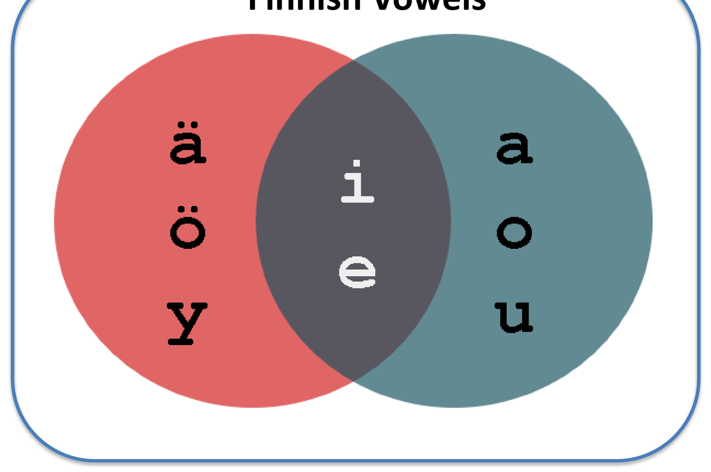

Turkish Vowels

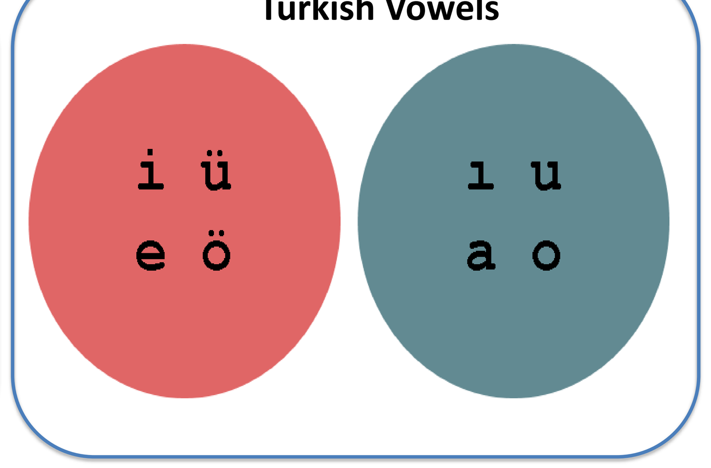

Finnish
…*teluttelemattomammuuksissansakaankopahan*
vs.
…*telmällistyttämättömyydellänsäkäänköhän*

Turkish (Left-to-Right, realized on suffixes)
Baş-lar vs. *Beşev-ler*

Fula (Right-to-Left, realized on roots)
mbeel-u → *mbɛɛl-ɔn* 'shadow'
peec-i → *pɛɛc-ɔn* (proper noun)

Early Acquisition

- Infants as young as seven months are sensitive to vowel harmony alternations in acoustic input (Mintz et al. 2006)
- Preferential listening task over unsegmented input (raw stream of continuous audio)
- Effect present for children whose normal language environment has *no* harmony

From Raw Input to Productive Grammar

- Phonological theory describes speaker's behavior involving vowel harmony
- But the child learner has no direct access to the latent parts involved in such a grammar
- **A learning model must exist to translate from raw input to the abstract input/output grammar**

Requirements of a Harmony Acquisition Model

- Requires only limited input data to learn
- Input is unsegmented speech
  - Stream of phonemes rather than neatly cut words
  - No frequency counts (need to handle high frequency exceptions)
  - However, Vowel/Consonant tiers are differentiable (Newport and Aslin 2004)
- Reasonable Cognitive Tools
  - Processing happens online rather than in batch
  - Posited calculations shuld be implementable by the learner

Distributional Hypothesis

- In a no-harmony language (e.g. English) we don't expect any particular vowel to be more or less likely to follow another
- In a vowel-harmony language one class of vowels is **mutually exclusive** to the other
- **The distributional signal of vowel harmony is a divergence from the base uniform distribution of vowel co-occurrences**
- Simply tracking conditional probabilities fails to capture this pattern
  - Frequency differences between vowels disrupt the signal
- Instead we normalize conditional probabilitiy by the frequency of the preceding vowel context

---

- We present a computational model of vowel harmony acquisition.
  - To date no model had been able to explain children's experimental and empirical performance with respect to vowel harmony (cf. Goldsmith and Riggle 2012)
  - Following previous experimental work the model learns only from limited unsegmented linguistic input
- The model accounts for a wide range of typological facts over varied cross-linguistic input

Lang.	1ary H?	Correct	2ary H?	Correct
Turkish	✓	8/8	✓	4/4
Finnish	✓	8/8	–	–
Uyghur	✓	7/8	–	–
Warlpiri	✓	3/3	–	–
German	–	5/5	–	–
English	–	6/6	–	–

Table showing performance of the vowel harmony learning model on various languages. Performance matches the target grammars of learners in almost all cases (97% of vowels correctly categorized with respect to harmony).

The small failure in Uyghur may be caused by orthographic factors, since the model was fed input as characters rather than phonemes.

English Vowel Pair PMI

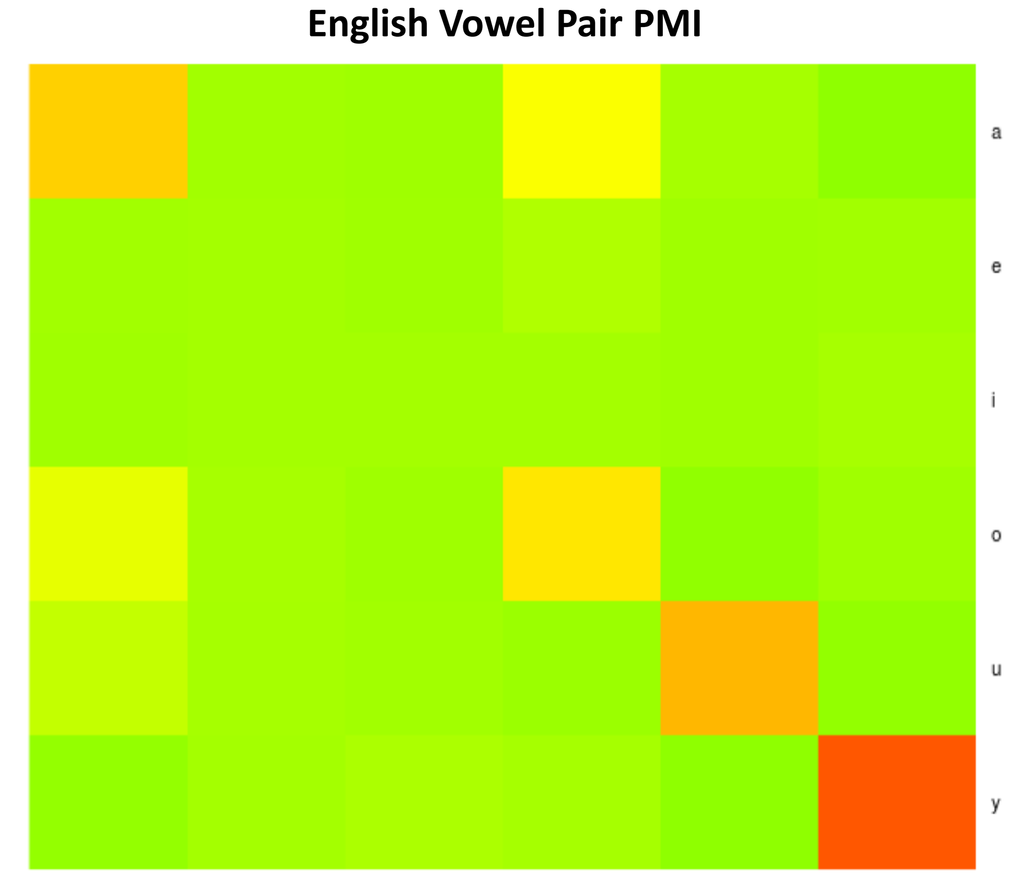

Finnish Vowel Pair PMI

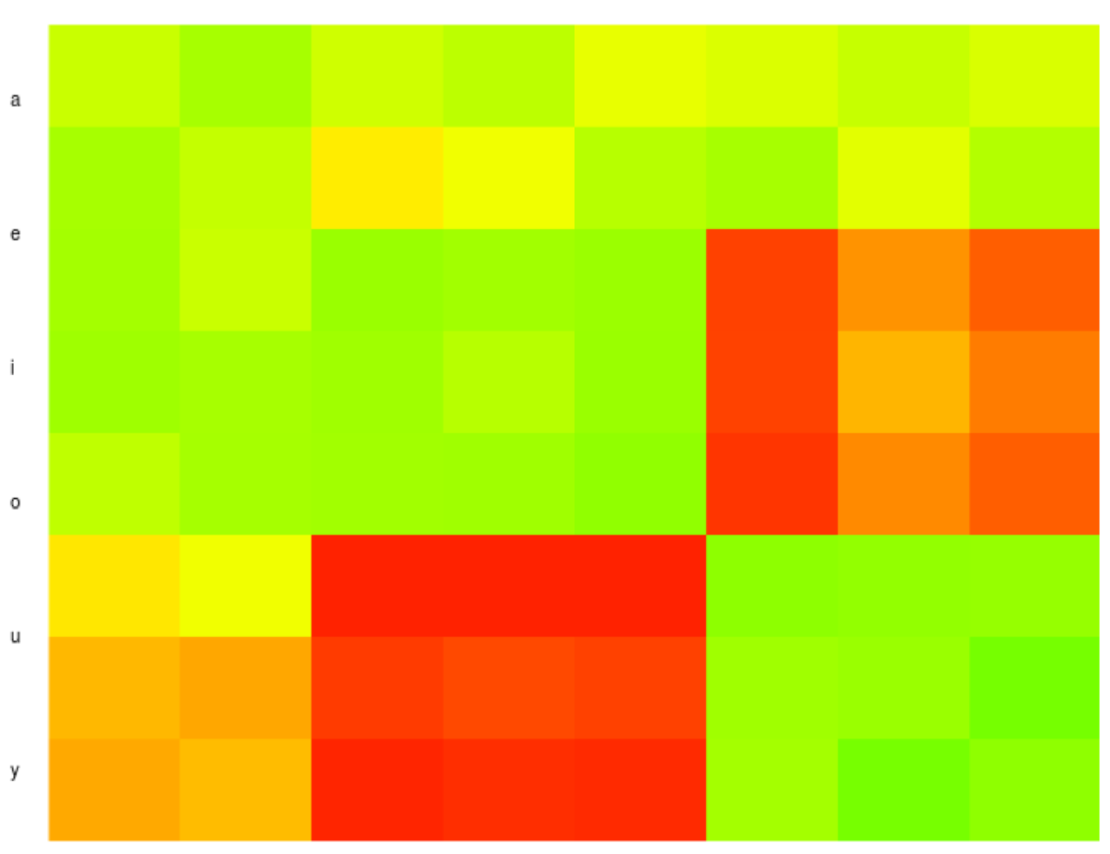

Turkish Vowel Pair PMI

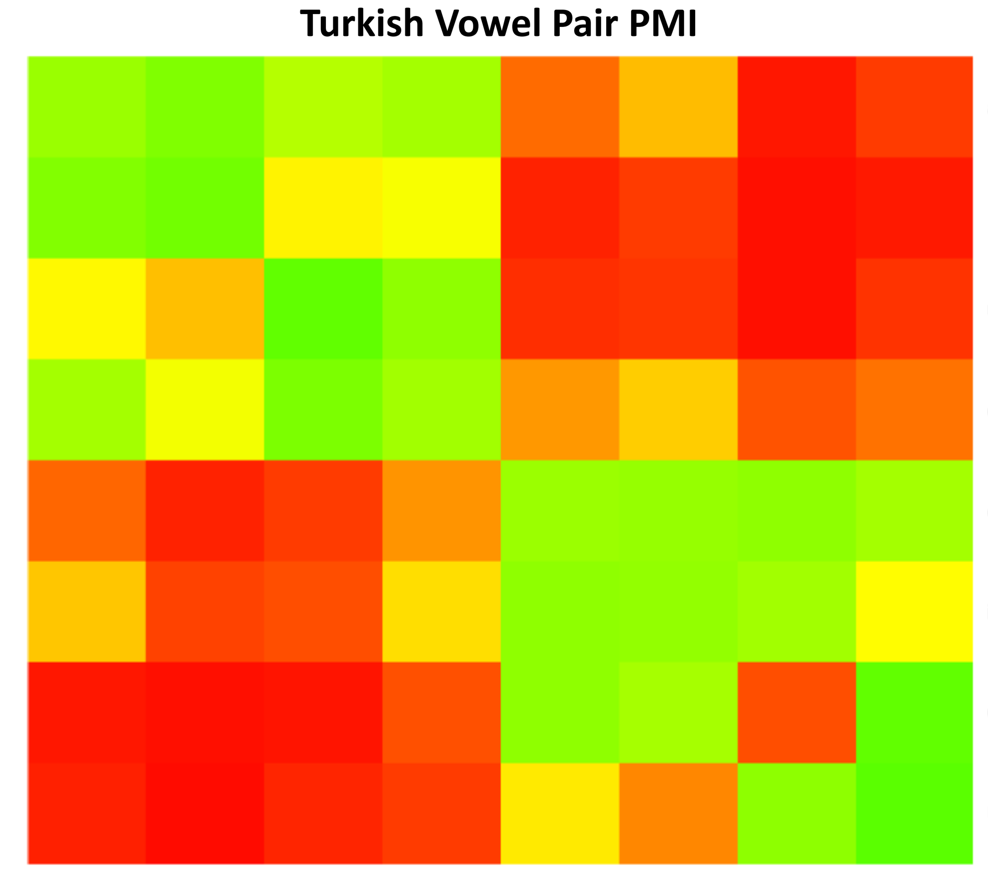

Heatmaps showing PMI for each pair within the vowel spaces of various test languages (English, Finnish, and Turkish)

- Red color indicates pairs which are distributionally less likely to co-occur
- Green pairs co-occur either at or above change

The distribution for English (as a non-harmony) language is close to uniform.

Finnish and Turkish show clear division of vowel distribution into latent harmonizing sets

- Note the two neutral vowels in Finnish "i" and "e" in the two top-most rows show no significant deviation from the uniform distribution.

Empirical Predictions and Questions for Future Research

- Is primary harmony within a language in fact acquired first as predicted by the model?
- Does harmony need to be a function over a single phonological feature (i.e. a natural class)?
- How would performance look given asymmetry within the distribution of only a single vowel?
- How does a learner differentiate input data which shows non-productive disharmony (e.g. Estonian) from a productive harmony process (e.g. Finnish)? (e.g. Yang 2005)

Model Implementation

1) Tabulate Vowel-Vowel Cooccurrence Matrix
- **Counting adjacencies over vowel-tier**
- **However signal and algorithm are robust to perform well under other counting schemes**

Utterence: "kababesisata"

Coocurrence matrix:

	A	E	I
A	2	0	1
E	1	0	0
I	0	1	0

Vowel frequencies: C(a)=4, C(e)=1, C(i)=1

Vowel probabilities: P(a) = 4/6, P(e) = 1/6, P(i) = 1/6

2) Convert to normalized conditional probabilities

$$normP(a|e) = \frac{C(a|e)}{C(a)P(e)}$$

3) Identify Neutral vs. Harmonizing Vowels
- **Threshold proportional to cardinality of vowel set**
- **If all vowels appear neutral this is a non-harmony language (e.g. English)**

$$Threshold = \frac{0.5}{\#Vowels}$$

For each pair of vowels v1, v2
If normP(v1|v2) < Threshold:
v1 is not neutral

4) Find featural partition

Cluster harmonizing vowels into two sets

K-means (k=2)

Division must be a natural class over a single phonological feature

5) Collapse over feature and repeat

Remove harmonizing feature from vowel space and repeat algorithm until all remaining vowels appear neutral.

This identifies any secondary harmony processes

Thank you

… to Charles Yang, Mitch Marcus, Ryan Budnick, and the audience at NECPhon10 for helpful feedback and advice

Code Available Open Source

https://github.com/scaplan/VowelHarmonyAcquisition

References

[1] Aslin, Richard N., Jenny R. Saffran, and Elissa L. Newport. "Computation of conditional probability statistics by 8-month-old infants." Psychological science 9.4 (1998): 321-324.
[2] Clements, George N. "Vowel harmony in nonlinear generative phonology." Bloomington: Indiana University Linguistics Club (1976).
[3] Goldsmith, John, and Jason Riggle. "Information theoretic approaches to phonological structure: the case of Finnish vowel harmony." Natural Language & Linguistic Theory 30.3 (2012): 859-896.
[4] Mintz, Toby, and Rachel Walker. "Infants' Sensitivity to Vowel Harmony and its Role in Word Segmentation." (2006).
[5] Yang, Charles. "On productivity." Linguistic variation yearbook 5.1 (2005): 265-302.