



*Socioindexical Variation and the
'Ladder of Abstraction'
in Phonology*

Benjamin Munson¹

Eden Kaiser²

¹Department of Speech-Language-Hearing Sciences

²Program in Linguistics

University of Minnesota



Up-Front Acknowledgments

- *Major collaborators on the experiments in this presentation: Di Schempp, Wendy Smith, Julie Johnson, Alysse Zittnan, Evan Figg, Celina Marnie*
- *This work was inspired by my teachers and colleagues, but especially: Janet Pierrehumbert*
- *Funding: McKnight Foundation, University of Minnesota Undergraduate Research Opportunities Program (UROP), University of Minnesota College of Liberal Arts*
- *Collaborators on pretty much everything else we work on: Jan Edwards, Mary Beckman, Kari Urberg-Carlson*
- *Useful comments from: Jen Hay, Paul Foulkes, Ghada Khattab, Katie Drager*
- *All remaining errors are: mine and mine alone*



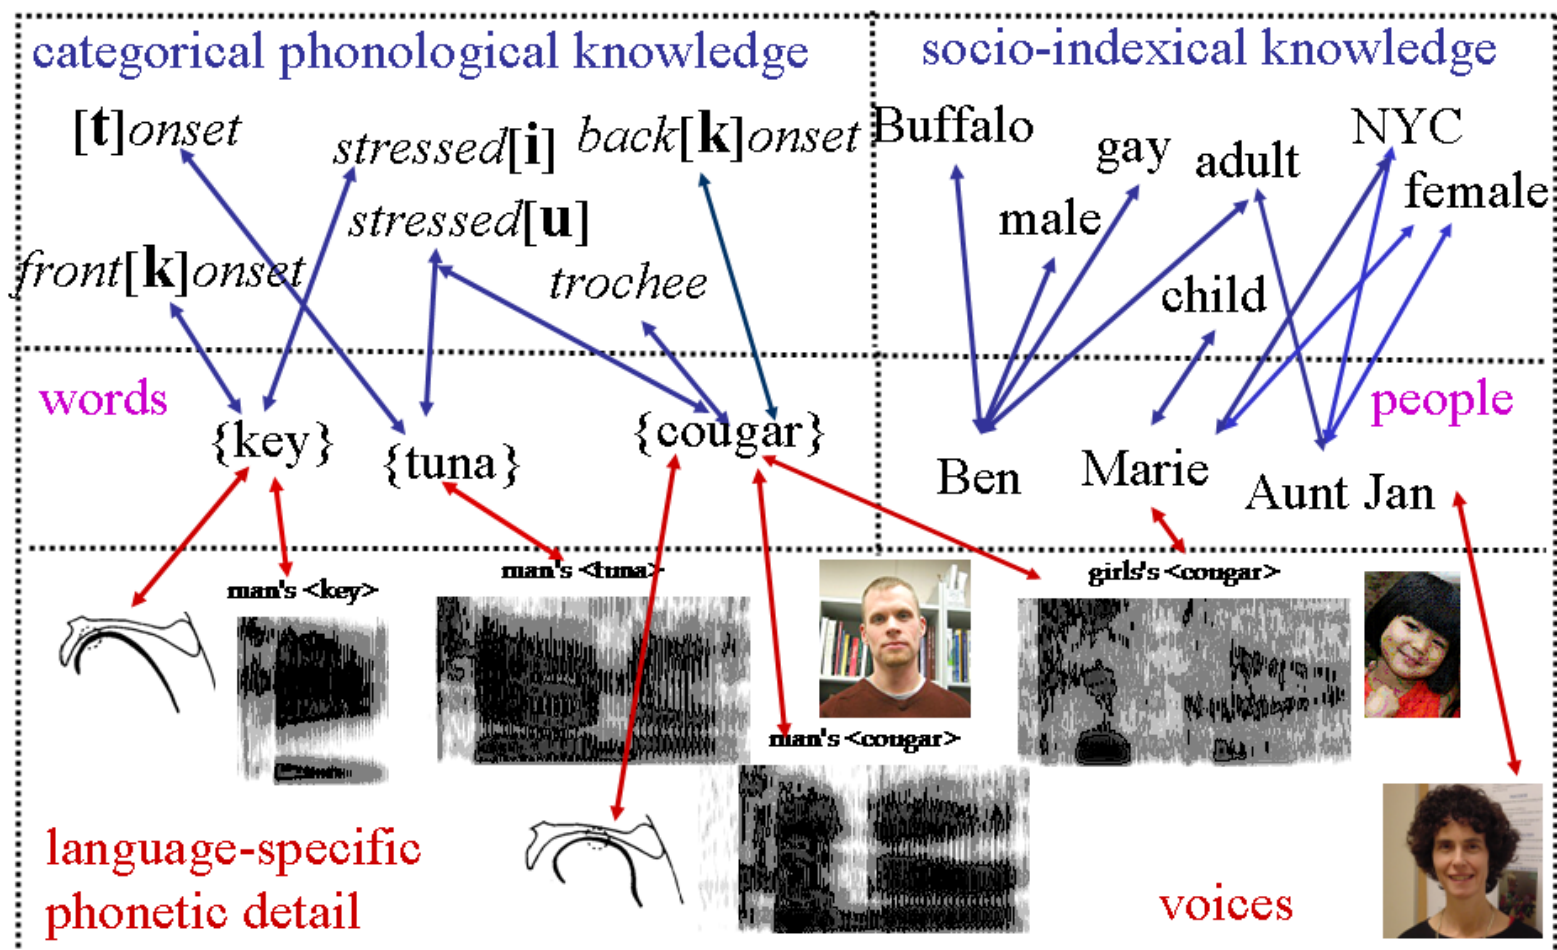
The Ladder of Abstraction

- Individuals' knowledge of the sound structure of language occurs at multiple levels of abstraction (Pierrehumbert, 2003; Beckman, Munson, and Edwards, 2007).
- Learning phonology involves developing progressively more abstract knowledge of the sound structure of language.



Figure: M. Beckman & J. Edwards

Levels of knowledge about speech sounds






The Ladder of Abstraction

- On the 'lexical-indexical' side, there are
- *Language-specific parametric phonetic representations:* detailed encodings of speech that has been heard, produced, and seen.
- Begins *in utero*
- Continues throughout life (Goldinger & Azuma, 2004; Sancier & Fowler, 1997; *inter alia*)



The Ladder of Abstraction

- *A Discretized Lexicon:*
word-form
representations that are
segmented only
insomuch as is enforced
by the nonlinear
relationship between
articulation and acoustics





The Ladder of Abstraction

- The development of discretized representations begins in the middle of the first year of life as the child begins to match the phonetics characteristics of the vocalizations of the adults in the native language



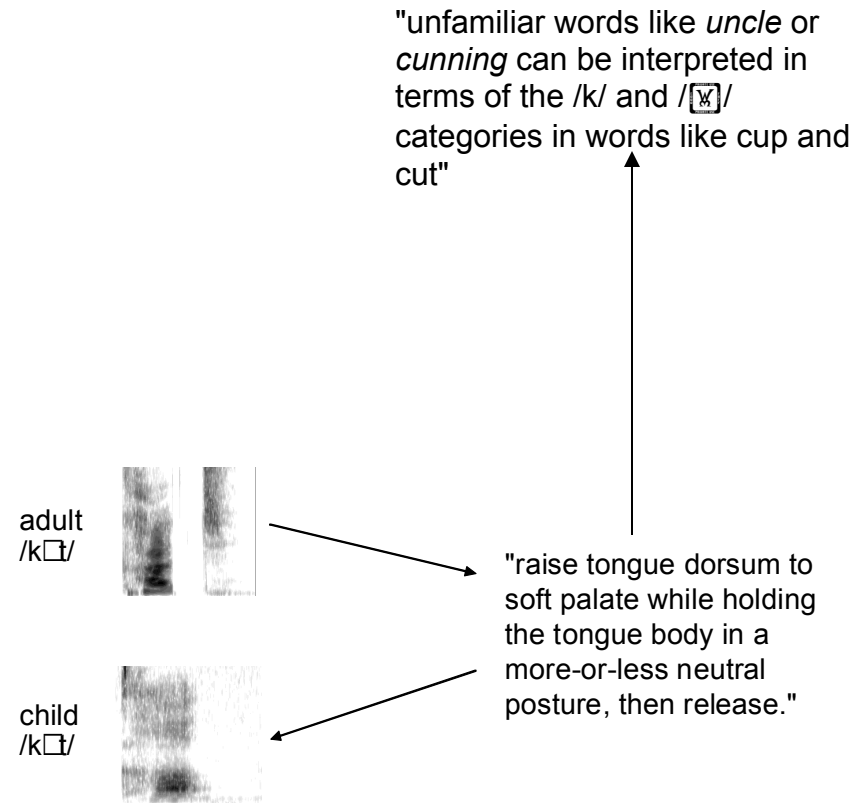
The Ladder of Abstraction

- *Generalizations about the sound structure of the language, which include*
 - The generalization that novel words can be parsed in terms of categories from already-known words
 - Generalizations about phonotactic constraints and morphophonological alternations



The Ladder of Abstraction

- These are projected from the lexicon, and they change as the lexicon changes and grows (Edwards, Beckman, & Munson, 2004)





What is Phonology?

A Progressive View

- Everything is 'phonological' in the sense of involving language specific knowledge...
 - Parametric phonetics differ from language to language
 - Each language has a unique combination of phones in the words in its lexicon
 - The phonotactic constraints and morphophonological alternations that emerge from the lexicon are a by-product of language-specific lexica
- ...derived via domain-neutral skills
 - Articulatory-acoustic learning
 - Statistical learning



What is Phonology?

A Conservative View

- 'Phonological' refers only to the most-abstract level of knowledge (i.e., phonotactic constraints and morphophonological alternations)
 - Halle 1985 is every laboratory phonologist's favorite straw man for this claim



The Ladder of Abstraction

- One hypothesis: Individuals make parallel sets of generalizations about the aspects of variation that are relevant for lexical contrast and those that relate to social-indexical variation



Socioindexicality and Phonology

- Knowledge of the sound structure of language is affected by knowledge of how pronunciation varies as a function of macro- and microsociological categories
- Example: the most-widely studied macrosociological category, gender.



Socioindexicality and Phonology

- Knowledge of (and beliefs about) gender-related variation affects...
 - Phoneme identification (Strand & Johnson, 1996; Munson, 2009 [this conference])
 - Spoken-word recognition (Strand, 2000 dissertation)



Socioindexicality and Phonology

- A conventional interpretation: this is evidence that memories for speech sounds are both highly phonetically detailed, and are indexed to attributes about the speakers that produced them (e.g., Hay et al., 2006)
- When speaker attributes are activated, they spread activation to the items that are associated with them.



Socioindexicality and Phonology

- If phonological knowledge involves the progressive abstraction from instance-based encodings, to what level of abstraction is information about speaker attributes retained?
- Is it retained at level at which phonotactic constraints and morphophonological alternations are made?
- Such a finding could convince even the most conservative phonologist that socioindexical variation is relevant for phonology.



An Experiment

- Listeners can make generalizations about phonotactics without explicit instruction based on a very short exposure to a novel language (Onishi, Chambers, & Fisher, 2002). These track the statistics of the novel language (Goldrick & Larson, 2008)
- These are not tied to the productions of a particular speaker (Onishi et al., 2002)
 - I.e., they appear to be 'grammatical' rather than ad hoc



An Experiment

- Are listeners' phonotactic generalizations about a novel language sensitive to differences in men's and women's production of this language?
 - A 'proof-of-concept' study on the learnability of socially stratified phonotactic constraints



Conducted by Benjamin Munson, Di Schempp, and Wendy Smith



An Experiment

- Listeners heard a series of 'CVCCVC nonwords and rated the clarity with which they had been produced
- #CV, VC, CV, and VC# sequences were legal (but generally low-frequency) English sequences
- CC sequences were low- or zero-frequency sequences in English.
 - /fk/, /fm/, /Sp/, /Sk/, /vb/, /vd/, /vg/



An Experiment

- Two sequences were trained with both men and women's voices, paired with men and women's faces



/vb/
/vg/
*/vd/



/vb/
/vg/
*/vd/



Images courtesy of Michael J. Tarr, Brown University, <http://www.tarrlab.org/>



An Experiment

- Two sequences were trained with either men or women's voices, paired with men and women's faces



/fk/
*/fm/
*/fp/



*/fk/
/fm/
*/fp/



Images courtesy of Michael J. Tarr, Brown University, <http://www.tarrlab.org/>



An Experiment

- Two sequences were trained with both men and women's voices, paired with men and women's faces



/Sp/
*/Sk/
*/Sm/



*/Sp/
/Sk/
/Sm/



Images courtesy of Michael J. Tarr, Brown University, <http://www.tarrlab.org/>



An Experiment

- Three men's and three women's productions were used during the training phase.
- Following a distracter task, participants listened to a new set of wordforms, spoken by four new talkers (two men, two women).
- They rated how good they would be as additions to the novel language. Judgments were elicited with a visual analog scale.





Image courtesy of Michael J. Tarr, Brown University, <http://www.tarrlab.org/>

Very
Good



Very
Bad

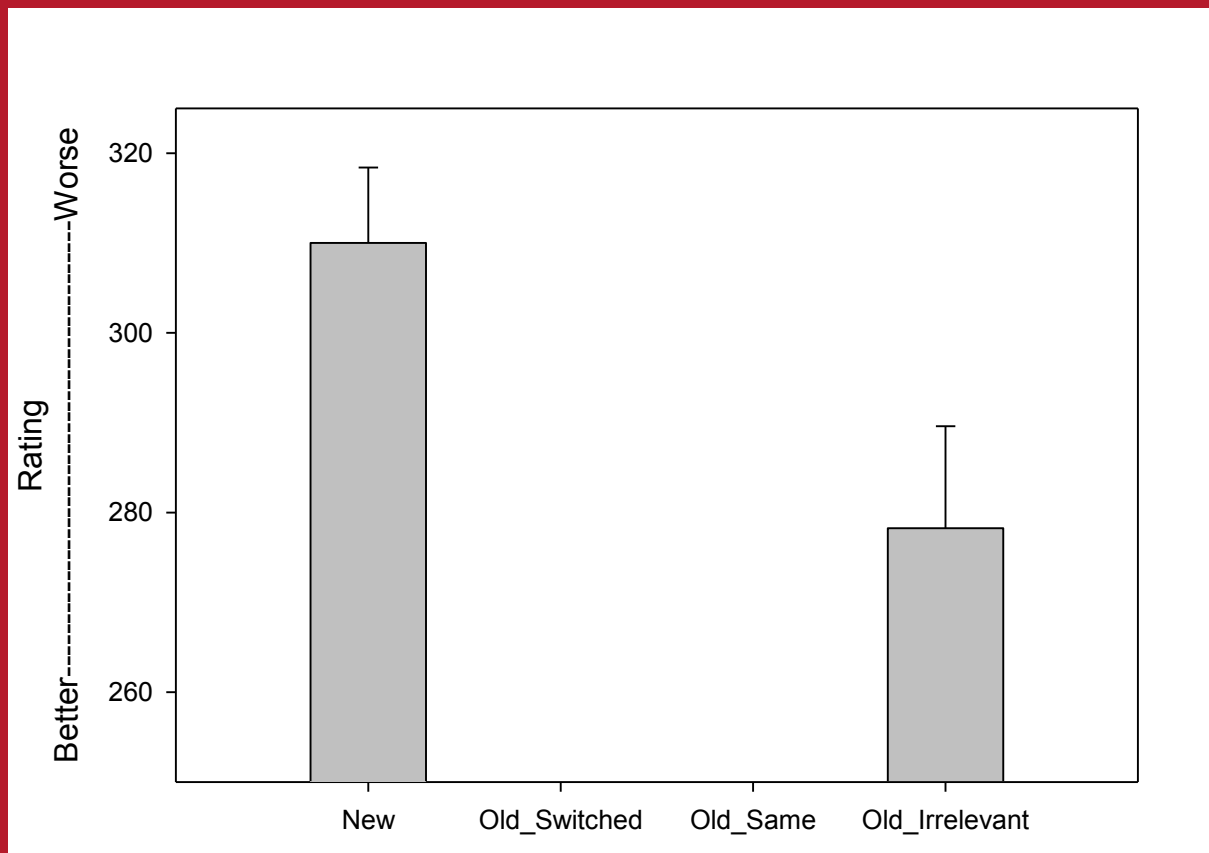


An Experiment

- Four types of CC sequences were tested
- 1. Sequences that had been trained in both men and women (*old_irrelevant*)
- 2. Sequences that were trained with only one sex, and were tested with the same sex with which they were trained (*old_same*)
- 3. Sequences that were trained with only one sex, and were tested with the other sex (*old_different*)
- 4. Completely untrained sequences (*new*)




An Experiment's Results



Words with untrained CC sequences were rated poorer than words with CC sequences which had been trained with both men and women



An Experiment's Results




$F[3,66]=8.59$,
 $p = 0.001$,
 $\eta^2_{partial} = 0.28$

...but old CC sequences that had been trained with only one sex were rated equally well regardless of the sex of the talker who produced them in the test phase



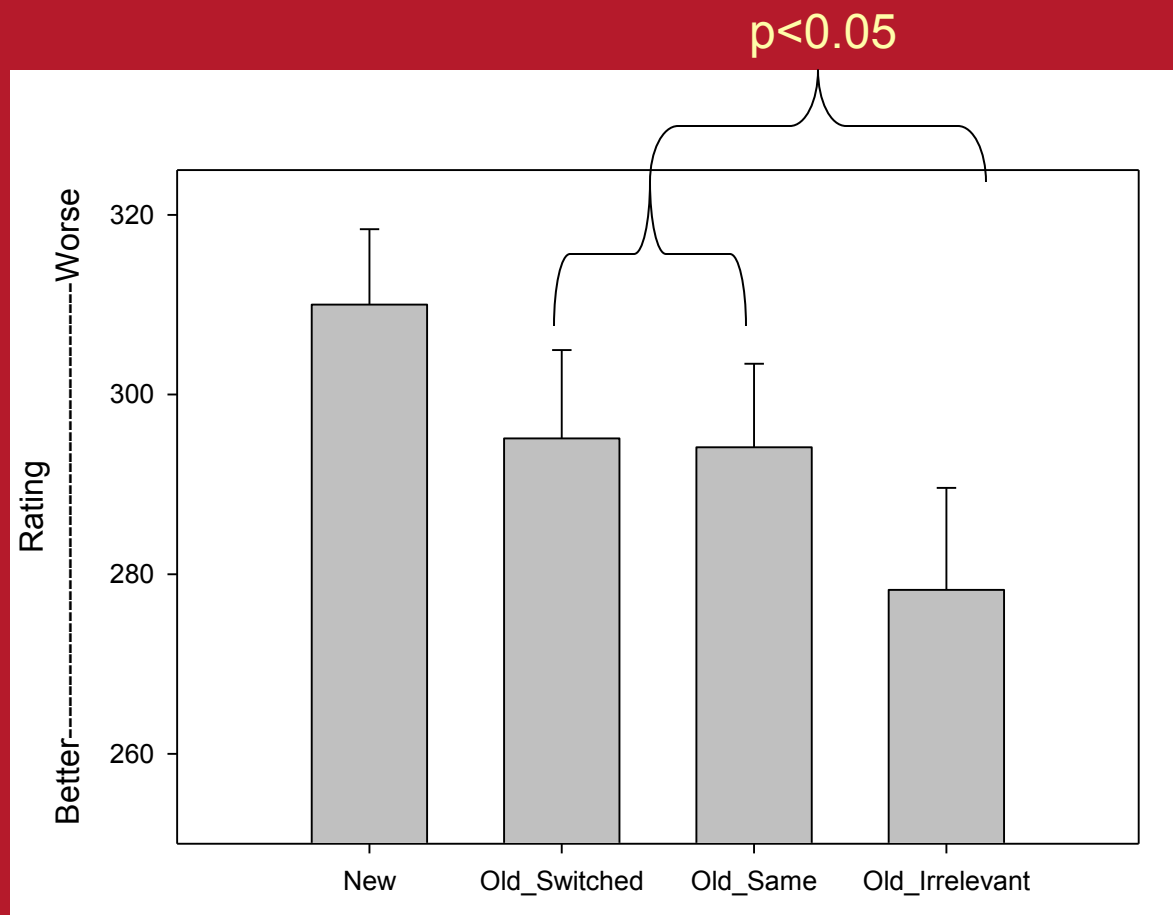
An Experiment's Results



- The most reliable difference was between untrained sequences and all other sequences




An Experiment's Results



- The sequences trained with only one sex were rated as marginally worse than the sequences trained with both sexes
- This stands to reason, given that only half as many talkers produced them

When plotted on the VAS line, the differences are quite small, though reliable across listeners





An Experiment's Discussion

- Listeners easily make generalizations about the likelihood of occurrence of sequences in an artificial language.
- These are gradient: sequences with more exemplars are rated as better than those with fewer.
- Generalizations are not linked to the gender of the talker who produces them.
 - Gender has fallen off the ladder of abstraction



An Experiment's Implications

- Null results are never definitive on their own.
- We are currently doing studies to make sure that the training that we did offered enough opportunities to make *any* implicit associations between speech-sound characteristics and gender.



An Experiment's Implications

- We are also examining whether listeners can learn associations between parametric phonetic characteristics and gender in this artificial-language learning task.



Definitely a
man



Definitely
a woman





A Call to Arms (or Brains?)

- Conclusion? No, no conclusion quite yet.
- Rather, this hopefully sets an agenda. A serious investigation of the extent to which socioindexical knowledge is relevant for phonology needs both a well-articulated model of levels of abstraction in phonology and both experiments and observational studies that can evaluate them.



Up-Front Acknowledgments

- *Major collaborators on the experiments in this presentation: **Di Schempp, Wendy Smith, Julie Johnson, Alysse Zittnan, Evan Figg, Celina Marnie***
- *This work was inspired by my teachers and colleagues, but especially: **Janet Pierrehumbert***
- *Funding: **McKnight Foundation, University of Minnesota Undergraduate Research Opportunities Program (UROP), University of Minnesota College of Liberal Arts***
- *Collaborators on pretty much everything else we work on: **Jan Edwards, Mary Beckman, Kari Urberg-Carlson***
- *Useful comments from: **Jen Hay, Paul Foulkes, Ghada Khattab, Katie Drager***
- *All remaining errors are: **mine and mine alone***



References