Modelling similarity-driven phonological generalisations

Deepthi Gopal[1], Henri Kauhanen[1 2] & Stephen Nichols[1]

[1]University of Manchester & [2]Universität Konstanz

deepthi.gopal@manchester.ac.uk, henri.kauhanen@manchester.ac.uk, stephen.nichols@manchester.ac.uk

We investigate:

i. guesses at the trajectory and actuation of a change based on a *narrow* apparent-time sample;
ii. the descriptive facts for a novel set of observations of the Turkish vowel space; mechanisms generalising a phonological rule to a class *and* producing an unequal but opposed effect in a different class;
iii. the amount of information we'd have to give the most minimal possible 'toy' iterative learner in order for it to evolve in a qualitatively convincing manner.

Some real-world data

In Turkish, an apparently categorical phonological process lowers the front mid vowels /e/ and /ø/ to [æ] and [œ] in pre-sonorant {r, l, m, n} contexts. (We focus here on the state of /e/, which is significantly higher-frequency overall and for which the data are more abundant.) **Data** (production study, n=10; all female speakers, age range 20–39, from major metropolitan areas in Turkey) show several interesting properties (for a detailed discussion, esp. of exceptionality, see Gopal & Nichols 2016 & in prep), & represent the first systematic investigation:

i. Clear bimodality in F1/F2 space: realisations of /e/ in pre-sonorant and other contexts are nonoverlapping.
ii. A further effect of coda sonority appears in pre-sonorant /e/: F1[r] > F1[l] > F1[n], F1[m]
iii. Pre-obstruent /e/s are generally higher (smaller F1, larger F2) than /e/s in unchecked syllables. Pre-obstruent /ø/s are generally *lower* than /ø/s in unchecked syllables.

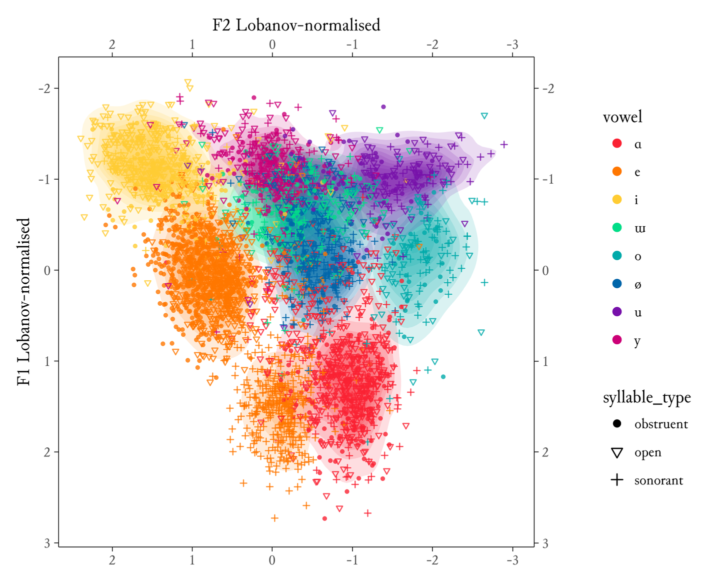

/e/ → [æ]:		
/sen/	[sæn]	'you'
/sen-in/	[se.nin]	'your'
/erdem/	[ær.dæm]	'virtue'
/gel-mek/	[gæl.mek]	'to come'
/ø/ → [œ]		
/dørt/	[dœrt]	'four'
/gør/	[gœr]	'see!'
/tørpy/	[tœr.py]	'file'
/tʃøp/	[tʃøp]	'rubbish'

Figure 1: (Lobanov-normalised) F1/F2 space across all subjects, by vowel and by coda type.

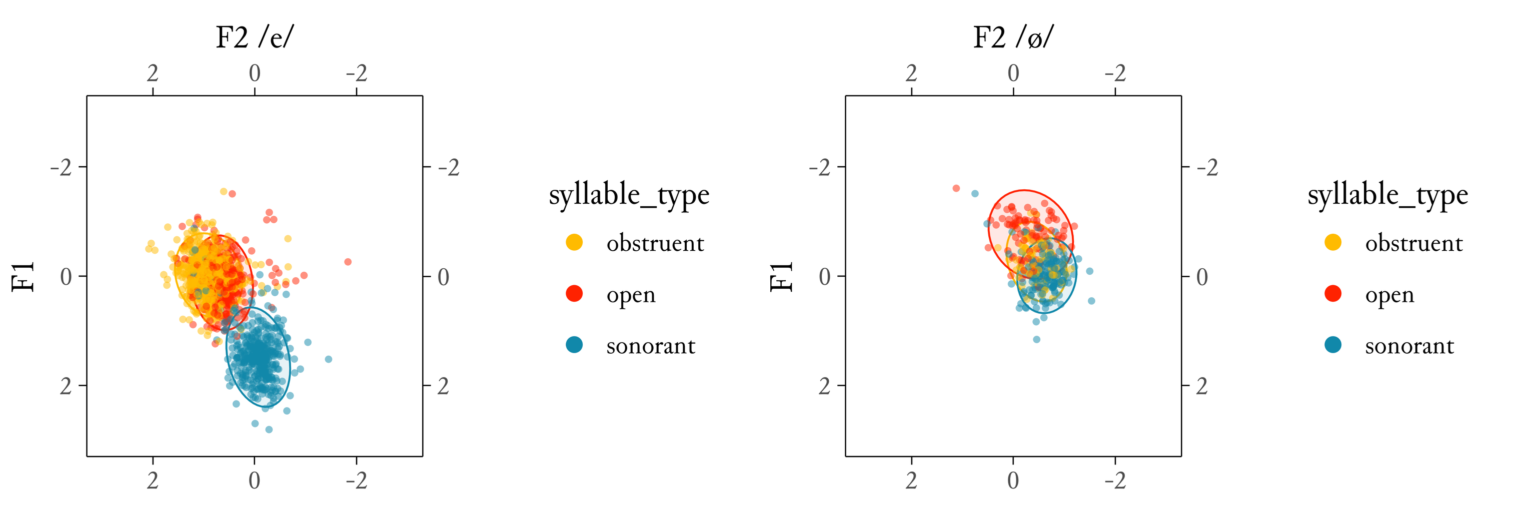

Figure 2: F1/F2 space, showing 95% confidence ellipses. *Left:* underlying /e/ (1529 tokens), by coda type. *Right:* underlying /ø/ (332 tokens), by coda type.

When we search for birth-year effects (with the necessary caveats emptor on sampling and x-axis sparseness):

i. pre-obstruent /e/-realisations begin to diverge noticeably from others with time (near-linearly);
ii. the rate of change for /e/ is quite small for other contexts; the pre-sonorant /ø/s are moving faster;
iii. for /e/, individual sonorant contexts' separation also generally time-invariant.
iv. Overall: /ø/ *is behind* /e/. Open syllable raising in /ø/ disappears in /e/; pre-sonorant /ø/-lowering is less drastic, and distributions still show overlap; the pre-obstruent /ø/ trajectory begins to diverge from the pre-sonorant one.

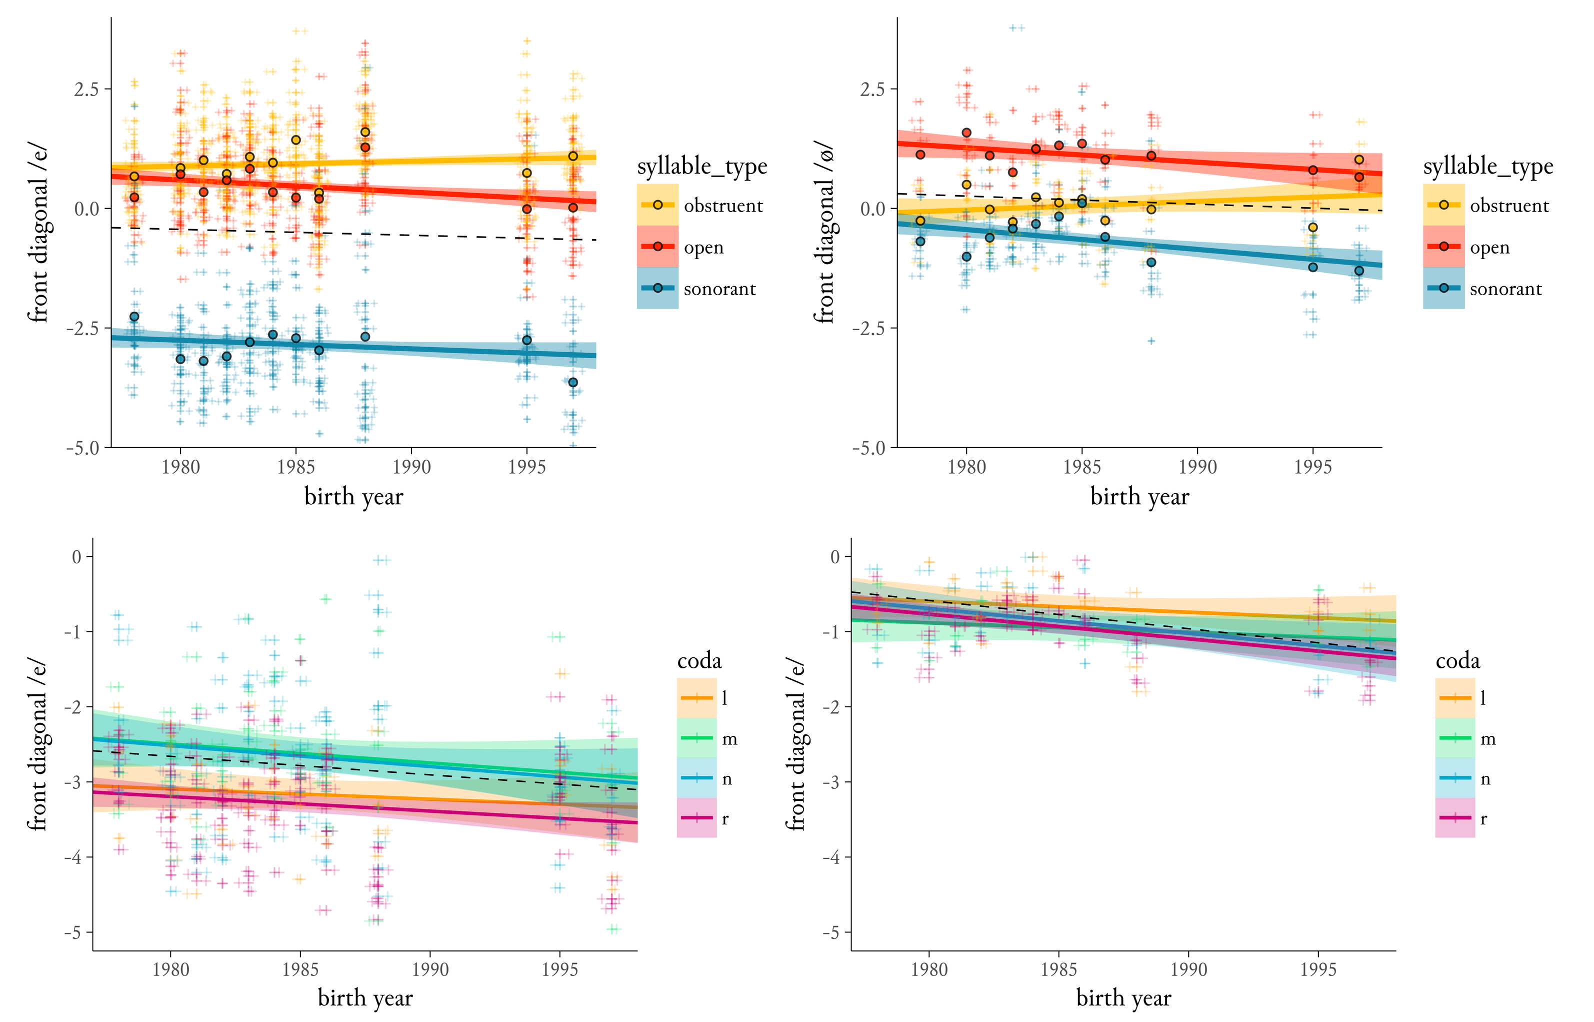

Figure 3: Trajectories across 20 years' apparent time, using F2-2∗F1 to model movement along the front diagonal of the vowel space. Dashed line tracks overall (cross-category) mean; generalised linear model fit. *Top:* left /e/, right /ø/. *Bottom:* Sonorants only, further split by individual coda; left /e/, right /ø/. NB: /ø/ tokens are not evenly distributed across coda environments – we avoid overinterpretation.

r is special: pre-/r/ tokens are consistently lower; for speakers who are far further behind in the change (see Gopal & Nichols 2016) nonetheless have pre-/r/ lowering. **Claim: mid-vowel lowering is ultimately *r-triggered*, and generalises to the sonorants.** Compare northeastern Swiss German varieties (Keel 1982; Janda & Joseph 2001, Mielke 2008), in which o → ɔ generalises from pre-r to pre-sonorant and other. **What's different here:** the further distribution by individual coda segment, and the pre-obstruent raising.

Suggestion: One way to do this: generalisation is driven by a pressure to equalise across similar environments (the set of sonorants is fairly self-similar), and a pressure to avoid dissimilar environments (obstruents and sonorants disagree). No-coda neutral – absent pressure to raise in open syllables, drift towards the mean.

Some imaginary-world data

Let's consider a learner (a speaker-listener loop) whose capacities are quite limited: the only quantity it can observe and manipulate is a parameter corresponding to vowel height $b \in [0, 1]$. On each iteration, this learner: *samples input distribution of /e/ heights → adjusts this distribution with some biases → produces output distribution*

Initialisation. Extrapolating from the production data: the 'initial state' involves some raising in unchecked syllables, and no other significant cross-context variation. Turkish-like context probabilities.

Biases. We provide our learner with:

i. a non-bias: default to the 'UR', or mean across all input tokens.
ii. a 'phonetic' bias: at each iteration, a fraction of the set of pre-*r* tokens is adjusted downward;
iii. an operational-phonological bias: for each environment, the learner wants to converge towards other similar environments, and diverge from dissimilar ones.

We (optionally; see fig. 5) introduce a further bias parameter as follows: the listener may privilege comparison against the *r* context (is more 'aware' that pre-*r* tokens constitute an independent set). Does this affect its behaviour? Converges towards extreme states faster, but qualitative behaviour isn't very different. Suggestion: a more 'realistic' learner may vary such a parameter as it perceives that context separation has increased.

An elementary similarity-biasing function

At each iteration i, the *next* height parameter $b_c[i+1]$ for context c is given by:

$$(1-\eta)UR + \eta \sum_{q \in \text{contexts}} p_q \cdot \text{sim}(c,q) \cdot b_c(1-b_c)(b_q - b_c)$$

where η determines conservativity of the learner ($\eta = 0$ blindly follows the across-context mean; $\eta = 1$ ignores it entirely), p_q is the probability that context q appears in the overall distribution of tokens, $\text{sim}(c,q)$ 'similarity score' $\in [-1, 1]$ for contexts c and q; $b_c(1-b_c)$ forces the floor and ceiling $[0, 1]$; $(b_q - b_c)$ is the difference in context heights, which we aim to maximise if dissimilar and minimise if similar.

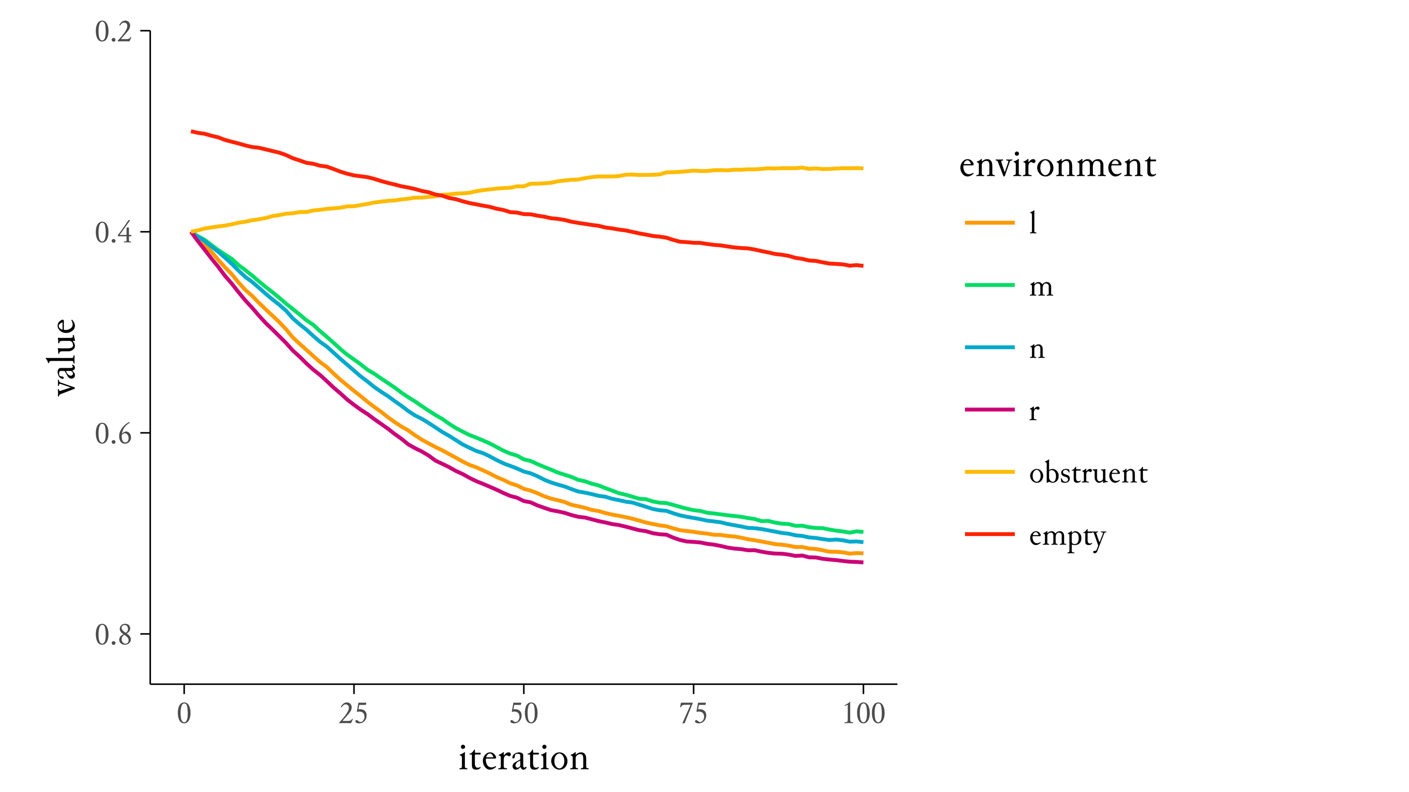

Figure 4: 100 iterations of our model, averaged over 25 runs. Obstruent environments have been collapsed; sonorant environments are shown individually.

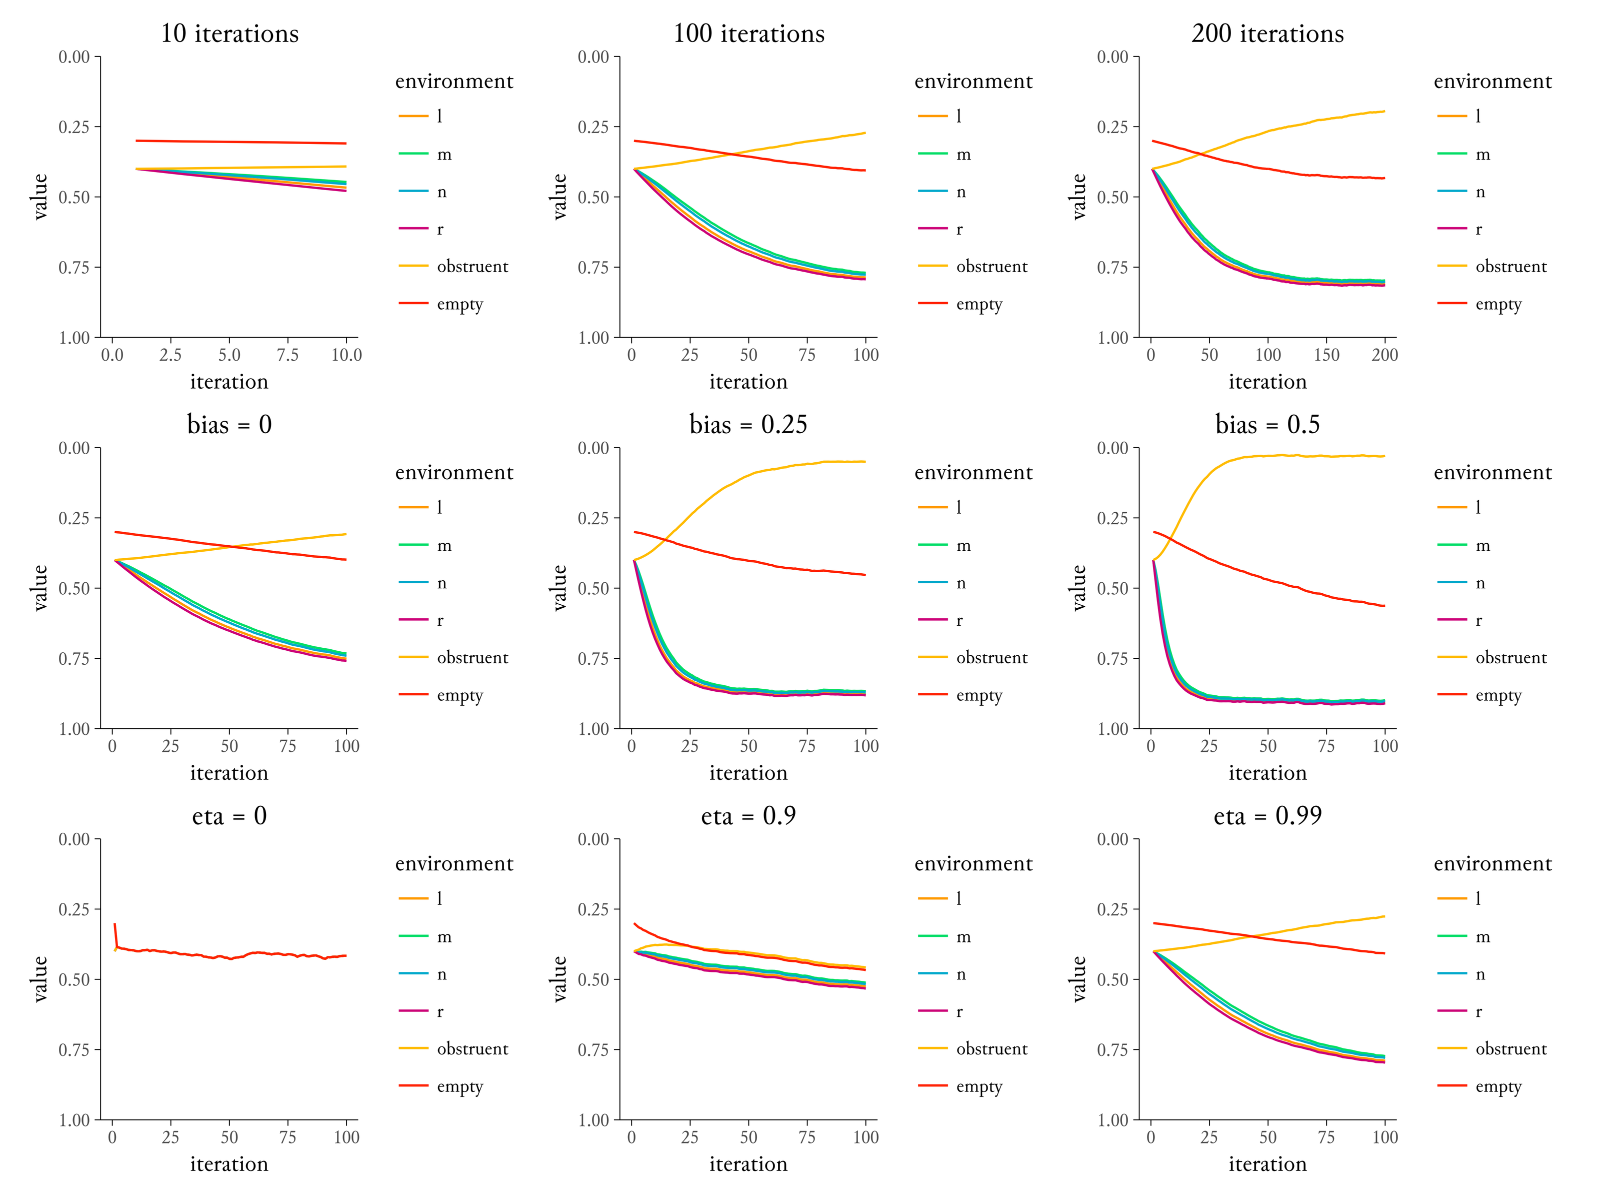

Figure 5: Model averaged over 25 runs: parameter variation. *Top row:* Evolution over 10, 100, 200 iterations. *Centre row:* Evolution for values of 'phonologised' bias towards evaluation wrt. r. *Bottom row:* Evolution for values of conservativity/towards-mean parameter η.

We:

i. have experimental evidence of a categorical lowering of /e/ (and, incipiently, /ø/) before sonorants in Turkish, and of an opposite effect of raising for obstruents;
ii. suggest aspects of the diachronic trajectory of this change that can be triangulated from a limited apparent-time sample;
iii. have presented a simple computational model of the speaker-listener loop based on assessed similarity that gives a qualitatively convincing trajectory vis-à-vis our data.

Remark. We see (in prep) inter-speaker variation on the status of coda voiced obstruents (low-frequency in Turkish – some generalise lowering to /v/, /z/ – model can account for this; see e. g. Mielke 2008, Cristia et al. 2013 on how this may support (featural?) similarity-, rather than class- generalisation.