From Abstract to Episodic:
Representations in Speech Production

Jonathan C. Paramore
jcparamo@ucsc.edu

University of Arizona
March 27, 2026

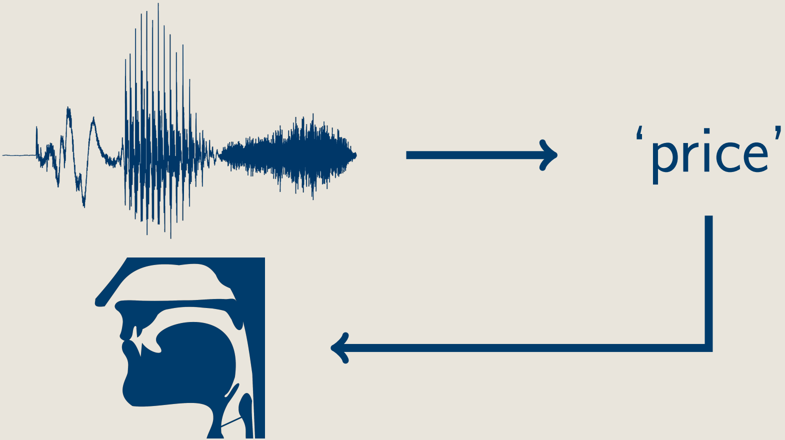

Perception and production require activating stored speech representations

What is the substance of those representations?

Representations as Episodic Memory Traces

- Individual experiences hearing and producing words retained as **exemplars** (Kruszewski, 1995; Bybee, 2001)
  - Phonetic detail is preserved (Goldinger, 1998; Clapp and Sumner, 2025)
  - Extralinguistic social information co-indexed with speech events (Johnson, 2006; Babel, 2012; Hay et al., 2019)

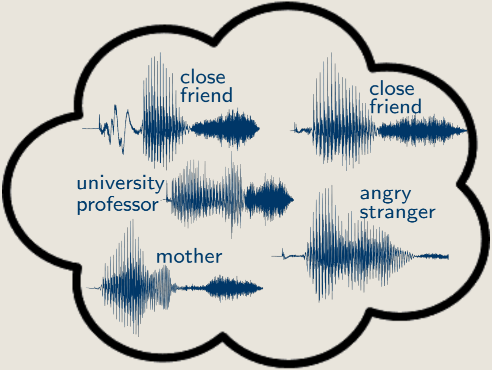

- Over time, complex network of exemplars builds up in memory
  - **Substance**: highly detailed, concrete experiences

Representations as Abstract Categories

- Evidence for abstract categories (Wedel, 2012; Pierrehumbert, 2016)
  - Why sounds across words and contexts are treated as members of the same category
  - Why regular sound changes affect the whole lexicon

- Generative approaches
  - Typical to propose highly abstract representations

    e.g., /t/ $\rightarrow$ la[ɾ]er 'latter'
  - **Covert representation**: at least one feature of the UR is never realized

- **Substance**: more general, context-free categories

This Talk

- Assumption: multiple levels of speech representation
  - Highly detailed **episodic memories**
  - Abstract **underlying representations** (URs)
- How abstract can URs be?
  - Two experiments exploring representation of vowels before nasal consonants (**pre-N vowels**) in Pakistani Punjabi
  - **Covert URs** are needed to account for pre-N vowels
- Are covert URs learnable?
  - A novel computational learner: **UR PHaSE Learner**

Background on Punjabi

- Native language of more than 78 million people in Pakistan
  - Bashir and Conners (2019)
- Approximately 33 million people in India

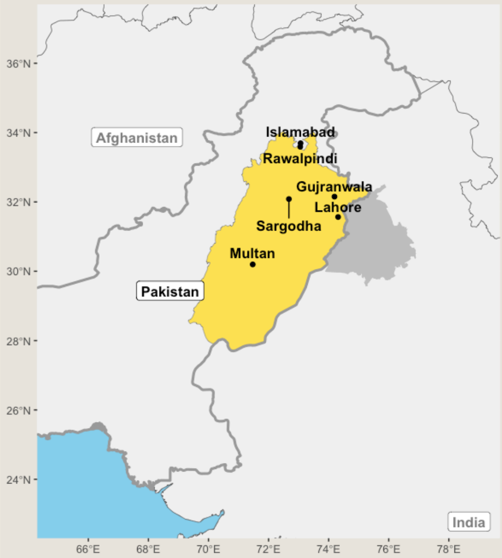

Background on Punjabi Phonology

- Long vowels **contrast in nasality**:

		Front			Central			Back		
		short	long	nasal	short	long	nasal	short	long	nasal
High	tense		iː	ĩː					uː	ũː
	lax	ɪ						ʊ		
Mid	tense		eː	ẽː					oː	õː
	lax		ɛː	ɛ̃ː	ə				ɔː	ɔ̃ː
Low	lax								ɑː	ɑ̃ː

Table 1: *Punjabi Vowel Inventory (Shackle, 2003)*

Background on Punjabi Phonology

- Nasality contrast neutralized before nasal consonants (Zahid and Hussain, 2012)

[sɑ̃ː]	'I was'	vs.	[sɑː]	'breath'
[sɑ̃ːɳ]	'grindstone'	vs.	*[sɑːɳ]	

Background on Punjabi Phonology

- Nasal vowels trigger regressive nasal harmony (Bhatia, 1993)
  - Vowels and glides (/ʋ/ and /j/) are targets
  - Other non-nasal consonants block transmission

/sɑː-ʋɑ̃ː/	→	[sɑ̃ːʋ̃ɑ̃ː]	'breath-PL'
/mətʃiː-jɑ̃ː/	→	[mətʃĩːj̃ɑ̃ː]	'fish-PL'
/ʃɑːk-ɑ̃ː/	→	[ʃɑːxɑ̃ː]	'branch-PL'

- Unclear from previous literature whether pre-N vowels trigger nasal harmony

[ɑːʋɑ̃ːm]?	or	[ɑ̃ːʋ̃ɑ̃ːm]?	'public'
[siːjɑ̃ːɳ]?		[sĩːj̃ɑ̃ːɳ]?	'recognition'

Experiment 1: How are pre-N vowels realized?

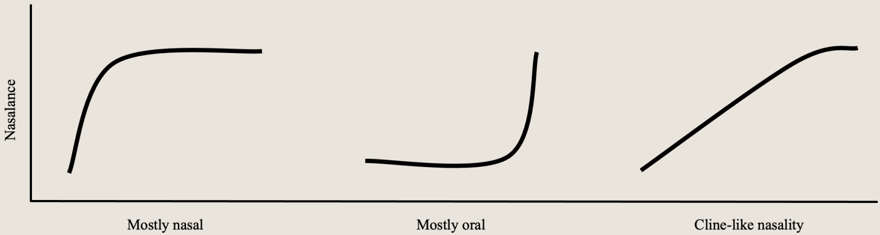

- **Nasalance**: amount of nasal airflow in the system as a proportion of the total amount of airflow

$$\frac{A^n}{A^n + A^o}$$

Experiment 1: Participants & Stimuli

- 20 native Punjabi speakers
  - 11 men and 9 women
  - Ages 22-79 (mean = 40)
- Stimuli: 63 monosyllabic tokens sourced from three conditions
  - All VːN sequences part of the same morpheme

oral Vː (26)		contrastive nasal Ṽː (20)		pre-N Vː (17)	
tʰɑː	'was'	tʰɑ̃ː	'room'	tʰɑːɳ	'piece of cloth'
seːk	'warmth'	sẽːk	'termite'		
doː	'two'			doːɳ	'cot ropes'
		pĩːg	'swing'	biːn	'musical instrument'

Table 2: Sample of words used in experiment 1

Experiment 1: Measurements

- Dual chamber oro-nasal airflow mask from Glottal Enterprises
- Measures oral and nasal airflow separately
- Outputs separate time-aligned waveforms for each cavity

Experiment 1: Procedure

- Recorded in a soundproof room at a university in Rawalpindi
- Words presented in Shahmukhi script in a randomized order:

  باں [bãː] ‘arm’

- Each word produced 4 times while wearing the mask
- Instructions given in Punjabi by a native speaker consultant

Experiment 1: Analysis

- Annotation
  - Hand-annotated in Praat
  - By me & 2 undergraduate RAs
- Measurement
  - Nasalance ($\frac{A^n}{A^n + A^o}$)
  - Measured at **11 equidistant timepoints** across each vowel
- Dataset
  - 19 of 20 participants included in the results
  - 5,062 total vowel tokens $=$
    67 words $\times$ 4 reps $\times$ 19 speakers $-$ 30 low-quality tokens

Experiment 1: Results

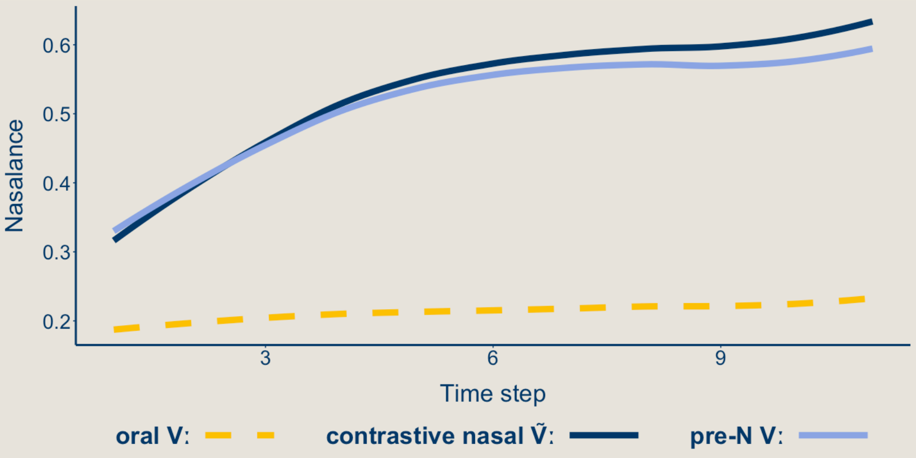

Experiment 1: Results

- **Linear mixed-effects** model revealed no significant effect of vowel type (contrastive nasal Ṽː vs. pre-N Vː) on nasalance

Comparison	β	95% CI	p
contrast. Ṽː − pre-N Vː	.016	[-.007, .039]	.17

- Lack of significance does not indicate equivalence

- Equivalence Test can confirm statistical equivalence
  1. Define a Region of Practical Equivalence (**RoPE**)
  2. Check whether the 95% CI falls entirely within the RoPE
     - **Yes**: evidence for equivalence
     - **No**: no evidence for equivalence

Experiment 1: Results

- Defining the **RoPE**
  1. Compute by-speaker nasalance mean and sd for contrastive nasal Ṽː tokens
  2. $RoPE = \pm 2sd$ of the speaker with the lowest variance

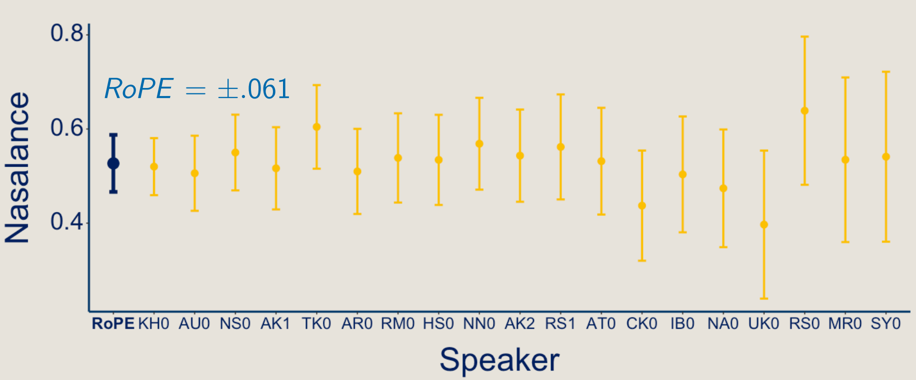

Experiment 1: Results

- Comparing the 95% CI to the $\pm .061$ *RoPE*

Comparison	β	95% CI	*p*
contrast. Ṽː – pre-N Vː	.016	**[-.007, .039]**	.17

- **Takeaway**: pre-N Vː and contrastive nasal Ṽː are statistically equivalent in nasalance

Experiment 2: Do pre-N vowels trigger harmony?

- Bhatia (1993) notes that **'nasal'** vowels trigger harmony

/sɑː-ʋãː/	→	[sãːʋ̃ãː]	'breath-PL'
/mətʃiː-jãː/	→	[mətʃĩːj̃ãː]	'fish-PL'
/ʃɑːk-ãː/	→	[ʃɑːxãː]	'branch-PL'

- Given similarity of pre-N Vː and contrastive nasal Ṽː
  - might also expect pre-N Vː to trigger harmony

[ãːʋ̃ãːm]	'public'
[sĩːj̃ãːɳ]	'recognition'

Experiment 2: Participants & Stimuli

- 16 native Punjabi speakers
  - 11 men and 5 women
  - Ages 18-43 (*mean* = 28)

- Stimuli
  - Same three conditions based on word-final vowel type
  - Di- and tri-syllabic words followed by genitive postposition, [də]
  - **VGV sequence** across final two syllables
  - All /Vː N/ sequences part of the same morpheme

oral Vː (10)		contrastive nas. Ṽː (8)		pre-N Vː (6)	
/pɑːʋeː də/	'cot leg'	/tʃɑːʋẽː də/	'pumicstone'	/ədʒəʋɛːn də/	'omum seed'
/tɑːʋuː də/	'paternal uncle'	/sɑː-ʋɑ̃ː də/	'breaths'	/ɑːʋɑːm də/	'public'
/sətɑːjiː də/	'twenty-seven'	/tiːʋĩː də/	'woman'	/geːjɑːɳ də/	'knowledge'

Experiment 2: Procedure & Analysis

- Procedure and analysis the same as experiment 1
- Removed data for 1 speaker due to poor recording quality
- 1,389 total tokens: 24 words x 4 reps x 15 speakers - 51 low-quality tokens

Experiment 2: Optional Harmony

- 44% of tokens with contrastive /Ṽː/ vowels lacked harmony

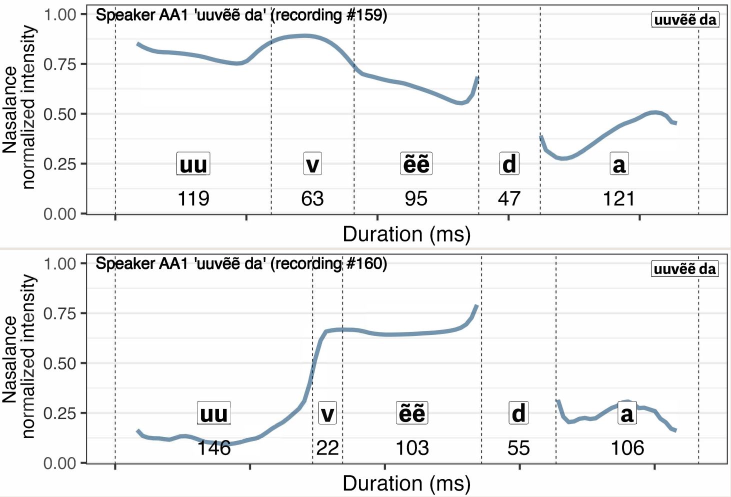

Experiment 2: Results

- 13 speakers showed a pattern of categorical pre-N nasalization

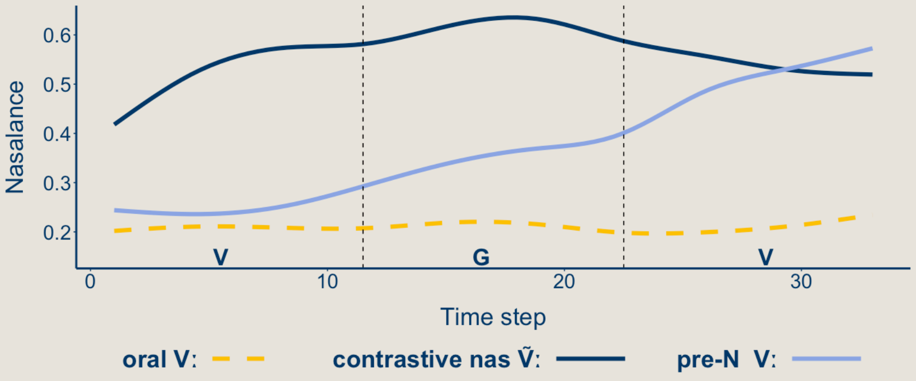

What do these two experiments tell us?

- Key result:
  - Pre-N Vː and contrastive nasal Ṽː produced with statistically **equivalent nasalance**
  - They **pattern differently in triggering harmony**
- Both harmony and pre-N nasalization are productive
  - Informal wug tests indicate patterns extend to novel forms
  - Not an artifact of historical change
- Puzzle: Why do phonetically similar vowels behave differently?

Generative Analysis: Abstract URs

- Generative analysis relies on abstract representations
  - **Covert URs**: URs with a feature never realized faithfully
  - Pre-N vowels underlyingly oral
  - Pre-N nasal assimilation covers up the oral representation
  - Nasal harmony only triggered by underlyingly nasal vowels

Generative Interpretation: Abstract URs

	a. /sɑːʋɑ̃ː/	b. /ɑːʋɑːm/
Nasal Harmony: **[-cons] $\rightarrow$ [+nas]/_$\mathbf{X}_0$ $_{\text{[-cons]}}\mathbf{V}_{\text{[+nas]}}$**	sɑ̃ːʋ̃ɑ̃ː	–
pre-N Assimilation: **V $\rightarrow$ [+nas]/_$\mathbf{C}_{\text{[+nas]}}$**	–	ɑːʋɑ̃ːm
	[sɑ̃ːʋ̃ɑ̃ː]	[ɑːʋɑ̃ːm]

Table 3: Pre-N assimilation counterfeeds nasal harmony in Punjabi

Alternatives to an abstract analysis

- The generative analysis is quite a strong claim
  - Posits that learners acquire representations they never experience
- Are there effective alternatives?
  - I argue that any surface-true analysis is inadequate

Alternative: Complex Co-Occurrence Constraint

- A more complex co-occurrence constraint transparently explains the harmony patterns (Albright, 2002)

/ɑːʋɑ̃ːm/	*[G̃Ṽːɴ]	HARMONY
→ a. ɑːʋɑ̃ːm		**
b. ɑ̃ːʋ̃ɑ̃ːm	*! W	L

- Two issues with this approach:
  1. **Computational**
  2. **Pathological Predictions**

Computational Costs

- Learning segmental co-occurrence constraints becomes intractable quickly (Hayes and Wilson, 2008, p.391-2)

$$*\begin{bmatrix} \pm Feat1 \\ \pm Feat2 \\ \vdots \end{bmatrix}\begin{bmatrix} \pm Feat3 \\ \pm Feat4 \\ \vdots \end{bmatrix}\cdots\begin{bmatrix} \pm Feat5 \\ \pm Feat6 \\ \vdots \end{bmatrix}$$

- Learning co-occurrence constraints with >2 natural classes requires strong restrictions on constraint types
  - Otherwise, set of possible constraints is too large

- Permitting constraints like *[G̃Ṽ:N] dramatically increases the number of possible constraints the learner must consider

Pathological Predictions

- Complex co-occurrence constraints risk overfitting the lexicon
  (Chomsky and Halle, 1965; Hayes and White, 2013; Wilson and Gallagher, 2018)
  - Grammar should rule out **principled** phonotactic Gaps: e.g., ill-formed *bnick* [bnɪk] in eng.
  - Grammar should not rule out **accidental** phonotactic Gaps: e.g., unattested but acceptable *blick* [blɪk] in eng.
- Permitting complex constraints like *[G̃Ṽː N] will generate many constraints that are surface-true by chance
  - More complex structures are more likely to be accidentally absent from the lexicon

Conclusion from Nasality in Punjabi

- No clear way to model the Punjabi nasality data without appealing to **covert URs**
- **How abstract can URs be?**
  - They can contain covert features that are never experienced by the language user

Abstraction & Learnability

- Nasality in Pakistani Punjabi motivates Covert URs
- But are covert URs **learnable**?

Abstraction raises two general learnability problems (Jarosz, 2019):

1. **Preference for minimally abstract representations**
   - All else equal, learner should minimize abstraction
   - Faithful /sãːɳ/→[sãːɳ] over unfaithful /sɑːɳ/→[sãːɳ]

2. **Search Space Explosion**
   - Potential URs for [sãːɳ] are infinite with unconstrained abstraction

/sãːɳ/	/sɑːɳ/	/siːɳ/
/sĩːɳ/	/sãːɳə/	/sãːɖ/
⋮	⋮	⋮

Previous Solutions

Several learning algorithms have been proposed that minimize UR abstraction (O'Hara, 2017; Paramore, 2025)

The issue of search space tractability is largely unsolved

- Wang and Hayes (2025): **UR abstraction must be limited** to levels that permit exhaustive search of the candidate space.

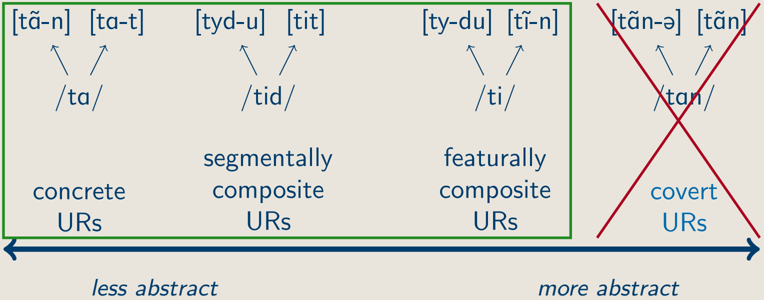

UR Progressive Hypothesis Space Expansion Learner

UR PHaSE Learner is a novel learning algorithm

- Considers UR candidates in batches, based on UR→SR mapping disparities
  - As opposed to considering all potential URs in parallel
- Stops searching for optimal solution once all observed forms are sufficiently likely
  - As opposed to exhaustively searching the space for the maximally optimal solution

UR PHaSE Learner

Four iterative learning steps

Step I: phonotactic learning
- Acquires phonotactic patterns

Step II: n-disparity UR Expectation Maximization learning
- Determines optimal UR for each morpheme
- UR set constrained to *n* disparities

Step III: Minimum Likelihood Threshold Evaluation
- Evaluates whether observed data is sufficiently likely under the current grammar
- Learning successful if likelihood of all SRs $> .95$

Step IV: UR Candidate Expansion
- Updates UR candidate space to include $n+1$ disparity URs

Step I: Phonotactic Learning

Input:

- 30 Punjabi forms from 5 categories provided to the learner
  - Morpheme boundaries labeled

i.	[sɑː]	‘breath’	ii.	[sɑ̃ː-ʋ̃ɑ̃ː]	‘breath-PL’
iii.	[gɑ̃ː]	‘cow’	iv.	[gɑ̃ː-ʋ̃ɑ̃ː]	‘cow-PL’
v.	[tɑʋɑ̃ːɳ]	‘penalty’			

- 11 constraints (see appendix for constraint definitions)
  - SPRD-L[NAS]: nasal harmony constraint
  - *VN: pre-N nasality assimilation
  - ID[NAS]/__V (Hauser and Hughto, 2020): handles opacity in parallel constraint evaluation frameworks

Goal: Calibrate constraint weights to acquire phonotactic patterns

Step I: Phonotactic Learning Results

Constraint	Type	initial w	final w
SPRD-L[NAS]	mark.	50.00	0.13
*NASOBS	mark.	50.00	100.00
*NASG	mark.	50.00	0.47
*NASV	mark.	50.00	0.00
*VN	mark.	50.00	100.00
*LOWRD	mark.	50.00	100.00

Table 4: Constraint weights during phonotactic learning

Step II: EM Learning with 0-disparity URs

1. Define the UR Hypothesis space for each morpheme

Morpheme	'breath'		'penalty'
SRs	[sɑː]	[sɑ̃ː]	[tɑːʋɑ̃ːɳ]
0-disparity URs	sɑː	sɑ̃ː	tɑːʋɑ̃ːɳ

Step II: EM Learning with 0-disparity URs

2. Use EM learning to search for the combination of URs and weights that maximize data likelihood (Jarosz, 2006a,b, 2009, 2010, 2015)

Morpheme	UR candidates	Prior $\mathbb{P}$
'breath'	/sɑː/	0.5
	/sɑ̃ː/	0.5
'penalty'	/tɑːʋɑ̃ːɳ/	1.0

E-step: Find UR probability distribution over potential URs that maximizes likelihood of the data.

M-step: Update constraint weights to maximize UR→SR mappings likelihood under posterior UR probabilities.

Step II: EM Learning with 0-disparity URs

Constraint	Type	initial w	final w
SPRD-L[NAS]	mark.	0.13	62.92
*NASOBS	mark.	100.00	100.00
*NASG	mark.	0.47	60.24
*NASV	mark.	0.00	23.16
*VN	mark.	100.00	100.00
ID[NAS]	faith.	0.00	83.44
IDFin[NAS]	faith.	0.00	89.46
ID[NAS]/_V	contfaith.	0.00	100.00
ID[RD]	faith.	0.00	6.29
*LOWRD	mark.	100.00	100.00
ID[VOC]	faith.	0.00	0.01

Morpheme	UR candidates	Prior ℙ	Posterior ℙ
'breath'	/sɑː/	0.5	0.5
	/sɑ̃ː/	0.5	0.5
'penalty'	/tɑːʋɑ̃ːɳ/	1.0	1.0

- 0-disparity UR space is too restrictive
  - Learner can't jointly model harmony in [sɑ̃ːʋ̃ɑ̃ː] and non-harmony in [tɑːʋɑ̃ːɳ]
  - Distributes probability across competing URs for 'breath'

Step III: Minimum Likelihood Evaluation

SR	Likelihood
[sɑː]	0.50
[sɑ̃ːʋ̃ɑ̃ː]	0.48
[tɑːʋɑ̃ːɳ]	1.00

Minimum Likelihood Threshold of $>.95$ **not met!**

Step IV: UR Candidate Expansion

- In an attempt to improve likelihood, learner expands UR set to include 1-disparity URs
  - If adjusting constraint weights is insufficient, the problem lies in the representations
- **Crucial**: Does not expand using all possible features
  - Candidate space would grow too large
- Exploits information from the EM learning stage
  - Eligible features must be associated with markedness constraints violated by at least one SR
    - → Markedness violations identify features active in the data
  - Flipping these features introduces new UR→SR mappings
    - → Creates new tradeoffs between markedness and faithfulness constraints

Step IV: UR Candidate Expansion

- Only constraints associated with the [NAS] feature are violated by at least one SR

Constraint	final w
SPRD-L[NAS]	62.92
*NASOBS	100.00
*NASG	60.24
*NASV	23.16
*VN	100.00
*LOWRD	100.00

Step IV: UR Candidate Expansion

- Expand UR space: create 1-disparity candidates by flipping [NAS] on eligible segments

Morpheme	'breath'		'penalty'
SRs	[saː]	[sãː]	[taːʋãːɳ]
0-disparity URs	saː	sãː	taːʋãːɳ
1-disparity URs			taːʋ̃ãːɳ, tãːʋãːɳ, taːʋaːɳ

Step II: EM Learning with 1-disparity URs

Constraint	Type	initial w	final w
SPRD-L[NAS]	mark.	0.13	48.45
*NASOBS	mark.	100.00	100.00
*NASG	mark.	0.47	48.45
*NASV	mark.	0.00	0.00
*VN	mark.	100.00	100.00
ID[NAS]	faith.	0.00	0.00
IDFin[NAS]	faith.	0.00	100.00
ID[NAS]/_V	contfaith.	0.00	100.00
ID[RD]	faith.	0.00	6.28
*LOWRD	mark.	100.00	100.00
ID[VOC]	faith.	0.00	0.01

Morpheme	UR candidates	Prior ℙ	Posterior ℙ
‘breath’	/sɑː/	0.5	1.0
	/sɑ̃ː/	0.5	0.0
‘penalty’	/tɑːʋɑ̃ːɳ/	0.25	0.0
	/tɑːʋɑːɳ/	0.25	1.0
	/tɑːʋ̃ɑ̃ːɳ/	0.25	0.0
	/tɑ̃ːʋ̃ɑ̃ːɳ/	0.25	0.0

- 1-disparity UR space resolves the learning problem
  - Successfully models harmony in [sɑ̃ːʋ̃ɑ̃ː]
  - Non-harmony in [tɑːʋɑ̃ːɳ] successfully blocked by ID[NAS]/__V

Step III: Minimum Likelihood Evaluation

SR	Likelihood
[sɑː]	1.00
[sɑ̃ːʋ̃ɑ̃ː]	0.96
[tɑːʋɑ̃ːɳ]	1.00

Minimum Likelihood Threshold satisfied!

Zooming Out

- Previous work indicates speech is stored at multiple levels of representation:
  - Phonetically/socially detailed episodic memories (e.g., Bybee 2001)
  - Abstract representational categories (e.g., Wedel 2012; Pierrehumbert 2016)
- This talk: How abstract can those categories be?
  - Some patterns require completely **covert** representations
  - With the right learning architecture, these representations are learnable

Future Directions: From Episodic to Abstract

- Abstract categories have been shown to emerge from episodic memories
  - Local interactions give rise to higher-order structure over time (Wedel, 2004, 2012)
- What drives learners to posit covert structure?
  - How can unobserved features enter the representation?

Future Directions: From Abstract to Episodic

- In speech production, how do the various levels of representation interact?

‘price’ → /pɹaɪs/ $\xrightarrow{\text{grammar}}$ [pʰɹ̥aɪ̆s]

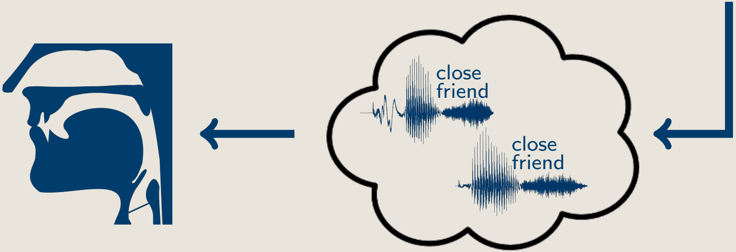

Future Directions: UR PHaSE Learner

- Does the learner generalize to other patterns requiring abstract URs? (O'Hara, 2017; Gilbert, 2023; Wang and Hayes, 2025)
- In principle, 2-disparity, 3-disparity, etc. URs are learnable
  - At deeper disparity levels, the UR hypothesis space could grow too large
  - Predicts an upper bound on learnable abstraction

Thank you!

References I

Albright, A. C. (2002). *The identification of bases in morphological paradigms*. PhD thesis, UCLA.

Babel, M. (2012). Evidence for phonetic and social selectivity in spontaneous phonetic imitation. *Journal of Phonetics*, 40:177–189.

Bashir, E. and Conners, T. J. (2019). *A descriptive grammar of Hindko, Panjabi, and Saraiki*. Mouton-CASL Grammar Series. De Gruyter Mouton.

Beddor, P. S. (2007). Nasals and nasalization: The relation between segmental and coarticulatory timing. In *Proceedings of the XVIth International Congress of Phonetic Sciences*, pages 249–254. Universität des Saarlandes, Saarbrüken, Germany.

Bhatia, T. K. (1993). *Punjabi*. Descriptive Grammars. Routledge.

Bybee, J. (2001). *Phonology and language use*. Cambridge: Cambridge University Press.

Chomsky, N. and Halle, M. (1965). Some controversial questions in phonological theory. *Journal of Linguistics*, 1:97–138.

Clapp, W. and Sumner, M. (2025). Talker-specificity beyond the lexicon: Recognition memory for spoken sentences. *Psychonomic Bulletin & Review*, 32:3201–3213.

Gilbert, M. (2023). Testing for underlying representations: Segments and clusters in Sevillian Spanish. *Natural Language & Linguistic Theory*, 42:493–531.

Goldinger, S. D. (1998). Echoes of echoes? an episodic theory of lexical access. *Psychological Review*, 105(2):251–279.

Hauser, I. and Hughto, C. (2020). Analyzing opacity with contextual faithfulness constraints. *Glossa: a journal of general linguistics*, 5(1):1–33.

Hay, J., Walker, A., Sanchez, K., and Thompson, K. (2019). Abstract social categories facilitate access to socially skewed words. *PLoS ONE*, 14(2):e0210793.

References II

Hayes, B. and White, J. (2013). Phonological naturalness and phonotactic learning. *Linguistic Inquiry*, 44:45–75.

Hayes, B. and Wilson, C. (2008). A maxiumum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry*, 39(3):379–440.

Jakubíček, M., Kilgarrif, A., Kovář, V., Rychlý, P., and Suchomel, V. (2013). The TenTen Corpus family. In *7th International Corpus Linguistics Conference CL*, pages 125–127.

Jarosz, G. (2006a). *Rich lexicons and restrictive grammars - Maximum likelihood learning in Optimality Theory*. PhD thesis, Johns Hopkins University.

Jarosz, G. (2006b). Richness of the base and probabilistic unsupervised learning in Optimality Theory. In *Proceedings of the Eighth Meeting of the ACL SPecial Interest Group on Computational Phonology*, pages 50–59.

Jarosz, G. (2009). Restrictiveness and phonological grammar and lexicon learning. In Elliot, M., Kirby, J., Sawada, O., Staraki, E., and Yoon, S., editors, *Proceedings of the 43rd Annual Meeting of the Chicago Linguistics Society*, volume 43, pages 125–134.

Jarosz, G. (2010). Implicational markedness and frequency in constraint-based computational models of phonological learning. *Journal of Child Language. Special Issue on Computational Models of Child Language Learning*, 37(3):565–606.

Jarosz, G. (2015). Expectation driven learning of phonology. University of Massachusetts manuscript.

Jarosz, G. (2019). Computational modeling of phonological learning. *Annual Review of Linguistics*, 5:67–90.

Johnson, K. (2006). Resonance in exemplar-based lexicon: The emergence of social identity and phonology. *Journal of Phonetics*, 34:485–499.

Kruszewski, M. (1883/1995). *Očrk Nauki O Jaszyke (An outline of linguistic science). Writings in general linguistics*, pages 43–173. John Benjamins.

References III

O'Hara, C. (2017). How abstract is more abstract? learning abstract underlying representations. *Phonology*, 34:325–345.

Paramore, J. C. (2025). Learning covert URs via disparity minimization. In *Eighth Annual Meeting of the Society for Computation in Linguistics (SCiL)*, volume 8.

Pierrehumbert, J. B. (2016). Phonological representation: Beyond abstract versus episodic. *Annual Review of Linguistics*, 2:33–52.

Shackle, C. (2003). *Panjabi*, chapter 16, pages 581–621. Routledge Language Family Series. Routledge.

Walker, R. (2003). *Reinterpreting transparency in nasal harmony*, pages 37–72. Amsterdam: John Benjamins.

Wang, Y. and Hayes, B. (2025). Learning phonological underlying representations: the role of abstractness. *Linguistic Inquiry*.

Wedel, A. B. (2004). *Self-organization and categorical behavior in phonology*. PhD thesis, University of California, Santa Cruz.

Wedel, A. B. (2012). Lexical contrast maintenance and the organization of sublexical contrast systems. *Language and Cognition*, 4(4):319–355.

Wilson, C. and Gallagher, G. (2018). Accidental gaps and surface-based phonotactic learning: A case study of South Bolivian Quechua. *Linguistic Inquiry*, 49(3):610–623.

Zahid, S. and Hussain, S. (2012). An acoustic study of vowel nasalization in Punjabi. *Language & Technology*, 61.

Appendix: 2 Lahori Speakers Nasality Pattern

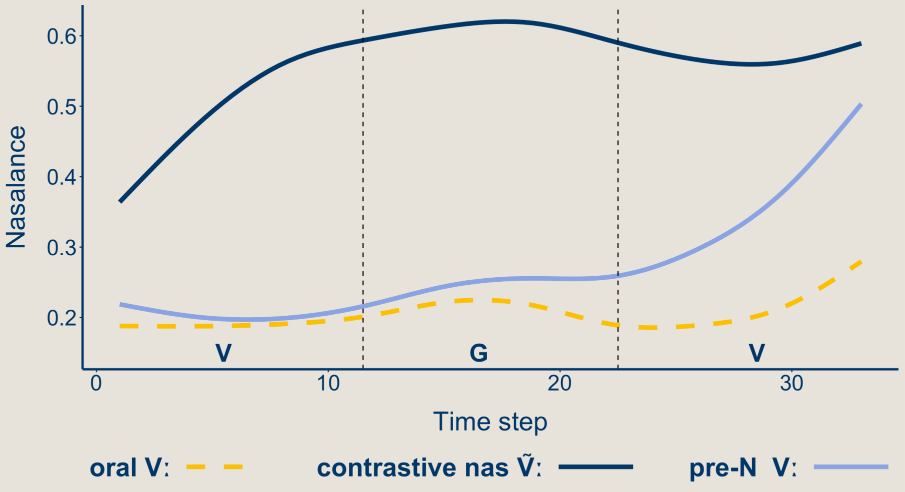

Pathological Predictions with Complex Constraints

- Corpus test: Pakistani Punjabi Phonotactics
  - **pnbTenTen corpus** (2017) (Jakubíček et al., 2013)
  - 2.8-million word corpus from 5,351 webpages
  - Analyzed 3-gram sequences over broad natural classes: N, L, O, Vː, Ṽː
- Key Finding:
  - **Principled** Gap: *[#NN]
  - **Accidental** Gap: *[LṼːL]
    - Words like [lɑːl] 'red' are abundant
    - Native speakers judge unattested [lɑ̃ːl] as acceptable

Alternative 2: A fixed nasalization window

- Harmony reflects a constant-sized nasal gesture (Beddor, 2007)
- Nasal gesture anchored to the rightmost nasal segment edge

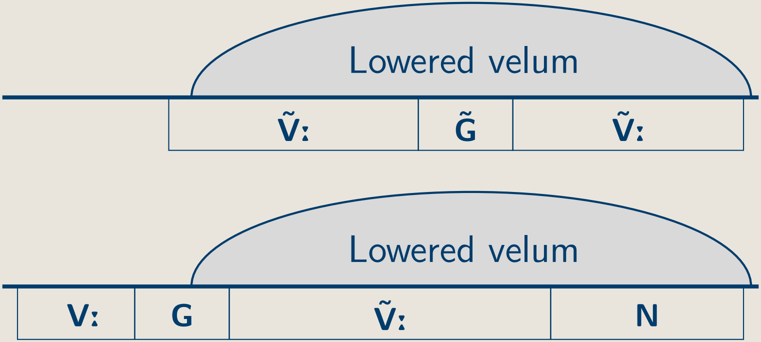

Adapted from Beddor (2007, p.249)

Alternative 2: A fixed nasalization window

Two problems with this analysis:

1. Fixed nasalization window incompatible with harmony optionality

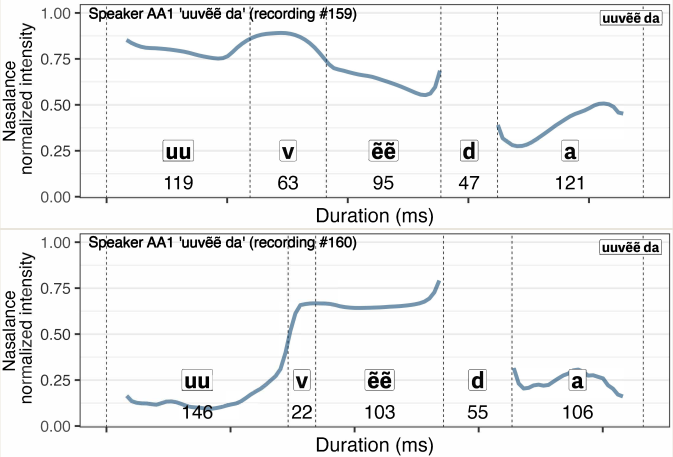

Alternative 2: A fixed nasalization window

Two problems with this analysis:

2. Incorrect predictions for contrastive nasal Ṽː=NV forms
   - Contrastive nasal Ṽː forms in exp2 also elicited with NV postpositions

contrastive nas. Ṽː=**CV** (8)		contrastive nas. Ṽː=**NV** (8)	
/tʃɑːʋẽː də/	‘pumicstone’ (GEN)	/tʃɑːʋẽː nũ/	‘pumicstone’ (ACC)
/sɑː-ʋɑ̃ː də/	‘breaths’ (GEN)	/sɑː-ʋɑ̃ː nẽ/	‘breaths’ (ERG)
/tiːʋĩː də/	‘woman’ (GEN)	/tiːʋĩː nẽ/	‘woman’ (ERG)

   - Nasalization window should shift right in presence of NV

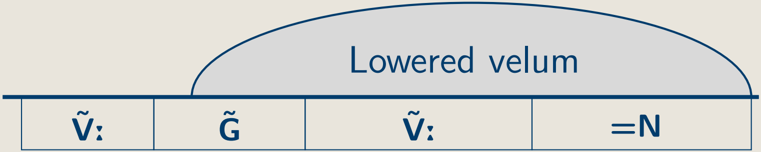

Alternative 2: A fixed nasalization window

- No shift in nasality observed

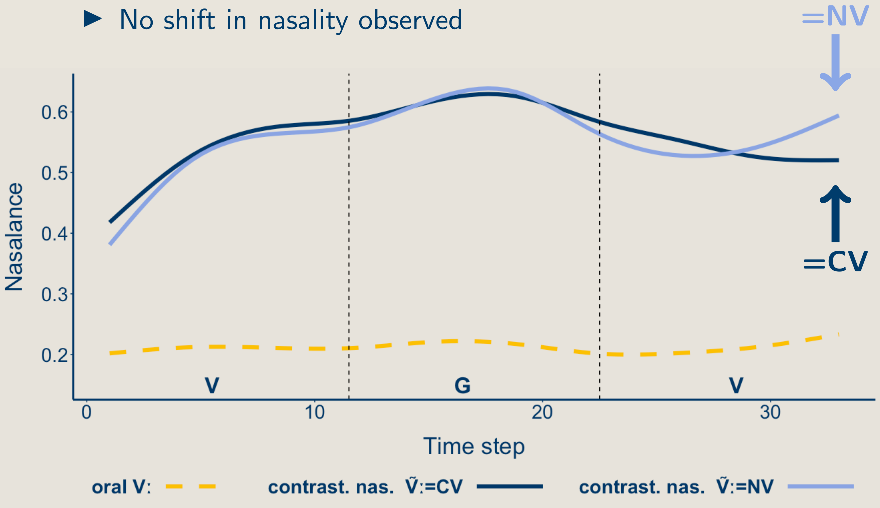

Appendix: Constraint Definitions I

i. SPRD-L[NAS] (cf. Walker, 2003, p.47)
For every occurrence of a [+NAS] feature in a prosodic word, if that [+NAS] feature is dominated by some segment, assign a violation for every segment to the left of that segment that does not dominate the [+NAS] feature.

ii. *NASOBS (Walker, 2003, p.51)
Assign a violation for every obstruent that dominates a [+NAS] feature.

iii. *NASG (Walker, 2003, p.51))
Assign a violation for every glide that dominates a [+NAS] feature.

iv. *NASV (Walker, 2003, p.51))
Assign a violation for every vowel that dominates a [+NAS] feature.

v. *VN
Assign a violation for every vowel that dominates a [-NAS] feature when directly preceding a nasal consonant.

vi. ID[NAS]
For every segment, *A*, assign a violation if the output value for the [NAS] feature dominated by *A* does not match the input value for the [NAS] feature dominated by *A*.

Appendix: Constraint Definitions II

vii. IDFIN[NAS]
For every segment, *A*, assign a violation if the output value for the [NAS] feature dominated by *A* does not match the input value for the [NAS] feature dominated by *A* in the final syllable of a prosodic word.

viii. ID[NAS]/__V
Let *A* be a segment that occurs before an oral vowel, V, in the input. Assign one violation if the output correspondent of *A* does not have the same specifications for [NAS] as *A*.

ix. ID[RD]
For every segment, *A*, assign a violation if the output value for the [RD] feature dominated by *A* does not match the input value for the [RD] feature dominated by *A*.

x. *LOWRD
Assign a violation for every vowel that dominates a [RD] feature and a [LOW] feature simultaneously.

xi. ID[VOC]
For every segment, *A*, assign a violation if the output value for the [VOC] feature dominated by *A* does not match the input value for the [VOC] feature dominated by *A*.