

Computational perspectives on syllable weight

Kaustubh Ghoshal, New York University, kg2776@nyu.edu

PhPhoNE at Yale, April 25, 2026

1 Learning weight-sensitive stress

- In Polish, primary stress always falls on the penult.

(1) Penultimate stress in Polish (Rubach and Booij, 1985; Newlin-Łukowicz, 2012)

a.	re po rtér	‘reporter.NOM’	b.	dvújka	‘two’
	re po rtéra	‘reporter.GEN’		listónof	‘mailman’
	re po rtéróvi	‘reporter.DAT’		tsudzɔzém ^j ets	‘foreigner’

- Compare this with Kabardian, where stress position is sensitive to the ‘weight’ of the final syllable. Syllables containing a long vowel (CVV) or a coda consonant (CVC) are ‘heavy’.

(2) Weight-sensitive stress in Kabardian: Stress the final syllable if it is heavy (2a), but the penult if the final is light (2b) (Gordon, 2017)

a.	təpʃég	‘plate’	b.	pá:sə	‘early’
	səbén	‘soap’		mófə	‘bear’
	sa:bí:	‘baby’		ʔedá:qə	‘rooster’
	na:nú:	‘kid’		χərzónə	‘good’

- Further, in Huallaga Quechua, only syllables with long vowels (CVV) are heavy.

(3) Weight-sensitive stress in Huallaga Quechua: Stress the final syllable if it is heavy (3a), but the penult if the final is light (3b) (Weber, 1983)

a.	uma:píta	‘from my head’	b.	ajwafá:	‘bear’
	ajwájkan	‘they are going’		ʈʂajʌará:	‘just a moment ago’
	manáray	‘not yet’		amá:	‘a prehistoric people’

- ◇ How does a learner of Kabardian or Huallaga Quechua figure out the conditions under which stress is final?

Kabardian (CVX)

$CV_\mu \quad CV_\mu C_\mu CV_\mu V_\mu$
Light Heavy Heavy

Huallaga Quechua (CVV)

$CV_\mu \quad CV_\mu C \quad CV_\mu V_\mu$
Light Light Heavy

- Analytically, heavy syllables can be represented as bimoraic. But how does the learner figure out which segments are moraic? This is a *hidden structure* problem.

1.1 Hidden structure

- For stress, making generalizations requires the assumption of higher-level structure like syllables or feet. But determining the structure may itself depend on the stress grammar.

(4) Input: *patáka*

- a. Parse 1: [(patá)ka] (Left-aligned iamb)
- b. Parse 2: [pa(táka)] (Right-aligned trochee)

- Foot type depends on the direction of alignment, and vice versa. This impasse is known as the ‘hidden structure’ problem (Tesar and Smolensky, 1998, 2000; Jarosz, 2013).
- Syllabification and syllable weight both present hidden structure problems. No reliable acoustic correlates for syllabification, numerous language-specific restrictions despite tendencies like the Sonority Sequencing Principle (Clements, 1990).

(5) Input: *pámpa*

- a. [pá_μm_μ.pa] (Stress on heavy CVC syllable, coda moraic)
- b. [pá_μm.pa] (Stress on light CVC syllable, coda non-moraic)
- c. [pá_μ.mpa] (Stress on light CV syllable, nasal-stop onset cluster)

1.2 Types of weight

Type	CVV	CVR	CVO	CV	Languages	N
CVX-heavy	H	H	H	L	Kabardian, Egyptian Arabic	42
CVR-heavy	H	H	L	L	Kwak’wala, Nuuchahnulth	3
CVV-heavy	H	L	L	L	Huallaga Quechua, Mongolian	35

Table 1: Weight-sensitive systems (Gordon, 2006)


- Not all types of weight are equally common
 - Subsequent phonetic studies have contradicted earlier descriptions for some CVV-heavy languages (Terzenbach, 2011; Lim, 2018).
 - CVR-heavy is severely underrepresented, with no phonetic support.
 - Stress sensitive to vowel quality called into question (Shih, 2018a,b; Rasin, 2018; Shih and De Lacy, 2019)
- (6) Frequency: CVX-heavy > CVV-heavy > CVR-heavy, vowel-quality-based systems
- Sensitivity to syllable *quantity* is more common than sensitivity to *quality*.
- Coarse-grained distinctions (consonant vs. vowel) more common than fine-grained (obstruent vs. sonorant, vowel quality).
- ◇ Is there a learnability explanation for this asymmetry? Coarse-grained distinctions show stronger statistical support in surface data.

1.3 Roadmap and key points

- Computational learning of phonological grammars: learning weight-sensitive stress
- Simulations with the UCLA Phonotactic Learner for different weight types:
 - ▷ *Feature-based constraint induction cannot learn all types of weight, makes the wrong predictions.*
- Alternatives using syllable structure improve learning:
 - ▷ *Learning weight requires solving for hidden structure, stress information can feed syllabification.*
- Discussion and next steps

2 Modeling phonological learning

- Certain patterns may be missing or underattested because they are difficult to learn (Stanton, 2016). Computational modeling offers an indirect, controlled way to study learnability.



- How much information is given to the learner? ~Approximates UG, less is better.
 - Constraints may be predefined with default weights
 - Input may encode structural information like syllable boundaries

2.1 Computation: Learning weight-sensitive stress

- Computational work on stress typically sidesteps the problems of syllabification and weight by encoding the input as strings of H and L (Apoussidou, 2007).
- Dresher (1999) provides an algorithm for determining syllable weight based on branching in the syllable. This assumes syllabification is universal, and has not been implemented.
- The UCLA Phonotactic Learner (Hayes and Wilson, 2008) induces constraints based on statistical distributions in the data.

- Works with realistic-sized data, and requires minimal assumptions: input strings of segments with corresponding feature values, the shape of constraints (linear n-grams of feature combinations), and optionally projections.
 - Claimed to work for weight-sensitive stress in Wargamay (CVV-heavy).
 - Can induce some hidden structure from surface data (Gouskova and Gallagher, 2020).
- ◇ Can the learner figure out syllable weight at the right edge?
- Weight in other positions expected to be harder, due to syllabification.

3 Learning stress constraints: The UCLA Phonotactic learner

- Tested the gain-based version of the learner (Wilson and Gallagher, 2018) on different types of syllable weight, using an artificial Latin dataset containing long and short vowels, and a wide range of syllable structures.
- Training data was created by adding weight-sensitive stress marking: stress the final syllable if it is heavy, else stress the penult.
- This was done in turn assuming three different types of heavy syllables: CVX, CVV, and CVR, resulting in three training sets, each with 86350 items (Table 2).
- In final position, syllable boundaries are unambiguous. Can such a feature-based induction model reliably learn weight-sensitive stress patterns?

Word	CVX-heavy	CVR-heavy	CVV-heavy
fors	fórs	fórs	fórs
detesto	detésto	detésto	detésto
arbitratus	arbitratús	arbitrátus	arbitrátus
perspectum	perspectúm	perspectúm	perspéctum
kausari:	kausarí:	kausarí:	kausarí:

Table 2: Sample words with stress marking in different training sets

- The learner was tasked with learning up to 50 constraints with maximum size of 3 (trigrams).




Figure 1: Projections for a word *kaarmageddoon*

3.1 Projections

- The learner was given a Default segmental projection with all features visible, and a vowel projection with limited features.
 - Constraints are based on linear adjacency. A constraint like *[-long][+wb] can be learnt on the Vowel projection, because the long vowel is directly adjacent to the word boundary. But this constraint would not be learnt on the other projections.
- The learner was also run with a CV projection including all segments but only limited features. In effect, the learner only sees templatic forms: #CVCVC#, #CVCVC#, etc.
 - Prosody is acquired before segmental phonotactics (Juszyk et al., 1993).
 - Stress generalizations are statistically more salient when segmental details are ignored.

<i>Projection</i>	<i>Critical features</i>	<i>Projected features</i>
Default run		
Default	none	all
Vowel	[+syll]	[syll, stress, wb]
CV run		
CV	none	[syll, stress, wb, (son)]
Vowel	[+syll]	[syll, stress, wb]

Table 3: Projections given to learner

3.2 Testing the learned grammar

- Testing dataset created by adding initial, penultimate and final stress to all words (244488 test forms). Test forms annotated as *legal* or *illegal* for each stress pattern.

Test form	CVX-heavy	CVR-heavy	CVV-heavy
detestó	illegal	illegal	illegal
detéstó	legal	legal	legal
perspectúm	legal	legal	illegal
perspéctum	illegal	illegal	legal
pérspectum	illegal	illegal	illegal

Table 4: Test forms with legality according to different patterns

- Learner assigns a harmony score to test forms, based on violations of constraints in its grammar and their weights.
 - If the learner has correctly learned the trained stress pattern, it should assign higher harmony scores to legal forms (for that pattern), than to illegal forms.
- (7) Evaluation Metric 0 (Separation): Difference between average harmony scores of legal and illegal forms (by group)

4 Results

4.1 Comparison of weight types: CVV-heavy

- (8) CVV-heavy: Stress the final syllable if it is CVV, else the penult.




Figure 2: Harmony scores by stress position and final syllable shape (CVV-heavy)

Constraint	Projection	Violations	Description
[+stress][][+syll]	Vowels	0	Stress falls within the final two syllables
[+wb][-stress][+wb]	Vowels	0	Monosyllables must be stressed
[+long,-stress][+wb]	Vowels	0	Final long vowels must be stressed
[-long,+stress][+wb]	Vowels	297	Final short vowels must be unstressed

Table 5: Intended constraints for CVV-heavy stress pattern

Constraint	Projection	Weight	Description
[+stress][][+syll]	Vowels	20	Stress falls within the final two syllables
[-stress][-long][+wb]	Vowels	20	Stress the penult if the final vowel is short
[+wb][-stress][+wb]	Vowels	20	Monosyllables must be stressed
[+stress][+stress]	Vowels	14.10	No stress clash
[+long]	Vowels	8.59	No long vowels
[-long,+stress][+wb]	Vowels	3.48	Final short vowels must be unstressed

Table 6: Constraints learned for CVV-heavy stress pattern

(9) Summary of CVV-heavy learning:

- a. Stress within final two syllables. Penultimate stress for final CV(C) words.
- b. No constraints learned for final long vowels, poor separation for CVV-final words.

4.2 Comparison of weight types: CVX-heavy

(10) CVX-heavy: Stress the final syllable if it is CVV or CVC, else the penult.

Constraint	Projection	Violations	Description
[+stress][][+syll]	Vowels	0	Stress falls within the final two syllables
[-long,+stress][+wb]	Default	33	No stress on word-final open short vowel
[-stress][+wb]	Vowels	42,451	Stress the final vowel

Table 7: Intended constraints for CVX-heavy stress pattern




Figure 3: Harmony scores by stress position and final syllable shape (CVX-heavy)

Constraint	Projection	Weight	Description
[+stress][][+syll]	Vowels	20	Stress falls within the final two syllables
[+stress][+stress]	Vowels	20	No stress clash
[+long]	Vowels	12.85	No long vowels
[-long,+stress][+wb]	Default	8.89	No stress on word-final open short vowel
[-stress][+wb]	Vowels	8.653	Stress the final vowel
[-stress][-long][+wb]	Vowels	7.774	Stress the penult if the final vowel is short

Table 8: Constraints learned for CVX-heavy stress pattern

(11) Summary of learning for CVX-heavy:

- Stress within final two syllables
- Stress dispreferred on word-final CV
- Stress preferred on word-final CVV
- For final-CVC words, final stress incorrectly penalized, poor separation

4.3 Evaluating learning quantitatively

(12) Evaluation Metric 1 (Categorization): Select X highest-scored forms, where X is the number of legal forms in the testing data. This approximates the forms categorized as legal by the learner.

▷ *How much of the learner's selection is actually legal?*

- (13) Evaluation Metric 2 (Prediction): For differently stressed forms of the same word, *what is the relative probability the learner assigns to the legally stressed form?* ($P \propto e^H$)

$$P(\text{legal stress for word}) \simeq \frac{P(\text{legal})}{P(\text{initial}) + P(\text{penultimate}) + P(\text{final})}$$

System	Projections	Probability(choose legal)			%legal
		CV	CVC	CVV	
CVV	Default	100	100	91.1	96.2
CVX	Default	99.8	92.9	100	68.4

Table 9: Evaluating the results of learning: CVV and CVX

4.4 Suppressing segmental information

- Segmental information was suppressed by restricting visible features, learner is effectively given templatic CVCV representations.
- Learning improves for both CVV and CVX on the prediction metric.




Figure 4: Effect of removing segmental information

Coding	Projection	Probability(choose legal)			%legal
		CV	CVC	CVV	
CVV	Default	100	100	90.5	93.5
CVV	CV	100	100	99.9	83.6
CVX	Default	99.8	92.9	100	68.4
CVX	CV	99.9	99.3	98.9	67.5

Table 10: Evaluation of results: Effect of removing segmental information

4.5 Coda sonority: CVR-heavy

(14) CVR-heavy: Stress the final syllable if it contains a long vowel or a sonorant coda, else stress the penult.

- For CVR-heavy, poor learning in the Default run.
- Learning improves when phonotactic information is suppressed.




Figure 5: CVR-heavy, with and without segmental information

System	Projections	Probability(choose legal)				%legal
		CV	CVO	CVR	CVV	
CVR	Default	90.5	87.0	44.2	63.7	57.5
CVR	CV	99.9	94.2	97.0	84	69.7

Table 11: CVR-heavy learning results

4.6 Interim Discussion

- The learner does best in situations where features on a single segment determine weight, like [+long] for CVV-heavy.
 - In H&W's Wargamay, *all* long vowels are stressed, so the constraint *[+long,-stress] has strong statistical support.
 - *Any* feature will do, no principled reason why [+long] and not, say, [+high] or [+back].
- CVX requires checking multiple segments, learner does not encode structure.
- CVX (following [-syll]) and CVR (following [-syll,+son]) are roughly equal in complexity.

- (15) a. Predicted frequency of occurrence: $CVV > CVX \simeq CVR$
b. Actual frequency of occurrence: $CVX > CVV \gg CVR$

- The MaxEnt/UCLA induction learner does not predict the right typology.
- The hidden structure problem of syllable weight cannot be solved by brute force generalization over linear representations, even in final position.

5 Syllable structure considerations

- ◇ Coarse-grained distinctions have more statistical support, are easier to learn (Pierrehumbert, 2001; Gouskova, 2025).
 - In syllabic terms, CVX is a more coarse-grained filter (rime has two segments).
 - CVV is more fine-grained (nucleus has two segments), proper subset of CVX.
- Does the learner fare better when syllabic information is encoded?

5.1 Moraic feature

- Every vowel and the first coda consonant was marked with [+mora].
- Learner was given moraic representations with only [+mora] segments, all other segments are effectively invisible.

<i>Projection</i>	<i>Criterial features</i>	<i>Projected features</i>
Default run		
Default-Mora	[+mora]	all
Vowel	[+syll]	[syll, stress, wb, longV]
CV run		
CV-Mora	[+mora]	[syll, stress, wb, longV, (son)]
Vowel	[+syll]	[syll, stress, wb, longV]

Table 12: Projections given to learner

System	Projections	Probability(choose legal)			%legal
		CV	CVC	CVV	
CVX	Default-Mora	99.9	92.9	100	83.4
CVX	CV-Mora	99.9	99.9	100	91.7
CVV	Default-Mora	100	100	100	99.9
CVV	CV-Mora	100	100	99.0	99.6
CVR	Default-Mora	99.8	86.4, 79.6	97.0	75.5
CVR	CV-Mora	99.9	97.6, 97.8	98.8	95.3

Table 13: Evaluating results of adding moraic features

- Learner performs better than before, especially when segmental information is suppressed.
- Learner posits constraints on vowel quality-coda interactions for stress like ‘do not stress a high vowel before a sonorant coda’ or ‘do not stress a low vowel before a coda’, potential for overgeneration.
- CVR results are slightly worse than CVV and CVX.

5.2 Marking syllable quantity on the vowel

- The vowel was marked with features [rime2], [nuc2] for whether the syllable containing it had more than one rime segment (VV or VC), and more than one nuclear segment (VV).
- The Vowel projection serves as a proxy for a syllable layer.
- Results show CVX is learned best, and CVR is worst, in line with the typology.

<i>Projection</i>	<i>Criterion features</i>	<i>Projected features</i>
CV	none	[syll, stress, wb, nuc2, (son)]
Vowel	[+syll]	[syll, stress, wb, rime2, nuc2]

Table 14: Projections given to learner

System	Projections	Probability(choose legal)			%legal
		CV	CVC	CVV	
CVX	CV	100	100	100	93.5
CVV	CV	100	100	99.8	90.0
CVR	CV	100	95.8, 99.3	98.3	77.3

Table 15: Evaluating the results of adding syllable structure features

- Constraints can target CVX with [+rime2], CVV and CVC are not disjoint.
- Learner posits constraints where shape of adjacent syllable affects stress like ‘Stress a CVX syllable if preceded/followed by a CV syllable’, potential overgeneration.

5.3 Summary

- The learner shows improved results when given syllable structure features, learning for different types is closer to attested typology.
- However, this featural proxy for syllable structure and stress is only a temporary solution.
 - Syllabic structure is not always directly reflected in vowel phonetics. And if these are indeed vocalic features, why are only these features privileged?
 - Stress affects the whole syllable including onsets and codas. But stress *position* does not seem to be sensitive to fine segmental details. This is an asymmetry not captured by featural constraints.
 - Too many segment types (moraic and non-moraic versions of all consonants, different vowel categories for CV, CVC, CVV) becomes computationally difficult for the learner.

- Where does the learner get the syllable structure feature values from?
 - Outside final position, this varies by language, and may require information on stress position to determine.

(16) Stress in Misantra Totonac (MacKay, 1991)

a.	kís.pa	‘corn kernel’	b.	ʃti.nì: tá:	‘ugly’
	sná.pu	‘gnat’		tʃíʃ.kùʃ.kán	‘her son’
	tá.ku	‘woman’		ìʃ.tá: ta	‘her father’

- This language allows [ʃt] in onsets (*ʃti.nì: tá:*), but stress suggests intervocalic [ʃt] is heterosyllabic. This helps generalize that all CVX syllables are stressed.
 - Otherwise, stress in *ì.ʃtá: ta* would unexpectedly fall on an open syllable.
- Syllabification and stress need to be in a feedback loop.
- Ongoing work: addressing this hidden structure problem using multilayered representations, allowing for encoding of syllabic structure.

6 Conclusion

- ◇ Learning weight-sensitive stress presents a hidden structure problem.
- ◇ Constraint induction models based on linear featural representations can best handle certain kinds of weight, and the predictions are inconsistent with the typology.
- ◇ Syllabic structure correctly predicts the most frequent type (CVX) makes the most coarse-grained distinction—quantity of the rime.
- ◇ Determining syllable structure and weight type at the same time as stress is a hidden structure problem and requires mutual feedback. This is part of my ongoing dissertation work.

References

- Apoussidou, Diana. 2007. The learnability of metrical phonology. Doctoral Dissertation, Universiteit van Amsterdam, Netherlands.
- Clements, G. N. 1990. The role of the sonority cycle in core syllabification. In *Papers in Laboratory Phonology*, ed. John Kingston and Mary E. Beckman, 283–333. Cambridge University Press, 1 edition.

- Dresher, Bezalel Elan. 1999. Charting the Learning Path: Cues to Parameter Setting. *Linguistic Inquiry* 30:27–67.
- Gordon, Matthew. 2006. *Syllable weight: phonetics, phonology, typology*. New York: Routledge.
- Gordon, Matthew. 2017. Syllable weight: A typological and theoretical overview. In *Current Issues in Linguistic Theory*, ed. Paul Newman, volume 338, 27–48. Amsterdam: John Benjamins Publishing Company.
- Gouskova, Maria. 2025. Phonological selection in small sublexicons. *Proceedings of Annual Meeting on Phonology 2023-2024* .
- Gouskova, Maria, and Gillian Gallagher. 2020. Inducing nonlocal constraints from baseline phonotactics. *Natural Language & Linguistic Theory* 38:77–116.
- Hayes, Bruce, and Colin Wilson. 2008. A Maximum Entropy Model of Phonotactics and Phonotactic Learning. *Linguistic Inquiry* 39:379–440.
- Jarosz, Gaja. 2013. Learning with hidden structure in Optimality Theory and Harmonic Grammar: beyond Robust Interpretive Parsing. *Phonology* 30:27–71.
- Juszyk, Peter W., Angela D. Friederici, Jeanine M.I. Wessels, Vigdis Y. Svenkerud, and Ann Marie Juszyk. 1993. Infants’ sensitivity to the sound patterns of native language words. *Journal of Memory and Language* 32:402–420.
- Lim, Keh Sheng. 2018. The Tonal and Intonational Phonology of Lhasa Tibetan. Doctoral Dissertation, Université d’Ottawa / University of Ottawa.
- Newlin-Łukowicz, Luiza. 2012. Polish stress: looking for phonetic evidence of a bidirectional system. *Phonology* 29:271–329.
- Pierrehumbert, Janet. 2001. Why phonological constraints are so coarse-grained. *Language and Cognitive Processes* 16:691–698.
- Rasin, Ezer. 2018. Modular Interactions in Phonology. Ph.D., Massachusetts Institute of Technology, Cambridge, MA.
- Rubach, Jerzy, and Geert E. Booij. 1985. A grid theory of stress in polish. *Lingua* 66:281–320.
- Shih, Shu-Hao. 2018a. Non-Moraic Schwa: Phonology and Phonetics. Doctoral Dissertation, Rutgers University, New Brunswick, NJ.
- Shih, Shu-hao. 2018b. On the existence of sonority-driven stress in Gujarati. *Phonology* 35:327–364.
- Shih, Shu-hao, and Paul De Lacy. 2019. Evidence for Sonority-Driven Stress. *Catalan Journal of Linguistics* 18:9.
- Stanton, Juliet. 2016. Learnability Shapes Typology: The Case of the Midpoint Pathology. *Language* 92:753–791.
- Terzenbach, Lauren M. 2011. Malayalam prominence and vowel duration : listener acceptability. MPhil thesis, The University of Texas at Austin.
- Tesar, Bruce, and Paul Smolensky. 1998. Learnability in Optimality Theory. *Linguistic Inquiry* 29:229–268.
- Tesar, Bruce, and Paul Smolensky. 2000. *Learnability in Optimality Theory*. The MIT Press.
- Weber, David John. 1983. A Grammar of Huallaga (Huanuco) Quechua. Ph.D., University of California, Los Angeles.
- Wilson, Colin, and Gillian Gallagher. 2018. Accidental Gaps and Surface-Based Phonotactic Learning: A Case Study of South Bolivian Quechua. *Linguistic Inquiry* 49:610–623.