

## Research Statement - Charlie O'Hara - November 12, 2021

A long-standing goal of phonological theory has been to develop a model of grammar that can divide the set of logically possible phonological patterns into those that are attested, and those that are unattested. However, in reality, the attestation of patterns is far more gradient—patterns range from very common to vanishingly rare. My research focuses on how the learnability of different patterns can impact their gradient attestation. By developing an integrated model of grammar and learnability, my work has cast light in three major areas of phonological theory. First, my integrated model is able to match observed patterns of gradient attestation in typology. Second, by splitting the responsibility of explaining typological gaps between learning and grammar, I show that simpler models of grammar can be used without losing the restrictiveness of a more complicated model of grammar. Third, learning offers a pipeline through which factors typically thought of as grammar-external, such as lexical frequency, can interact with the grammar to affect typology, without needing to be explicitly encoded in the grammar itself.

My work shows the ways in which simple models of grammar, learning, and the lexicon interact to bias the learner away from certain patterns and towards others. I analyze the formal properties of such a bias, through mathematical and computational methods, but I also evaluate these formal predictions by collecting and interpreting new empirical data: from typological surveys, lexical corpus development and analysis, and longitudinal child production analysis. I believe that computational methods and corpus analysis can be used to find new generalizations in understudied languages, and have done corpus work on a variety of languages of the Americas (e.g. Klamath, Navajo) and Finno-Ugric languages (e.g. Finnish, Estonian).

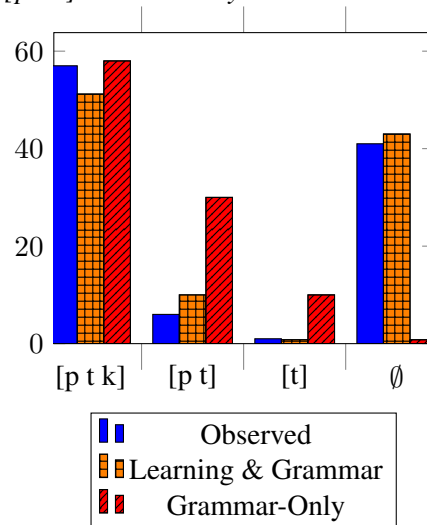
### Soft Typology—Learning of Grammar Skews Typology

In my dissertation *Soft Biases in Phonology: Learnability meets Grammar* (O'Hara, 2021), and a paper resubmitted to *Natural Language and Linguistic Theory* (O'Hara, resubmitted), I investigate the gradient typology of positional restrictions of place of articulation in stops. Many languages allow a three-way contrast of stops word-initially ([p t k]), but many show a restricted set of stops word-finally: Apalai (Cariban, Brazil) allows no stops with supralaryngeal place of articulation word-finally ([\*p \*t \*k]); Finnish (Uralic, Finland) allows only coronal place of articulation word-finally ([\*p t \*k]); Kiowa (Tanoan, United States) allows coronal and bilabial stops, but not dorsal word-finally ([p t \*k]); and Cebuano (Austronesian, Phillipines) allows all three word-finally.

I performed a large-scale typological survey of 94 genetically diverse languages that have a three way place contrast word-initially, and found a significant skew in the rate of attestation of these patterns. Languages are much more likely to allow either all or none of the stops available word-initially in word-final position, than they are to allow only some of the stops word-finally. I show that a model of learning and grammar closely matches this observed gradient typology, but previous grammar-only models of gradient typology diverge significantly (see Figure 1).

**Grammar-External Factors Influence Typology** A gradient perspective on typology can reveal that typically grammar-external factors such as the lexical distribution of a language can have a significant effect on the relative attestation of different patterns. In grammar-only theories of typology, it is difficult to encode for

Figure 1: *Soft typology of languages with [p t k] word initially*



these sorts of language-specific properties to have a major effect on the space of possible grammars without making use of language-specific constraints. In an integrated model of learning and typology, learning offers a pipeline through which grammar-external factors can project onto the typology. These language-specific properties impact the distribution of data that a learner observes when trying to learn a specific language, impacting the types of patterns that would be easy or hard to learn. So far, I have focused on lexical frequency as such a language-specific feature. In my dissertation, I identify four former principles that explain how different lexical frequency distributions affect the learnability of phonological forms and patterns containing them. These four principles can subvert more general differences in learnability across patterns. For example, in chapter four my dissertation, currently in preparation for journal submission, and several conference presentations (O’Hara, 2018a,b), I look at how different lexical frequency distributions affect the learnability of particular patterns of place of articulation contrast in initial vs. final voiceless stops. Of particular interest, Finnish (and other closely related Finno-Ugric languages) exhibit the extremely rare pattern of banning all but coronal stops word finally. I argued that the cross-typological rarity of this pattern was based on the general difficulty of learning this pattern—but Finnish learners have been able to maintain this pattern for five centuries at least. I show that looking across the space of possible lexical distributions, the coronal-only pattern represented by Finnish is only likely to be learned under very specific lexical conditions, predicting that this pattern should be very rare. The Finno-Ugric languages exhibit precisely the specific lexical conditions that condition learnability of the coronal-only pattern (high frequency of final [t], but low frequency of final [k]), explaining why this family is able to exhibit the pattern.

In ongoing work (O’Hara, 2020c,a), I examine a different phonological domain—contour tone licensing. Some languages, like Thai (Kra-Dai, Thailand) restrict contour tones only to syllables that do not have obstruent codas, regardless of vowel length, but other languages, like Navajo (Athabaskan, United States) restrict contour tones to syllables with long vowels, regardless of what appears in the coda of the syllable. This can be seen by looking at short sonorant final syllables, and long obstruent final syllables in each language: in Thai CVR syllables can have contour tones ([lǎŋ] ‘back’), but CV:O syllables cannot (\*[lǎ:k]); whereas Navajo shows the opposite pattern (CV:O: [t̤ɛ̄ɹ̄ɹ̄n̄:ɬon] ‘they shot at him’; CVR: \*[p̄ik̄h̄in]). I show that the lexical frequencies in these languages differ in a way associated with the relative markedness of the different types of contour toned syllables. Learning simulations show that languages with more short syllables than obstruent-final syllables are more likely to ban contour tones in obstruent final syllables like Thai; whereas languages with more obstruent-final syllables are more likely to ban contour tones in short syllables, like Navajo.

In future work, I plan to expand these learning models to investigate other factors that may affect soft typology such as asymmetries in channel bias (errors in transmission of phonetic forms from teacher to learner). Previous work (Jun, 1995) has suggested that language-specific aspects of the phonetic realization of syllable duration appear to correlate with patterns of contour tone licensing, with contour tones more likely to be licensed on syllables of longer duration in the language. I plan to test how these differences in phonetic realization would impact the learnability of patterns, by implementing a model of misperception. Contour tones are less likely to be perceived accurately on shorter duration syllables, so learners would effectively observe fewer tokens of contour tones on such syllables. As a result, asymmetries in misperception would influence learning and typology through the same mechanics as asymmetries in lexical frequencies.

**Simplifying the phonological component through emergent biases** My work focuses on the interaction of three aspects of the phonological component: the grammar, the learning algorithm, and the lexicon. Previous work has tried to capture typological restrictiveness by adding restrictions to one of these components, but I have shown through multiple projects that the interaction of the three components has a more powerful effect than previously noticed, allowing for the same typology to be captured without stipulating overt restrictions on these three systems.

THE GRAMMAR— Overgeneration has long been an argument against certain models of constraint based grammars, but in a gradient model of typology, learning provides a way to explain the unattestedness of grammatically possible patterns. In my dissertation, I show that the structure of the grammar and the learning process interact to make certain unattested patterns harder-to-learn than all attested patterns. An example of this kind is the pattern that bans dorsal stops in initial and final position, but allows coronal and labial stops in both positions. Extending this line of work, I have shown how learning and a gradient theory of typology can weigh in on debates between theories of probabilistic grammar. Maximum Entropy Harmonic Grammar (MaxEnt, Goldwater & Johnson 2003) has been shown to be simpler and more computationally tractable than competing models such as Noisy Harmonic Grammar (Flemming, 2017), but MaxEnt is conjectured to have an overgeneration not present in competing models, particularly because harmonically bounded candidates can receive probability only in MaxEnt (Jesney, 2007; Anttila & Magri, 2018). For example, MaxEnt can model an unattested, and unnatural pattern where onset consonants variably delete, along with a more natural pattern that variably epenthesiszes onset consonants, even when the set of constraints would only predict the latter in other models of grammar. In my talk at the *Annual Meeting on Phonology 2019* (O’Hara, 2019) and ongoing work, I show that grammars where harmonically bounded candidates (like onset deletion) receive significant amounts of probability are remarkably difficult to learn. As a result, the typological predictions of learning and MaxEnt together do not differ substantially from other models of grammar. Thus, even though MaxEnt alone restricts the space of possible grammars less than other more complex models, in an integrated model of learning and grammar that restriction can be found to be emergent from the simpler model of grammar.

Learning can also allow us to compare grammatical theories that appear to make similar typological predictions from the grammar alone. In an ongoing project with Caitlin Smith (Smith & O’Hara, in revision), we investigate the learnability of height harmony processes under different grammatical frameworks. Many languages show a pattern described as stepwise harmony, where low-mid vowels become high-mid /ɛCi/→[eCi], but high-mid vowels become high /eCi/→[iCi], before high vowel triggers (e.g. in Nzɛbi, Bantu: Gabon). On the other hand, no languages show a “saltatory” pattern, where low vowels become high, but mid vowels do not raise /aCi/→[iCi], but /eCi/→[eCi]. Both of these patterns can be generated under a variety of different grammatical frameworks both using features (Gnanadesikan, 1997; Hayes & White, 2015), and using gestures (Smith, 2018). We show that with gestural representations, the well-attested stepwise pattern is easier to learn than the unattested saltation pattern, but with featural representations, the saltation pattern is easy to learn. By comparing how the learning algorithm interacts with a gestural model of grammar compared to a featural model, we can better evaluate the predictions of both models. Further, in ongoing work also collaborating with Eric Rosen and Paul Smolensky, we examine how models of deep learning can offer insight about the nature of phonological representation. In Smith *et al.* (2021), we show that gestural representations seem to emerge from recurrent neural networks trained to learn the mapping between abstract morphemes and articulatory trajectories via an attention mechanism. In ongoing work in this area, we have developed an online phonological learner that combines lexical and grammatical learning, learning the gestural parameters of vowels simultaneously with grammatical aspects of the grammar such as which segments act as triggers or blockers. We plan to extend this work into learning the underlying representations in the lexicon as well.

THE LEXICON— Finding a single underlying representation (UR) for an alternating morpheme has long been a difficult challenge for phonological learners. The language Klamath (Penutian: Southern Oregon) demonstrates a pattern of alternation where in some stems [i] appears only when it is phonotactically necessary to break up strings of consonants—i.e. [w-<sup>h</sup>e:wɪtk<sup>h</sup>]-[n-<sup>h</sup>e:w\_<sup>h</sup>a](‘broken’-‘splinters on’). There are two concrete possible underlying forms for this alternating vowel—either it is underlyingly /i/ and deletes in the second word, or it is underlyingly absent /∅/ and [i] is epenthesiszed in the first word. Any other underlying representation would be *abstract*, because it does not surface faithfully in any context. In a naive model of lexical learning separated from grammar learning, there is no clear way to distinguish between the

many possible abstract URs: is it underlyingly /e/, /ɪ/, etc.? Previous work has therefore proposed an overt restriction banning abstract URs from the hypothesis space, in order to make the learning problem more tractable (Kiparsky, 1968; Albright, 2002; Allen & Becker, 2015). In my study of this problem, published in *Phonology* (O’Hara, 2017), I hand-transcribed a dictionary of the Klamath language Barker (1963) into a digital database of the Klamath lexicon (O’Hara, 2014), and found that neither concrete UR could serve for this alternation. I developed a MaxEnt learner of URs that simultaneously learns the lexicon and the grammar. In my learning simulations, I found that learners chose the abstract UR /e/ as the UR for this alternation, consistent with analytic principles like feature economy (Clements, 2003), without any overt biases built into the learner. The learning of the grammar impacted the learning of URs, such that it created a metric across the set of abstract URs that were previously indistinguishable as hypotheses, separating /e/ from the infinite space of abstract URs. Through an integrated model of grammar and lexicon learning, we can find that realistic abstract URs can be learned, but much of the abstract UR space is still untenable, without encoding any direct stipulations banning abstract URs from the hypothesis space.

THE LEARNER—My work also has shown how the interaction of the learner with the grammar and the shape of the lexicon itself results in emergent learning behavior that obviates the need for adding complications to the learning algorithm. This can be seen in my work on learning lexical idiosyncrasies and frequency matching behavior (at my poster at the *Society for Computation in Linguistics (SCiL 2020)* (O’Hara, 2020b) and ongoing work). When learning languages where a process applies at different rates for different lexical items, human language learners tend to apply that process to nonce words at a rate roughly matching the general statistical pattern in the lexicon (Hayes & Londe, 2006; Hayes *et al.*, 2009). For example, in Hungarian, the dative suffix usually appears as /nɛk/ after front vowels, though in 8% of monosyllabic stems with front unrounded vowels, the suffix appears as [nɔk], the form usually reserved for stems with back vowels. Hungarian speakers tested on nonce monosyllabic stems with front unrounded vowels use the back [nɔk] 7% of the time, closely matching the statistical pattern in the lexicon. Nonce word behavior resembles hold-out testing in statistical learning—by testing a learner on forms they were not trained on, we can observe whether the learner properly generalized. Recent work using batch MaxEnt learners (Zymet, 2018; Hughto *et al.*, 2019) show that typical assumptions about batch MaxEnt learning fail to frequency-match nonce words, so they suggest overtly biasing the learners so that they are more likely to move general constraints more than specific constraints to better capture frequency matching behavior. I show that in an online MaxEnt learner, where learners are exposed to one form at a time rather in a large batch of forms, there is an emergent bias to update general constraints more than lexically specific constraints, simply because they are relevant to a larger set of the data. I show that the size of the lexicon has a major effect when determining how closely a learner frequency matches the general statistical trend—with realistically sized lexicons frequency matching behavior is observed. By looking at the interaction of the constraint set, the lexicon shape and size, and the learner, I show that a learner can capture frequency matching without stipulating an overt bias on the learning of different types of constraints. In future work, I intend to expand this avenue of research by further studying the behavior of online learners with realistic data and attempting to characterize their emergent biases in formal mathematical terms.

## References

- ALBRIGHT, ADAM. 2002. The lexical bases of morphological well-formedness. *In*: BENDJABALLAH, SABRINA, DRESSLER, PFEIFFER, & VOEIKOVA (eds), *Morphology 2000: Selected papers from the 9th Morphology Meeting, Vienna*. Benjamins.
- ALLEN, BLAKE, & BECKER, MICHAEL. 2015. *Learning alternations from surface forms with sublexical phonology*. Ms, University of British Columbia & Stony Brook University. Available (May 2017) at <http://ling.auf.net/lingbuzz/002503>.
- ANTTILA, ARTO, & MAGRI, GIORGIO. 2018. *Comparing SHG and ME from the perspective of equiprobable mappings*. Talk given at Analyzing Typological Structure, Stanford University.

- BARKER, M.A.R. 1963. *Klamath Dictionary*. Berkeley and Los Angeles: University of California Press.
- CLEMENTS, G. N. 2003. Feature economy in sound systems. *Phonology*, **20**, 287–333.
- FLEMMING, EDWARD. 2017. *Stochastic Harmonic Grammars as Random Utility Models*. Poster presented to the 2017 Annual Meeting on Phonology, NYU.
- GNANADESIKAN, AMALIA. 1997. *Phonology with ternary scales*. Ph.D. thesis, University of Massachusetts Amherst.
- GOLDWATER, SHARON, & JOHNSON, MARK. 2003. Learning OT constraint rankings using a Maximum Entropy model. In: *Proceedings of the Workshop on Variation within Optimality Theory*. Stockholm University.
- HAYES, BRUCE, & LONDE, ZSUZSA. 2006. Stochastic phonological knowledge: The case of Hungarian vowel harmony. *Phonology*, **23**, 59–104.
- HAYES, BRUCE, & WHITE, JAMES. 2015. Saltation and the P-map. *Phonology*, **32**, 267–302.
- HAYES, BRUCE, ZURAW, KIE, SIPTAR, PETER, & LONDE, ZSUZSA. 2009. Natural and unnatural constraints in Hungarian vowel harmony. *Language*, **85**, 822–863.
- HUGHTO, CORAL, LAMONT, ANDREW, PRICKETT, BRANDON, & JAROSZ, GAJA. 2019. Learning Exceptionality and Variation with Lexically Scaled MaxEnt. In: *Proceedings of the Society for Computation in Linguistics*, vol. 2.
- JESNEY, KAREN. 2007. *The locus of variation in weighted constraint grammars*. Handout for poster presented at the Workshop on Variation, Gradience and Frequency in Phonology, Stanford University.
- JUN, JONGHO. 1995. *Perceptual and articulatory factors in place assimilation: an Optimality Theoretic approach*. Ph.D. thesis, University of California, Los Angeles.
- KIPARSKY, PAUL. 1968. *How abstract is phonology?* Indiana University Linguistics Club.
- O’HARA, CHARLIE. 2014. *Searchable Klamath Dictionary v. 0.5*. Electronic Database.
- O’HARA, CHARLIE. 2017. How abstract is too abstract: Learning abstract underlying representations. *Phonology*, **34**(2), 325–345.
- O’HARA, CHARLIE. 2018a (September). *Rare Hard-To-Learn Patterns Stably Learned Due To Language-Specific Lexical Frequencies*. Talk given at Analyzing Typological Structure: From Categorical to Probabilistic Phonology. Stanford University.
- O’HARA, CHARLIE. 2018b (October). *The Sweet Spot Effect: Rare Phonotactic Patterns Require Specific Lexical Frequencies*. Poster presented at Annual Meeting on Phonology 2018. UCSD.
- O’HARA, CHARLIE. 2019 (October). *Learning Prevents MaxEnt from Giving Probability to Harmonically Bounded Candidates*. Talk given at Annual Meeting on Phonology 2019.
- O’HARA, CHARLIE. 2020a (January). *The Effect of Learnability on Constraint Weighting: Case Study from Contour Tone Licensing*. Poster Presented at LSA Annual Meeting 2020.
- O’HARA, CHARLIE. 2020b (January). *Frequency Matching Behavior in On-line MaxEnt Learners*. Poster presented at the Society for Computation in Linguistics (SCiL 2020).
- O’HARA, CHARLIE. 2020c. Language-Specific Factors Influence Learnability: Case Study from Contour Tone Licensing. *Pages 235–245 of: ASATRYAN, MARIAM, SONG, YIXIAO, & WHITMAL, AYANA (eds), Proceedings of NELS 50*, vol. 2.
- O’HARA, CHARLIE. 2021. *Soft Biases in Phonology: Learnability meets grammar*. Ph.D. thesis, University of Southern California.
- O’HARA, CHARLIE. resubmitted. *Emergent Learning Bias and the Underattestation of Simple Patterns*. Ms. University of Southern California.
- SMITH, CAITLIN. 2018. *Harmony in Gestural Phonology*. Ph.D. thesis, University of Southern California, Los Angeles, CA.
- SMITH, CAITLIN, & O’HARA, CHARLIE. in revision. Learnability of derivationally opaque processes in the Gestural Harmony Model. *Phonology*.
- SMITH, CAITLIN, O’HARA, CHARLIE, ROSEN, ERIC, & SMOLENSKY, PAUL. 2021. Emergent Gestural Scores in a Recurrent Neural Network Model of Vowel Harmony. *Page Article 7 of: Proceedings of the Society for Computation in Linguistics*, vol. 4.
- ZYMET, JESSE. 2018. *Lexical propensities in phonology: corpus and experimental evidence, grammar, and learning*. Ph.D. thesis, University of California Los Angeles.