The Logical Structure of Phonological Generalizations

Jeffrey Heinz
Stony Brook University

ABSTRACT. Logical languages, which combine formal logics with representations, are argued to have many benefits for the description and analysis of phonological systems. They inform questions relating to locality, typology, psychology, as well as learning and acquisition, and offer insights unavailable with traditional rule-based or constraint-based phonological formalisms. This article reviews the main insights and recent developments offered by unearthing the logical structure of phonological generalizations.

Keywords: computation, logic, regular, subregular, constraints, processes

1. Introduction

In this article, I argue that the logical structure of phonological generalizations is not only "regular", but also "less than regular" in a particularly "local" way. These concepts provide a way to understand the extensive variation cross-linguistically, how these patterns can be acquired from examples, and the important role played by representation in grammar. I argue that formalizing these insights directly with logic will lead to better theories of phonology than theories based on global optimization (e.g. Prince and Smolensky 2004), or theories based on serial rule application (e.g. Chomsky and Halle 1968). The primary reason is that the hypothesis that phonological generalizations can be characterized with weak logics over particular representations is both sufficiently expressive to account for variation observed cross-linguistically and simultaneously sufficiently restrictive to provide an account for how such generalizations can be learned with limited resources. While this article overlaps with arguments I have made elsewhere (Heinz 2011a,b, 2018), it also incorporates recent research results since those pieces were written over a decade ago.

2. Background

I begin with what I regard to be the insight made in the 20th century which is foundational to generative phonology. In my view, it is that the best explanation of the systematic variation in the pronunciation of morphemes is to posit a single underlying mental representation of the phonetic form of each morpheme and to derive its pronounced variants with context-sensitive transformations. Readers are referred to Kenstowicz and Kisseberth (1979, chapter 6) and Odden (2014, chapter 5) for arguments for this position.

To illustrate, consider the following Finnish words in Table 1 (Odden 2014:86). The partitive singular suffix shows an [a/æ] alternation, but let us focus on the [e/i] alternation which occurs at the ends of some nouns, including the ones meaning 'river' and 'door.'

Table 1. Finnish Words

Nominative Singular	Partitive Singular	
aamu	aamua	'morning'
kello	kelloa	'clock'
kylmæ	kylmææ	'cold'
kømpelø	kømpeløæ	'clumsy'
æiti	æitiæ	'mother'
tukki	tukkia	'log'
yoki	yokea	'river'
ovi	ovea	'door'

A standard phonological analysis would posit that the long-term memory representation of the pronunciation of the morphemes ‘river’ and ‘door’ would be /yoke/ and /ove/, respectively and the phonological generalization that unround, mid, front vowels raise in word-final position. Formally, that transformation is expressed with the SPE-style rewrite rule e → [+high] / __ #. It is also expressed in Optimality Theory with the ranking *e# >> IDENT(HIGH), where, of all the faithful constraints whose violation could resolve a violation of *e#, IDENT(HIGH) is ranked the lowest. Despite their differences, both theories adopt the fundamental insight in the 20th century.

Any theory which agrees with the position that the best explanation for the systematic variation in the pronunciation of morphemes is due to a lawful transformation from underlying representation to the surface representation must address three questions.

1. What is the nature of the more abstract, underlying, lexical representations?
2. What is the nature of the more concrete, surface representations?
3. What is the nature of the transformation from underlying forms to surface forms?

While different theories of phonology may disagree on the answers to these questions, *they agree on the questions being asked*.

It is a truism that grammars are productive, and make predictions regarding unseen forms. It is also a truism that different grammars may generate the same transformations, like the SPE-style rewrite rule and the OT grammar above for Finnish word-final /e/ raising. Both these grammars map the URs /yoke, ove/ to SRs [yoki, ovi], respectively, and both would map possible URs like /manile/ to [manili]. Each (UR, SR) pair is a point in a space, and the set of these points – a string map — constitutes what I call the extension of the grammar, in the same way that there is a set of points that satisfy an equation for a line, such as y = 2x + 1.

Another important truism is that grammars may have properties largely independent of grammatical particulars. For example, lines are lines regardless of whether they are described with Cartesian or Polar coordinates. One foundation in computer science is that string sets and string maps can be defined equally well with different kinds of grammatical formalisms (Kleene 1956; Elgot and Mezei 1965; Scott and Rabin 1959). These ideas are not uncommon in phonology. For example, Tesar’s (2014) definition of output-driven maps does not depend on the grammatical formalism, and Baković and Blumenfeld (2024) provide an analysis of the ways string maps may interact. In the next section, I want to draw attention to an important property of phonological generalizations, regardless of how they are formalized: they are *regular*.

3. The regularity of phonology

In phonology, there are several computations that grammars engage in. For a given constraint C and representation R, we may want to know whether R violates C. We may also want to know how many times. Similarly, for a given grammar G and underlying representation R, we want to know which surface representation(s) G transforms R into? Determining the answers to these questions uses resources, and knowing what kinds of resources are necessary and sufficient to answer these questions tells us something about the given constraint or grammar.

Technically, a set of strings is *regular*, if it can be defined with a regular expression, a finite-state acceptor, or with Monadic Second Order logic over string models, and string maps are similarly defined (Kozen 1997; Engelfriet and Hoogeboom 2001). However, unless one is familiar with these terms, no knowledge is gained because one unfamiliar term has been replaced with other unfamiliar ones. Therefore, I will provide the following intuition for what it means for a grammar to be regular. *A grammar is regular provided the memory required to generate the output is bounded by a constant, regardless of the size of the input*. Bear in mind that this is not a technical definition, and this attempt to convey what *regular* means in terms of memory is not perfect. Nonetheless, I believe it suffices to get the crucial idea across. Figure 1 visualizes what it means for a computation to be regular. In the graph on the left, no matter how large the input is, the amount of memory required to process it never exceeds a fixed, constant threshold. On the other hand, in the graph on the right, as the input size increases, the amount of memory required to process the input, in the worst case, increases without bound.

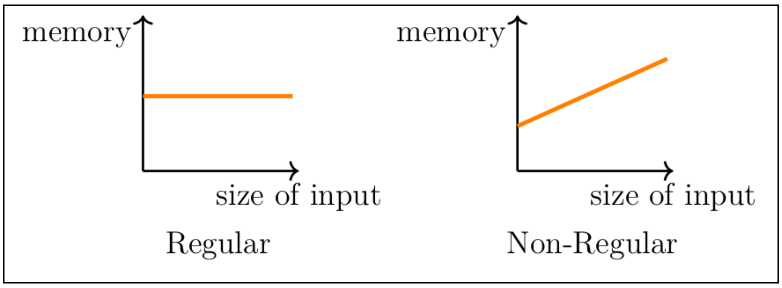

Figure 1. Regular and Non-regular patterns

To give a phonologically-motivated example, consider two logically-possible vowel harmony patterns. In Progressive Harmony, vowels agree in backness with the first vowel in the underlying representation. In Majority Rules Harmony, vowels agree in backness with the majority of vowels in the underlying representation. These generalizations are exemplified in Table 2 with hypothetical URs.

Table 2. Progressive and Majority Rules Vowel Harmony

UR	Progressive	Majority Rules
nokelu	nokolu	nokolu
nokeli	nokolu	nikeli
pidugo	pidige	pudugo
pidugomemi	pidigememi	pidigememi

It can be shown that Progressive Harmony is regular but Majority Rules Harmony is not (Riggle 2004; Heinz and Lai 2013). Intuitively for Progressive Harmony, when computing the SR for any UR, once the backness of the first vowel is known, producing the SR is just a matter of making the other vowels in the word match in backness. The only memory required is knowing the backness of the first vowel, which does not change as words become longer. On the other hand, in Majority Rules Harmony, when computing the SR for any UR, one must know whether the front or back vowels constitute a majority in the word. There is no single vowel that determines this. Instead, one needs to store in memory the difference between the number of front and back vowels in a word, and as words get longer, this difference can increase without bound in the worst case.

What makes this particular example interesting are three facts. First, Progressive Harmony is attested, but Majority Rules Harmony is not (Baković 2000). Second, human subjects fail to learn Majority Rules Harmony in artificial grammar learning experiments, unlike Progressive Harmony (Finley 2008, 2011). And third, as mentioned, Majority Rules Harmony is not regular, but Progressive Harmony is. Are these facts coincidental? In my opinion, the non-regular nature of Majority Rules Harmony accounts for the difference between it and Progressive Harmony with respect to attestedness and learnability in artificial grammar learning experiments.

This explanation is not available in Optimality Theory. There exists a CON and ranking over it which generates Majority Rules: AGREE(BACK)>>IDENT(BACK) (Frank and Satta 1998; Riggle 2004; Gerdemann and Hulden, 2012). Some believe changing CON helps address this issue. However, there are many results that pin the blame not on CON, but on global optimization itself. This is because limiting CON to simpler constraints does not appear to reign in global optimization (Hao 2019, 2024; Lamont 2021, 2022). In summary, global optimization allows one to solve problems that are non-regular; that is, which require memory to grow as the input size grows. It does this because it allows structures in different parts of the word to interact with each other in all kinds of non-local ways.

On the other hand, there is evidence that phonological generalizations are regular. This evidence originates with Johnson (1972) and Kaplan and Kay (1994), who showed how to translate any ordered sequence of SPE-style rewrite rules into a finite-state transducer which maps URs to SRs. Since virtually any phonological grammar can be expressed as an ordered, finite sequence of SPE-style rewrite rules, no matter how inelegantly, this means "being regular" is a property of the string maps that any phonological grammar defines. Furthermore, constraints on well-formed surface representations are regular since the set of outputs produced by finite- state transducers are also regular (Scott and Rabin 1959). A similar argument holds for the constraints on well-formed underlying representations.

An analogy may be useful. Consider grammars which build shapes. Some grammars can build all kinds of shapes: circles, pentagons, and quadrilaterals. Some grammars can only build quadrilaterals. Regular grammars are like the grammars that only build quadrilaterals. Examination of phonological generalizations shows they always have four sides. Optimality Theory by its very nature constructs non-quadrilaterals even when CON is sharply limited so it cannot explain this finding. Further inspection of phonological generalizations, reveal they are in fact squares. This last point will be explained in the remainder of this article while also introducing how logic and representation provide theories of phonology which help determine their shape and character.

4. Constraints and transformations with First Order logic

There is ample precedent for using logic and model theory to describe and analyze phonological generalizations (Coleman and Local 1991; Scobbie et al. 1996; Potts and Pullum 2002; Graf 2010). The basic idea is that phonological representations can be formalized as relational structures, in which there are a number of events, each which exhibit certain properties, and which are related to each other by particular kinds of relations. The properties and relations which constitute the structure are the atomic predicates in a logical language. This section illustrates these ideas using first order logic, broad classes of phonetic properties, and the successor relation, which relates events which occur one after another in time.

As a first example, I will illustrate how to define the constraint *e#, which is plausibly a markedness constraint driving word-final /e/ raising in Finnish. This constraint is defined below.

(1) $*\text{e\#} \coloneqq \neg\,(\exists x)\,[\text{vocalic}(x) \wedge \text{mid}(x) \wedge \text{front}(x) \wedge \text{unround}(x) \wedge (\exists y)\,[x \lhd y \wedge \#(y)]]$

To understand this sentence of First Order (FO) logic, it can be useful to think of variables like x and y as positions (events) in a sequence. A predicate property(x) means position x has that property. The relation $x \lhd y$ means position y is the next position after x and is read "y is the successor of x." The symbol $\wedge$ means AND; symbol $\vee$ means OR; symbol $\neg$ means NOT; and symbol $\exists$ means EXISTS. Putting it all together, the above sentence reads "There does not exist a position x which is a mid, front, unrounded vowel whose successor position y is the word boundary." More simply, this translates to "There are no word-final [e] vowels."

Figure 2 shows the relational structure for the form [ove]. There are five positions (events), each of which corresponds either to a word boundary or to one of the segments in the word. These are related by the successor relation, which arranges them in time.

Evaluating the constraint *e# is straightforward. We must check whether there is a position x satisfying the properties of a mid, front, unrounded vowel, followed by a position y satisfying the property of being a word boundary. There is such a position (position 4 in Figure 2). Thus, the logical sentence *e# in (1) evaluates to False when presented with the relational structure in Figure 2. This means that structure violates the constraint.

In contrast, if one imagined the relational structure for [ovi], it would be nearly identical

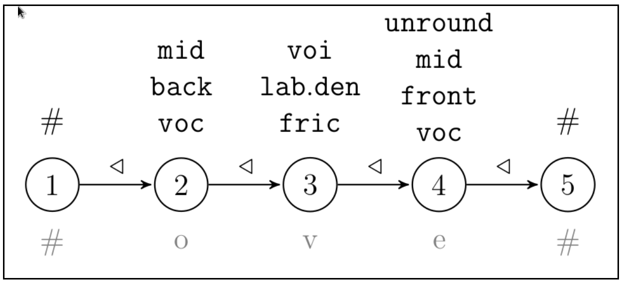

Figure 2. The relational structure of the word [ove]

to the one shown in Figure 2. The only difference would be that position 4 would satisfy a predicate *high* and not satisfy the predicate *mid*. When presented with this structure, the sentence *e# would evaluate to True because there are no positions corresponding to mid, front, unrounded vowels succeeded by word boundaries.

In this way, this logical language provides a precise way of evaluating constraints. Furthermore, since FO logic is a fragment of Monadic Second Order (MSO) logic, all constraints defined with this logical language, including conjunctions of such constraints, are regular. Good references on finite model theory and logic include Keisler and Robbin (1996), Enderton (2001), Libkin (2004), Hedman (2004), and Pratt-Hartmann (2023).

Logic can also be used to transform one structure into another (Courcelle 1994; Courcelle and Engelfriet 2012) to map a UR to a SR, for example. Such a transformation can be specified with a collection of sentences which collectively define properties of the positions in the output. These sentences have the form shown below, and should be understood as saying “Position x has property P in the output only if corresponding position x in the input satisfies predicate Q.”

(2) $\varphi P(x) \coloneqq Q(x)$

The symbol φ in the left-hand side is only there to signify that the predicate φP(x) may be distinct from predicate P(x). The former indicates whether position x has property P in the output structure and the latter indicates whether position x has property P in the input structure.

To illustrate, I will now provide a collection of sentences specifying the transformation of word-final /e/ raising in Finnish. First, most of the underlying structure surfaces intact. Only the properties *high* and *mid* are different between the underlying and surface structures. Therefore, any property P which is distinct from both *high* and *mid* (e.g. *vocalic*) is the same in the surface form as it is in the underlying form. This faithfulness of property P can be expressed with a sentence like the one below.

(3) $\varphi P(x) \coloneqq P(x)$

This means specifically that a position x will have the property P in the output only if the position x in the input has property P.

Next, the successor relation is itself unchanged. The sentence below asserts that position y is the successor of position x in the output only if y is the successor of position x in the input.

(4) $\varphi(x \triangleleft y) \coloneqq x \triangleleft y$

What remains is to specify which positions have the properties *high* and *mid*, respectively. A position x in the output bears the property *high* if its corresponding position x in the input is *high* OR if it is a word-final /e/. Similarly, a position x in the output bears the property *mid* if its corresponding position x in the input is *mid* AND it is NOT a word-final /e/.

(5) φhigh(x) ≔ high(x) ∨ (e(x) ∧ final(x))
(6) φmid(x) ≔ mid(x) ∧ ¬ (e(x) ∧ final(x))

The formulas above make use of the predicates e(x) and final(x). For completeness, these predicates are defined below.

(7) e(x) ≔ vocalic(x) ∧ mid(x) ∧ front(x) ∧ unround(x)
(8) final(x) ≔ (∃y) [x⊲ y ∧ #(y)]

This collection of sentences defines a logical transduction which, given some input relational structure, provides an effective procedure for constructing an output relational structure. I have omitted some details, and readers are referred to Strother-Garcia (2018, 2019), Dolatian (2020), and Chandlee and Jardine (2021) for more complete introductions.

To conclude this section, I return to the question of why anyone should care that we can describe constraints and transformation with logical languages. The first reason is that the use of First Order logic guarantees that the memory resources are bounded by a constant; that is, that the process we are describing is regular. This is because any logic equivalent to some fragment of MSO is at most regular, and First Order logic is a fragment of MSO.

The second reason is that choosing a logical formalism and a representation fixes a theory. The choice of logic and the choice of representation, which combine to yield a logical language, parameterize a space of possible theories. Logic and representation provide an "encyclopedia of categories" in Humboldt's (1836/1999) sense, against which the "encyclopedia of types" – that is the observed phonological patterning – can be compared against.

For example, the theory above combines FO logic with representations of words using the successor relation. A well-known aspect of this theory is that it is insufficiently expressive to account for long-distance constraints or processes that are found in phonology. To express such phenomena, one either has to use a more powerful logic such as MSO logic, or alter the representation of words to include either the notion of general precedence (Heinz, 2010; Rogers et al., 2013) or the notion of the phonological tier (Heinz et al. 2011; McMullin 2016; Rogers and Lambert 2019a; Lambert 2023). In this way, phonological phenomena can be classified according to the logical languages by which they can be expressed.

A third reason to be interested is that weaker logics facilitate learning and acquisition. It is well known that the class of constraints and processes defined with MSO logic are not learnable exactly from positive examples, even in the absence of noise. Nor are they even approximately learnable under typical probabilistic assumptions. On the other hand, constraints and processes which can be expressed with logics weaker than FO logic are known to be feasibly learnable (Valiant 1984; Strother-Garcia et al. 2016; Lambert et al. 2021).

In the example above, we used a FO logical language for Finnish. Evidently it was sufficient, but are there weaker logics that are still sufficient? In the next section I argue that weaker logics, with the right representations, are not only sufficient for expressing the diversity of phonological generalizations found in natural languages, but they are also sufficiently restrictive for successful learning to take place.

5. The strengths of weaker logics

This section considers logics weaker than FO. For constraints, I consider a Propositional logic and two smaller fragments of it. For transformations, I consider a Quantifier-Free logic.

5.1 Constraints

This section introduces a propositional logic for defining constraints over relational structures (Rogers and Lambert 2019a). Propositional logic has no variables like x nor quantifiers like ∃. It only uses connectives like ∧, ∨ and ¬ to combine the atomic propositions. The atomic propositions themselves are *factors*, which are connected pieces of structures. A relational structure S satisfies an atomic proposition F if and only if S *contains* F.

For example, the factor e# is defined to be the structure shown in Figure 3. This factor is an atomic proposition in the logical language. The relational structure for [ove] shown in Figure 2 *contains* the factor e# and therefore that structure satisfies the propositional sentence $\varphi \coloneqq$ e#. On the other hand, the relational structure for [ove] shown in Figure 2 does not satisfy the sentence $\varphi' \coloneqq \neg$ e# because only structures which *do not contain* the factor e# satisfy φ'.

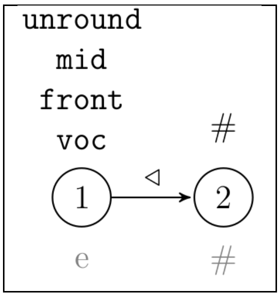

Figure 3. The factor e#

More generally, sentences of the form $\neg F_1 \wedge \neg F_2 \ . \ . \ . \wedge \neg F_n$ mean that well-formed structures do not contain any factor F_i for $1 \leq i \leq n$. If sentences of propositional logic are restricted to sentences of this form, we obtain a logic known as the Conjunction of Negative Literals (CNL). If these conjunctive sentences also admit positive literals – that is conjunctions like $F_1 \wedge F_2$ which would require structures to contain factors F_1 and F_2 – then they constitute a logic known as the Conjunction of Positive and Negative Literals (CPNL). Such sentences reduce well-formedness to simply checking whether local regions of the structure contain a forbidden or required factor.

It is of interest to know to what extent sentences of propositional logic are able to describe phonological well-formedness conditions in natural languages. An early result is that with the right representations, both local and long-distance segmental phonotactic patterns can be so described (Heinz 2010).

More recently, Rogers and Lambert (2019b) examine the StressTyp2 database (Goedemans et al. 2015), which provides non-lexical stress patterns for over 700 languages, for which 106 distinct patterns have been encoded as finite-state acceptors. They show that 98/106 patterns can be described with CPNL. Six more require implication: "F_1 implies $\neg F_2$" and two more appear to require MSO logic because they contain "hidden alternation pattern that requires an odd number of syllables to occur in certain spans of the word." These become CPNL if secondary stress is perceptible, as has been argued in similar cases (Becker 2022). Lambert (to appear) provides a more refined analysis of these stress patterns.

The above results employed linear representations, but non-linear representations, such as syllabic and autosegmental structures, have been studied too. These representations are also relational structures, and constraints can be expressed with sentences of CNL logic over factors that refer to different levels of structure. Strother-Garcia (2018, 2019) shows that constraints given by ONSET, NOCODA, and the Sonority Sequencing Principle follow from forbidding particular factors using syllabic representations. She also shows that the basic CV typology, and extensions thereof, can be obtained with sentences of CNL. Similarly, Jardine (2016, 2017) shows that well-studied patterns of tonal association can all be expressed with sentences of CNL over autosegmental representations, including position-specific plateaus, position-specific contours, melody constraints, and quality dependent plateaus.

These results are important for a couple of reasons. First, they show that global optimization of ranked constraints is not necessary to characterize well-formedness conditions. Second, they show that the computational complexity is less than what was previously known. In particular, the fact that many patterns can be expressed with CNL means the well-formedness conditions are localized: it is sufficient to attend to individual factors to determine well-formedness.

The examples so far are inviolable, language-specific constraints. This itself is not an issue if they can be learned. Chandlee et al. (2019) show that the space of factors forms a partial order, and design the Bottom-Up Factor Inference Algorithm (BUFIA) to search this space for forbidden factors without statistics. This algorithm is actually a family of algorithms whose implementation depends on the particular representations involved. Li (2024) implements BUFIA for autosegmental representations and applies it to a curated list of monomorphemic forms from a Hausa dictionary. It returns seven forbidden factors, three of which were previously identified by linguists, two of which are more specific versions of constraints previously reported (because the dictionary contains forms violating those previously reported constraints). The last two are previously unreported constraints. In another study, Swanson et al. (2024) implement BUFIA for learning segmental phonotactics and examine the local and long-distance phonotactics found in Quechua. Wilson & Gallagher (2018) had argued that phonotactic learning over featural representations necessitates statistical methods such as Maximum Entropy (Hayes & Wilson, 2008). On the contrary, Swanson et al. (2024) shows that BUFIA learns constraints over featural representations for Quechua, just as well as, if not better than, the maximum entropy learning algorithm Wilson and Gallagher employ. The reason for BUFIA's success is not a mystery: the space of possible constraints is structured and restricted in a way that facilitates successful, feasible learning.

To summarize, restricted forms of propositional logic provide good typological coverage, as well as learning algorithms with good theoretical and empirical results. Finally, I turn to the question of whether weaker logics for transformations have the same strengths.

5.2 Processes

It is not known to me how to define processes in terms of a propositional logic over factors as was the case above. However, Quantifier-Free (QF) logic is a fragment of FO logic that, as we will see, ensures that transformations it specifies involve only local computations.

A formula of FO logic is quantifier-free if the right-hand side does not include any quantification. Compare the two formulas below, each of which presents a definition for whether a position x in the output structure has property P.

(9) $\varphi P(x) \coloneqq Q(x) \wedge (\exists y)\,[R(y)]$

(10) $\varphi P(x) \coloneqq Q(x) \wedge R(x)$

The first formula uses a quantifier on the right-hand side, but the second does not. For the former, to decide whether position x in the output structure has property P, one has to examine the whole structure for a position y satisfying R, whereas for the latter, all the necessary information is local to position x.

Lindell and Chandlee (2016) show that Quantifier-Free transductions over string representations are Input Strictly Local (ISL) transformations, which Chandlee (2014) and Chandlee and Heinz (2018) showed is a property belonging to approximately 95% of the individual processes in P-Base (Mielke 2008, v.1.95), including local substitution, deletion, epenthesis, and synchronic metathesis. Furthermore, many opaque transformations without any special modification are also ISL (Chandlee et al. 2018).

Chandlee et al. (2014) show ISL transformations are feasibly learnable from positive examples given a non-negative integer parameter k. ISL transformations can also be generalized to operate on a phonological tier to account for long-distance harmony and spreading processes (McMullin 2016; Burness and McMullin 2019; Burness et al. 2021; Lambert and Heinz 2023). Given an arbitrary finite-state transducer, one can decide whether it is (tier) ISL or not (Lambert and Heinz 2023).

It is of course of interest to see to what extent QF transformations inform our understanding of non-linear representations. Strother-Garcia (2018) shows that the process of syllabification in Imdlawn Tashlhiyt Berber (ITB; Dell and Elmedlaoui 1985; Prince and Smolensky 1993) is Quantifier Free. She concludes "…syllabification in ITB can be represented by a QF [logical] transduction, a formalism restricted to substantially lower computational complexity than [traditional] phonological grammars…Establishing that ITB syllabification is QF highlights an insight not apparent from [those traditional] grammatical formalisms…" (152). Similarly, Dolatian (2020) examines the phonology-morphology interface in light of Quantifier Free logical transductions. He concludes "the bulk of the morphology-phonology interface requires local computation, not global computation." (iii).

The use of logical transformations over relational structures is also a valuable way to compare competing phonological representations to identify significant differences. When stronger logics are necessary to translate between representations, then the differences between them are more significant. Strother-Garcia (2019) shows that different syllabic representations can be obtained via QF translations, an indicator that their differences are relatively minor. In a similar vein, (Oakden 2020) compares different tonal representations (Yip 1989; Bao, 1990) and demonstrates their inter-translatability with quantifier free logic. Both these researchers conclude these different representations are better thought of as notational variants instead of as competing proposals. Jardine et al. (2021) shows that autosegmental representations and Q-theoretic representations (Shih and Inkelas 2018) are also intertranslatable with quantifier-free logic and shows that Q-theory is mostly the same as autosegmental representations, claims to the contrary notwithstanding. On the other hand, (Nelson 2022) examines the phonetics-phonology interface, and shows how to convert coupling graphs in Articulatory Phonology (Browman and Goldstein 1992) to familiar segmental representations and vice versa with FO logic. It remains to be seen whether a QF translation exists.

In short, like the propositionally-defined constraints earlier, QF transformations exhibit good typological coverage and theoretical learning results. They also strongly implicate local computation as an important characteristic of phonological generalizations, as opposed to global optimization. Finally, they provide a new technique for comparing representations proposed for phonology or its interfaces.

If MSO is a logic which can only build quadrilaterals, this section has presented logics that can only build squares. The fact that so many phonological generalizations are square-like is notable, and it is not something one would anticipate given a rule-based theories like SPE or theories based on global optimization like OT.

6. Conclusion

The logical structure of phonological generalizations reveals that many, if not all, phonological generalizations are local with the right representations and that the computations are mostly, if not all, *subregular*. Unlike in Optimality Theory, the constraints defined here are inviolable constraints and language-specific. That is not a problem provided they can be learned from examples. Learning algorithms exploiting these properties exist and continue to be developed. The logical structure of phonological generalizations reveals what learners must attend to, and in this way explains the kinds of phonological generalizations which are learned, and provides explanations for the phonological generalizations we do and do not observe.

There are many reasons to pursue phonological description and analysis with logical systems and relational structures. They can express phonological generalizations precisely, accurately, and completely. They are easy to learn with only a little practice. They provide an "encyclopedia of categories" which can be used to examine typologies and compare phonological theories. They inform questions regarding memory, processing, and learning. They can be weighted to compute probabilities, count violations, and handle optionality, among other uses (Droste and Gastin 2009). Developments in algebraic methods have made it easier to determine which logics over strings can express given constraints or processes (Lambert 2022, to appear; Lambert and Heinz, 2023, 2024). Research continues to show that weaker logics admit learnability results that are unavailable to more expressive logics. Finally, logic will still be here in hundreds of years. Thus, even researchers whose primary focus is the documentation of understudied languages benefit from describing grammars using logical languages as it provides a measure of longevity unequaled by other grammatical formalisms.

There are many avenues for future research. Many more phonological representations can be studied within phonology and its interfaces, and compared more carefully (Danis to appear). Similarly, more logics including fixed point logics, Boolean Monadic Recursive Schemes (Chandlee and Jardine, 2019, 2021), and logics without negation for the expression of natural classes (Nelson 2022b), can be further studied and developed. Finally, there are many open learning problems to be investigated for grammars expressed with logical languages. Learning lexicons, grammars, exceptions, and variation and learning transformations over non-linear representations are all current and exciting areas to pursue.

References

Baković, Eric. 2000. *Harmony, dominance and control.* Ph.D. dissertation, Rutgers University.

Baković, Eric and Lev Blumenfeld. 2024. A formal typology of process interactions. *Phonological Data and Analysis* 6:3.1–43.

Bao, Zhi-ming (1990). On the Nature of Tone. PhD dissertation, MIT.

Becker, Michael (2022). Cairene Arabic stress is local. *Radical: A Journal of Phonology* 4. 211–247.

Browman, Claire and Louis Goldstein (1992). Articulatory phonology: An overview. *Phonetica* 155–180.

Burness, Phillip and Kevin McMullin (2019). Efficient learning of output tier-based strictly 2-local functions. In *Proceedings of the 16th Meeting on the Mathematics of Language*. Toronto, Canada: Association for Computational Linguistics, 78–90.

Burness, Phillip A., Kevin J. McMullin and Jane Chandlee (2021). Long-distance phonological processes as tier-based strictly local functions. *Glossa: a Journal of General Linguistics* 6:1.

Chandlee, Jane (2014). Strictly Local Phonological Processes. Ph.D. dissertation, The University of Delaware.

Chandlee, Jane, Rémi Eyraud and Jeffrey Heinz (2014). Learning strictly local subsequential functions. *Transactions of the Association for Computational Linguistics* 2. 491–503.

Chandlee, Jane, Remi Eyraud, Jeffrey Heinz, Adam Jardine and Jonathan Rawski (2019). Learning with partially ordered representations. In *Proceedings of the 16th Meeting on the Mathematics of Language*. Toronto, Canada: 91–101.

Chandlee, Jane and Jeffrey Heinz (2018). Strict locality and phonological maps. *Linguistic Inquiry* 49:1. 23–60.

Chandlee, Jane, Jeffrey Heinz and Adam Jardine (2018). Input strictly local opaque maps. *Phonology* 35:2. 171–205.

Chandlee, Jane and Adam Jardine (2019). Quantifier-free least fixed point functions for phonology. In *Proceedings of the 16th Meeting on the Mathematics of Language*. Toronto, Canada: Association for Computational Linguistics, 50–62.

Chandlee, Jane and Adam Jardine (2021). Computational universals in linguistic theory: Using recursive programs for phonological analysis. *Language* 93. 485–519.

Chomsky, Noam and Morris Halle (1968). *The Sound Pattern of English*. New York: Harper and Row.

Coleman, John and John Local (1991). The “no crossing constraint” in autosegmental phonology. *Linguistics and Philosophy* 14:3. 295–338.

Courcelle, Bruno (1994). Monadic second-order definable graph transductions: a survey. *Theoretical Computer Science* 126. 53–75.

Courcelle, Bruno and Joost Engelfriet (2012). *Graph Structure and Monadic Second-Order Logic, a Language Theoretic Approach*. Cambridge: Cambridge University Press.

Danis, Nick (to appear). Logical transductions are not sufficient for notational equivalence. In *Proceedings of the 2024 Annual Meeting of Phonology*.

Dell, Francois and Mohamed Elmedlaoui (1985). Syllabic consonants and syllabification in Imdlawn Tashlhiyt Berber. *Journal of African Languages and* Linguistics 7. 105–130.

Dolatian, Hossep (2020). Computational Locality of Cyclic Phonology in Armenian. Ph.D. dissertation, Stony Brook University.

Droste, Manfred and Paul Gastin (2009). Weighted automata and weighted logics. Springer, *Monographs in Theoretical Computer Science*, chapter 5.

Elgot, C. C. and J. E. Mezei (1965). On relations defined by generalized finite automata. *IBM Journal of Research and Development* 9:1. 47–68.
Enderton, Herbert B. (2001). *A Mathematical Introduction to Logic*. Academic Press, 2nd edition.
Engelfriet, Joost and Hendrik Jan Hoogeboom (2001). MSO definable string transductions and two-way finite-state transducers. *ACM Transactions on Computational Logic* 2:2. 216–254.
Finley, Sara (2008). The formal and cognitive restrictions on vowel harmony. PhD dissertation, Johns Hopkins University, Baltimore, MD.
Finley, Sara (2011). The privileged status of locality in consonant harmony. *Journal of Memory and Language* 65. 74–83.
Frank, Robert and Giorgo Satta (1998). Optimality Theory and the generative complexity of constraint violability. *Computational Linguistics* 24:2. 307–315.
Gerdemann, Dale and Mans Hulden (2012). Practical finite state optimality theory. *In Proceedings of the 10th International Workshop on Finite State Methods and Natural Language Processing*. Donostia–San SebastiÃ¡n: Association for Computational Linguistics, 10–19.
Goedemans, R. W. N., Jeffrey Heinz and Harry van der Hulst (2015). http://st2.ullet.net/. Retrieved 30 January 2018. Stresstyp2.
Graf, Thomas (2010). Logics of Phonological Reasoning. MA dissertation, University of California, Los Angeles.
Hao, Sophie (2024). Universal generation for Optimality Theory is PSPACE-Complete. *Computational Linguistics*. 1–35.
Hao, Yiding (2019). Finite-state optimality theory: non-rationality of harmonic serialism. *Journal of Language Modelling* 7:2. 49–99.
Hayes, Bruce and Colin Wilson (2008). A maximum entropy model of phonotactics and phonotactic learning. *Linguistic Inquiry* 39. 379–440.
Hedman, Shawn (2004). *A First Course in Logic*. Oxford: Oxford University Press.
Heinz, Jeffrey (2010). Learning long-distance phonotactics. *Linguistic Inquiry* 41:4. 623–661.
Heinz, Jeffrey (2011a). Computational phonology part I: Foundations. *Language and Linguistics Compass* 5:4. 140–152.
Heinz, Jeffrey (2011b). Computational phonology part II: Grammars, learning, and the future. *Language and Linguistics Compass* 5:4. 153–168.
Heinz, Jeffrey (2018). The computational nature of phonological generalizations. In Larry Hyman and Frans Plank (eds.) *Phonological Typology*, De Gruyter Mouton, Phonetics and Phonology, chapter 5, 126–195.
Heinz, Jeffrey and Regine Lai (2013). Vowel harmony and subsequentiality. In Andras Kornai and Marco Kuhlmann (eds.) *Proceedings of the 13th Meeting on the Mathematics of Language*. Sofia, Bulgaria: 52–63.
Heinz, Jeffrey, Chetan Rawal and Herbert G. Tanner (2011). Tier-based strictly local constraints for phonology. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*. Portland, Oregon: 58–64.
Humboldt, Wilhem von (1999). *On Language.* In Cambridge Texts in the History of Philosophy. Edited by Michael Losonsky. Translated by Peter Heath. Originally published 1836. Cambridge: Cambridge University Press.
Jardine, Adam (2016). Locality and non-linear representations in tonal phonology. PhD dissertation, University of Delaware.
Jardine, Adam (2017). The local nature of tone-association patterns. *Phonology* 34. 385–405.
Jardine, Adam, Nick Danis and Luca Iacoponi (2021). A formal investigation of Q-theory in comparison to autosegmental representations. *Linguistic Inquiry* 52:2. 333–358.
Johnson, C. Douglas (1972). *Formal Aspects of Phonological Description.* The Hague: Mouton.
Kaplan, Ronald and Martin Kay (1994). Regular models of phonological rule systems. *Computational Linguistics* 20:3. 331–378.
Keisler, H. Jerome and Joel Robbin (1996). *Mathematical Logic and Computability*. New York: McGraw-Hill.
Kenstowicz, Michael and Charles Kisseberth (1979). *Generative Phonology*. New York: Academic Press, Inc.
Kleene, S.C. (1956). Representation of events in nerve nets. In C.E. Shannon and J. McCarthy (eds.) *Automata Studies*, Princeton University. Press, 3–40.
Kozen, Dexter (1997). *Automata and Computability*. New York: Springer.
Lambert, Dakotah (2022). Unifying Classification Schemes for Languages and Processes with Attention to Locality and Relativizations Thereof. PhD dissertation, Stony Brook University.
Lambert, Dakotah (2023). Relativized adjacency. *Journal of Logic Language and Information* 32. 707–731.
Lambert, Dakotah (to appear). Multitier phonotactics with logic and algebra. *Phonology.*
Lambert, Dakotah and Jeffrey Heinz (2023). An algebraic characterization of total input strictly local functions. In *Proceedings of the Society for Computation in Linguistics.* Volume 6. 25–34.
Lambert, Dakotah and Jeffrey Heinz (2024). Algebraic reanalysis of phonological processes described as output-oriented. In *Proceedings of the Society for Computation in Linguistics*. Volume 7, 129–138.

Lambert, Dakotah, Jonathan Rawski and Jeffrey Heinz (2021). Typology emerges from simplicity in representations and learning. *Journal of Language Modelling* 9:1. 151–194.

Lamont, Andrew (2021). Optimizing over subsequences generates context-sensitive languages. *Transactions of the Association for Computational Linguistics* 9. 528–537.

Lamont, Andrew (2022). Optimality theory implements complex functions with simple constraints. *Phonology* 38:4. 729–740.

Li, Han (2024). Learning tonotactic patterns over autosegmental representations. Qualifying Paper, Stony Brook University.

Libkin, Leonid (2004). Elements of Finite Model Theory. Berlin: Springer.

Lindell, Steven and Jane Chandlee (2016). A logical characterization of input strictly local functions. Presented at the Fourth Workshop on Natural Language and Computer Science, in affiliation with LICS 2016.

McMullin, Kevin (2016). Tier-Based Locality in Long-Distance Phonotactics: Learnability and Typology. Ph.D. dissertation, University of British Columbia.

Mielke, Jeff (2008). *The Emergence of Distinctive Features*. Oxford: Oxford University Press.

Nelson, Scott (2022a). Are representations in articulatory and generative phonology so different? Poster presented at the 2022 Annual Meeting on Phonology.

Nelson, Scott (2022b). A model theoretic perspective on phonological feature systems. In Allyson Ettinger, Tim Hunter, and Brandon Prickett (eds.) *Proceedings of the Society for Computation in Linguistics* 2022. 1–10.

Oakden, Chris (2020). Notational equivalence in tonal geometry. *Phonology* 37:2. 257–296.

Odden, David (2014). *Introducing Phonology*. Cambridge: Cambridge University Press, 2nd edition.

Potts, Christopher and Geoffrey Pullum (2002). Model theory and the content of OT constraints. *Phonology* 19. 361–393.

Pratt-Hartmann, Ian (2023). *Fragments of First-Order Logic*. Oxford: Oxford University Press.

Prince, Alan and Paul Smolensky (1993). Optimality Theory: Constraint Interaction in Generative Grammar. Technical Report 2, Rutgers University Center for Cognitive Science.

Prince, Alan and Paul Smolensky (2004). *Optimality Theory: Constraint Interaction in Generative Grammar*. Malden, MA: Blackwell Publishing.

Riggle, Jason (2004). Generation, Recognition, and Learning in Finite State Optimality Theory. Ph.D. dissertation, University of California, Los Angeles.

Rogers, James, Jeffrey Heinz, Margaret Fero, Jeremy Hurst, Dakotah Lambert and Sean Wibel (2013). Cognitive and sub-regular complexity. In Glyn Morrill and Mark-Jan Nederhof (eds.) Formal Grammar. Springer, volume 8036 of Lecture Notes in Computer Science, 90–108.

Rogers, James and Dakotah Lambert (2019a). Some classes of sets of structures definable without quantifiers. In *Proceedings of the 16th Meeting on the Mathematics of Language*. Toronto, Canada: Association for Computational Linguistics, 63–77.

Rogers, James and Dakotah Lambert (2019b). Extracting subregular constraints from regular stringsets. *Journal of Language Modelling* 7:2. 143–176.

Scobbie, James M., John S. Coleman and Steven Bird (1996). Key aspects of declarative phonology. In Jacques Durand and Bernard Laks (eds.) *Current Trends in Phonology: Models and Methods*, Manchester, UK: European Studies Research Institute, volume 2, 685–709. University of Salford.

Scott, Dana and Michael Rabin (1959). Finite automata and their decision problems. *IBM Journal of Research and Development* 5:2. 114–125.

Shih, Stephanie S and Sharon Inkelas (2018). Autosegmental aims in surface-optimizing phonology. *Linguistic Inquiry* 50:1. 137–196.

Strother-Garcia, Kristina (2018). Imdlawn Tashlhiyt Berber syllabification is quantifier-free. In *Proceedings of the Society for Computation in Linguistics*. Volume 1. Article 16.

Strother-Garcia, Kristina (2019). Using Model Theory in Phonology: A Novel Characterizaton of Syllable Structure and Syllabification. Ph.D. dissertation, University of Delaware.

Strother-Garcia, Kristina, Jeffrey Heinz and Hyun Jin Hwangbo (2016). Using model theory for grammatical inference: a case study from phonology. In Sicco Verwer, Menno van Zaanen, and Rick Smetsers (eds.) *Proceedings of The 13th International Conference on Grammatical Inference*. volume 57 of JMLR: Workshop and Conference Proceedings, 66–78.

Swanson, Logan, Jeffrey Heinz and Jon Rawski (2024). Phonotactic learning and constraint selection without statistics. Poster presented to the 2024 Annual Meeting of Phonology.

Tesar, Bruce (2014). *Output-driven Phonology*. Cambridge: Cambridge University Press.

Valiant, L.G. (1984). A theory of the learnable. *Communications of the ACM* 27. 1134–1142.

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

Yip, Moira (1989). Contour tones. *Phonology* 6. 149–174.