

A Framework for Learning Phonological Maps as Logical Transductions

Tatevik Yolyan

Department of Linguistics / Rutgers University / New Brunswick, NJ
tatevik.yolyan@rutgers.edu

Abstract

This paper presents an approach to learning a subclass of phonological maps that can be expressed with non-recursive quantifier-free logical transductions within the framework of Boolean Monadic Recursive Schemes (BMRS). Building on previous work on phonological learning with partially-ordered hypothesis spaces, this paper shows that similar structures can be used in learning phonological maps when these maps are represented by logical transducers over unconventional string models. A primary contribution of this paper is a demonstration of how the BMRS framework can be used to learn *phonological generalizations*.

1 Introduction

Input Strictly Local (ISL) functions characterize phonological maps in which sounds undergo change based on local information in the input (Chandlee, 2014; Chandlee and Heinz, 2018). In other words, every ISL function has a parameter k such that for every x in the input string, the information needed to determine the output of x is within some window of size k around x . Consider the postnasal voicing map below from Zoque, in which voiceless stops become voiced when they are immediately preceded by a nasal (Wonderly, 1951). The input-output pair of words in Figure 1 illustrates why this map is ISL-2.

pama	‘clothing’	mbama	‘my clothing’
tatah	‘father’	ndatah	‘my father’
kama	‘cornfield’	ngama	‘my (corn)field’
hayah	‘husband’	nhayah	‘my husband’

$$\begin{array}{c} k = 2 \\ / \boxed{m \quad p} \quad a \quad m \quad a \quad / \\ \downarrow \\ [m \quad \boxed{b} \quad a \quad m \quad a \quad] \end{array}$$

Figure 1: Postnasal voicing in Zoque is ISL-2.

The ISL class of functions is learnable from possible data; for every ISL function f , it is possible to learn a finite state transducer which computes f from a characteristic sample $S \subseteq f$ of input-output pairs of strings (Chandlee, 2014; Chandlee et al., 2014). The purpose of this paper is to revisit learning from a model-theoretic perspective, where the goal is to learn ISL functions as *logical* transductions from pairs of input-output string *models*. While finite state transducers traditionally operate over string representations, logical transducers operate over more enriched representations. This paper uses unconventional string models which represent strings with phonological features (Strother-Garcia et al., 2016; Chandlee et al., 2019).

One of the main ideas presented in this paper is that we can learn logical transductions from input-output pairs of string models by using partially-ordered hypothesis spaces. Similar spaces have been used for various learning problems within phonology (Rawski, 2021; Chandlee et al., 2019; Tesar, 2013; Heinz et al., 2012; Heinz, 2010). A common theme among them is that a partially-ordered hypothesis space encodes entailment relations that are relevant for learning. These entailments allow a large space to be pruned efficiently with a small number of data points. This paper shows that a similar approach can be applied to the problem of learning logical transducers. This paper culminates with a demonstration of how a logical transducer for postnasal voicing can be learned using a very small number of input-output pairs.

A further contributions of this work is that it is closely related to previous work on model-theoretic learning of phonotactic constraints from Chandlee et al. (2019) and Rawski (2021). More specifically, this paper shows that the problem of learning the environment that triggers a particular feature to undergo change employs the same structure as learning banned k -factors in a grammar. In the case of postnasal voicing, for example, Zoque has

a phonotactic constraint $*[+nas][-voi,-cont]$ which says a voiceless stop cannot be immediately preceded by a nasal. This constraint also expresses the *environment* where voicing takes place: if an underlying form violates $*[+nas][-voi,-cont]$, the second sound undergoes voicing in order to repair the violation. The partially-ordered hypothesis space used to learn the surface constraint $*[+nas][-voi,-cont]$ can therefore be adapted to learning the environment where a voiceless sound becomes voiced. As such, the procedure presented here is a means of *lifting* phonotactic learning to learning phonological maps via logical transductions.

A broader contribution of this paper is a demonstration of how this model-theoretic approach can be used to learn phonological generalizations from a sample of underlying and surface form pairs. To that end, the style of this paper is mainly expository, and the learning problem is abstract rather than formally-defined. This paper moreover focuses on length-preserving and order-preserving ISL functions because they are a natural starting point for motivating the larger research program of model-theoretic phonology. The method introduced here can be extended to more complex function classes, as well as more complex representations.

The structure of this paper is as follows. Section 2 provides the relevant background and definitions of string models, logical transductions, subfactors, and partially-ordered sets (posets). Section 3 discusses previous work on phonological learning with posets, with emphasis on the Bottom-Up Feature Inference Algorithm (BUFIA; Chandlee et al., 2019). Section 4 presents the learning goal and procedure, and Section 4.4 illustrates these ideas using postnasal voicing as a case study. Future extensions of this work are discussion in Section 5.

2 Preliminaries

Given an alphabet Σ , a string/word is a finite concatenation of symbols in Σ . The set of all such strings, along with the empty string λ , is denoted Σ^* . For each $w \in \Sigma^*$, we let w_i denote the i^{th} character in w , and $|w|$ denote the length of w . This section presents definitions of relational structures and string models which are used to represent words in Σ^* , and ultimately functions over Σ^* . These structures are then enhanced with phonological features over which phonological maps are defined.

2.1 String Models

Definition 1. A *signature* is a collection of relation symbols $\{R_i\}_{i \leq m}$, and function symbols $\{f_i\}_{i \leq n}$.¹

Definition 2. Given a signature \mathcal{S} , an *\mathcal{S} -structure*

$$\mathcal{M} = \langle D, \{R_i^{\mathcal{M}}\}_{i \leq m}, \{f_i^{\mathcal{M}}\}_{i \leq n} \rangle$$

consists of a domain D , a relation $R_i^{\mathcal{M}} \subseteq D^k$ for each k -ary relation symbol R_i , and a function $f_i^{\mathcal{M}} : D^k \rightarrow D$ for every k -ary function symbol f_i . $\text{Struc}(\mathcal{S})$ denotes the set of \mathcal{S} -structures.

Structures are built over a signature by specifying a domain and interpretations for the relation and functions symbols. Every word in Σ^* is associated with a particular kind of structure called a *string model* built over a signature consisting of unary predicates $\{\bowtie, \times, P_\sigma\}_{\sigma \in \Sigma}$ and successor and predecessor functions. An alphabet and a signature over that alphabet are often designated by the same symbol; the string models built over an alphabet Σ are called Σ -structures.

This paper looks at a specific type of string model which we refer to here as *monadic* string models, following Chandlee and Lindell (forth.).

Definition 3. For an alphabet Σ and word $w \in \Sigma^*$, a *monadic string model* for w is a Σ -structure

$$\mathcal{M}(w) := \langle D, \bowtie, \times, P_\sigma, s, p \rangle_{\sigma \in \Sigma}$$

where $D = \{0, \dots, |w| + 1\}$ is a set of indices; each $P_\sigma : D \rightarrow \{\top, \perp\}$ is a monadic predicate s.t. $P_\sigma(i) = \top$ iff $w_i = \sigma$; $\bowtie(i) = \top$ iff $i = 0$; $\times(i) = \top$ iff $i = |w| + 1$; $s, p : D \rightarrow D$ are successor and predecessor functions over D ; for each $i \in \{1, \dots, |w|\}$, exactly one σ is s.t. $P_\sigma(i) = \top$.

The predicates P_σ indicate which symbol of the alphabet appears at an index, the symbols \bowtie and \times indicate the left and right edges of the word, and p and s capture the linear ordering of the symbols. The monadic model for the string ‘baa’ over the alphabet $\{a, b\}$ is illustrated in Figure 2.

String models can be further enriched with other relations/functions to represent tiers (Lambert and Rogers, 2020; Chandlee and Jardine, 2019a; Jardine, 2017), syllable structure (Strother-Garcia, 2019, 2018; Strother-Garcia et al., 2016), prosodic structure (Dolatian, 2020), and syntactic trees (Rogers, 2003). String models have been used

¹The concept of ‘signature’ is also commonly referred to in model theory literature as ‘vocabulary’ (e.g. Libkin, 2004) or ‘language’ (e.g. Marker, 2002).

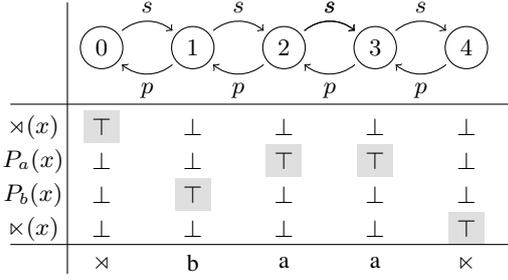


Figure 2: $\mathcal{M}(baa)$: monadic model of the string ‘baa’.

to characterize subregular classes of phonological grammars (Lambert and Rogers, 2020; Rogers and Lambert, 2019; Rogers and Pullum, 2011) and to develop inference algorithms for these grammars (Rawski, 2021; Chandlee et al., 2019; Strother-Garcia et al., 2016; Vu et al., 2018).

2.2 Logical Transductions

The focus of this paper is on phonological *maps*. We represent the underlying and surface forms of words with string models, and represent phonological maps as *relations* between string models. Consider a string function $f : \Sigma^* \rightarrow \Gamma^*$. For every input string $w \in \Sigma^*$ and corresponding Σ -structure $\mathcal{M}(w)$, there is an output string $f(w)$ and corresponding Γ -structure $\mathcal{M}(f(w))$. We call $\Sigma = \langle P_\sigma, \times, \times, p, s \rangle_{\sigma \in \Sigma}$ the *input signature* and $\Gamma = \langle P_\gamma, \times, \times, p, s \rangle_{\gamma \in \Gamma}$ the *output signature*, and capture the relationship between input and output string models with a map that transforms Σ -structures into Γ -structures by means of logical formulas. In particular, the predicates $\{P_\gamma\}_{\gamma \in \Gamma}$ in the output Γ -structure are associated with logical formulas over the predicates $\{P_\sigma\}_{\sigma \in \Sigma}$ in the input Σ -structure. These maps are called *logical transductions* (Courcelle, 1994; Engelfriet and Hoogboom, 2001; Courcelle and Engelfriet, 2012). Logical transductions over monadic string models are quantifier-free (Chandlee and Jardine, 2019b, 2021; Chandlee and Lindell, forthcoming).

As an example, consider the simple function f expressed by the rewrite rule $a \rightarrow b/b_$, applied over the alphabet $\Sigma = \{a, b\}$. The logical transducer which expresses f transforms a Σ -structure $\langle D, \times, \times, P_a, P_b, p, s \rangle$ into a Σ -structure $\langle D, \times, \times, P'_a, P'_b, p, s \rangle$, where P'_a and P'_b are expressed as logical formulas over the *signature of the input model*, as in (1). The predicates P_a and P_b express when an index in the input string model carries an ‘a’ or ‘b’ symbol, respectively. The predicates P'_a and P'_b express when an index in the

input	\times	b	a	a	\times
	0	1	2	3	4
$\times(x)$	\top	\perp	\perp	\perp	\perp
$P_a(x)$	\perp	\perp	\top	\top	\perp
$P_b(x)$	\perp	\top	\perp	\perp	\perp
$\times(x)$	\perp	\perp	\perp	\perp	\top
$P'_a(x)$	\perp	\perp	\perp	\top	\perp
$P'_b(x)$	\perp	\top	\top	\perp	\perp
output	\times	b	b	a	\times

Figure 3: Logical transducer in ((1)) over $\mathcal{M}(baa)$.

output string model carries an ‘a’ or ‘b’. Since f is a length-preserving function, the input and output models have the same domain, linear ordering, and boundary predicates \times and \times .

- (1) *Logical transducer which expresses $a \rightarrow b/b_$*
- $$P'_a(x) = P_a(x) \wedge \neg P_b(p(x))$$
- $$P'_b(x) = P_b(x) \vee (P_a(x) \wedge P_b(p(x)))$$

The equation for $P'_a(x)$ in (1) expresses the following: the index x will carry an ‘a’ in the output model iff it carries an ‘a’ in the input model, and the index preceding it does not carry a ‘b’ in the input model. Stated more simply, x will output as an ‘a’ iff it is an ‘a’ that is not preceded by a ‘b’ in the input. Similarly, $P'_b(x)$ says: x will carry a ‘b’ in the output iff it is either a ‘b’ in the input, or it is an ‘a’ that is preceded by a ‘b’. The string transformation $baa \mapsto bba$ is presented in Figure 3. In the input model $\mathcal{M}(baa)$, the logical formula for $P'_a(2)$ evaluates to \perp and the logical formula for $P'_b(2)$ evaluates to \top . These facts together capture the following relationship between the input and output models: the ‘a’ at index 2 of the input model becomes a ‘b’ in the output model.

More recent work within model-theoretic phonology uses the framework of Boolean Monadic Recursive Schemes (BMRS) to represent string functions. Logical transductions over monadic string models within the BMRS framework have been used to represent phonological maps (Chandlee and Jardine, 2021; Jardine and Oakden, 2023), model process interaction (Oakden, 2021), and provide logical characterizations of the expressivity of phonological maps (Bhaskar et al., 2020, 2023; Yolyan, 2025; Chandlee and Lindell, forthcoming). This framework expresses the predicates in the output signature using an *if...then...else* syntax; logical transductions within the BMRS framework are referred to as *pro-*

grams. The logical transducer in (1) can equivalently be expressed as the BMRS program in (2).

(2) *Logical transducer with BMRS syntax*

$$\begin{aligned} P'_a(x) &= \text{if } P_b(px) \text{ then } \perp \text{ else } P_a(x) \\ P'_b(x) &= \text{if } P_a(x) \text{ then } P_b(px) \text{ else } P_b(x) \end{aligned}$$

The equation for $P'_b(x)$ in (2) expresses the following: if x is an ‘a’ in the input, then it will output as ‘b’ iff it is preceded by a ‘b’; otherwise, it will output as an ‘b’ iff it is a ‘b’ in the input. Chandlee and Jardine (2021) discuss how the if...then...else syntax of BMRS can be used to represent meaningful phonological generalizations such as licensing/blocking structures and elsewhere conditions. Section 4.1 of this paper discusses how this syntax is also valuable for expressing the generalizations involved in learning phonological maps. Thus, while these logical transductions can be expressed using propositional logic operators, the syntax of BMRS programs is useful for refining the learning problem.

2.3 Feature Models

The string models discussed so far have the requirement that the predicates $\{\times, \times, P_\sigma\}_{\sigma \in \Sigma}$ partition the domain. Phonological maps, however, target particular features or bundles of features (Jakobson et al., 1952; Clements and Hume, 1995). In order to model phonological words and maps, we use *unconventional* string models (Strother-Garcia et al., 2016), in which more than one predicate can hold at each index.² In the case where the predicates represent phonological features, we will refer to these unconventional string models as ‘feature models’. In this case, the monadic predicates range over an alphabet of phonological features \mathbb{F} , rather than an alphabet of characters. For simplicity, we use the same notation for each feature $F \in \mathbb{F}$ as the associated monadic predicate. That is, $F(x) = \top$ means the sound at index x of the model has feature $[+F]$ and $F(x) = \perp$ means it has feature $[-F]$.³ Moreover, because the requirement that only one predicate hold at each index is dropped for feature models, the predicates INITIAL/FINAL or MIN/MAX can be used in place of the boundary symbols \times/\times . For simplicity of demonstration, we consider here the limited collection of sounds /p,b,t,d,a,h/ over the four binary phonologi-

²String models in which exactly one predicate holds at each index are referred to as ‘conventional models’ in previous research because that is the convention used in computer science literature (e.g. Buchi, 1960).

cal features [sonorant], [coronal], [continuant], and [voice]. Each of these sounds corresponds with a unique combination of these four feature specifications, given in Figure 4. The sound /d/ for example, has the specification $[-\text{son}, +\text{cor}, -\text{cont}, +\text{voi}]$. Moreover, because postnasal voicing does not distinguish between nasal sounds, the feature [nasal] will be used to identify *any* nasal sound.

	p	b	t	d	a	h	N
[son]	-	-	-	-	+	-	+
[cor]	-	-	+	+	-	-	-
[voi]	-	+	-	+	+	-	+
[cont]	-	-	-	-	+	+	-
[nas]	-	-	-	-	-	-	+

Figure 4: Limited inventory of sounds and features

A logical transducer over feature models specifies output predicates $F'(x)$ for every feature $F \in \mathbb{F}$. The logical transducer which expresses the postnasal voicing map in Zoque is given in (3). A sample transduction of /mpama/ \rightarrow [mbama] ‘my clothing’ is presented in Figure 5. In order to represent feature models more compactly, we use feature matrices. The input feature model for /mpama/ in Figure 5, for example, can be represented with the shorthand notation in (4) which encodes exactly the same information as the input model.

(3) *Logical transducer for Zoque postnasal voicing*

$$\begin{aligned} [\text{voi}]'(x) &= [\text{voi}](x) \vee \\ &\quad (\neg[\text{cont}](x) \wedge [\text{nas}](px)) \\ F'(x) &= F(x) \quad \text{for all other } F \in \mathbb{F} \end{aligned}$$

³This convention is not the only way to model phonological maps within this framework. We may instead allow \mathbb{F} to contain both $[+F]$ and $[-F]$ as predicates. Moreover, Chandlee and Jardine (2021, pg.42) show that feature models can be adapted for non-binary feature systems as well. Further discussion of feature systems from a model-theoretic perspective can also be found in Nelson (2022).

input	m	p	a	m	a
	0	1	2	3	4
Initial(x)	\top	\perp	\perp	\perp	\perp
[son](x)	\top	\perp	\top	\top	\top
[cor](x)	\perp	\perp	\perp	\perp	\perp
[voi](x)	\top	\perp	\top	\top	\top
[cont](x)	\perp	\perp	\top	\perp	\top
[nas](x)	\top	\perp	\perp	\top	\perp
Final(x)	\perp	\perp	\perp	\perp	\top
[voi]'(x)	\top	\top	\top	\top	\perp
output	m	b	a	m	a

Figure 5: Logical transducer in (3) over $\mathcal{M}(mpama)$

(4) *Shorthand notation for $\mathcal{M}(mpama)$*

+son	-son	+son	+son	+son
-cor	-cor	-cor	-cor	-cor
+voi	-voi	+voi	+voi	+voi
-cont	-cont	+cont	-cont	+cont
+nas	-nas	-nas	+nas	-nas

2.4 Subfactors

The substructures which are relevant to this paper are *subfactors* of monadic feature models. This section does not delve into formal definitions of subfactors in general; discussions and definitions of subfactors can be found in Rogers and Lambert (2019); Rawski (2021); Chandlee et al. (2019); Strother-Garcia et al. (2016). The relevant result is that subfactors of monadic *string* models corresponds with substring models. The substrings of $w \in \Sigma^*$ are the strings $v \in \Sigma^*$ for which there exist $x, y \in \Sigma^*$ such that $w = xvy$. For example, the substrings of ‘bac’ are $\{\lambda, b, a, c, ba, ac, bac\}$. The *subfactors* of $\mathcal{M}(w)$ are the models $\mathcal{M}(v)$ such that v is a substring of w .

The subfactors of *unconventional* string models, however, are more complex than substring models because any number of predicates can be true at each index. Figure 6 presents the subfactors of a feature model with a domain that has cardinality 1. The feature [coronal] is left out of this figure in order to keep the space manageable. An example of a subfactor of [-nas,-son,-cont,-voi] is a model with no specification for [voice]: [-nas,-son,-cont]. These structures no longer represents a particular sound, but instead collections of sounds. The subfactor [-nas,-son,-cont] corresponds with the set of sounds $\{/p/, /b/, /t/, /d/\}$. Although the space of subfactors for feature models is very large, they allow for a formal representation of relevant linguistic generalizations because each subfactor in Figure 6 represents a natural class.

For any model \mathcal{M} , the set of subfactors is denoted $Subfact(\mathcal{M})$. A subfactor is called a k -factor if the domain has cardinality k . The set of k -factors of \mathcal{M} is denoted $Subfact_k(\mathcal{M})$. The subfactor relation between models is denoted \sqsubseteq .

2.5 Partially-Ordered Sets

A structure (X, \leq) is a partially-ordered set (poset) iff \leq is reflexive, antisymmetric, and transitive. Figure 6 is an example of a poset. Chandlee et al. (2019) show that this is true in general; for a model \mathcal{M} , the structure $(Subfact(\mathcal{M}), \sqsubseteq)$ is a poset.

Two types of posets which are particularly relevant for the discussion of learning in this paper

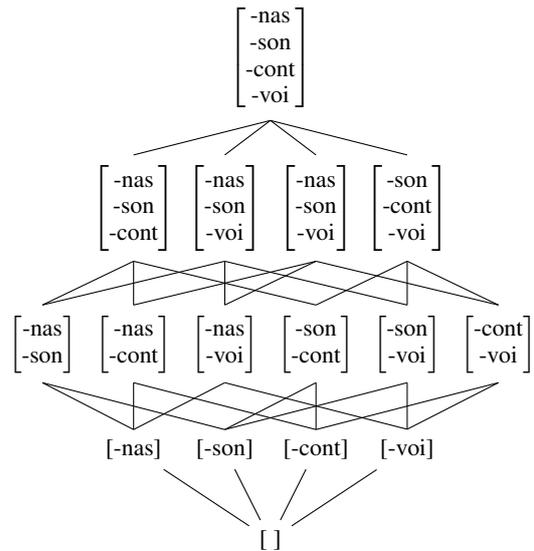


Figure 6: Poset of subfactors for feature models.

are filters and ideals, defined in Definitions 4 and 5, respectively. For every $x \in X$ of a poset, the subset of objects above x forms a filter, and the set of objects below x forms an ideal. The remainder of this paper shows how filters and ideals are used to generate and prune a hypothesis space.

Definition 4 (Filters). For a poset (X, \leq) , a filter is a non-empty set $F \subseteq X$ such that:

- (a) for every $x \in F$, if $x \leq y$ then $y \in F$
- (b) for every $x, y \in F$, there is some $z \in F$ such that $z \leq x$ and $z \leq y$

The **principle filter** generated by $x \in X$ is the set

$$\uparrow x := \{y \in X | x \leq y\}$$

Definition 5 (Ideals). For a poset (X, \leq) , an ideal is a non-empty set $I \subseteq X$ such that:

- (a) for every $x \in I$, if $y \leq x$ then $y \in I$
- (b) for every $x, y \in I$, there is some $z \in I$ such that $x \leq z$ and $y \leq z$

The **principle ideal** generated by $x \in X$ is the set

$$\downarrow x := \{y \in X | y \leq x\}$$

Posets have been used as a hypothesis space for various learning problems within phonology (Rawski, 2021; Chandlee et al., 2019; Tesar, 2013; Heinz et al., 2012; Heinz, 2010). A partially-ordered hypothesis space encodes entailment relations that are relevant for learning, which allow a large space to be pruned efficiently with a small number of data points. The next section presents two such examples to highlight this point.

3 Phonological Learning with Partially-Ordered Structures

Tesar (2013) uses a stress-length system to illustrate a learning algorithm for underlying forms with a partially-ordered hypothesis space.⁴ In this system every vowel in a word can be stressed or unstressed, and long or short. Consider the surface form [paká:], where the second vowel is both stressed and long. The space of possible underlying forms of [paká:] is a poset under the *relative similarity* relation, defined as follows: $x \leq y$ iff y is more similar to [paká:] than x . This poset encodes entailment relations relevant to learning. If the map /paká/ → [paká:] leads to inconsistency, then /paká/ is not a possible underlying form. This information entails that anything *less similar to* [paká:] than /paká/ will lead to inconsistency (in an output-driven map). Thus, being inconsistent is *downward entailing*, as illustrated in Figure 7 with the candidates marked \times .

Although the poset contains 16 possible candidates for the underlying form of [paká:] (Figure 8), downward entailment of inconsistency allows the space to be pruned efficiently. If the learner determines that /paká/ cannot be the underlying form, the entire substructure highlighted in gray (i.e. the principle ideal \downarrow /paká/) in Figure 8 can be removed from the hypothesis space. As Tesar (2013) explains, the benefit of the structured hypothesis space is computational efficiency:

[...]the possible optimality of a node entails upward, while the definite non-optimality of a node entails downward. This makes it possible to draw conclusions about whole sublattices of candidates based on the evaluation of only a single candidate. [...] This converts exponential search into linear search: the number of possible underlying forms is exponential in the number of features, but the number of forms to actually be

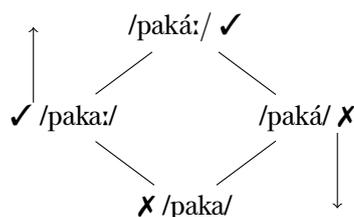


Figure 7: Entailments for learning underlying forms

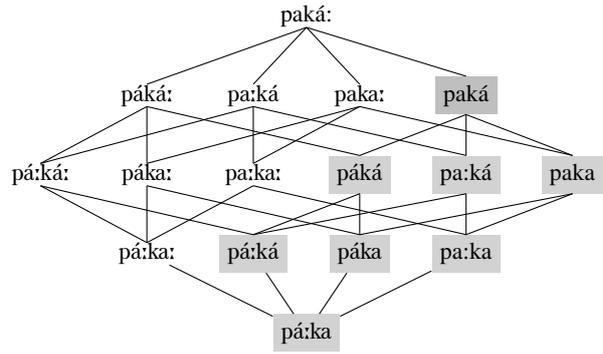


Figure 8: Removing \downarrow /paká/ from hypothesis space.

tested is linear in the number of features. (Tesar, 2013, pg.287-289)

Similar observations are made by Chandlee et al. (2019) for the problem of learning phonotactic grammars. In this case, subfactors form a poset of potential banned structures in a grammar. Ungrammaticality is upward entailing; if a structure is ungrammatical, then all its superfactors must be ungrammatical. For example, if [+nas][-voi,-cont] is an ungrammatical surface form (as in the case of Zoque), then [+nas][-nas,-voi,-cont] must also be ungrammatical. On the other hand, grammaticality is downward entailing; if a structure is grammatical, then all its subfactors must also be grammatical. These entailments are illustrated in Figure 9, where ungrammatical structures are represented by \times and grammatical structures are represented by \checkmark .

Similar to the previous example, these entailments are used to prune the hypothesis space. In this case, if a structure is ungrammatical, all the structures above it (i.e. the principle filter) will be removed from the hypothesis space. This process is discussed in the next section.

⁴This work is specifically for output-driven maps, which are outside the scope of this paper.

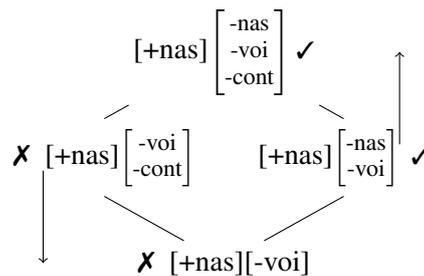


Figure 9: Entailments for learning grammars

3.1 Bottom-Up Feature Inference Algorithm (BUFIA)

The learning algorithm introduced by Chandlee et al. (2019) is called the Bottom-Up Factor Inference Algorithm (BUFIA). The goal of BUFIA is to learn the *banned* k -factors of a language from a positive sample (i.e. collection of grammatical words from the language). BUFIA runs as follows. There is a finite sample S containing words from a language. The goal is to learn the *grammar* of the language, which is represented as a set of k -factors which are not permitted in the language. Because there are many ways to express the ungrammatical k -factors, a further goal is to find the most general description. For this reason, the algorithm searches through the hypothesis space *bottom-up*. If some k -factor x in the space is not found in any surface form in S , then it is added to the set of constraints. Because ungrammaticality is upward entailing, all of x 's superfactors must also be ungrammatical. Thus, the entire filter $\uparrow x$ is removed from the hypothesis space since the grammaticality of these structures no longer needs to be considered. In this way, only the minimal (most general) ungrammatical k -factors are part of the resulting grammar.

BUFIA is illustrated here with the example of postnasal voicing in Zoque. For brevity, the features [sonorant] and [coronal] are left out of the discussion because they are ultimately not necessary to identify the minimal ungrammatical 2-factor in Zoque.⁵ The surface forms from Section 1 provide evidence that all the subfactors of [+nas][-voi,-cont] are grammatical. These subfactors are presented in Figure 10 with N, V, and C for the features [nasal], [voi], and [cont], respectively. The form [tatah] ‘father’ provides evidence that [][-voi,-cont] is a grammatical structure. The form [nhayah] ‘my husband’ provides evidence that [+nas][-voi] is grammatical. The form [ndatah] ‘my father’ provides evidence that [+nas][-cont] is grammatical. Moreover, all the subfactors of these three 2-factors are grammatical. The 2-factor [+nas][-voi,-cont] is the *minimal* structure in the poset that is not found in any surface form of Zoque; all the structures below it are grammatical, while all the structures above it ungrammatical.

⁵The [-voi,-son] /h/ in [nhayah] ‘my husband’ does not have a voiced counterpart in Zoque, and therefore does not become voiced after a nasal (Wonderly, 1951). Thus, [nhayah] provides evidence that [+nas][-voi,-son] (and every subfactor of it) is grammatical. The feature [cont] is therefore necessary to distinguish the sounds which can be preceded by a

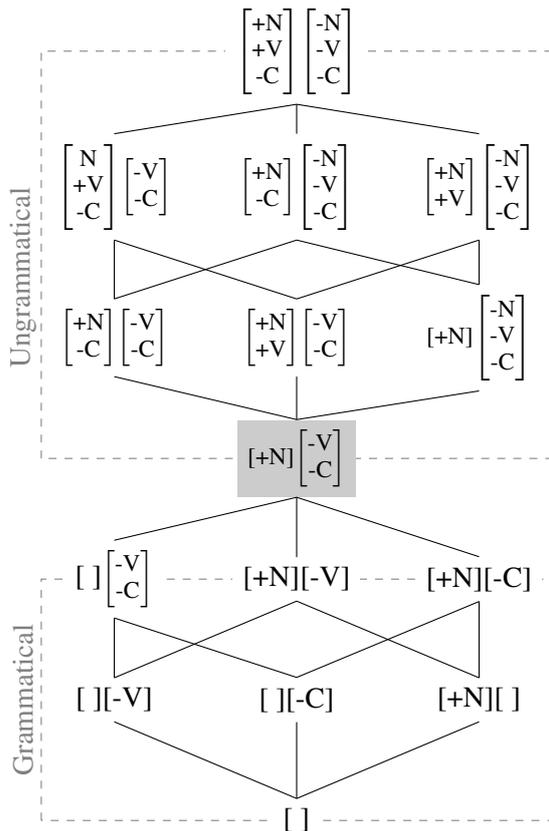


Figure 10: Minimal ungrammatical structure in Zoque.

These data points are revisited in the next section to show how the same structures and principles can be used to learn the postnasal voicing map. Further discussion of BUFIA and abduction principles used in learning can be found in Rawski (2021). Recent applications of BUFIA can also be found in Li (2025) and Payne (2024).

As a final note, the learning problem for BUFIA (Def. 9 of Chandlee et al., 2019) is not stated in terms of finding the *correct* grammar in the limit, as in Gold (1967), but instead finding the most *general* grammar that is consistent with the sample of surface forms. The goal of learning phonological maps is similar: generate a logical transducer that expresses the most general phonological map that is consistent with the sample of input-output pairs. A formal statement of the learning problem in this paper is left for future work.

nasal from the sounds which cannot. Moreover, since the ungrammaticality of [+nas][-voi,-cont] entails the ungrammaticality of both [+nas][-voi,-son,-cont] and [+nas][-voi,-cor,-cont], [son] and [cor] are not necessary to represent the ungrammatical structures in Zoque because a more general subfactor suffices.

4 Learning Logical Transducers

Chandlee and Lindell (forthcoming) show that non-recursive quantifier-free logical transductions over monadic models are the logical characterization of finite-to-one ISL functions. A string function $f : \Sigma^* \rightarrow \Gamma^*$ is finite-to-one iff for every $w \in \Gamma^*$, the set $\{u \in \Sigma^* \mid f(u) = w\}$ is finite. In other words, every surface form has finitely many underlying forms that map to it under the application of f . The output of the learning procedure introduced here is a quantifier-free BMRS program. This means we assume that we are always learning finite-to-one phonological maps. We moreover restrict the maps we consider to those that order-preserving and length-preserving. That is, epenthesis, deletion, and metathesis are not considered here. Future extensions to non-length-preserving maps are discussed in Section 5.

The examples from Tesar (2013) and Chandlee et al. (2019) illustrate how two very different learning goals can be pursued with partially-ordered hypothesis spaces, where entailment relations encoded in the structure make it possible for the learner to remove large substructures with a single data point. The remainder of this paper shows that the same observation holds of the problem of learning environments where features undergo change, and ultimately logical transducers which express a phonological map. The main idea expressed in this section is that learning a phonological map amounts to learning several phonotactic grammars simultaneously. In that sense, the learning procedure here is a way of *lifting* the ideas in Chandlee et al. (2019) to maps. This idea is formalized with the normal form for logical transducers proposed in Definition 6 and the main result in Theorem 7.

4.1 Normal Form for Logical Transductions

Consider a signature $\mathcal{S} = \{P_1, \dots, P_n, s, p\}$, where each P_i is a unary predicate. A *term* over this signature is recursively defined as follows: a variable is a term; for every term T , $p(T)$ and $s(T)$ are terms. The *atoms* over this signature are defined as

$$\text{atoms}(\mathcal{S}) := \{P_i(T) \mid T \text{ is a term}\}$$

Definition 6. Let $\{P'_\sigma\}_{\sigma \in \Sigma}$ be a collection of formulas defined over an input signature Σ . Each P'_σ is in **normal form** if it is expressed as

$$P'_\sigma(x) = \text{if } P_\sigma(x) \text{ then } \neg\phi_\sigma(x) \text{ else } \psi_\sigma(x)$$

where $\phi_\sigma(x)$ and $\psi_\sigma(x)$ are either \top , \perp , or disjunctive normal form expressions over $\text{Atoms}(\Sigma)$.⁶

Although normal form programs use the `if...then...else` syntax, the ‘then’ and ‘else’ parts of the equations are expressed with propositional operators.⁷ This brings into question why we would bother using the `if...then...else` syntax at all. The reason for this representations is that the syntax compartmentalizes each equation in a program into two components: the ‘then’ expression and the ‘else’ expression. We consider first two examples which illustrate the information the ‘then’ and ‘else’ parts of these equations encode.

The normal form of the program in (2) is presented in (5). Recall that this program expresses the same function as the rewrite rule $a \rightarrow b/b_$. The four pieces of information contained within this normal form program are presented in (6). The expression $\phi_a(x)$ is a logical description of the environment where an ‘a’ input becomes a non-‘a’ output. $\psi_b(x)$ is a logical description of the environment where a non-‘b’ input becomes a ‘b’ output. Similarly, $\psi_a(x) = \perp$ and $\phi_b(x) = \perp$ encode the facts that a non-‘a’ input never becomes an ‘a’, and a ‘b’ input never becomes a non-‘b’, respectively. The normal form in (5) rearranges the atoms of (2) into a syntactic form that explicitly references environments where change takes place.

(5) *Logical transducer from (2) in normal form*

$$\begin{aligned} P_a^*(x) &= \text{if } P_a(x) \text{ then } \neg P_b(px) \text{ else } \perp \\ P_b^*(x) &= \text{if } P_b(x) \text{ then } \neg\perp \\ &\quad \text{else } (P_a(x) \wedge P_b(px)) \end{aligned}$$

(6) *Information encoded in the normal form in (5)*

$$\begin{aligned} \phi_a(x) &= P_b(px) & \psi_a(x) &= \perp \\ \phi_b(x) &= \perp & \psi_b(x) &= P_a(x) \wedge P_b(px) \end{aligned}$$

A similar translation can be done for the postnasal voicing transducer in (3), presented in (7). The highlighted expression $[nas](px) \wedge \neg[cont](x)$ is a logical description of the environment in which a [-voi] sound in the input model becomes [+voi] in the output (i.e. $\psi_{[voi]}(x)$).

(7) *Logical transducer from (3) in normal form*

$$\begin{aligned} [voi]^*(x) &= \text{if } [voi](x) \text{ then } \neg\perp \\ &\quad \text{else } ([nas](px) \wedge \neg[cont](x)) \\ F^*(x) &= \text{if } F(x) \text{ then } \neg\perp \text{ else } \perp \end{aligned}$$

⁶We assume that the input and output alphabets are the same. This means that we assume the underlying and surface forms have the same feature inventory. This assumption is not necessary, and is only used for simplicity.

⁷Every propositional operator can be expressed using `if...then...else`. For example, $p \vee q = \text{if } p \text{ then } \top \text{ else } q$.

The intuition for the normal form in Definition 6 is that every string function can be uniquely identified with $\{\phi_\sigma, \psi_\sigma\}_{\sigma \in \Sigma}$, which encode information about the changes the function induces, and the environments those changes take place. With respect to learning, the significance of the normal form is that it refines the learning problem in the following way: in order to learn a program $\{P'_\sigma\}_{\sigma \in \Sigma}$, it is sufficient to learn $\{\phi_\sigma, \psi_\sigma\}_{\sigma \in \Sigma}$. This fact is the immediate result of Theorem 7.

Theorem 7. *Every program $\{P'_\sigma\}_{\sigma \in \Sigma}$ over an input signature $\Sigma = \langle P_\sigma, s, p \rangle_{\sigma \in \Sigma}$ is logically equivalent to a program $\{P^*_\sigma\}_{\sigma \in \Sigma}$ where each P^*_σ is in normal form.*

The proof of Theorem 7 is presented in the appendix. The main idea is that for any output predicate P'_σ , there is a systematic way to construct an equivalent normal form predicate P^*_σ by generating a truth table of all the atoms in the definition of P'_σ , and determining ϕ_σ and ψ_σ from that truth table.

For phonological maps, learning a program over feature models amounts to learning $\{\phi_F, \psi_F\}_{F \in \mathbb{F}}$, where ϕ_F (ψ_F) expresses the environment where a [+F] ([-F]) sound in the underlying form becomes [-F] ([+F]) in the surface form. In this way, learning programs in this normal form amount to learning *phonological generalizations* regarding environments where features undergo change. The observation that connects this learning goal to the examples presented in Section 3 is that the space of possible expressions for ϕ_F and ψ_F can also be modeled as a poset, with relevant entailment relations that allow the space to be pruned efficiently.

Consider the case of postnasal voicing and the goal of learning the environment where a [-voice] input becomes [+voice] ($\psi_{[voi]}$). A similar diagram to Figures 7 and 9 is presented in Figure 11 for the goal of learning the environment where voicing takes place. Environments where voicing takes place are upward entailing, while environments where voicing does *not* take place are downward

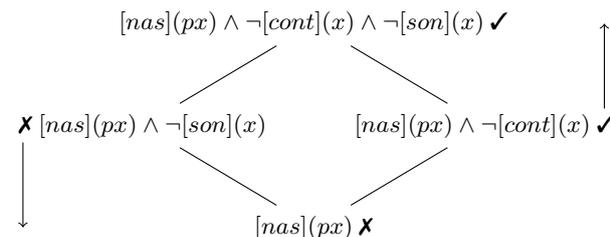


Figure 11: Entailments for learning environments

entailing. For example, if being a [-cont] that is immediately preceded by a [+nas] is sufficient to trigger voicing, then being a [-cont,-son] that is immediately preceded by a [+nas] is also sufficient. If an input-output pair indicates that some environment does *not* trigger voicing, then its entire principle ideal can be removed from the hypothesis space, similar to the example of learning underlying forms in Figure 8.

Each logical expression in Figure 11 corresponds with a description of a 2-factor. The expression $[nas](px) \wedge \neg[cont](x)$, for example, describes the 2-factor [+nas][-cont]. The subfactors which are used to learn grammars in BUFIA can therefore be adapted to learning environments. The distinction is that when we are learning *maps* rather than surface constraints, some index of the subfactor is the target of the map. Every k -factor we consider will therefore have an extra predicate target which identifies the unique index of the structure that is the target of the feature change in question. Figure 12 presents the 2-factor [+nas][-voi,-cont] with an additional predicate target which picks out the index that the voicing process targets. As a shorthand, we underline the subfactor that is the target of the map, as (8).

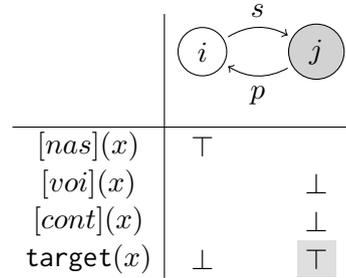


Figure 12: 2-factor targeted by postnasal voicing

(8) *Shorthand notation for 2-factor in Figure 12*

$$[+nas] \left[\begin{array}{l} \underline{-cont} \\ -son \end{array} \right]$$

4.2 Variables and Initialization

In the case of learning grammars, there is a single hypothesis space consisting of all the possible subfactors. In the case of learning maps, there are several hypothesis spaces: one for each of the learning targets $\{\phi_F, \psi_F\}_{F \in \mathbb{F}}$. Thus, learning a program can be viewed as learning several grammars simultaneously. This section shows how principle ideals are used to construct each of the hypothesis spaces.

For every $X \in \{\phi_F, \psi_F\}_{F \in \mathbb{F}}$ we need two sets

of k -factors: $V(X)$ and $\widehat{V}(X)$. $V(X)$ contains the relevant k -factors in which the targeted change takes place. $V(\phi_F)$, for example, contains all the k -factors in which an index x is such that $F(x) = \top$ in the input model but $F'(x) = \perp$ in the output model. $\widehat{V}(\phi_F)$ contains all the k -factors in which an index x is such that $F(x) = \top$ in the input model and $F'(x) = \top$ in the output model. In other words, $V(\phi_F)$ contains all the k -factors in which a [+F] sound becomes [-F], and $\widehat{V}(\phi_F)$ contains all the k -factors in which [+F] sounds do not undergo change. In this way, the collections V and \widehat{V} encode positive and negative evidence for the environments where features undergo change. The hypothesis spaces are then defined in Definition 8.

Definition 8. For every $X \in \{\phi_F, \psi_F\}_{F \in \mathbb{F}}$, the corresponding *hypothesis space* is the set

$$\mathcal{H}(X) := \bigcup_{x \in V(X)} \downarrow x - \bigcup_{x \in \widehat{V}(X)} \downarrow x$$

subject to the following two restrictions:

- (i) Every $\mathcal{M} \in \mathcal{H}(X)$ has exactly one index i in the domain of \mathcal{M} such that $\text{target}(i) = \top$.
- (ii) If X is ϕ_F for some $F \in \mathbb{F}$, then for every $\mathcal{M} \in \mathcal{H}(X)$, $F(i) = \top$ must hold for the unique i such that $\text{target}(i) = \top$. If X is ψ_F , then $F(i) = \perp$ must hold.

The posets defined in Definition 8 takes all the k -factors in which the relevant change was observed in some pair of input-output models, and removes the subfactors in which the relevant change was *not* observed. The restrictions in (i) and (ii) then ensure that only the subfactors which represent relevant environments are included in the space. The restriction in (i) ensures that every structure in the hypothesis space represents an *environment*. The restriction in (ii) ensures that the hypothesis space includes only the *relevant* environments. If the goal is to learn the environment where a [-voi] sound becomes [+voi], then the hypothesis space should only include environments where the target of the change is [-voi]. This means that $\mathcal{H}(\psi_{[-voi]})$ will only include k -factors where the unique index that satisfies the predicate target is such that $[-voi](x) = \perp$.

The learning procedure starts with an initial hypothesis, and generates a new hypothesis with each input-output pair of models. Before any data have been observed, $V(X)$ and $\widehat{V}(X)$ are empty for

each $X \in \{\phi_F, \psi_F\}$. This means that the initial hypothesis space for each X will also be empty by definition. Consequently the set of minimal elements, will be empty. In this case, the corresponding logical formula for X must be \perp . The initial variables, hypothesis, and corresponding BMRS program are summarized in (9). The consequence of V and \widehat{V} being empty is that the initial (null) hypothesis corresponds with a BMRS program that represents the identity map.

(9) Initial Variables and Hypothesis

Variables; for all $X \in \{\phi_F, \psi_F\}_{F \in \mathbb{F}}$

$$V(X) = \emptyset$$

$$\widehat{V}(X) = \emptyset$$

$$\text{min}(\mathcal{H}(X)) = \emptyset$$

Individual learning goals

$$\{\phi_F(x) = \perp; \psi_F(x) = \perp\}_{F \in \mathbb{F}}$$

Initial Hypothesis as a BMRS program

$$\{F'(x) = \text{if } F(x) \text{ then } \neg \perp \text{ else } \perp\}_{F \in \mathbb{F}}$$

$$\equiv \{F'(x) = F(x)\}_{F \in \mathbb{F}}$$

4.3 Updating the Hypothesis

When an input-output pair provides evidence for a feature change, the corresponding hypothesis space is updated. For example, if a pair of models $(\mathcal{M}, \mathcal{M}')$ is such that some index x is [-F] in the input model but [+F] in the output model, the set $V(\psi_F)$ will be updated with the k -factor(s) of \mathcal{M} which contain x as the target. If, on the other hand, some index x is [-F] in the input and remains [-F] in the output, the set $\widehat{V}(\psi_F)$ will be updated with the relevant k -factors. Each time V or \widehat{V} is updated, the set of minimal elements must also be updated. Thus, there are two possible types of updates.

Consider first the case in which a k -factor x is added to $V(X)$ for some X . In this case, the hypothesis space $\mathcal{H}(X)$ grows, and potentially the set of minimal elements grow. This means that when a new k -factor x is added to $V(X)$, the update function maintain all previous minimal elements and only determines whether to add new ones. The procedure for determining which elements to add to $M(X)$ is as follows: we traverse the structure $\downarrow x$ bottom-up in order to find the minimal subfactors of x which are neither subfactors of any current minimal elements, nor subfactors of any k -factor in $\widehat{V}(X)$. Because both k and the number of features we consider are fixed, the structure $\downarrow x$ is always of fixed size. Moreover, we only need to know the set of factors in $M(X)$ and $\widehat{V}(X)$ in order to determine the new set of minimal elements.

In the case where a new k -factor x is added to $\widehat{V}(X)$, the update function must remove any element $m \in M(X)$ such that $m \sqsubseteq x$, and potentially add new minimal elements. In this case, there are two possibilities: if $x \in V(X)$, then m is removed from the set of minimal elements; otherwise m is replaced by new minimal elements that are immediately above m in the poset. In this case, we only need information about $V(X)$ and $M(X)$ in order to update the set of minimal elements.

Both types of updates therefore amount to traversing a restricted space in order to find new minimal elements. Although the hypothesis space can get very large, for each $X \in \{\phi_F, \psi_F\}$, only the sets $M(X)$, $V(X)$, and $\widehat{V}(X)$ need to be maintained in memory in order to update the hypothesis. The next section illustrates the update process with the example of postnasal voicing.

4.4 Case Study: Postnasal Voicing in Zoque

Consider first the input-output pair of models in Figure 5 representing $/mpama/ \rightarrow [mbama]$. The input-output models for this transformation were presented in Figure 5. The sound at index 1 is [-voi] in the input model and [+voi] in the output model. This means that the set $V(\psi_{[voi]})$ must be updated with the 2-factors that contain the sound at index 1 as a target. There are two possible 2-factors in the input model which contain the sound at index 1. These are presented in Figure 13. The updated hypothesis space is presented in Figure 14. The corresponding updates are summarized in (10).

	m	p
[nas](x)	⊥	⊥
[voi](x)	⊥	⊥
[cont](x)	⊥	⊥
target(x)	⊥	⊤

	p	a
[nas](x)	⊥	⊥
[voi](x)	⊥	⊤
[cont](x)	⊥	⊤
target(x)	⊤	⊥

Figure 13: 2-factors added to $V(\psi_{[voi]})$

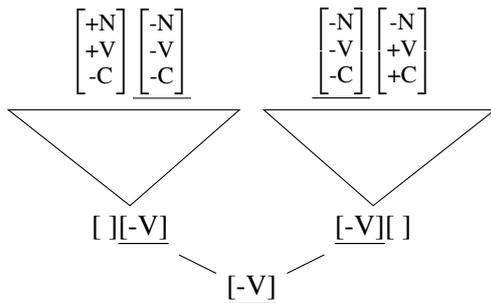


Figure 14: $\mathcal{H}(\psi_{[voi]})$ given (/mpama/, [mbama]).

Since being [-voi] is presumed to be true when learning $\psi_{[voi]}$ (by restriction (ii) of Definition 8), the minimal element is [-V] rather than \emptyset . This element corresponds with $\psi_{[voi]}(x) = \top$. We could also set $\psi_{[voi]} = \neg[voi](x)$; this would give a logically equivalent program and corresponding map, but with unnecessary redundancy. The new hypothesis therefore says that a [-voi] sound *always* becomes [+voi]. This hypothesis corresponds with the *most general* hypothesis that accounts for the input-output sample observed.

(10) Hypothesis update given $/mpama/ \rightarrow [mbama]$

$$\begin{aligned} \min(\mathcal{H}(\psi_{[voi]})) &= \{[-V]\} \\ \psi_{[voi]}(x) &= \top \end{aligned}$$

as a BMRS program

$$[voi]^*(x) = \text{if } [voi](x) \text{ then } \neg\perp \text{ else } \top$$

as a phonological map

$$[-voi] \rightarrow [+voi]$$

We consider next the case where the underlying and surface forms are the same. Consider the surface form [tatah] ‘father’. The hypothesis in (10) predicts that we should not see this surface form; we would instead expect to see [dadah]. This surface form therefore provides evidence that the previous hypothesis is incorrect. In this case, the pair of 2-factors in Figure 15 are added to $\widehat{V}(\psi_{[voi]})$. The updated hypothesis space is presented in Figure 16. The corresponding updates are summarized in (11)

	a	t
[nas](x)	⊥	⊥
[voi](x)	⊤	⊥
[cont](x)	⊤	⊥
target(x)	⊥	⊤

	t	a
[nas](x)	⊥	⊥
[voi](x)	⊥	⊤
[cont](x)	⊥	⊤
target(x)	⊤	⊥

Figure 15: 2-factors added to $\widehat{V}(\psi_{[voi]})$

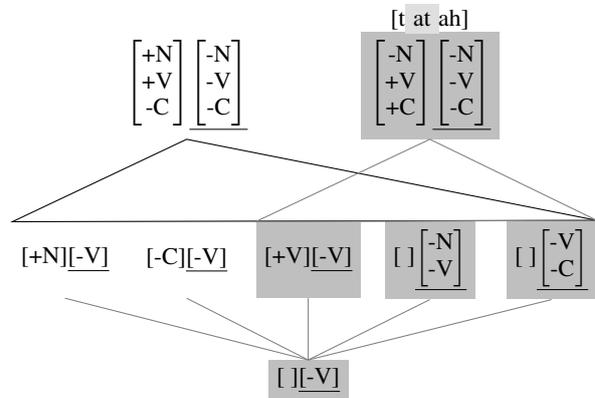


Figure 16: Updated $\mathcal{H}(\psi_{[voi]})$ given (/tatah/, [tatah]).

are also ISL (Chandlee, 2014). An interesting extension of this work would be to develop a learning procedure for input-output index correspondences. Consider for example a simple deletion map $a \rightarrow \lambda/b_b$, where an ‘a’ is deleted whenever it is between two ‘b’s. Figure 18 presents the input-output correspondences for the string transformation $ababa \mapsto abba$.

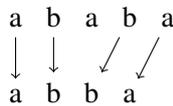


Figure 18: Input-output correspondences

Recent work from Li (2025) uses BUFIA to learn tonotactic patterns, where the structures in the hypothesis space encode autosegmental representations. An interesting question to pursue further is whether a similar hypothesis space can be used to represent input-output correspondences. Doing so would amount to learning the environments where features undergo change and the input-output correspondence between segments in the underlying and surface forms simultaneously.

Adding output structures to the hypothesis space would also make it possible to learn a larger class of transductions. Theorem 7 stated that every non-recursive BMRS programs can be expressed as an equivalent normal form program. However, it seems intuitive that this result should extend to programs in general because every map can be uniquely identified in terms of environments where features undergo change; this is after all how SPE-style rewrite rules express a function. However, such a result requires reasoning about recursion in BMRS, which uses a fixed point semantics (Bhaskar et al., 2020), and is outside the scope of this paper. A valuable result for future work is an extension of Theorem 7 to more general programs.

Ultimately, the purpose of this paper was to present a starting point for a model-theoretic approach to learning phonological maps, and the merits of such an approach. The ideas and procedure presented here can be extended to different representations and classes of functions. In order to extend this work to more complex representations, the next step for this research is an implementation of the learning procedure presented here. Doing so would also make it possible to obtain empirical results.

References

- Siddharth Bhaskar, Jane Chandlee, and Adam Jardine. 2023. [Rational functions via recursive schemes](#). *Preprint*, arXiv:2302.03074.
- Siddharth Bhaskar, Jane Chandlee, Adam Jardine, and Christopher Oakden. 2020. [Boolean monadic recursive schemes as a logical characterization of the subsequential functions](#). In *Language and Automata Theory and Applications (LATA)*, Lecture Notes in Computer Science, pages 157–169. Springer.
- Richard Buchi. 1960. [Weak second-order arithmetic and finite automata](#). *Mathematical Logic Quarterly*, 6:66–92.
- Jane Chandlee. 2014. *Strictly Local Phonological Processes*. Ph.D. thesis, University of Delaware.
- Jane Chandlee, Rémi Eyraud, and Jeffrey Heinz. 2014. [Learning strictly local subsequential functions](#). *Transactions of the Association for Computational Linguistics*, 2:491–504.
- Jane Chandlee, 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*, pages 91–101, Toronto, Canada. Association for Computational Linguistics.
- Jane Chandlee and Jeffrey Heinz. 2018. [Strict locality and phonological maps](#). *Linguistic Inquiry*, 49(1):23–60.
- Jane Chandlee and Adam Jardine. 2019a. [Autosegmental input strictly local functions](#). In *Transactions of the Association for Computational Linguistics*, volume 7.
- Jane Chandlee and Adam Jardine. 2019b. [Quantifier-free least fixed point functions for phonology](#). In *Proceedings of the 16th Meeting on the Mathematics of Language (MOL)*, pages 50–62. Association for Computational Linguistics.
- Jane Chandlee and Adam Jardine. 2021. [Computational universals in linguistic theory: Using recursive programs for phonological analysis](#). *Language*, 97:485–519.
- Jane Chandlee and Steven Lindell. forthcoming. [Logical perspectives on strictly local transformations](#). In *Doing Computational Phonology*. Oxford University Press.
- Nick Clements and Elizabeth Hume. 1995. [The internal organisation of speech sounds](#). In John Goldsmith, editor, *The Handbook of Phonological Theory*, chapter 7, pages 245–307. Blackwell.
- Bruno Courcelle. 1994. [Monadic second-order definable graph transductions: a survey](#). *Theoretical Computer Science*, 126(1):53–75.

- Bruno Courcelle and Joost Engelfriet. 2012. [Monadic second-order transductions](#). In *Graph Structure and Monadic Second-Order Logic: A Language-Theoretic Approach*, chapter 7, pages 505–577. Cambridge University Press.
- Hossep Dolatian. 2020. *Computational locality of cyclic phonology in Armenian*. Ph.D. thesis, Stony Brook University.
- Joost Engelfriet and Hendrik Jan Hoogeboom. 2001. [Mso definable string transductions and two-way finite-state transducers](#). *ACM Transactions on Computational Logic (TOCL)*, 2(2):216–254.
- Mark Gold. 1967. Language identification in the limit. *Information and Control*, 10:447–474.
- Jeffrey Heinz. 2010. [String extension learning](#). In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 897–906, Uppsala, Sweden. Association for Computational Linguistics.
- Jeffrey Heinz, Anna Kasprzik, and Timo Kötzing. 2012. [Learning in the limit with lattice-structured hypothesis spaces](#). *Theoretical Computer Science*, 457:111–127.
- Roman Jakobson, Gunnar Fant, and Morris Halle. 1952. *Preliminaries to Speech Analysis*. MIT Press.
- Adam Jardine. 2017. [The local nature of tone-association patterns](#). *Phonology*, 34:385–405.
- Adam Jardine and Chris Oakden. 2023. [Computing process-specific constraints](#). *Linguistic Inquiry*, pages 1–9.
- Dakotah Lambert and James Rogers. 2020. Tier-based strictly local stringsets: Perspectives from model and automata theory. In *Proceedings of the Society for Computation in Linguistics 2020*, pages 159–166. Association for Computational Linguistics.
- Han Li. 2025. [Learning tonotactic patterns over autosegmental](#). *Proceedings of the 2023 and 2024 Annual Meetings on Phonology*.
- Leonid Libkin. 2004. *Elements of Finite Model Theory*. Springer.
- David Marker. 2002. *Model Theory: An Introduction*. Springer.
- Scott Nelson. 2022. A model theoretic perspective on phonological feature systems. In *Proceedings of the Society for Computation in Linguistics*, volume 5.
- Chris Oakden. 2021. *Modeling phonological interactions using recursive schemes*. Ph.D. thesis, Rutgers University.
- Sarah Payne. 2024. [A generalized algorithm for learning positive and negative grammars with unconventional string models](#). In *Proceedings of the Society for Computation in Linguistics (SCiL)*, pages 75–85.
- Jonathan Rawski. 2021. *Structure and Learning in Natural Language*. Ph.D. thesis, Stony Brook University.
- James Rogers. 2003. [Syntactic structures as multi-dimensional trees](#). *Research on Language and Computation*, 1:265–305.
- James Rogers and Dakotah Lambert. 2019. [Some classes of sets of structures definable without quantifiers](#). In *Proceedings of the 16th Meeting on the Mathematics of Language*, pages 63–77. Association for Computational Linguistics.
- James Rogers and Geoffrey Pullum. 2011. [Aural pattern recognition experiments and the subregular hierarchy](#). *Journal of Logic, Language and Information*, 20:329–342.
- Kristina Strother-Garcia. 2018. [Imdlawn Tashlhiyt Berber syllabification is quantifier-free](#). In *Proceedings of the Society for Computation in Linguistics*, page 145–153.
- Kristina Strother-Garcia. 2019. *Using model theory in phonology: a novel characterization of syllable structure and syllabification*. Ph.D. thesis, University of Delaware.
- Kristina Strother-Garcia, Jeffrey Heinz, and Hyun Jin Hwangbo. 2016. Using model theory for grammatical inference: a case study from phonology. In *Proceedings of The 13th International Conference on Grammatical Inference*, volume 57, pages 66–78. PMLR.
- Bruce Tesar. 2013. *Output-Driven Phonology: Theory and Learning*. Cambridge University Press.
- Mai Vu, Ashkan Zehfroosh, Kristina Strother-Garcia, Michael Sebok, Jeffrey Heinz, and Herbert G. Tanner. 2018. [Statistical relational learning with unconventional string models](#). *Frontier in Robotics and AI*, 5.
- William L. Wonderly. 1951. Zoque II: Phonemes and Morphophonemics. *International Journal of American Linguistics*, 17:105–123.
- Tatevik Yolyan. 2025. [A logical characterization of weak determinism as simultaneous application](#). *Journal of Logic, Language and Information*, 34(1):89–153.

A Appendix

Proof of Theorem 7. Let $\{P'_\sigma\}_{\sigma \in \Sigma}$ be a non-recursive BMRS program over the input signature $\Sigma = \{P_\sigma, p, s\}_{\sigma \in \Sigma}$. Fix $\sigma \in \Sigma$, and let $\text{atoms}(P'_\sigma(x))$ be the set of Boolean expressions in $\text{atoms}(\Sigma)$ which appear in the equation $P'_\sigma(x)$. We can then construct a truth table with columns $\mathbb{C} = \text{atoms}(P'_\sigma(x)) \cup \{P'_\sigma(x)\}$. For every $\alpha(x) \in \mathbb{C}$ and row r , we say $r \models \alpha(x)$ if and only if $\alpha(x)$ evaluates to \top in row r of the table. For every row r , we define a conjunction consisting of all the information contained in row r as follows

$$\Psi_r(x) := \bigwedge_{\alpha(x) \in \text{Atoms}(P'_\sigma(x))} \begin{cases} \alpha(x) & \text{if } r \models \alpha(x) \\ \neg\alpha(x) & \text{otherwise} \end{cases}$$

We show that there is a normal form expression P_σ^* such that P'_σ is logically equivalent to P_σ^* .

Case 1. Consider first the case where $P_\sigma(x) \notin \mathbb{C}$. In other words, the equation $P'_\sigma(x)$ is not defined in terms of the input predicate $P_\sigma(x)$. Set $\psi_\sigma(x)$ and $\phi_\sigma(x)$ as follows.

$$\begin{aligned} \phi_\sigma(x) &:= \bigvee_{\{r:r \not\models P'_\sigma(x)\}} \Psi_r(x) \\ \psi_\sigma(x) &:= \bigvee_{\{r:r \models P'_\sigma(x)\}} \Psi_r(x) \end{aligned}$$

Let $P_\sigma^*(x) = \text{if } P_\sigma(x) \text{ then } \neg\phi_\sigma(x) \text{ else } \psi_\sigma(x)$. By construction, $\phi_\sigma(x) = \neg\psi_\sigma(x)$, and therefore $P_\sigma^*(x) \equiv \psi_\sigma(x)$. Because $\psi_\sigma(x)$ exhaustively contains all the possible situations in which $P'_\sigma(x)$ evaluates to \top , $\psi_\sigma(x) \equiv P'_\sigma(x)$. Thus, P_σ^* is equivalent to P'_σ .

Case 2. $P_\sigma(x) \in \text{atoms}(P'_\sigma(x))$. Define $\phi_\sigma(x)$ and $\psi_\sigma(x)$ as follows.

$$\phi_\sigma(x) := \bigvee_{\{r:r \models P_\sigma(x), r \not\models P'_\sigma(x)\}} \Psi_r(x) \quad (1)$$

$$\psi_\sigma(x) := \bigvee_{\{r:r \not\models P_\sigma(x), r \models P'_\sigma(x)\}} \Psi_r(x) \quad (2)$$

Set $P^*(x) = \text{if } P_\sigma(x) \text{ then } \neg\phi_\sigma(x) \text{ else } \psi_\sigma(x)$. We show that $P_\sigma^*(x)$ evaluates to \top iff $P'_\sigma(x)$ evaluates to \top . Consider first the case where $P_\sigma(x)$ evaluates to \top . Then $P_\sigma^*(x)$ evaluates to \top iff $\phi_\sigma(x)$ evaluates to \perp . By the definition in (1), this means Ψ_r evaluates to \perp for every row r such that $r \models P_\sigma(x)$ and $r \not\models P'_\sigma(x)$. In other words, none of the conditions under which $P'_\sigma(x)$ can evaluate to \perp are satisfied.

Thus, $\phi_\sigma(x)$ evaluates to \perp iff $P'_\sigma(x)$ evaluates to \top . Similarly in the case where $P_\sigma(x)$ evaluates to \perp , $P_\sigma^*(x)$ evaluates to \top iff $\psi_\sigma(x)$ evaluates to \top . By the definition in (2), $\psi_\sigma(x)$ evaluates to \top iff $\Psi_r(x)$ evaluates to \top for *some* row r . In other words, one of the conditions under which $P'_\sigma(x)$ can evaluate to \top is satisfied. Thus, $\psi_\sigma(x)$ evaluates to \top iff $P'_\sigma(x)$ evaluates to \top . \square