

# Typology emerges from simplicity in representations and learning

*Dakotah Lambert*<sup>1</sup>, *Jonathan Rawski*<sup>1</sup>, and *Jeffrey Heinz*<sup>1</sup>

Department of Linguistics and Institute for Advanced Computational Science,  
Stony Brook University

## ABSTRACT

We derive well-understood and well-studied Subregular classes of formal languages purely from the computational perspective of algorithmic learning problems. We parameterise the learning problem along dimensions of representation and inference strategy. Of special interest are those classes of languages whose learning algorithms are necessarily not prohibitively expensive in space and time, since learners are often exposed to adverse conditions and sparse data. Learned natural language patterns are expected to be most like the patterns in these classes, an expectation supported by previous typological and linguistic research in phonology. A second result is that the learning algorithms presented here are completely agnostic to choice of linguistic representation. In the case of the Subregular classes, the results fall out from traditional model-theoretic treatments of words and strings. The same learning algorithms, however, can be applied to model-theoretic treatments of other linguistic representations such as syntactic trees or autosegmental graphs, which opens a useful direction for future research.

*Keywords: model theory, Subregularity, grammatical inference, formal language theory, phonology, learning complexity*

1

## INTRODUCTION

We present an analysis supporting the view that the computational simplicity of learning mechanisms have considerable impact on the types of patterns found in natural languages.

First, we derive well-understood and well-studied Subregular classes of formal languages purely from the computational perspective of algorithmic learning problems. We present a family of four learning algorithms, generalizing Heinz (2010b). We show that the simplest classes of languages in these hierarchies are precisely the ones whose learning algorithms use the least computational resources. In fact, these are the only ones that are not prohibitively expensive to learn. A reasonable expectation is that learned natural language patterns would be most similar to patterns in the simplest of these classes, an expectation supported by previous typological and linguistic research in the domain of phonology.

A second result is that we introduce *linear-time* learning algorithms for some Subregular classes. As we explain, these algorithms are helpful in certain cases and not so helpful in others, depending on the extent to which the target patterns interact with other constraints. At issue is that a set of data points which may be helpful in identifying one constraint do not occur because they also happen to violate another. A virtue of this analysis is that we can identify precisely the situations where the linear-time learning algorithms can be applied.

A third result is that the learning algorithms presented here are completely agnostic to choice of linguistic representation. These learning algorithms essentially parameterise the learning problem in two ways: the structural knowledge salient to the learner (the representation), and the way the learner collects and combines this structural information to derive sets of acceptable and unacceptable linguistic structures. In the case of the Subregular classes of formal languages, the results fall out from traditional model-theoretic treatments of words and strings on the representational side and how the combinatorics of the grammars relate to kinds of logical languages on the other side.

Because the algorithms are agnostic to the representations, the same learning algorithms can be applied to model-theoretic treatments of other linguistic representations such as syntactic tree structures or autosegmental graphs. Of course, the real-life learning problem is complicated by the fact that language learners do not have direct access to linguistic structures like trees. Nonetheless the generality of these learning algorithms means the real-life learning problems may be reduced to these algorithms coupled with appropriate parsing mecha-

nisms.

### 1.1 *Priors in language learning*

Language acquisition succeeds despite sparse, underdetermined, Zipf-distributed input, compounded by a lack of invariance in the signal – the so-called poverty of stimulus (Yang 2013). This holds across all domains of language, from phonological to syntactic induction.

It is uncontroversial that *some* bias or innate component restricts a learner’s hypothesis space regardless of its strategy to solve this induction problem, often referred to as Universal Grammar (Nowak *et al.* 2002). The question is its nature. How is it rich, and how is it poor?

Data-driven statistical learning does not change this basic calculus. One reason is that children often learn language in ways that defy adult distributions (Legate and Yang 2002). Another is that induction from a data distribution without a prior may only recapitulate the training data (Fodor and Pylyshyn 1988; Mitchell 1982, 2017), and cannot generalize. Without a lens in which linguistic experience is viewed, even the input distribution cannot be recovered, simply because distributions are based on the structure of their parameters (Lappin and Shieber 2007). Consequently, the nontrivial open question central to learnability research in linguistics instead concerns the characteristics of this additional prior knowledge or bias such that learners *generalize* from limited experience (Rawski and Heinz 2019).

### 1.2 *Regular and Subregular patterns*

Recent typological and experimental work highlights the Regular region of languages as a sufficient structural bound on computational expressivity for phonological and morphological grammars. This Regular characterization has been extended to syntactic distributions when the data structure characterizing the computational trace is formulated as a tree rather than strings, which enforce syntactic membership in the Mildly Context-Sensitive class of languages (Kobele 2011; Graf 2011). However, the Regular class is not learnable from under various learning scenarios including identification in the limit from positive data, and the Probably Approximately Correct (PAC) framework (Gold 1967; Valiant 1984; de la Higuera 2010). Additionally, the range of distributions present in phonology and morphology that sit in the Regular region do not require the full complexity of Regular

power (Heinz 2018; Chandlee 2017).

For these reasons, phonological constraints are hypothesised to inhabit structured subclasses of the Regular languages, lumped under the term Subregular (Heinz 2010a, 2018). Various connections between logic, formal languages and automata defining these classes have been explored in great detail. These characterizations build on two classical results in formal language theory: Büchi's monadic second-order characterization of the Regular languages (1960), and the first-order characterization by McNaughton and Papert (1971) of the Star-Free languages, which are also characterized by aperiodic deterministic finite-state automata (Schützenberger 1965). Refinements of these results from logical, automata-theoretic, and algebraic viewpoints have defined the Local and Piecewise hierarchies (Rogers *et al.* 2012). Linguistically, these refinements have garnered interest since the morphological and phonological typology correlates with these refinements, favouring the weakest subclasses in the Subregular hierarchy. Experimental work also favours this characterization (Finley 2008; Lai 2015; McMullin and Hansson 2019). Our learning algorithms can be applied to model-theoretic treatments of other linguistic representations such as syntactic trees or autosegmental graphs, which opens a useful direction for future research.

### 1.3

#### *Outline*

This paper proceeds as follows. Section 2 defines a general model-theoretic treatment of linguistic representation, and several types of linguistic structures based on different model signatures. Section 3 defines a typology of online learning algorithms and derives the Subregular language classes, hierarchically organised by space complexity. Section 4 characterizes this space of algorithms according to time complexity, and picks out of the least space-intensive Subregular classes those that can be learned in linear time. Section 5 characterizes interactions of constraints defined in and between these classes. Section 6 discusses model signatures for other linguistic representations. Section 7 describes related work. Section 8 concludes with future directions.

## MODEL THEORIES

Our learning algorithms can be applied to model-theoretic treatments of other linguistic representations such as syntactic trees or autosegmental graphs, which opens a useful direction for future research. This section will introduce the structures that the learning algorithms will work over. We will first discuss a general notion of structural information, and use it to derive a notion of *substructures*. This will allow us to describe several representations of words. Structural information is defined relationally in terms of model theory. Finite model theory provides a unified ontology and a vocabulary for representing many kinds of objects, by considering them as *relational structures* (see Libkin 2004 for a thorough introduction). This allows flexible but precise definitions of the structural information in an object, by explicitly defining its parts and the relations between them. This makes model-theoretic representations a powerful tool for analyzing the information characterizing a certain structure.

The discussion of this section is organized around different notions of order: successor, precedence, and relativized successor. The successor and precedence orders give rise to the Local and Piecewise branches of the Subregular hierarchy, and the relativized successor gives rise to the Tier-Based Local branch. We assume some familiarity with these classes. Because this presentation focuses on deriving these Subregular classes from a model-theoretic and learning perspective, we postpone most references to these classes and related work to Section 7.

A relational structure in general is a set of domain elements,  $\mathcal{D}$ , which is augmented with a set of relations of arbitrary arity,  $R_i \subseteq \mathcal{D}^{n_i}$ . The relations provide information about the domain elements. The *model signature*  $\mathcal{M} = \langle \mathcal{D}; R_i \rangle$  collects these parts and defines the nature of the structure in terms of the information in the model. Let  $w$  be a string over some alphabet  $\Sigma$ . Then a model for a word  $w$  is a structure:

$$\mathcal{M}_\Sigma^{R_i}(w) := \langle \mathcal{D}_w; R_i, \sigma_w \rangle_{\sigma \in \Sigma}$$

where  $\mathcal{D}_w$  is isomorphic to an initial segment  $\langle 1, \dots, |w| \rangle$  of the nonzero natural numbers and represents the positions in  $w$ , and each  $\sigma_w$  is a unary relation that holds for all and only those positions at which  $\sigma$  occurs. Note that it is assumed that the set  $\{\sigma_w\}_{\sigma \in \Sigma}$  is a partition of

$\mathcal{D}_w$ .<sup>1</sup> Without loss of generality, consider an alphabet  $\Sigma = \{s, \int, \acute{a}, \grave{a}\}$ , which represent two types of sibilants and a vowel with either low or high tone. Strings are combinations of these symbols at certain events, like the word ‘sásàjá’.

The remaining  $R_i$  are the other salient relations, which are used to define order in a particular structure. One model signature for strings, called the *precedence model*, is given as

$$\mathcal{M}^<(w) = \langle \mathcal{D}_w; <_w, s_w, \int_w, \acute{a}_w, \grave{a}_w \rangle.$$

This model says that for every symbol  $\sigma$  in alphabet  $\Sigma$ , there is a unary relation  $R_\sigma$  in  $\mathcal{R}$  that can be thought of as a labelling relation for that symbol. For our set  $\Sigma = \{s, \int, \acute{a}, \grave{a}\}$ ,  $\mathcal{R}$  includes the unary relations  $R_s$ ,  $R_\int$ ,  $R_{\acute{a}}$ , and  $R_{\grave{a}}$ . It also defines a binary relation  $(x < y)$ , the general precedence relation on the domain  $\mathcal{D}$ . A visual of the word model for ‘sásàjá’ under this signature is given in Figure 1.


The general precedence relation describes a notion of structural information purely in terms of whether a node precedes another one. While the information that, say, the last element in a string comes after the first is immediately accessible from the model, this distinction collapses the notions of immediate and general structural adjacency. Building on this precedence relation we can derive different types of

<sup>1</sup> One can convert a model in which multiple unary relations may apply to a given domain element into a partitioned normal form by simply replacing these unary relations with their powerset.

Figure 1:

The general precedence model of ‘sásàjá’, along with the 3-factor ‘àáá’.

Each edge defined by the relation is pictured, while the thick solid edges designate those that form the window from which this 3-factor is derived



relational structure. These refine the model of a word to describe immediate, relativized, or multiply-relativized adjacency.

Perhaps we would like to consider only immediately adjacent elements. Rather than a general precedence relation  $<$ , we may consider an immediate precedence, or successor, relation  $\triangleleft$ . The standard successor relation is the transitive reduction of the precedence relation and is first-order definable from the latter as follows:

$$x \triangleleft y := x < y \wedge (\forall z)[x < z \Rightarrow y \leq z].$$

This relation gives a different word model, where elements are arranged according to immediate adjacency, commonly called the *successor model*. The signature for this model is given as

$$\mathcal{M}^{\triangleleft}(w) = \langle \mathcal{D}_w; \triangleleft_w, s_w, \int_w, \acute{a}_w, \grave{a}_w \rangle.$$

A visual of the successor word model for the word ‘sásàfá’ is given in Figure 2.

The general precedence relation can alternatively be refined to discuss a form of immediate adjacency relativized to certain unary relations in the signature. In particular, we can form relations between subsets of the alphabet, commonly called a *tier-alphabet*. For example, we may want to discuss the relations between only the sibilant elements present in a word, to the exclusion of all others. Similarly to how the successor relation is derived, we can restrict the precedence relation to the intended tier-alphabet  $\tau$  and first-order define a similar tier-successor relation  $\triangleleft^\tau$ :

$$x \triangleleft^\tau y := \tau(x) \wedge \tau(y) \wedge x < y \wedge (\forall z)[(\tau(z) \wedge x < z) \Rightarrow y \leq z].$$

Figure 3 depicts the relationships among these ordering relations.

Adjusting the model signature appropriately, shown below, we get a tier-based notion of structure, shown visually in Figure 4.




Figure 2:  
The immediate successor model of ‘sásàfá’, along with its 3-factor ‘ásà’




Figure 3:  
Relationships between the generative precedence relation and others first-order definable from it

$$\mathcal{M}^{\triangleleft^{\{s, \text{f}\}}}(w) = \langle \mathcal{D}_w; \triangleleft_w^{\{s, \text{f}\}}, s_w, \text{f}_w, \acute{a}_w, \grave{a}_w \rangle.$$

Because the unary relations partition the domain elements, we can create a tier-adjacency relation for each element of the powerset of these relations. This merely amounts to adding tier-adjacency relations to the model signature to create a multi-tier signature. A model of the multi-tier relations is shown in Figure 5.

$$\mathcal{M}^{\triangleleft^{\{s, \text{f}\}}, \triangleleft^{\{\acute{a}, \grave{a}\}}}(w) = \langle \mathcal{D}_w; \triangleleft_w^{\{s, \text{f}\}}, \triangleleft_w^{\{\acute{a}, \grave{a}\}}, s_w, \text{f}_w, \acute{a}_w, \grave{a}_w \rangle.$$


These four model signatures are by no means the only relational word models that may be considered. However, for the purposes of this paper we restrict ourselves to these signatures. Additionally, the definability of these signatures from other signatures leads to a general ability to define a notion of substructure, which we cover below.

### 2.1 Windows and factors

Now that we have a general model-theoretic notion of structure, we would like a way to define certain parts of each structure, each of which is a structure in itself defined by the signature. Here we generalize the method of Lambert and Rogers (2020) in defining these restrictions on models.

In order to pick out subparts of a word model, we first pick out sets of elements that will define the substructure. Given a homogeneous

Figure 4:  
The tier-successor model of ‘sásàfá’ relativized over the set  $\tau = \{s, \text{f}\}$ , along with its only 3-factor ‘sfs’








Figure 5:  
The multi-tier-successor model of ‘sásàfá’ relativized over the sets  $\{s, f\}$  and  $\{\hat{a}, \acute{a}\}$ , along with its only two 3-factors, ‘sfs’ and ‘ààá’

relation  $R$  of arity  $a$ , the set

$$\mathcal{W}_a^R(m) := \left\{ \left\{ \langle x_i^i, x_{i+1}^{i+1} \rangle : 1 \leq i < a \right\} : \langle x_1, \dots, x_a \rangle \in R_m \right\}$$

is the set of  $a$ -windows over  $R$  in the context of the model  $m$ . These are merely directed acyclic graphs (represented by their edge sets alone) constructed from the relations in  $R$ , such that each instance of a given domain element in the tuple is represented by a distinct node in the window, rather than merging all instances into a single node. Concretely, if 1 were a domain element and  $\langle 1, 1 \rangle$  an element of the relation, the corresponding 2-window would have two distinct nodes, both labelled by an index and the domain element 1:  $\langle 1^1, 1^2 \rangle$ . The set of windows of length greater than  $a$  is defined inductively by

$$\begin{aligned} \mathcal{W}_{k+1}^R(m) := & \left\{ A \cup \left\{ \langle x_{a-1}^{j_{a-1}}, x_a^{k+1} \rangle \right\} : \right. \\ & A \in \mathcal{W}_k^R(m), \langle x_1, \dots, x_a \rangle \in R_m, \\ & \left. \left\{ \langle x_i^{j_i}, x_{i+1}^{j_{i+1}} \rangle : 1 \leq i < a - 1 \right\} \subseteq A, \langle x_{a-1}^{j_{a-1}}, x_a^{j_a} \rangle \notin A \right\}. \end{aligned}$$

This means that for each  $k$ -window, we find a linear subgraph (a path) that maps to the initial  $a - 1$  domain elements of one of the  $a$ -tuples that comprise  $R$  and add an edge from the final node of this path to a newly constructed node representing the final domain element from that tuple. The conditions are arranged in such a way that each iteration actually adds a new step to the path rather than simply repeating an older step, while still allowing cycles to be taken arbitrarily many times. Each of these larger windows can then be thought of as a graph of positions that are formed from a set of overlapping  $a$ -windows, which in turn are merely representations of tuples in the relation  $R$ . However, we may also wish to discuss a window which is

of shorter length than the arity of the relation that defines it. To do so, we simply state that any connected subgraph of a window is itself a window.

For a given window  $x$  of a word model  $m$ , we define *the factor at  $x$*  (written  $\llbracket x \rrbracket_m$ ) as the restriction of  $m$  to the domain elements that occur in  $x$ . This lets us define the set of all  $k$ -factors of  $m$  as follows:

$$\mathcal{F}_k^R(m) := \{\llbracket x \rrbracket_m : x \in \mathcal{W}_k^R(m)\}.$$

Note that a window is distinct from a factor in that the former is a graph of positions while the latter describes a word model whose domain consists of only a certain set of positions.

As example, consider the tier successor model of the word ‘sàsàfá’ as above. Consider a 3-window  $x$  which contains all and only the domain elements  $\{1, 3, 5\}$ . Here the restriction of the word model that defines this 3-factor is

$$\llbracket x \rrbracket_m = m \upharpoonright x = \langle \{1, 3, 5\}; \{\langle 1, 3 \rangle, \langle 3, 5 \rangle\}, \{1, 3\}, \{5\}, \emptyset, \emptyset \rangle.$$

Similar examples can be seen above in Figures 1–5. Various parallels emerge. The precedence word model contains a strict superset of the factors of every other word model we have considered. The tier-based and multi-tier-based word models have ‘sjs’ as a 3-factor, but the immediate successor model does not. On the other hand, ‘àsà’ is a 3-factor of only the precedence and immediate successor models. Only the precedence and multi-tier successor models have both ‘sjs’ and ‘ààà’ (a sequence of High-Low-High tone vowels) as 3-factors.

## 2.2

### *Anchored word models*

The word models considered up to this point do not encode domain boundaries explicitly. However, many prior treatments, including that of Lambert and Rogers (2020), explicitly assume such boundaries. One approach that has not been explicitly considered in this prior work is a model whose string yield is biinfinite. Here left and right boundary symbols (labelled  $\times$  and  $\times$ , respectively) exist in the model and both participate in and are self-related under any ordering relations. This approach naturally captures words shorter than  $k$  symbols in its concept of a  $k$ -factor, without having to consider a union of smaller factor widths. The successor model for ‘sàsàfá’ is shown along with each of its 3-factors in Figure 6.

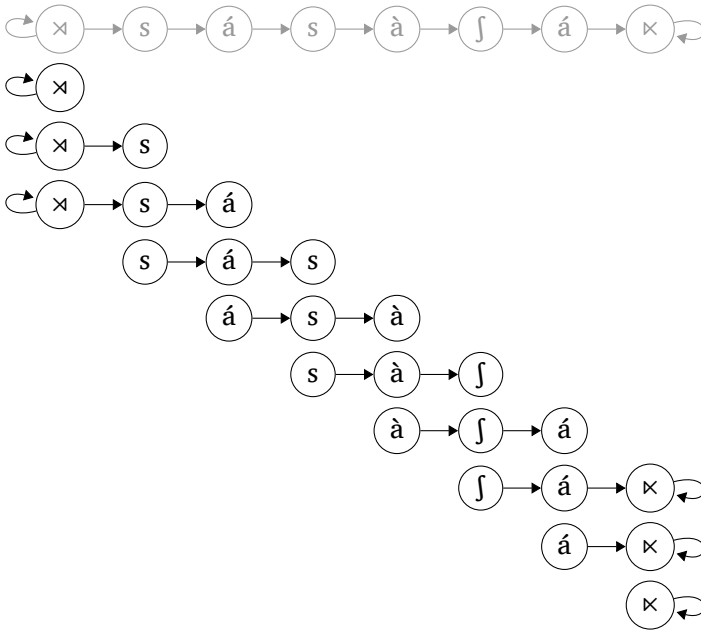


Figure 6:  
The anchored word model under successor for ‘sásàjà’ along with each of its 3-factors. Note that every factor that includes a boundary symbol has an infinite yield. Those factors shorter than 3 symbols are formed from windows of length 3 that repeat the boundary symbols

The learning algorithms that we consider in this work are not bound to any particular model signature. Thus we may consider the standard word models as shown in, for example, Figure 1, or we might consider these anchored word models.

### 3 SPACE COMPLEXITY AND THE SUBREGULAR GRID

This section will examine four learning algorithms in the style of Gold (1967), where stringsets are learned in the limit from distribution-free positive texts. Indeed, we will only be considering subclasses of a style of learning algorithm presented by Heinz (2010b, expanded upon by Heinz *et al.* 2012). We show that of these subclasses, some require substantially more space than others to properly account for the distinctions that must be made in the course of learning, and we argue that this alone would cause linguistic typology to tend toward the simpler, less space-intensive classes.

First we briefly discuss some background from learning theory. Let  $L$  be a set of strings drawn from  $\Sigma^*$  and let  $L_{\odot}$  represent  $L$  with an adjoined element  $\odot$ . An online learner is a function  $\varphi : \mathcal{G} \times L_{\odot} \rightarrow \mathcal{G}$ ,

where  $\mathcal{G}$  is some kind of grammar representation, a mechanism by which one can decide whether a given string is in  $L$ . In other words, an online learner begins with some guess as to what the grammar might be and updates this guess for each input word. Let  $\mathcal{L} : \mathcal{G} \rightarrow \mathcal{P}(\Sigma^*)$  be the function that maps a grammar to its extensions, the set of strings it represents. Two grammars  $G_1$  and  $G_2$  are equivalent ( $G_1 \equiv G_2$ ) iff they are extensionally equal, that is,  $\mathcal{L}(G_1) = \mathcal{L}(G_2)$ .

A text for  $L$  is a function  $t : \mathbb{N} \rightarrow L_\circ$ , a sequence of strings drawn from  $L$  or pauses in which data does not appear. Following traditional mathematical notation for sequences, we use  $t_n$  to represent  $t(n)$ . If  $\emptyset$  represents an initial guess at what the grammar might be, then the recursively-defined sequence

$$a_n(t) := \begin{cases} \emptyset & \text{if } n = 0 \\ \varphi(a_{n-1}, t_n) & \text{otherwise.} \end{cases}$$

represents the learning trajectory over a given text. Then given a text  $t$  for a language  $L$ , we say that a learning algorithm  $\varphi$  *converges* on  $t$  iff there is some  $i \in \mathbb{N}$  such that for all  $j > i$  it holds that  $a_j(t) \equiv a_i(t)$ . If for every possible text  $t$  over  $L$  it is the case that  $\varphi$  converges on  $t$  and  $\mathcal{L}(\lim_{n \rightarrow \infty} a_n(t)) = L$ , then we say  $\varphi$  converges on  $L$ . As a second lift, if for every stringset  $L$  in a class  $\mathbb{L}$  it is the case that  $\varphi$  converges on  $L$ , then we say  $\varphi$  converges on  $\mathbb{L}$ .

### 3.1 *String extension learning*

Heinz (2010b, expanded by Heinz *et al.* 2012) defined a general notion of learning from gathered substructures. Originally treated only as a batch learner, the online definition is trivial to derive. Given a function  $f : \mathcal{M} \rightarrow \mathcal{S}$  that extracts informational content from a word model, where  $\mathcal{S}$  represents some notion of structural content, along with a combinator  $\oplus : \mathcal{G} \times \mathcal{S} \rightarrow \mathcal{G}$  that somehow informs the grammar of these structures, we define

$$\varphi(G, w) := \begin{cases} G & \text{if } w = \circ \\ G \oplus f(\mathcal{M}(w)) & \text{otherwise.} \end{cases}$$

In a simple case,  $\mathcal{G}$  and  $\mathcal{S}$  will be the same type, and  $\oplus$  will simply be set union, but this is not a necessary requirement.

3.1.1 Learning with factors

This simple case is exemplified by a learner that makes distinctions only between permitted and nonpermitted factors. This learner is parameterised by a factor width  $k$ . We have  $\mathcal{G} = \mathcal{S} = \mathcal{P}(\Sigma^k)$  and  $G \oplus S = G \cup S$ . The information extraction function is

$$f(m) := \mathcal{F}_k(m).$$

Since the grammar only needs to maintain a single merged set of attested factors, the space complexity for this class of learner is  $\mathcal{O}(|\Sigma|^k)$ . This will be referred to as Algorithm II. A variant, which will be Algorithm I, will be discussed in Section 3.1.4.

3.1.2 Learning with sets

The primary difference between learning with factors and learning with sets thereof is the grammar augmentation combinator. Rather than set union,  $G \oplus S = G \cup S$ , we have set insertion,  $G \oplus S = G \cup \{S\}$ . This of course means that  $\mathcal{G}$  and  $\mathcal{S}$  are no longer equal, with  $\mathcal{G}$  being the powerset  $\mathcal{P}(\mathcal{S})$ , adding a layer of structure. Since a given grammar is in this case a set of sets of factors, with this larger grammar the space complexity is  $\mathcal{O}(2^{|\Sigma|^k})$ . These set-based classes can make more distinctions than the purely factor-based classes, but this power comes at a cost. This is Algorithm III.

3.1.3 Learning with multisets

A set is simply a structure that contains for each possible element a Boolean value describing whether or not that element is included. Given the natural isomorphism between the Booleans and the subset of  $\mathbb{N}$  consisting of 0 and 1, one might consider a natural expansion of this structure which denotes number of occurrences saturating not at 1 but at some arbitrary value  $t$ . We can learn classes in which well-formedness is characterized by the saturating multisets of factors in a word as follows. With  $\mathcal{S} = \mathcal{P}(\Sigma^k \times \mathbb{N}_t)$  and  $\mathcal{G} = \mathcal{P}(\mathcal{S})$ , we can maintain from the set-based learner the augmentation combinator where  $G \oplus S = G \cup \{S\}$ . However, the function that extracts informational content must be modified to include the  $t$ -counts associated with a given factor as follows

$$f(m) := \langle \llbracket x \rrbracket_m : x \in \mathcal{W}_k(m) \rangle_t.$$

The notation  $\{\dots\}_t$  represents a multiset that saturates at a count of  $t$ . Note that this parallels the window-based definition of factors in a model, except that a saturating multiset is formed rather than merely a set. The space complexity here is much like that of Algorithm III, except that the base of the exponent is changed to correspond with the number of values each factor may be associated with:  $\mathcal{O}((t+1)^{|\Sigma|^k})$ . This is Algorithm IV. Using this algorithm with  $t = 1$  is equivalent in every way to Algorithm III, so in fact there are only three algorithms under discussion. That said, we will retain this separation for the current discussion.

#### 3.1.4 Learning with factors, revisited

A variant of Algorithm II ignores all input words longer than  $k$  symbols. The only difference then is the information extraction function

$$f(m) := \begin{cases} m & \text{if } |m| \leq k \\ \emptyset & \text{otherwise.} \end{cases}$$

This is Algorithm I. Notably, using the anchored word models with this algorithm produces only finite languages. In contrast to the other algorithms, translating such models into unanchored ones provides an increase in expressive power.

#### 3.1.5 Illustration

Considering a standard unanchored word model, with the algorithmic parameters  $k$  and  $t$  both set to 2, Table 1 represents the outputs of these four learning algorithms after seeing the single word ‘aaaab’. Notably, this word is not short enough to inform Algorithm I of anything. Also, despite the fact that ‘aa’ occurs as a substring three distinct times, Algorithm IV saturates at a count of 2 under these assumed parameters.

Table 1:  
Encountering the single word ‘aaaab’ with each learner

| Algorithm | Resulting Grammar                            |
|-----------|--|
| I         | $\emptyset$                                  |
| II        | $\{\text{aa}, \text{ab}\}$                   |
| III       | $\{\{\text{aa}, \text{ab}\}\}$               |
| IV        | $\{\{\{\text{aa}, 2\}, \{\text{ab}, 1\}\}\}$ |

These learning algorithms are model-agnostic. As long as there exists some way to extract windows or factors (*i.e.*, substructures) from a model, the algorithms will work with that. When allowed to range over selected model signatures, the classes learned by each algorithm are shown in Figure 7. Note that Algorithm I learns only strong classes (those in which domain boundaries are unreferenced), while the others do not have this restriction. In the Piecewise case, where the model signature contains the general precedence relation ( $<$ ), the strong classes are equivalent to the general classes and this distinction is irrelevant.

We note that the amount of space required to store the grammar is fairly large for any of these algorithms. But Algorithms III and IV require exponentially more space than I and II. These space requirements are shown in Figure 8, where it is apparent that even on merely a binary alphabet, the Locally 5-Testable class of languages, for example, requires more storage space than there are synapses in an average human being (Azevedo *et al.* 2009; Herculano-Houzel and Lent 2005).

Due to the enormous space requirements in terms of alphabet size of the higher-numbered algorithms, it seems in general unfeasible to learn patterns that lie strictly in their corresponding classes. That is, from learnability alone we would expect the typology of patterns in natural language to lie primarily within the region spanned by Algorithms I and II. This region is highlighted in Figure 7 by a dotted line. Further, we would expect any attested patterns to require relatively small values for the factor width parameter  $k$ , since that is the exponent of these singly and doubly exponential space complexities.

The Strictly Local (SL) class (McNaughton and Papert 1971) is learned by gathering the factors of simple adjacency. Under such a model, there exists at most a single window of size  $k$  at any given point. Thus for each index in the word, we can simply insert the contents of this single window into the grammar. Including the time it takes to insert a factor into a set, the class is learnable in  $\mathcal{O}(nk \log|\Sigma|)$  time for input of size  $n$ , and since  $\Sigma$  and  $k$  are assumed constant this amounts to linear

Figure 7:  
The Subregular classes. Learning difficulty increases along the vertical axis. The horizontal axis is categorical, describing the type of substructure. The dotted line indicates the predicted region of phonological typology

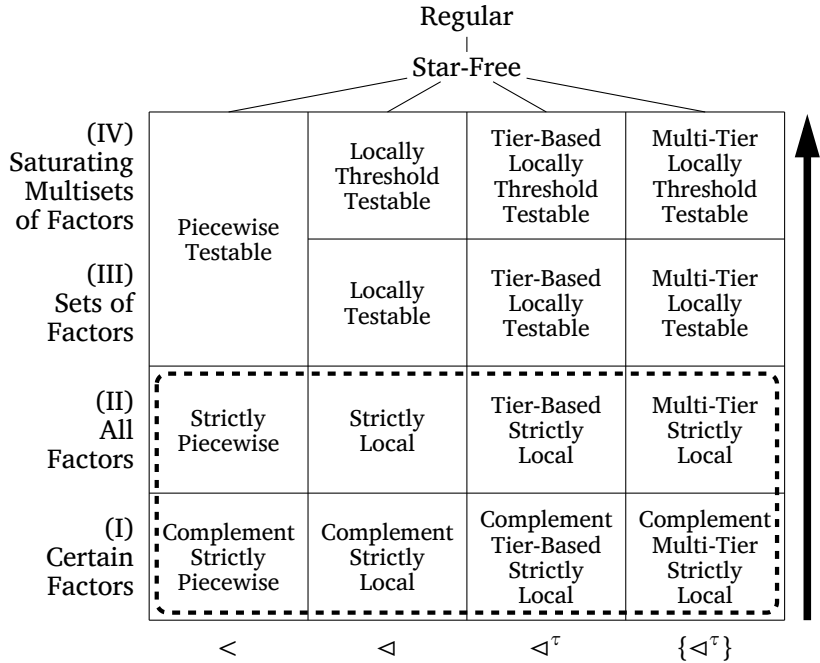



Figure 8:  
While gathering factors requires space exponential in terms of factor width, the requirements are doubly exponential for any of the larger structures we might employ. Here the space requirements are shown for just a binary alphabet





time. As discussed, this Algorithm II learner also uses constant space that is but singly exponential in the width of the factors.

For the Tier-Based Strictly Local (TSL) class (Heinz *et al.* 2011, see also Lambert and Rogers 2020), if the tier alphabet  $\tau$  is known, then this approach applies directly to the projection of the word to  $\tau$ . But generally we assume that  $\tau$  is not known, and one might initially assume that a learner might need to construct grammars for all possibilities, which would result in increased resource requirements, be that in terms of time, space, or both. Per Jardine and McMullin (2017), maintaining the factors of width bounded above by  $k + 1$  is sufficient to determine the value of  $\tau$ . But their approach seems to require a batch approach, first deciding the value of the  $\tau$  parameter and then processing the (projections of) the input as for the Strictly Local class. But it turns out that, due to (inverse-)projection closure and the fact that in the Gold framework we assume a complete text, we can guarantee that any substring whose projection will appear on the tier will itself appear as a substring in some word. Since we still need to determine the value of  $\tau$ , we do still require the factors of width bounded above by  $k + 1$ , but nothing more. The exact learning algorithm used for  $SL_{k+1}$  will produce a grammar for  $TSL_k$ , and only the interpretation of the result is changed (Lambert 2020). These same properties hold true for the relativized variants of the Locally Testable (LT) (McNaughton and Papert 1971) and Locally Threshold Testable (LTT) (Beauquier and Pin 1989) classes as well, where the corresponding adjacency-based learners suffice to learn the relativized-adjacency classes (Lambert 2020).

The Strictly Piecewise (SP) class (Rogers *et al.* 2010, see also Haines 1969) is similar in that, one might expect a time complexity on the order of  $\mathcal{O}(n^k)$  to find all of the subsequences of each word. Heinz and Rogers (2013) show that in fact a factored approach can use simply  $\mathcal{O}(n|\Sigma|^k)$ , but we can reduce this even further by taking advantage of this same property. Given a complete text, every attested subsequence will eventually occur as a substring due to the SP stringsets' closure under deletion. Again then, the same learning algorithm used for  $SL_k$  will produce a grammar for  $SP_k$  as well, where the difference lies only in interpretation.

Given this ability to learn the SP, SL, and TSL classes in linear time and in space only singly exponential in factor width, we can modify

Figure 7 to indicate the boundary between the classes that are learnable within these resource bounds and those that are not. This boundary is indicated by a thick line in Figure 9, which also uses dashed lines to indicate where one algorithm may be used for multiple distinct classes.

Considering only the SP class of stringsets, there are at least three online learning algorithms of various complexities:

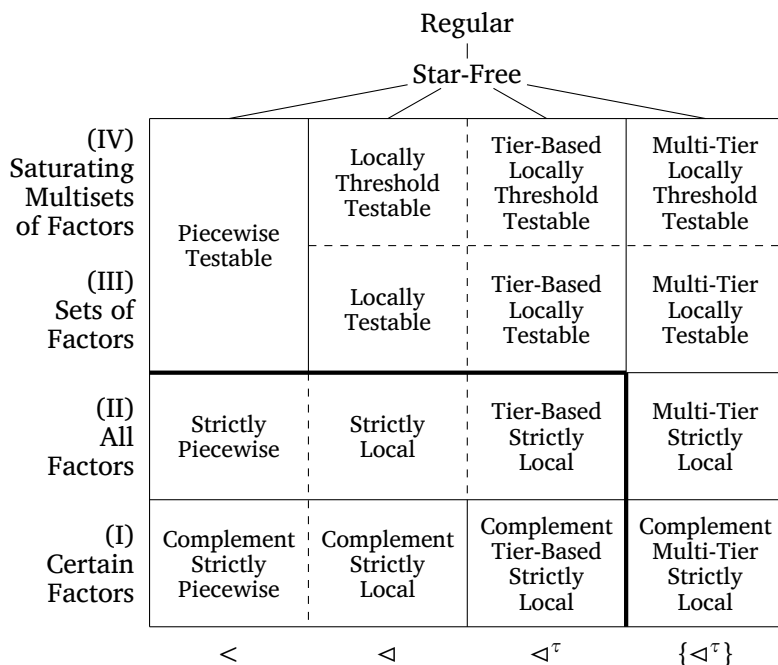
- Gather all factors under general precedence of each word.
  - $\mathcal{O}(n^k k \log|\Sigma|)$  time
  - $\mathcal{O}(|\Sigma|^k)$  space
  - Learns an SP least upper bound (lub) of the source text.
- Use factored learning as per Heinz and Rogers (2013).
  - $\mathcal{O}(n|\Sigma|^k)$  time
  - $\mathcal{O}(|\Sigma|^k)$  space
  - Learns an SP lub of the source text.
- Gather all adjacency factors of each word.
  - $\mathcal{O}(nk \log|\Sigma|)$  time
  - $\mathcal{O}(|\Sigma|^k)$  space
  - Only works if every permitted subsequence eventually occurs adjacently, which holds for SP targets.

One caveat is that these optimizations of the learning algorithms for the SP and TSL classes rely on certain properties of the input stringset. The nonoptimized variants are guaranteed to learn a smallest in-class superset of the input stringset, a property which is lost in this optimization. For example, a long distance sibilant harmony constraint will not be learned by the optimized SP learner if the text is drawn from a language that exemplifies both this constraint and a CV syllable structure, even though it would be learned by the nonoptimized variant. This prompts a question regarding the learnability consequences of constraint interaction.

## 5 LEARNING INTERACTIONS

Most natural languages are describable not by a single Subregular class but by an interaction of constraints from multiple such classes.

Figure 9: Dashed lines indicate that the constraints on either side can be learned by exactly the same algorithm. The thick solid line denotes the tier between linear-time  $(|\Sigma|^k)$ -space learning and resource-intensive learning



The interaction of constraints from different classes might influence the learnability of each constraint individually, in which case time or space tradeoffs might be necessary.

For example, we might consider the default-to-opposite stress pattern of Chuvash (Krueger 1961), where primary stress falls on the rightmost heavy syllable if there is one, or on the initial syllable otherwise. One way of describing this invokes the conjunction of two constraints from two different classes, namely an SP constraint detailing a lack of:

- A heavy syllable anywhere after a stressed syllable,
- A stressed light syllable after any other syllable, or
- Two stressed syllables in the same word,

and a coSP constraint that states that every word contains some stressed syllable.

The requirement that some stress must occur does not affect which substrings may appear in a word, and so the SP constraint may be learned by any of the three algorithms that have been discussed

for that class so far, including the optimized substring-based learner. Further, the precedence restrictions do not prevent seeing the two words (light and heavy stressed monosyllables) required to learn the stress requirement. In other words, these constraints interact in such a way that learning is not hindered. This is not always the case.

Consider now the sibilant harmony of Samala (Applegate 1972), in which ‘s’ and an ‘j’ may not appear in the same word. Since this constraint acts on the segment level rather than the syllable level, we might assume that it is isolated from any kind of stress constraint. But other segment-level constraints will certainly have the possibility of interaction. For example, imposing a CV syllable pattern restricts the substrings that may occur, in such a way that using an SL learner to infer the SP constraints is not a possibility. This means that one has to decide between among the other possible SP learning algorithms, where time or space tradeoffs must be made.

In contrast, a tone plateauing constraint like that which occurs in Luganda (Hyman and Katamba 1993) is  $SP_3$ , which means that it could be learned directly alongside this sibilant harmony constraint without fear of interaction effects. Note that the word ‘sásàǰá’ that has been our running example violates both the harmony constraint and the tone plateauing constraint.

Given our space-based learnability considerations, we would assume that Algorithms III or IV are not practically learnable and would likely be unattested. In other words, we would expect linguistic typology to inhabit only the lower regions of the hierarchy, or at least be biased heavily toward this region. Rogers and Lambert (2019b) provide strong evidence that this is in fact the case when it comes to stress patterns. Their exhaustive analysis of the more than one hundred stress patterns in the StressTyp2 database (Goedemans *et al.* 2015) showed that each of these can be described as the interaction of constraints that can be learned by Algorithms I and II.

## 6 MODEL-THEORETIC REPRESENTATIONS OF NONLINEAR STRUCTURES

This model-theoretic formulation provides a distinct advantage when applied to various linguistic objects. It allows one to characterize the content of a particular linguistic representation, and in so doing, im-

mediately guarantee that there are learning algorithms which can describe various constraints over those representations. This is important, because work describing nonlinear structures in syntax and phonology has proceeded in an ad-hoc way, by first defining constraints, and working backwards to the representations, often without any learning algorithms at all, or ones relativized to that structure.

The previous sections used various model signatures that characterized information based on a string data structure. This is because the Subregular classes that were the central motivation for this paper are defined over strings, or model signatures based on strings, in the work of Büchi, Thomas, and others. The constructions considered to this point are not restricted to simple string models. Without modification, the algorithms may be applied to any relational model at all. They in fact apply to any structure that can be characterized as a graph. In this sense, strings are a special case, but the distinctions that each of the four learning algorithms pick out carry over onto these more general factors as well. In this section we discuss some other linguistically-motivated models that one might consider.

### 6.1 *Autosegmental graphs*

An example of a nonlinear structure where the graph perspective is clearly relevant to linguistic research concerns autosegmental representations in phonology. Graphs were proposed to handle a variety of prosodic phenomena for which the string-based perspective was inadequate. Phonological processes affecting domains larger than two adjacent segments, such as tonal alternations in tonal languages, have temporal properties that do not always map consistently onto discrete vowel segments in one-to-one fashion (Goldsmith 1976; Williams 1976). Goldsmith introduced a model of the phonological word where tonal features formed an independent string from the segmental string, called a tier. Segments on the two strings are linked via many-to-one relations, turning the structure into a graph.

In practice, encoding these adjustments into a word model involves adding more relational structure. Jardine (2017a,b, 2019) uses a binary relation  $\alpha(x, y)$  to encode the association relation between autosegmental tiers. Augmenting the successor model signature used

throughout this paper gives a signature as

$$\mathcal{M}^{\alpha, \triangleleft}(w) = \langle \mathcal{D}_w; \alpha_w, \triangleleft_w, s_w, \int_w, a_w, H_w, L_w \rangle.$$

Here the domain is increased to accommodate the new autosegments, and the successor relation holds between elements on both tiers. The unary relations encoding vowels with tonal features have been split, into a relation ‘a’ for vowel information, and distinct ‘H’ and ‘L’ relations for tonal information. Under this signature, a word model for the example ‘sàsàjà’ is given in Figure 10.

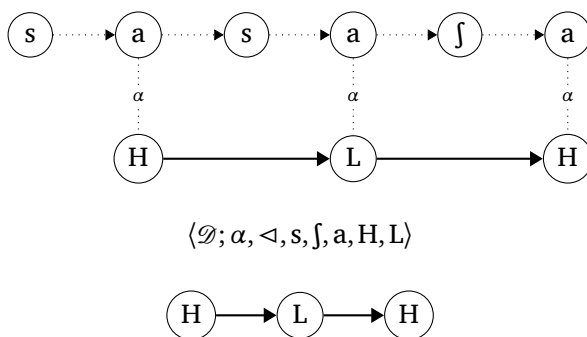
Our notion of a factor is exactly a notion of a subgraph. The previous section showcased how this word violates a constraint on tone plateauing. The autosegmental model makes this information immediately accessible, by encoding the ‘HLH’ structure as its own subgraph, shown on the bottom of Figure 10. Thus the permissibility of tone sequences is liberated from the segmental elements that carry them.

## 6.2

### Tree models

< The model-theoretic framework also allows describing tree structures (Rogers 1998), and opens the door to study parallels between phonological and syntactic constraints (Graf 2014). Rogers (2003) describes a model-theoretic characterization of trees of arbitrary dimensionality. In this framework, we specify the domain  $\mathcal{D}$  as a Gorn tree domain (Gorn 1967). This is a hereditarily prefix closed set  $D$  of node addresses, that is to say, for every  $d \in D$  with  $d = \alpha i$ , it holds that  $\alpha \in D$ , and if  $i \neq 0$  then  $\alpha(i - 1) \in D$  as well. In this view, a string may be called a one-dimensional or unary-branching tree, since it has one axis along which its nodes are ordered. In a standard tree, on the other hand, the set of nodes is ordered as above by two relations,

Figure 10:  
The autosegmental successor model of ‘sàsàjà’, along with its 3-factor ‘HLH’. The  $\alpha$  relation is shown without tips because it is symmetric



“dominance” and “immediate right-of”. Suppose  $s$  is the mother of two nodes  $t$  and  $u$  in some standard tree, and also assume that  $t$  precedes  $u$ . Then we might say that  $s$  dominates the string  $tu$ .

While a Gorn tree domain as written encodes these dominance and adjacency relations implicitly, we may explicitly write them out model-theoretically so that a signature for a  $\Sigma$ -labelled two-dimensional tree  $T$  is  $\mathcal{M}^{\triangleleft^\downarrow, \triangleleft^\rightarrow} = \langle D; \triangleleft^\downarrow, \triangleleft^\rightarrow, R_\sigma \rangle_{\sigma \in \Sigma}$  where  $\triangleleft^\downarrow$  is the immediate dominance relation and  $\triangleleft^\rightarrow$  is the immediate right-of relation. Model signatures that include the transitive closures of each of these relations have also been studied. Additionally, the anchored word models considered above for strings lift naturally to trees, where a root node is an anchor and each leaf is a separate anchor, or there is a single additional node which serves as the anchor for every leaf.

Recent work in syntax has synthesised the model-theoretic approach to trees with insights from the Subregular approach to phonology. For instance, Graf and Shafiei (2019) hypothesise that the TSL class is sufficient to characterize syntactic constraints.

7 FURTHER READING

The Subregular classes considered here have been widely studied for decades. McNaughton and Papert (1971) introduce the Local hierarchy, with Beauquier and Pin (1989) adding the Locally Threshold Testable class. The Piecewise branch of the hierarchy stems from Si-




Figure 11:  
A tree model. Nodes are organised by immediate dominance (black tip) and immediate right-of (white tip) relations. Labelling relations are omitted to show Gorn addresses. All edges are shown, with a particular factor noted with solid thick lines

mon (1975), with the Strictly Piecewise class only being integrated into the hierarchy in 2010 by Rogers *et al.*. (Languages closed under subsequence had been discussed by Haines 1969, though not in connection with other Subregular classes.) The Tier-Based Strictly Local class was introduced by Heinz *et al.* (2011) and extended in various ways by De Santo and Graf (2019), Lambert and Rogers (2020), and Lambert (2020). Recent work in syntax has synthesized the model-theoretic perspective on trees with insights from the Subregular program (Graf and Shafiei 2019; Graf 2020, 2014)

Provided a finite-state automaton, Caron (1998, see also Caron 2000) describe algorithms that decide whether the corresponding language is Locally or Piecewise Testable. An efficient algorithm for deciding SL is described by Edlefsen *et al.* (2008). Algorithms that extract SL and SP factors from a given language (and thus can also be used to decide class membership) are due to Rogers and Lambert (2019a), and these were extended to the TSL class by Lambert and Rogers (2020).

While this paper has so far focused on constraints, this work is easily extended to consider mappings between structures, expressed mathematically as Regular functions (Courcelle 1994; Courcelle and Engelfriet 2012; Filiot 2015; Engelfriet and Hoogetboom 2001). The notion of strict locality has been generalized to functions and shown to be relevant for natural language phonology and morphology (Chandlee 2014, 2017). These local functions have been model-theoretically characterized and extended to consider nonlinear structures in phonology (Chandlee and Jardine 2019; Strother-Garcia 2019). Relativizing input representations to consider multi-arity functions allows a notion of strictly local transducers expressed using multi-tape automata (Rawski and Dolatian 2020; Dolatian and Rawski 2020). Expressed as functions, these Subregular characterizations have been extended to consider continuous functions over vector spaces, and learning algorithms operating over them (Rawski 2019; Nelson *et al.* 2020).

There exist other learning algorithms alongside the string extension learners of Heinz (2010b) and Heinz *et al.* (2012). Garcia *et al.* (1990) demonstrate the learnability of SL. Heinz and Rogers (2013) provide learning algorithms for the SL and SP classes as well as their Testable correlates. Other approaches have directly incorporated phonological features into the models (Vu *et al.* 2018; Chandlee *et al.* 2019). Learning of TSL classes has been discussed by Jardine



and Heinz (2016) and Jardine and McMullin (2017), while online learners for this class and the remaining single-tier-based hierarchy were proposed by Lambert (2020).

8

## CONCLUSION

This paper showed how the nature of phonological typology emerges from simple representations and inference strategies. We discussed the nature of these representations in model-theoretic terms, forming a general notion of structural information (factors) that characterizes virtually any linguistic representation, from strings, to trees, to graphs. We also discussed a series of learning algorithms that work over any form of these factors, and are organised into a hierarchy of space complexity based on the distinctions they make with respect to structural information. We then derived the full hierarchical range of Subregular formal language classes from the product of these different representations and inference strategies. Consideration of time complexity further parameterises this hierarchy, drawing equivalences and distinctions amongst the classes with respect to learning. We find that the scope of phonological typology is strongly biased into the range defined by the simplest learning algorithms and representations.

The relevance of these results for linguistic theory are clear. A learner, faced with dramatically sparse data, favours grammar induction strategies that limit the amount of necessary distinctions between structural forms. The requirement for learners to structure and limit their hypothesis spaces plays off the distinctions learners make and the representations they make them over. The results here, as well as typological and experimental evidence, suggest that a learner may fix a learning algorithm and allow representational primitives to vary. From this perspective, the requirement of parsing from a linguistic input to a particular linguistic form is of the utmost importance. Linguistic learning can be relativized over various representations, be they strings or graphs for phonology, to trees for syntax. In this way, natural language typology, considered through an algorithmic lens, can be shown to emerge from the interaction of simple learning algorithms and simple but wide-ranging notions of representation.

## REFERENCES

- Richard Brian APPLGATE (1972), *Ineseño Chumash Grammar*, Ph.D. thesis, University of California, Berkeley.
- Frederico Augusto Casarsa AZEVEDO, Ludmila Ribeiro Bezerra CARVALHO, Lea Tenenholz GRINBERG, José Marcelo FARFEL, Renata Eloah de Lucena FERRETTI, Renata Elaine Paraizo LEITE, Wilson Jacob FILHO, Roberto LENT, and Suzana HERCULANO-HOUZEL (2009), Equal Numbers of Neuronal and Nonneuronal Cells Make the Human Brain an Isometrically Scaled-Up Primate Brain, *The Journal of Comparative Neurology*, 513:532–541, doi:10.1002/cne.21974.
- Danièle BEAUQUIER and Jean-Eric PIN (1989), Factors of Words, in Giorgio AUSIELLO, Mariangiola DEZANI-CIANCAGLINI, and Simonetta RONCHI DELLA ROCCA, editors, *Automata, Languages and Programming: 16th International Colloquium*, volume 372 of *Lecture Notes in Computer Science*, pp. 63–79, Springer Berlin / Heidelberg, doi:10.1007/BFb0035752.
- Julius Richard BÜCHI (1960), Weak Second-Order Arithmetic and Finite Automata, *Zeitschrift für mathematische Logik und Grundlagen der Mathematik*, 6(1–6):66–92, doi:10.1002/malq.19600060105.
- Pascal CARON (1998), LANGAGE: A Maple Package for Automaton Characterization of Regular Languages, in Derick WOOD and Sheng YU, editors, *Automata Implementation*, volume 1436 of *Lecture Notes in Computer Science*, pp. 46–55, Springer Berlin / Heidelberg, doi:10.1007/BFb0031380.
- Pascal CARON (2000), Families of Locally Testable Languages, *Theoretical Computer Science*, 242(1–2):361–376, doi:10.1016/S0304-3975(98)00332-6.
- Jane CHANDLEE (2014), *Strictly Local Phonological Processes*, Ph.D. thesis, University of Delaware, [https://chandlee.sites.haverford.edu/wp-content/uploads/2015/05/Chandlee\\_dissertation\\_2014.pdf](https://chandlee.sites.haverford.edu/wp-content/uploads/2015/05/Chandlee_dissertation_2014.pdf).
- Jane CHANDLEE (2017), Computational Locality in Morphological Maps, *Morphology*, 27(4):599–641, doi:10.1007/s11525-017-9316-9.
- Jane CHANDLEE, Rémi 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*, pp. 91–101, Association for Computational Linguistics, Toronto, Canada, doi:10.18653/v1/W19-5708.
- Jane CHANDLEE and Adam JARDINE (2019), Autosegmental Input Strictly Local Functions, *Transactions of the Association for Computational Linguistics*, 7:157–168, doi:10.1162/tacl\_a\_00260.
- Bruno COURCELLE (1994), Monadic Second-Order Definable Graph Transductions: A Survey, *Theoretical Computer Science*, 126(1):53–75.
- Bruno COURCELLE and Joost ENGELFRIET (2012), *Graph Structure and Monadic Second-Order Logic: A Language-Theoretic Approach*, volume 138, Cambridge University Press.

- Colin DE LA HIGUERA (2010), *Grammatical Inference: Learning Automata and Grammars*, Cambridge University Press, doi:10.1017/CBO9781139194655.
- Aniello DE SANTO and Thomas GRAF (2019), Structure Sensitive Tier Projection: Applications and Formal Properties, in Raffaella BERNARDI, Greg KOBELE, and Sylvain POGODALLA, editors, *Formal Grammar 2019*, volume 11668 of *Lecture Notes in Computer Science*, pp. 35–50, Springer Verlag, doi:10.1007/978-3-662-59648-7\_3.
- Hossep DOLATIAN and Jonathan RAWSKI (2020), Multi-Input Strictly Local Functions for Templatic Morphology, in *Proceedings of the Society for Computation in Linguistics*, volume 3, pp. 282–296, New Orleans, Louisiana.
- Matt EDLEFSEN, Dylan LEEMAN, Nathan MYERS, Nathaniel SMITH, Molly VISSCHER, and David WELLCOME (2008), Deciding Strictly Local (SL) Languages, in Jon BREITENBUCHER, editor, *Proceedings of the 2008 Midstates Conference for Undergraduate Research in Computer Science and Mathematics*, pp. 66–73.
- Joost ENGELFRIET and Hendrik Jan HOOGEBOOM (2001), MSO Definable String Transductions and Two-Way Finite-State Transducers, *ACM Transactions on Computational Logic*, 2(2):216–254, doi:10.1145/371316.371512.
- Emmanuel FILIOT (2015), Logic-Automata Connections for Transformations, in *Logic and Its Applications*, volume 8923 of *Lecture Notes in Computer Science*, pp. 30–57, Springer Berlin / Heidelberg, doi:10.1007/978-3-662-45824-2\_3.
- Sara FINLEY (2008), *Formal and Cognitive Restrictions on Vowel Harmony*, Ph.D. thesis, Johns Hopkins University.
- Jerry Alan FODOR and Zenon Walter PYLYSHYN (1988), Connectionism and Cognitive Architecture: A Critical Analysis, *Cognition*, 28(1–2):3–71, doi:10.1016/0010-0277(88)90031-5.
- Pedro GARCIA, Enrique VIDAL, and José ONCINA (1990), Learning Locally Testable Languages in the Strict Sense, in *Proceedings of the 1st International Workshop on Algorithmic Learning Theory*, pp. 325–338, Tokyo, Japan, <https://grfia.dlsi.ua.es/repositori/grfia/pubs/111/alt1990.pdf>.
- R. W. N. GOEDEMANS, Jeffrey HEINZ, and Harry VAN DER HULST (2015), StressTyp2, <http://st2.ullet.net/>.
- Edward Mark GOLD (1967), Language Identification in the Limit, *Information and Control*, 10(5):447–474, doi:10.1016/S0019-9958(67)91165-5.
- John Anton GOLDSMITH (1976), *Autosegmental Phonology*, Ph.D. thesis, Swarthmore College, <https://dspace.mit.edu/bitstream/handle/1721.1/16388/03188555-MIT.pdf>.
- Saul GORN (1967), Explicit Definitions and Linguistic Dominoes, in John HART and Satoru TAKASU, editors, *Systems and Computer Science*, pp. 77–115, University of Toronto Press, Toronto, doi:10.3138/9781487592769.

- Thomas GRAF (2011), Closure Properties of the Minimalist Derivation Tree Languages, in Sylvain POGODALLA and Jean-Philippe PROST, editors, *Logical Aspects of Computational Linguistics*, volume 6736 of *Lecture Notes in Computer Science*, pp. 96–111, Springer Berlin / Heidelberg, doi:10.1007/978-3-642-22221-4\_7.
- Thomas GRAF (2014), Beyond the Apparent: Cognitive Parallels Between Syntax and Phonology, in Carson T. SCHÜTZE and Linnaea STOCKALL, editors, *Connectedness: Papers by and for Sarah VanWagenen*, volume 18 of *UCLA Working Papers in Linguistics*, pp. 161–174, <http://phonetics.linguistics.ucla.edu/wpl/issues/wpl18/papers/graf.pdf>.
- Thomas GRAF (2020), Curbing Feature Coding: Strictly Local Feature Assignment, in *Proceedings of the Society for Computation in Linguistics*, volume 3, pp. 362–371, New Orleans, Louisiana, doi:10.7275/f7y5-xz32.
- Thomas GRAF and Nazila SHAFIEI (2019), C-Command Dependencies as TSL String Constraints, in *Proceedings of the Society for Computation in Linguistics*, volume 2, pp. 205–215, New York City, New York, doi:10.7275/4rrx-x488.
- Leonard H. HAINES (1969), On Free Monoids Partially Ordered by Embedding, *Journal of Combinatorial Theory*, 6(1):94–98, doi:10.1016/s0021-9800(69)80111-0.
- Jeffrey HEINZ (2010a), Learning Long-Distance Phonotactics, *Linguistic Inquiry*, 41(4):623–661, doi:10.1162/ling\_a\_00015.
- Jeffrey HEINZ (2010b), String Extension Learning, in *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pp. 897–906, Association for Computational Linguistics, Uppsala, Sweden, <https://www.aclweb.org/anthology/P10-1092>.
- Jeffrey HEINZ (2018), The Computational Nature of Phonological Generalizations, in Larry HYMAN and Frank PLANK, editors, *Phonological Typology*, volume 23 of *Phonetics and Phonology*, chapter 5, pp. 126–195, Mouton de Gruyter, doi:10.1515/9783110451931-005.
- Jeffrey HEINZ, Anna KASPRZIK, and Timo KÖTZING (2012), Learning in the Limit with Lattice-Structured Hypothesis Spaces, *Theoretical Computer Science*, 457:111–127, doi:10.1016/j.tcs.2012.07.017.
- Jeffrey HEINZ, 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: Short Papers*, volume 2, pp. 58–64, Association for Computational Linguistics, Portland, Oregon, <https://aclweb.org/anthology/P11-2011>.
- Jeffrey HEINZ and James ROGERS (2013), Learning Subregular Classes of Languages with Factored Deterministic Automata, in *Proceedings of the 13th Meeting on the Mathematics of Language*, pp. 64–71, Association for

- Computational Linguistics, Sofia, Bulgaria,  
<https://www.aclweb.org/anthology/W13-3007>.
- Suzana HERCULANO-HOUZEL and Roberto LENT (2005), Isotropic Fractionator: A Simple, Rapid Method for the Quantification of Total Cell and Neuron Numbers in the Brain, *Journal of Neuroscience*, 25(10):2518–2521, doi:10.1523/JNEUROSCI.4526-04.2005.
- Larry M. HYMAN and Francis X. KATAMBA (1993), A New Approach to Tone in Luganda, *Language*, 69(1):34–67, doi:10.2307/416415.
- Adam JARDINE (2017a), The Local Nature of Tone-Association Patterns, *Phonology*, 34(2):385–405, doi:10.1017/s0952675717000185.
- Adam JARDINE (2017b), On the Logical Complexity of Autosegmental Representations, in *Proceedings of the 15th Meeting on the Mathematics of Language*, pp. 22–35, Association for Computational Linguistics, London, UK, doi:10.18653/v1/W17-3403.
- Adam JARDINE (2019), The Expressivity of Autosegmental Grammars, *Journal of Logic, Language and Information*, 28(1):9–54, ISSN 1572-9583, doi:10.1007/s10849-018-9270-x.
- Adam JARDINE and Jeffrey HEINZ (2016), Learning Tier-Based Strictly 2-Local Languages, *Transactions of the Association for Computation in Linguistics*, 4:87–98, doi:10.1162/tacl\_a\_00085.
- Adam JARDINE and Kevin MCMULLIN (2017), Efficient Learning of Tier-Based Strictly k-Local Languages, in Frank DREWES, Carlos MARTÍN-VIDE, and Bianca TRUTHE, editors, *Language and Automata Theory and Applications: 11th International Conference*, volume 10168 of *Lecture Notes in Computer Science*, pp. 64–76, Springer, Cham, doi:10.1007/978-3-319-53733-7\_4.
- Gregory M. KOBELE (2011), Minimalist Tree Languages are Closed Under Intersection with Recognizable Tree Languages, in Sylvain POGODALLA and Jean-Philippe PROST, editors, *Logical Aspects of Computational Linguistics*, volume 6736 of *Lecture Notes in Computer Science*, pp. 129–144, Springer Berlin / Heidelberg, doi:10.1007/978-3-642-22221-4\_9.
- John Richard KRUEGER (1961), *Chuvash Manual: Introduction, Grammar, Reader, and Vocabulary*, volume 7 of *Uralic and Altaic Series*, Indiana University, Bloomington.
- Regine LAI (2015), Learnable vs. Unlearnable Harmony Patterns, *Linguistic Inquiry*, 46(3):425–451, doi:10.1162/LING\_a\_00188.
- Dakotah LAMBERT (2020), Grammar Interpretations and Learning TSL Online, unpublished manuscript.
- 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*, volume 3, pp. 330–337, New Orleans, Louisiana, doi:10.7275/2n1j-pj39.

- Shalom LAPPIN and Stuart Merrill SHIEBER (2007), Machine Learning Theory and Practice as a Source of Insight into Universal Grammar, *Journal of Linguistics*, 43(2):393–427, doi:10.1017/S0022226707004628.
- Julie Anne LEGATE and Charles D. YANG (2002), Empirical Re-assessment of Stimulus Poverty Arguments, *The Linguistic Review*, 18(1–2):151–162, doi:10.1515/tlr.19.1-2.151.
- Leonid LIBKIN (2004), *Elements of Finite Model Theory*, Texts in Theoretical Computer Science, Springer Berlin / Heidelberg, doi:10.1007/978-3-662-07003-1.
- Kevin McMULLIN and Gunnar Ólafur HANSSON (2019), Inductive Learning of Locality Relations in Segmental Phonology, *Laboratory Phonology: Journal of the Association for Laboratory Phonology*, 10(1):1–53, doi:10.5334/labphon.150, article 14.
- Robert MCNAUGHTON and Seymour A. PAPERT (1971), *Counter-Free Automata*, MIT Press.
- Tom Michael MITCHELL (1982), Generalization as Search, *Artificial Intelligence*, 18(2):203–226, doi:10.1016/0004-3702(82)90040-6.
- Tom Michael MITCHELL (2017), Key Ideas in Machine Learning, in *Machine Learning: Second Edition*, McGraw Hill, <http://www.cs.cmu.edu/~tom/mlbook/keyIdeas.pdf>, forthcoming.
- Max NELSON, Hossep DOLATIAN, Jonathan RAWSKI, and Brandon PRICKETT (2020), Probing RNN Encoder-Decoder Generalizations of Subregular Functions Using Reduplication, in *Proceedings of the Society for Computation in Linguistics*, volume 3, pp. 31–42, New Orleans, Louisiana, doi:10.7275/xd0r-pg04.
- Martin Andreas NOWAK, Natalia L. KOMAROVA, and Partha NIYOGI (2002), Computational and evolutionary aspects of language, *Nature*, 417(6889):611–617, doi:10.1038/nature00771.
- Jonathan RAWSKI (2019), Tensor Product Representations of Subregular Formal Languages, in *Proceedings of the International Joint Conference on Artificial Intelligence workshop on Neural-Symbolic Learning and Reasoning*, pp. 36–42.
- Jonathan RAWSKI and Hossep DOLATIAN (2020), Multi-Input Strictly Local Functions for Tonal Phonology, in *Proceedings of the Society for Computation in Linguistics*, volume 3, pp. 245–260, New Orleans, Louisiana.
- Jonathan RAWSKI and Jeffrey HEINZ (2019), No Free Lunch in Linguistics or Machine Learning: REsponse to Pater, *Language*, 93(1):e125–e135, doi:10.1353/lan.2019.0004.
- James ROGERS (1998), *A Descriptive Approach to Language-Theoretic Complexity*, (Monograph.) Studies in Logic, Language, and Information, CSLI/FoLLI.

- James ROGERS (2003), Syntactic Structures as Multi-dimensional Trees, *Research on Language and Computation*, 1(3–4):265–305, doi:10.1023/A:1024695608419.
- James ROGERS, Jeff HEINZ, Margaret FERO, Jeremy HURST, Dakotah LAMBERT, and Sean WIBEL (2012), Cognitive and Sub-Regular Complexity, in Glyn MORRILL and Mark-Jan NEDERHOF, editors, *Formal Grammar 2012*, volume 8036 of *Lecture Notes in Computer Science*, pp. 90–108, Springer Verlag, doi:10.1007/978-3-642-39998-5\_6.
- James ROGERS, Jeffrey HEINZ, Gil BAILEY, Matt EDLEFSEN, Molly VISSCHER, David WELLCOME, and Sean WIBEL (2010), On Languages Piecewise Testable in the Strict Sense, in Christian EBERT, Gerhard JÄGER, and Jens MICHAELIS, editors, *The Mathematics of Language: Revised Selected Papers from the 10th and 11th Biennial Conference on the Mathematics of Language*, volume 6149 of *LNCS/LNAI*, pp. 255–265, FoLLI/Springer, doi:10.1007/978-3-642-14322-9\_19.
- James ROGERS and Dakotah LAMBERT (2019a), Extracting Subregular Constraints from Regular Stringsets, *Journal of Language Modelling*, 7(2):143–176, doi:10.15398/jlm.v7i2.209.
- James ROGERS and Dakotah LAMBERT (2019b), Some Classes of Sets of Structures Definable Without Quantifiers, in *Proceedings of the 16th Meeting on the Mathematics of Language*, pp. 63–77, Association for Computational Linguistics, Toronto, Canada, doi:10.18653/v1/W19-5706.
- Marcel-Paul SCHÜTZENBERGER (1965), On Finite Monoids Having Only Trivial Subgroups, *Information and Control*, 8(2):190–194, doi:10.1016/s0019-9958(65)90108-7.
- Imre SIMON (1975), Piecewise Testable Events, in Helmut BRAKHAGE, editor, *Automata Theory and Formal Languages*, volume 33 of *Lecture Notes in Computer Science*, pp. 214–222, Springer Verlag, Berlin, doi:10.1007/3-540-07407-4\_23.
- Kristina STROTHER-GARCIA (2019), *Using Model Theory in Phonology: A Novel Characterization of Syllable Structure and Syllabification*, Ph.D. thesis, University of Delaware.
- Wolfgang THOMAS (1982), Classifying Regular Events in Symbolic Logic, *Journal of Computer and Systems Sciences*, 25:360–376, doi:10.1016/0022-0000(82)90016-2.
- Leslie Gabriel VALIANT (1984), A Theory of the Learnable, *Communications of the ACM*, 27(11):1134–1142, doi:10.1145/1968.1972.
- Mai Ha VU, Ashkan ZEHFROOSH, Kristina STROTHER-GARCIA, Michael SEBOK, Jeffrey HEINZ, and Herbert G. TANNER (2018), Statistical Relational Learning with Unconventional String Models, *Frontiers in Robotics and AI*, 5(76):1–26, doi:10.3389/frobt.2018.00076.
- Edwin Samuel WILLIAMS (1976), Underlying Tone in Margi and Igbo, *Linguistic Inquiry*, 7(3):463–484.

*Lambert et al.*

Charles YANG (2013), Who's Afraid of George Kingsley Zipf? Or: Do Children and Chimps Have Language?, *Significance*, 10(6):29–34,  
doi:10.1111/j.1740-9713.2013.00708.x.

*This work is licensed under the Creative Commons Attribution 3.0 Unported License.*

<http://creativecommons.org/licenses/by/3.0/>

