

Sample-efficient Linguistic Generalizations through Program Synthesis: Experiments with Phonology Problems

Saujas Vaduguru¹ Aalok Sathe² Monojit Choudhury³ Dipti Misra Sharma¹

¹ IIT Hyderabad ² MIT BCS*

³ Microsoft Research India

¹ {saujas.vaduguru@research., dipti@}iit.ac.in

² aalok.sathe@{mit.edu, richmond.edu}

³ monojitc@microsoft.com

Abstract

Neural models excel at extracting statistical patterns from large amounts of data, but struggle to learn patterns or reason about language from only a few examples. In this paper, we ask: Can we learn explicit rules that generalize well from only a few examples? We explore this question using program synthesis. We develop a synthesis model to learn phonology rules as programs in a domain-specific language. We test the ability of our models to generalize from few training examples using our new dataset of problems from the Linguistics Olympiad, a challenging set of tasks that require strong linguistic reasoning ability. In addition to being highly sample-efficient, our approach generates human-readable programs, and allows control over the generalizability of the learnt programs.

1 Introduction

In the last few years, the application of deep neural models has allowed rapid progress in NLP. Tasks in phonology and morphology have been no exception to this, with neural encoder-decoder models achieving strong results in recent shared tasks in phonology (Gorman et al., 2020) and morphology (Vylomova et al., 2020). However, the neural models that perform well on these tasks make use of hundreds, if not thousands of training examples for each language. Additionally, the patterns that neural models identify are not interpretable. In this paper, we explore the problem of learning interpretable phonological and morphological rules from only a small number of examples, a task that humans are able to perform.

Consider the example of verb forms in the language Mandarin presented in Table 1. How would a neural model tasked with filling the two blank cells do? The data comes from a language that

to <i>V</i>	to be <i>Ved</i>
mappasuŋ	dipasuŋ
mattunu	ditunu
?	ditimbe
?	dipande

Table 1: Verb forms in Mandarin (McCoy, 2018)

is not represented in large-scale text datasets that could allow the model to harness pretraining, and the number of samples presented here is likely not sufficient for the neural model to learn the task.

However, a human would fare much better at this task even if they didn’t know Mandarin. Identifying rules and patterns in a different language is a principal concern of a descriptive linguist (Brown and Ogilvie, 2010). Even people who aren’t trained in linguistics would be able to solve such a task, as evidenced by contestants in the Linguistics Olympiads¹, and general-audience puzzle books (Bellos, 2020). In addition to being able to solve the task, humans would be able to express their solution explicitly in terms of rules, that is to say, a *program* that maps inputs to outputs.

Program synthesis (Gulwani et al., 2017) is a method that can be used to learn programs that map an input to an output in a *domain-specific language* (DSL). It has been shown to be a highly sample-efficient technique to learn interpretable rules by specifying the assumptions of the task in the DSL (Gulwani, 2011).

This raises the questions (i) Can program synthesis be used to learn linguistic rules from only a few examples? (ii) If so, what kind of rules can be learnt? (iii) What kind of operations need to explicitly be defined in the DSL to allow it to model linguistic rules? (iv) What knowledge must be im-

*Work done while at the University of Richmond

¹<https://www.ioling.org/>

PLICITLY provided with these operations to allow the model to choose rules that generalize well?

In this work, we use program synthesis to learn phonological rules for solving Linguistics Olympiad problems, where only the minimal number of examples necessary to generalize are given (Şahin et al., 2020). We present a program synthesis model and a DSL for learning phonological rules, and curate a set of Linguistics Olympiad problems for evaluation.

We perform experiments and comparisons to baselines, and find that program synthesis does significantly better than our baseline approaches. We also present some observations about the ability of our system to find rules that generalize well, and discuss examples of where it fails.

2 Program synthesis

Program synthesis is “the task of automatically finding programs from the underlying programming language that satisfy (user) intent expressed in some form of constraints” (Gulwani et al., 2017). This method allows us to specify domain-specific assumptions as a language, and use generic synthesis approaches like FlashMeta (Polozov and Gulwani, 2015) to synthesize programs.

The ability to explicitly encode domain-specific assumptions gives program synthesis broad applicability to various tasks. In this paper, we explore applying it to the task of learning phonological rules. Whereas previous work on rule-learning has focused on learning rules of a specific type (Brill, 1992; Johnson, 1984), the DSL in program synthesis allows learning rules of different types, and in different rule formalisms.

In this work, we explore learning rules similar to rewrite rules (Chomsky and Halle, 1968) that are used extensively to describe phonology. Sequences of rules are learnt using a noisy disjunctive synthesis algorithm NDSyn (Iyer et al., 2019) extended to learn *stateful multi-pass rules* (Sarathi et al., 2021).

2.1 Phonological rules as programs

The synthesis task we solve is to learn a program in a domain-specific language (DSL) for string transduction, that is, to transform a given sequence of input tokens $i \in \mathcal{I}^*$ to a sequence of output tokens $o \in \mathcal{O}^*$, where \mathcal{I} is the set of input tokens, and \mathcal{O} is the set of output tokens. Each token is a symbol accompanied by a feature set, a set of key-value pairs that maps feature names to boolean values.

We learn programs for token-level examples, which transform an input token in its context to output tokens. The program is a sequence of rules which are applied to each token in an input string to produce the output string. The rules learnt are similar to rewrite rules, of the form

$$\phi_{-l} \cdots \phi_{-2} \phi_{-1} X \phi_1 \phi_2 \cdots \phi_r \rightarrow T$$

where (i) $X : \mathcal{I} \rightarrow \mathbb{B}$ is a boolean predicate that determines input tokens to which the rule is applied (ii) $\phi_i : \mathcal{I} \rightarrow \mathbb{B}$ is a boolean predicate applied to the i^{th} character relative to X , and the predicates ϕ collectively determine the context in which the rule is applied (iii) $T : \mathcal{I} \rightarrow \mathcal{O}^*$ is a function that maps an input token to a sequence of output tokens.

X and ϕ belong to a set of predicates \mathcal{P} , and T is a function belonging to a set of transformation functions \mathcal{T} . \mathcal{P} and \mathcal{T} are specified by the DSL.

We allow the model to synthesize programs that apply multiple rules to a single token by synthesizing rules in passes and maintaining state from one pass to the next. This allows the system to learn stateful multi-pass rules (Sarathi et al., 2021).

2.2 Domain-specific language

The domain-specific language (DSL) is the declarative language which defines the allowable string transformation operations. The DSL is defined by a set of operators, a grammar which determines how they can be combined, and a semantics which determines what each operator does. By defining operators to capture domain-specific phenomena, we can reduce the space of programs to be searched to include those programs that capture distinctions relevant to the domain. This allows us to explicitly encode knowledge of the domain into the system.

Operators in the DSL also have a score associated with each operator that allows for setting domain-specific preferences for certain kinds of programs. We can combine scores for each operator in a program to compute a ranking score that we can use to identify the most preferred program among candidates. The ranking score can capture implicit preferences like shorter programs, more/less general programs, certain classes of transformations, etc.

The DSL defines the predicates \mathcal{P} and set of transformations \mathcal{T} that can be applied to a particular token. The predicates and transformations in the DSL we use, along with the description of their semantics, can be found in Tables 2 and 3.

Predicate	
$\text{IsToken}(w, s, i)$	Is x equal to the token s ? This allows us to evaluate matches with specific tokens.
$\text{Is}(w, f, i)$	Is f true for x ? This allows us to generalize beyond single tokens and use features that apply to multiple tokens.
$\text{TransformationApplied}(w, t, i)$	Has the transformation t has been applied to x in a previous pass? This allows us to reference previous passes in learning rules for the current pass.
$\text{Not}(p)$	Negates the predicate p .

Table 2: Predicates that are used for synthesis. The predicates are applied to a token x that is at an offset i from the current token in the word w . The offset may be positive to refer to tokens after the current token, zero to refer to the current token, or negative to refer to tokens before the current token.

Transformation	
$\text{ReplaceBy}(x, s_1, s_2)$	If x is s_1 , it is replaced with s_2 . This allows the system to learn conditional substitutions.
$\text{ReplaceAnyBy}(x, s)$	x is replaced with s . This allows the system to learn unconditional substitutions.
$\text{Insert}(x, S)$	This inserts a sequence of tokens S after x at the end of the pass. It allows for the insertion of variable-length affixes.
$\text{Delete}(x)$	This deletes x from the word at the end of the pass.
$\text{CopyReplace}(x, i)$ $\text{CopyInsert}(x, i)$	These are analogues of the ReplaceBy and Insert transformations where the token which is added is the same as the token at an offset i from x . They allow the system to learn phonological transformations such as assimilation and gemination.
$\text{Identity}(x)$	This returns x unchanged. It allows the system where a transformation applies under certain conditions, but does not under others.

Table 3: Transformations that are used for synthesis. The transformations are applied to a token x in the word w . The offset i for the Copy transformations may be positive to refer to tokens after the current token, zero to refer to the current token, or negative to refer to tokens before the current token.

```
output := Map(disjunction, input_tokens)
disjunction := Else(rule, disjunction)
rule := transformation
      | IfThen(predicate, rule);
```

Figure 1: IfThen-Else statements in the DSL

Sequences of rules are learnt as disjunctions of IfThen operators, and are applied to each token of the input using a Map operator (Figure 1). The conjunction of predicates X and ϕ that define the context are learnt by nesting IfThen operators.

A transformation produces an token that is tagged with the transformation that is applied. This allows for maintaining state across passes.

The operators in our DSL are quite generic and can be applied to other string transformations as well. In addition to designing our DSL for string transformation tasks, we allow for phonological information to be specified as features, which are a set of key-value pairs that map attributes to boolean values. While we restrict our investigation to fea-

tures based only on the symbols in the input, more complex features based on meaning and linguistic categories can be provided to a system that works on learning rules for more complex domains like morphology or syntax. We leave this investigation for future work.

2.3 Synthesis algorithm

We use an extension (Sarathi et al., 2021) of the NDSyn algorithm (Iyer et al., 2019) that can synthesize stateful multi-pass rules. Iyer et al. (2019) describe an algorithm for selecting disjunctions of rules, and use the FlashMeta algorithm as the rule synthesis component. Sarathi et al. (2021) extend the approach proposed by Iyer et al. (2019) for disjunctive synthesis to the task of grapheme-to-phoneme (G2P) conversion in Hindi and Tamil. They propose the idea of learning transformations on token aligned examples, and use language-specific predicates and transformations to learn rules for G2P conversion. We use a similar approach, and use a different set of predicates and

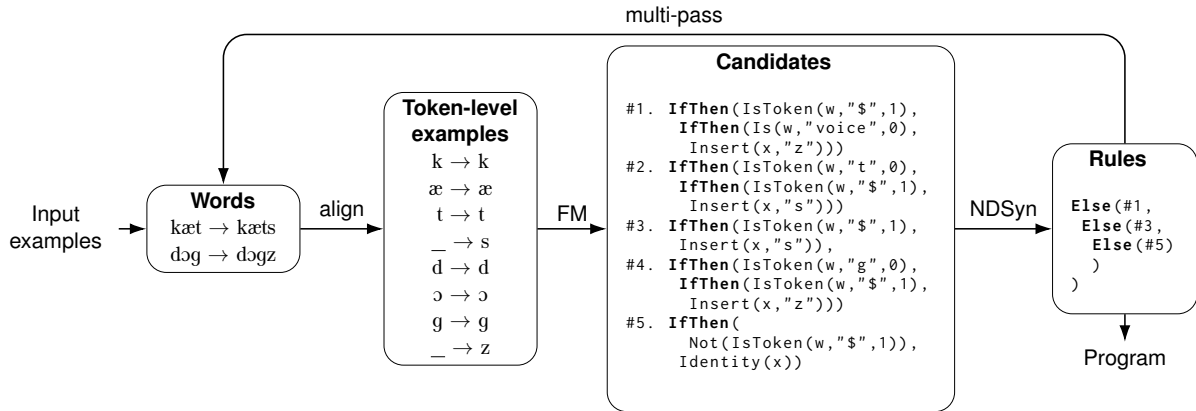


Figure 2: An illustration of the synthesis algorithm. FM is FlashMeta, which synthesizes rules which are combined into a disjunction of rules by NDSyn. Here, rule #1 is chosen over #4 since it uses the more general concept of the voice feature as opposed to a specific token, and thus has a higher ranking score.

transformations that are language-agnostic. Figure 2 sketches the working of the algorithm.

The NDSyn algorithm is an algorithm for learning disjunctions of rules, of the form shown in Figure 1. Given a set of examples, it first generates a set of candidate rules using the FlashMeta synthesis algorithm (Polozov and Gulwani, 2015). This algorithm searches for a program in the DSL that satisfies a set of examples by recursively breaking down the search problem into smaller sub-problems. Given an operator, and the input-output constraints it must satisfy, it infers constraints on each of the arguments to the operator, allowing it to recursively search for programs that satisfy these constraints on each of the arguments. For example, given the Is predicate and a set of examples where the predicate is true or false, the algorithm infers constraints on the arguments the token s and offset i such that the set of examples is satisfied. The working of FlashMeta is illustrated with an example in Figure 3. We use the implementation of the FlashMeta algorithm available as part of the PROSE² framework.

From the set of candidate rules, NDSyn selects a subset of rules with a high ranking score that correctly answers the most examples as well incorrectly answers the least³. Additional details about the algorithm are provided in Appendix A.

The synthesis of multi-pass rules proceeds in passes. In each pass, a set of token-aligned examples is provided as input to the NDSyn algorithm. The resulting rules are then applied to all the exam-

ples, and those that are not solved are passed as the set of examples to NDSyn in the next pass. This proceeds until all the examples are solved, or for a maximum number of passes.

3 Dataset

To test the ability of our program synthesis system to learn linguistic rules from only a few examples, we require a task with a small number of training examples, and a number of test examples which measure how well the model generalises to unseen data. Additionally, to ensure a fair evaluation, the test examples should be chosen such that the samples in the training data provide sufficient evidence to correctly solve the test examples.

To this end, we use problems from the Linguistics Olympiad. The Linguistics Olympiad is an umbrella term describing contests for high school students across the globe. Students are tasked with solving linguistics problems—a genre of composition that presents linguistic facts and phenomena in enigmatic form (Derzhanski and Payne, 2010). These problems typically have 2 parts: the *data* and the *assignments*.

The data consists of examples where the solver is presented with the application rules to some linguistic forms (words, phrases, sentences) and the forms derived by applying the rules to these forms. The data typically consists of 20-50 forms, the minimal number of examples required to infer the correct rules is presented (Şahin et al., 2020).

The assignments provide other linguistic forms, and the solver is tasked with applying the rules inferred from the data to these forms. The forms in the assignments are carefully selected by the

²<https://www.microsoft.com/en-us/research/group/prose/>

³A rule will not produce any answer to examples that don't satisfy the context constraints of the rule.

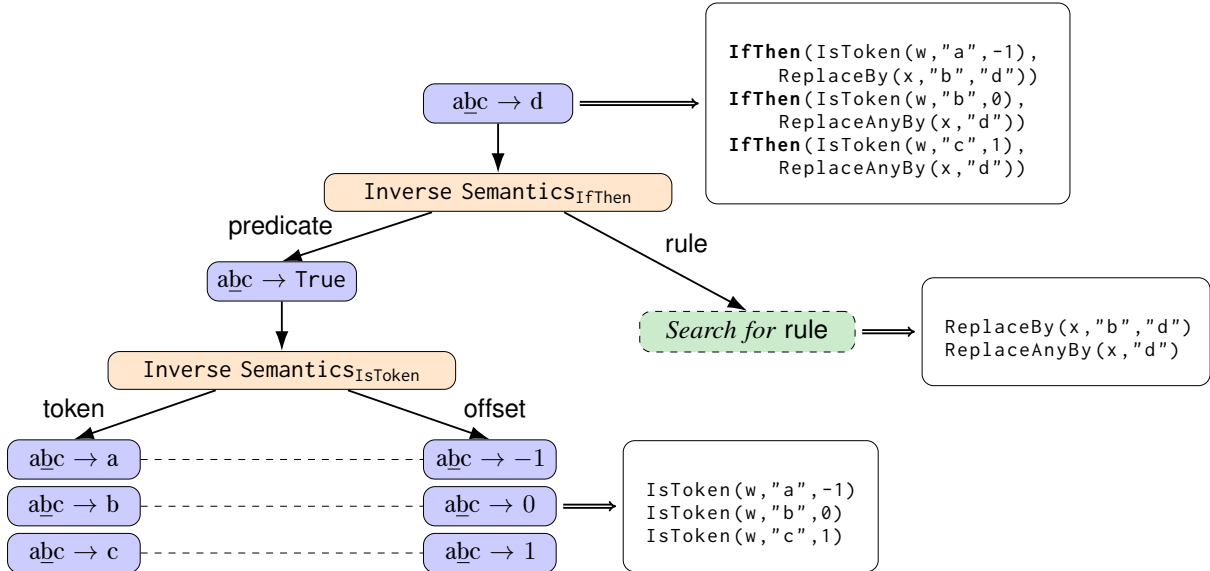


Figure 3: An illustration of the search performed by the FlashMeta algorithm. The blue boxes show the specification that an operator must satisfy in terms of input-output examples, with the input token underlined in the context of the word. The Inverse Semantics of an operator is a function that is used to infer the specification for each argument of the operator based on the semantics of the operator. This may be a single specification (as for predicate) or a disjunction of specifications (as for token and offset). The algorithm then recursively searches for programs to satisfy the specification for each argument, and combines the results of the search to obtain a program. The search for the rule in an IfThen statement proceeds similarly to the search for a predicate. Examples of programs that are inferred from a specification are indicated with \implies . A dashed line between inferred specifications indicates that the specifications are inferred jointly.

designer to test whether the solver has correctly inferred the rules, including making generalizations to unseen data. This allows us to see how much of the intended solution has been learnt by the solver by examining responses to the assignments.

The small number of training examples (data) tests the generalization ability and sample efficiency of the system, and presents a challenging learning problem for the system. The careful selection of test examples (assignment) lets us use them to measure how well the model learns these generalizations.

We present a dataset of 34 linguistics problems, collected from various publicly accessible sources. These problems are based on phonology, and some aspects of the morphology of languages, as well as the orthographic properties of languages. These problems are chosen such that the underlying rules depend only on the given word forms, and not on inherent properties of the word like grammatical gender or animacy. The problems involve (1) inferring phonological rules in morphological inflection (Table 4a) (2) inferring phonological changes between multiple related languages (Table 4b) (3) converting between the orthographic

form of a language and the corresponding phonological form (Table 4c) (4) marking the phonological stress on a given word (Table 4d). We refer to each of these categories of problems as morphophonology, multilingual, transliteration, and stress respectively. We further describe the dataset in Appendix B⁴.

3.1 Structure of the problems

Each problem is presented in the form of a matrix M . Each row of the matrix contains data pertaining to a single word/linguistic form, and each column contains the same form of different words, i.e., an inflectional or derivational paradigm, the word form in a particular language, the word in a particular script, or the stress values for each phoneme in a word. A test sample in this case is presented as a particular cell M_{ij} in the table that has to be filled. The model has to use the data from other words in the same row ($M_{i:}$) and the words in the column ($M_{:j}$) to predict the form of the word in M_{ij} .

In addition to the data in the table, each problem contains some additional information about the symbols used to represent the words. This addi-

⁴The dataset is available [here](#).

base form	negative form	Turkish	Tatar	Listuguj	Pronunciation	Aleut	Stress
joy	kas joya:ya'	bandır	mandır	<i>g'p'ta'q</i>	gəbədə:x	tatul	01000
bi:law	kas bika'law	yelken	cilkän	<i>epsaqtejg</i>	epsaxteck	nətyəlqin	000010000
tipoyసు:da	?	?	osta	<i>emtoqwatg</i>	?	sawat	?
?	kas wurula:la'	bilezik	?	?	əmteskəm	qalpuqal	000010000

(a) Movima negation

(b) Turkish and Tatar

(c) Micmac orthography

(d) Aleut stress

Table 4: A few examples from different types of Linguistics Olympiad problems. ‘?’ represents a cell in the table that is part of the test set.

tional information is meant to aid the solver understand the meaning of a symbol they may not have seen before. We manually encode this information in the feature set associated with each token for synthesis. Where applicable, we also add consonant/vowel distinctions in the given features, since this is a basic distinction assumed in the solutions to many Olympiad problems.

We use the assignments that accompany every problem as the test set, ensuring that the correct answer can be inferred based on the given data.

3.2 Dataset statistics

The dataset we present is highly multilingual. The 34 problems contain samples from 38 languages, drawn from across 19 language families. There are 15 morphophonology problems, 7 multilingual problems, 6 stress, and 6 transliteration problems. The set contains 1452 training words with an average of 43 words per problem, and 319 test words with an average of 9 per problem. Each problem has a matrix that has between 7 and 43 rows, with an average of 23. The number of columns ranges from 2 to 6, with most problems having 2.

4 Experiments

4.1 Baselines

Given that we model our task as string transduction, we compare with the following transduction models used as baselines in shared tasks on G2P conversion (Gorman et al., 2020) and morphological reinflection (Vylomova et al., 2020).

Neural: We use LSTM-based sequence-to-sequence models with attention as well as Transformer models as implemented by Wu (2020). For each problem, we train a single neural model that takes the source and target column numbers, and the source word, and predicts the target word.

WFST: We use models similar to the pair n -gram models (Novak et al., 2016), with the implementation similar to that used by Lee et al. (2020). We

train a model for each pair of columns in a problem. For each test example M_{ij} , we find the column with the smallest index j' such that $M_{ij'}$ is non-empty and use $M_{ij'}$ as the source string to infer M_{ij} .

Additional details of baselines are provided in Appendix C.

4.2 Program synthesis experiments

As discussed in Section 3.1, the examples in a problem are in a matrix, and we synthesize programs to transform entries in one column to entries in another. Given a problem matrix M , we refer to a program to transform an entry in column i to an entry in column j as $M_{:i} \rightarrow M_{:j}$. To obtain token-level examples, we use the Smith-Waterman alignment algorithm (Smith et al., 1981), which favours contiguous sequences in aligned strings.

We train three variants of our synthesis system with different scores for the Is and IsToken operators. The first one, NOFEATURE, does not use features, or the Is predicate. The second one, TOKEN, assigns a higher score to IsToken and prefers more specific rules that reference tokens. The third one, FEATURE, assigns a higher score to Is and prefers more general rules that reference features instead of tokens. All other aspects of the model remain the same across variants.

Morphophonology and multilingual problems:

For every pair of columns (s, t) in the problem matrix M , we synthesize the program $M_{:s} \rightarrow M_{:t}$. To predict the form of a test sample M_{ij} , we find a column k such that the program $M_{:k} \rightarrow M_{:j}$ has the best ranking score, and evaluate it on M_{ik} .

Transliteration problems: Given a problem matrix M , we construct a new matrix M' for each pair of columns (s, t) such that all entries in M' are in the same script. We align word pairs (M_{is}, M_{it}) using the Phonetisaurus many-to-many alignment tool (Jiampojamarn et al., 2007), and build a simple mapping f for each source token to the target token with which it is most frequently aligned. We fill in M'_{is} by applying f to each token of M_{is} and

Model	All		Morphophonology		Multilingual		Transliteration		Stress
	EXACT	CHRF	EXACT	CHRF	EXACT	CHRF	EXACT	CHRF	EXACT
NOFEATURE	26.8%	0.64	30.1%	0.72	42.1%	0.59	12.0%	0.51	15.4%
TOKEN	32.7%	0.63	37.5%	0.68	45.3%	0.60	16.4%	0.52	22.2%
FEATURE	30.9%	0.51	38.6%	0.56	39.9%	0.42	9.5%	0.49	23.0%
LSTM	8.2%	0.44	9.2%	0.49	5.7%	0.45	2.1%	0.31	15.0%
Transformer	5.4%	0.42	2.3%	0.39	9.2%	0.50	1.7%	0.42	12.6%
WFST	20.9%	0.56	16.3%	0.47	38.7%	0.63	29.7%	0.71	2.8%

Table 5: Metrics for all problems, and for problems of each type. The CHRF score for stress problems is not calculated, and not used to determine the overall CHRF score.

$M'_{it} = M_{it}$. We then find a program $M'_{:s} \rightarrow M'_{:t}$.

Stress problems: For these problems, we do not perform any alignment, since the training pairs are already token aligned. The synthesis system learns to transform the source string to the sequence of stress values.

4.3 Metrics

We calculate two metrics: exact match accuracy, and CHRF score (Popović, 2015). The exact match accuracy measures the fraction of examples the synthesis system gets fully correct.

$$\text{EXACT} = \frac{\#\{\text{correctly predicted test samples}\}}{\#\{\text{test samples}\}}$$

The CHRF score is calculated only at the token level, and measures the n -gram overlaps between the predicted answer and the true answer, and allows us to measure partially correct answers. We do not calculate the CHRF score for stress problems as n -gram overlap is not a meaningful measure of performance for these problems.

4.4 Results

Table 5 summarizes the results of our experiments. We report the average of each metric across problems for all problems and by category.

We find that neural models that don't have specific inductive biases for the kind of tasks we present here are not able to perform well with this amount of data. The synthesis models do better than the WFST baseline overall, and on all types of problems except transliteration. This could be due to the simple map computed from alignments before program synthesis causing errors that the rule learning process cannot correct.

5 Analysis

We examine two aspects of the program synthesis models we propose. The first is the way it uses the

explicit knowledge in the DSL and implicit knowledge provided as the ranking score to generalize. We then consider specific examples of problems, and show examples of where our models succeed and fail in learning different types of patterns.

Model	100%	$\geq 75\%$	$\geq 50\%$
NOFEATURE	3	5	7
TOKEN	3	6	10
FEATURE	3	6	11
WFST	1	2	7

Table 6: Number of problems where the model achieves different thresholds of the EXACT score.

5.1 Features aid generalization

Since the test examples are chosen to test specific rules, solving more test examples correctly is indicative of the number of rules inferred correctly. In Table 6, we see that providing the model with features allows it to infer more general rules, solving a greater fraction of more problems. We see that allowing the model to use features increases its performance, and having it prefer more general rules involving features lets it do even better.

5.2 Correct programs are short

In Figure 4 we see that the number of rules in a problem⁵ tends to be higher when the model gets the problem wrong, than when it gets it right. This indicates that when the model finds many specific rules, it overfits to the training data, and fails to generalize well. This holds true for all the variants, as seen in the downward slope of the lines.

We also find that allowing and encouraging a model to use features leads to shorter programs. The average length of a program synthesized by

⁵To account for some problems having more columns than others (and hence more rules), we find the average number of rules for each pair of columns.




Figure 4: Number of rules plotted against EXACT score

NOFEATURES is 30.5 rules, while it is 25.8 for TOKEN, and 20.7 for FEATURE. This suggests that explicit access to features, and implicit preference for them leads to fewer, more general rules.

5.3 Using features

Some problems provide additional information about certain sounds. For example, a problem based on the alternation retroflexes in Warlpiri words (Laughren, 2011) explicitly identifies retroflex sounds in the problem statement. In this case, a program produced by our FEATURE system is able to use these features, and isolate the focus of the problem by learning rules such as

```
IfThen(Not(Is(w, "retroflex", 0)),
      Identity(x))
```

The system learns a concise solution, and is able to generalize using features rather than learning separate rules for individual sounds.

In the case of inflecting a Mandarin verb (McCoy, 2018), the FEATURE system uses a feature to find a more general rule than is the case. To capture the rule that the prefix *di-* changes to *mas-* when the root starts with *s*, the model synthesizes

```
IfThen(Is(w, "fricative", 1),
      ReplaceBy(x, "i", "s"))
```

However, since *s* is the only fricative in the data, this rule is equivalent to a rule specific to *s*. This rule also covers examples where the root starts with *s*, and causes the model to miss the more general rule of a voiceless sound at the beginning of the root to be copied to the end of the prefix. It identifies this rule only for roots starting with *p* as

```
IfThen(IsToken(w, "p", 1),
      CopyReplace(x, w, 1))
```

The TOKEN system does not synthesize these

rules based on features, and instead chooses rules specific to each initial character in the root.

Since the DSL allows for substituting one token with one other, or inserting multiple tokens, the system has to use multiple rules to substitute one token with multiple tokens. In the case of Mandarin, we see one way it does this, by performing multiple substitutions (to transform *di-* to *mas-* it replaces *d* and *i* with *a* and *s* respectively, and then inserts *m*).

5.4 Multi-pass rules

In a problem on Somali verb forms (Somers, 2016), we see a different way of handling multi-token substitutions by using multi-pass rules to create a complex rule using simpler elements. The problem requires being able to convert verbs from 1st person to 3rd person singular. The solution includes a rule where a single token (*I*) is replaced with (*sh*). The learned program uses two passes to capture this rule through sequential application of two rules: first `ReplaceBy(x, "l", "h")`, followed by

```
IfThen(TransformationApplied(w,
  "{ReplaceBy, h}", 1),
  Insert(x, "s"))
```

5.5 Selecting spans of the input

In a problem involving reduplication in Tarangan (Warner, 2019), all variants fail to capture any synthesis rules. Reduplication in Tarangan involves copying one or two syllables in the source word to produce the target word. However, the DSL we use does not have any predicates or transformations that allow the system to reference a span of multiple tokens (which would form a syllable) in the input. Therefore, it fails to model reduplication.

5.6 Global constraints

Since we provide the synthesis model with token-level examples, it does not have access to word-level information. This results in poor performance on stress problems, as stress depends on the entire word. Consider the example of Chickasaw stress (Vaduguru, 2019). It correctly learns the rule

```
IfThen(Is(w, "long", 0),
      ReplaceAnyBy(x, "1"))
```

that stresses any long vowel in the word. However, since it cannot check if the word has a long vowel that has already been stressed, it is not able to correctly model the case when the word doesn't have a long vowel. This results in some samples being marked with stress at two locations, one where

the rule for long vowels applies, and one where the rule for words without long vowels applies.

6 Related work

Gildea and Jurafsky (1996) also study the problem of learning phonological rules from data, and explicitly controlling generalization behaviour. We pursue a similar goal, but in a few-shot setting.

Barke et al. (2019) and Ellis et al. (2015) study program synthesis applied to linguistic rule learning. They make much stronger assumptions about the data (the existence of an underlying form, and the availability of additional information like IPA features). We take a different approach, and study program synthesis models that can work only on the tokens in the word (like NOFEATURE), and also explore the effect of providing features in these cases. We also test our approach on a more varied set of problems that involves aspects of morphology, transliteration, multilinguality, and stress.

Şahin et al. (2020) also present a set of Linguistics Olympiad problems as a test of the metalinguistic reasoning abilities of NLP models. While problems in their set involve finding phonological rules, they also require the knowledge of syntax and semantics that are out of the scope of our study. We present a set of problems that only requires reasoning about surface word forms, and without requiring the meanings.

7 Conclusion

In this paper, we explore the problem of learning linguistic rules from only a few training examples. We approach this using program synthesis, and demonstrate that it is a powerful and flexible technique for learning phonology rules in Olympiad problems. These problems are designed to be challenging tasks that require learning rules from a minimal number of examples. These problems also allow us to specifically test for generalization.

We compare our approach to various baselines, and find that it is capable of learning phonological rules that generalize much better than existing approaches. We show that using the DSL, we can explicitly control the structure of rules, and using the ranking score, we can provide the model with implicit preferences for certain kinds of rules.

Having demonstrated the potential of program synthesis as a learning technique that can work with very little data and provide human-readable models,

we hope to apply it to learning more complex types of linguistic rules in the future.

In addition to being a way to learn rules from data, the ability to explicitly control the generalization behaviour of the model allows for the use of program synthesis to understand the kinds of learning biases and operations that are required to model various linguistic processes. We leave this exploration to future work.

Acknowledgements

We would like to thank Partho Sarthi for invaluable help with PROSE and NDSyn. We would also like to thank the authors of the ProLinguist paper for their assistance. Finally, we would like to thank the anonymous reviewers for their feedback.

References

- Shraddha Barke, Rose Kunkel, Nadia Polikarpova, Eric Meinhardt, Eric Bakovic, and Leon Bergen. 2019. [Constraint-based learning of phonological processes](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6176–6186, Hong Kong, China. Association for Computational Linguistics.
- A. Bellos. 2020. *The Language Lover’s Puzzle Book: Lexical perplexities and cracking conundrums from across the globe*. Guardian Faber Publishing.
- Eric Brill. 1992. [A simple rule-based part of speech tagger](#). In *Speech and Natural Language: Proceedings of a Workshop Held at Harriman, New York, February 23-26, 1992*.
- Keith Brown and Sarah Ogilvie. 2010. *Concise encyclopedia of languages of the world*. Elsevier.
- Noam Chomsky and Morris Halle. 1968. The sound pattern of english.
- Ivan Derzhanski and Thomas Payne. 2010. *The Linguistic Olympiads: academic competitions in linguistics for secondary school students*, page 213–226. Cambridge University Press.
- Kevin Ellis, Armando Solar-Lezama, and Josh Tenenbaum. 2015. [Unsupervised learning by program synthesis](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 973–981. Curran Associates, Inc.
- Daniel Gildea and Daniel Jurafsky. 1996. [Learning bias and phonological-rule induction](#). *Computational Linguistics*, 22(4):497–530.

- Kyle Gorman, Lucas F.E. Ashby, Aaron Goyzueta, Arya McCarthy, Shijie Wu, and Daniel You. 2020. [The SIGMORPHON 2020 shared task on multilingual grapheme-to-phoneme conversion](#). In *Proceedings of the 17th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 40–50, Online. Association for Computational Linguistics.
- Sumit Gulwani. 2011. [Automating string processing in spreadsheets using input-output examples](#). *SIGPLAN Not.*, 46(1):317–330.
- Sumit Gulwani, Oleksandr Polozov, and Rishabh Singh. 2017. [Program synthesis](#). *Foundations and Trends® in Programming Languages*, 4(1-2):1–119.
- Arun Iyer, Manohar Jonnalagedda, Suresh Parthasarathy, Arjun Radhakrishna, and Sriram K Rajamani. 2019. Synthesis and machine learning for heterogeneous extraction. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 301–315.
- Sittichai Jiampojarn, Grzegorz Kondrak, and Tarek Sherif. 2007. [Applying many-to-many alignments and hidden Markov models to letter-to-phoneme conversion](#). In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*, pages 372–379, Rochester, New York. Association for Computational Linguistics.
- Mark Johnson. 1984. [A discovery procedure for certain phonological rules](#). In *10th International Conference on Computational Linguistics and 22nd Annual Meeting of the Association for Computational Linguistics*, pages 344–347, Stanford, California, USA. Association for Computational Linguistics.
- Mary Laughren. 2011. [Stopping and flapping in warlpiri](#). In Dragomir Radev and Patrick Littell, editors, *North American Computational Linguistics Olympiad 2011: Invitational Round*. North American Computational Linguistics Olympiad.
- Jackson L. Lee, Lucas F.E. Ashby, M. Elizabeth Garza, Yeonju Lee-Sikka, Sean Miller, Alan Wong, Arya D. McCarthy, and Kyle Gorman. 2020. [Massively multilingual pronunciation modeling with WikiPron](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4223–4228, Marseille, France. European Language Resources Association.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*.
- Tom McCoy. 2018. [Better left unsaid](#). In Patrick Littell, Tom McCoy, Dragomir Radev, and Ali Sharman, editors, *North American Computational Linguistics Olympiad 2018: Invitational Round*. North American Computational Linguistics Olympiad.
- Josef Robert Novak, Nobuaki Minematsu, and Keikichi Hirose. 2016. [Phonetisaurus: Exploring grapheme-to-phoneme conversion with joint n-gram models in the wfst framework](#). *Natural Language Engineering*, 22(6):907–938.
- Oleksandr Polozov and Sumit Gulwani. 2015. [Flashmeta: A framework for inductive program synthesis](#). *SIGPLAN Not.*, 50(10):107–126.
- Maja Popović. 2015. [chrF: character n-gram F-score for automatic MT evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Gözde Gül Şahin, Yova Kementchedjiev, Phillip Rust, and Iryna Gurevych. 2020. [PuzzLing Machines: A Challenge on Learning From Small Data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1241–1254, Online. Association for Computational Linguistics.
- Partho Sarthi, Monojit Choudhury, Arun Iyer, Suresh Parthasarathy, Arjun Radhakrishna, and Sriram Rajamani. 2021. ProLinguist: Program Synthesis for Linguistics and NLP. *IJCAI Workshop on Neuro-Symbolic Natural Language Inference*.
- Temple F Smith, Michael S Waterman, et al. 1981. Identification of common molecular subsequences. *Journal of molecular biology*, 147(1):195–197.
- Harold Somers. 2016. [Changing the subject](#). In Andrew Lamont and Dragomir Radev, editors, *North American Computational Linguistics Olympiad 2016: Invitational Round*. North American Computational Linguistics Olympiad.
- Saujas Vaduguru. 2019. [Chickasaw stress](#). In Shardul Chiplunkar and Saujas Vaduguru, editors, *Panini Linguistics Olympiad 2019*. Panini Linguistics Olympiad.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Ekaterina Vylomova, Jennifer White, Elizabeth Salesky, Sabrina J. Mielke, Shijie Wu, Edoardo Maria Ponti, Rowan Hall Maudslay, Ran Zmigrod, Josef Valvoda, Svetlana Toldova, Francis Tyers, Elena Klyachko, Ilya Yegorov, Natalia Krizhanovskiy, Paula Czarnowska, Irene Nikkarinen, Andrew Krizhanovskiy, Tiago Pimentel, Lucas Torroba Hennigen, Christo Kirov, Garrett Nicolai, Adina Williams, Antonios Anastasopoulos, Hilaria Cruz, Eleanor Chodroff, Ryan Cotterell, Miikka Silfverberg, and Mans Hulden. 2020. [SIGMORPHON 2020 shared task 0: Typologically diverse morphological inflection](#). In *Proceedings of the 17th SIGMORPHON Workshop on Computational*

Research in Phonetics, Phonology, and Morphology, pages 1–39, Online. Association for Computational Linguistics.

Elysia Warner. 2019. **Tarangan**. In Samuel Ahmed, Bozhidar Bozhanov, Ivan Derzhanski (technical editor), Hugh Dobbs, Dmitry Gerasimov, Shinjini Ghosh, Ksenia Gilyarova, Stanislav Gurevich, Gabrijela Hladnik, Boris Iomdin, Bruno L’Astorina, Tae Hun Lee (editor-in chief), Tom McCoy, André Nikulin, Miina Norvik, Tung-Le Pan, Aleksejs Peguševs, Alexander Piperski, Maria Rubinstein, Daniel Rucki, Artūrs Semēņuks, Nathan Somers, Milena Veneva, and Elysia Warner, editors, *International Linguistics Olympiad 2019*. International Linguistics Olympiad.

Shijie Wu. 2020. Neural transducer. <https://github.com/shijie-wu/neural-transducer/>.

A NDSyn algorithm

We use the NDSyn algorithm to learn disjunctions of rules. We apply NDSyn in multiple passes to allow the model to learn multi-pass rules.

At each pass, the algorithm learns rules to perform token-level transformations that are applied to each element of the input sequence. The token-level examples are passed to NDSyn, which learns the if-then-else statements that constitute a set of rules. This is done by first generating a set of candidate rules by randomly sampling a token-level example and synthesizing a set of rules that satisfy the example. Then, rules are selected to cover the token-level examples.

Rules that satisfy a randomly sampled example are learnt using the FlashMeta program synthesis algorithm (Polozov and Gulwani, 2015). The synthesis task is given by the DSL operator P and the specification of constraints \mathcal{X} that the synthesized program must satisfy. In our application, this specification is in the form of token-level examples, and the DSL operators are the predicates and transformations defined in the paper. The algorithm recursively decomposes the synthesis problem (P, \mathcal{X}) into smaller tasks (P_i, \mathcal{X}_i) for each argument P_i to the operator. \mathcal{X}_i is inferred using the inverse semantics of the operator P_i , which is encoded as a *witness function*. The inverse semantics provides the possible values for the arguments of an operator, given the output of the operator. We refer the reader to the paper by Polozov and Gulwani (2015) for a full description of the synthesis algorithm.

After the candidates are generated, they are ranked according to a ranking score of each program. The ranking score for an operator in a program is computed as a function of the scores of

its arguments. The arguments may be other operators, offsets, or other constants (like tokens or features). The score for an operator in the argument is computed recursively. The score for an offset favours smaller numbers and local rules by decreasing the score for larger offsets. The score for other constants is chosen to be a small negative constant. The scores for the arguments are added up, along with a small negative penalty to favour shorter programs, to obtain the final score for the operator.

This ranking score selects for programs that are shorter, and favours either choosing more general by giving the `Is` predicate a higher score (FEATURE) or more specific rules by giving the `IsToken` predicate a higher score (TOKEN). The top k programs according to the ranking function are chosen as candidates for the next step.

To choose the final set of rules from the candidates generated using the FlashMeta algorithm, we use a *set covering* algorithm that chooses the rules that correctly answer the most number of examples while also incorrectly answering the least. These rules are applied to each example, and the output tokens are tagged with the transformation that is applied. These outputs are then the input to the next pass of the algorithm.

B Dataset

We select problems from various Linguistics Olympiads to create our dataset. We include publicly available problems that have appeared in Olympiads before. We choose problems that only involve rules based on the symbols in the data, and not based on knowledge of notions such as gender, tense, case, or semantic role. These problems are based on the phonology of a particular language, and include aspects of morphology and orthography, and maybe also the phonology of a different language. In some cases where a single Olympiad problem involves multiple components that can be solved independent of each other, we include them as separate problems in our dataset.

We put the data and assignments in a matrix, as described in Section 3.1. We separate tokens in a word by a space while transcribing the problems from their source PDFs. We do not separate diacritics as different tokens, and include them as part of the same token. For each token in the Roman script, we add the boolean features `vowel` and `consonant`, and manually tag the tokens according to whether

they are a vowel or consonant.

We store the problems in JSON files with details about the languages, the families to which the languages belong, the data matrix, the notes used to create the features, and the feature sets for each token.

C Baselines

C.1 Neural

Following Şahin et al. (2020), we use small neural models for sequence-to-sequence tasks. We train a single neural model for each task, and provide the column numbers as tags in addition to the source sequence. We find that the single model approach works better than training a model for each pair of columns.

LSTM: We use LSTM models with soft attention (Luong et al., 2015), with embeddings of size 64, hidden layers of size 128, a 2-layer encoder and a single layer decoder. We apply a dropout of 0.3 for all layers. We train the model for 100 epochs using the Adam optimizer with a learning rate of 10^{-3} , learning rate reduction on plateau, and a batch size of 2. We clip the gradient norm to 5.

Transformer: We use Transformer models (Vaswani et al., 2017) with embeddings of size 128, hidden layers of size 256, a 2-layer encoder and a 2-layer decoder. We apply a dropout of 0.3 for all layers. We train the model for 2000 steps using the Adam optimizer with a learning rate of 10^{-3} , warmup of 400 steps, learning rate reduction on plateau, and a batch size of 2. We use a label smoothing value of 0.1, and clip the gradient norm to 1.

We use the implementations provided at <https://github.com/shijie-wu/neural-transducer/> for all neural models.

C.2 WFST

We use the implementation the WFST models available at <https://github.com/sigmorphon/2020/tree/master/task1/baselines/fst> for the WFST models. We train a model for each pair of columns. We report the results for models of order 5, which were found to perform the best on the test data (highest EXACT score) among models of order 3 to 9.