

# Neural Induction of Finite-State Transducers

Michael Ginn<sup>1</sup> and Alexis Palmer<sup>1</sup> and Mans Hulden<sup>2</sup>

<sup>1</sup>University of Colorado <sup>2</sup>New College of Florida  
michael.ginn@colorado.edu

## Abstract

Finite-State Transducers (FSTs) are effective models for string-to-string rewriting tasks, often providing the efficiency necessary for high-performance applications, but constructing transducers by hand is difficult. In this work, we propose a **novel method for automatically constructing unweighted FSTs** following the hidden state geometry learned by a recurrent neural network. We evaluate our methods on real-world datasets for morphological inflection, grapheme-to-phoneme prediction, and historical normalization, showing that the constructed FSTs are highly accurate and robust for many datasets, **substantially outperforming classical transducer learning algorithms by up to 87% accuracy on held-out test sets.**

## 1 Introduction

*Finite-State Transducers (FSTs)* are effective models for string-to-string tasks such as spellchecking and autocomplete (Ouyang et al., 2017), morphological inflection (Žiga Golob et al., 2012), grapheme-to-phoneme conversion (G2P, Manohar et al., 2022), and transliteration (Hellsten et al., 2017). Neural networks have largely supplanted FSTs in research settings, thanks to their improved accuracy, tolerance for noise, and ease of construction from data (e.g. Kann and Schütze, 2016). However, FSTs are typically more efficient and less resource-intensive than neural networks, making them an appropriate choice for high-performance applications such as mobile keyboards (Wolf-Sonkin et al., 2019) and embedded systems (Wang et al., 2012).

Construction of FSTs is difficult, requiring domain knowledge and significant human effort—for example, Beemer et al. (2020) observed that it took approximately forty hours for an expert to handcraft a single transducer for basic morphological inflection. Prior research has proposed grammar induction algorithms to automatically learn transduc-




Figure 1: There is a theoretical and empirical correspondence between recurrent neural network’s continuous state space (left) and the finite-state space of an automaton (right). In the RNN state space, individual activation values may form clusters that correspond to finite states in the FSA (blue). Transitions between individual activations can be aggregated to form the transitions of the FSA (purple).

ers from datasets with pairs of inputs and outputs, but these algorithms generally struggle to generalize robustly (Gildea and Jurafsky, 1995) or require an oracle to verify arbitrary examples (Vilar, 1996).

In this work, we propose a novel approach for **unweighted transducer induction**<sup>1</sup> which utilizes neural networks as intermediate learners. Neural networks are highly effective at learning from data, even in the presence of noise. Furthermore, there is a clear theoretical correspondence between the continuous, vector-valued hidden state space of a recurrent neural network (RNN) and the discrete state space of a transducer. Our method is inspired by the seminal work of Giles et al. (1992); Omlin and Giles (1996), which sought to understand RNN behavior by forming finite-state automata (FSAs) from clusters of hidden state vectors. Our work is novel in three key ways

1. We use real-world, noisy datasets instead of small toy grammars.

<sup>1</sup>Weighted transducer induction is a distinct, and typically easier problem, and has a number of solutions discussed in §7

2. We extract FSTs instead of FSAs.
3. We propose architectural modifications for the intermediate neural networks that facilitate transducer extraction. Specifically, we design a custom transduction training objective and use a spectral penalty to encourage the RNN to model finite-state-esque dynamics.

We evaluate our method on real-world datasets for morphological inflection, grapheme-to-phoneme conversion, and historical normalization. **We find that our method is highly accurate on inflection**, but struggles on grapheme-to-phoneme and normalization data due to bidirectional dependencies. Nonetheless, our method produces more accurate transducers than existing methods for nearly every dataset in the study. For inflection, the extracted transducers often achieve **similar accuracy to transducers hand-crafted by human experts**. Our code will be released on GitHub.<sup>2</sup>

## 2 Background

A finite-state transducer is defined as a tuple  $A = (\Sigma, \Pi, Q, q_0, \delta)$ , where:

- $\Sigma$  is the *input alphabet* of the automaton, a finite non-empty set of input symbols.
- $\Pi$  is the *output alphabet* of the automaton (often the same as the input alphabet).
- $Q$  is a finite non-empty set of states.
- $q_0 \in Q$  is the *initial state*.
- $\delta : Q \times (\Sigma \cup \{\epsilon\}) \rightarrow Q \times (\Pi \cup \{\epsilon\})$  is the *transition function*, which maps a given state and input symbol (or the empty string) to a new state and output symbol (or empty string).

In this work, we focus on **regular transducers**, which are deterministic/sequential, i.e. a given input string only has one possible output. An example FST is given in Figure 2, which inflects the word ‘cat’ to the plural ‘cats’. For this FST,  $\Sigma = \{c, a, t\}$ ,  $\Pi = \{c, a, t, s\}$ ,  $Q = \{0, 1, 2, 3, 4\}$ , and  $q_0 = 0$ .

Transducers are closely related to *finite-state automata (FSAs)*, which recognize a set of strings. Automata have a single alphabet of symbols, and each transition is associated with a single symbol

<sup>2</sup><https://github.com/michaelpginn/fst-distillation>




Figure 2: Example FST that rewrites "cat" to "cats".

rather than an input and output symbol. Every automaton can be formulated as a transducer that performs the identity function, mapping each allowed string to itself. Likewise, every transducer can be formulated as an automaton that recognizes a valid sequence of (input symbol : output symbol) tuples. We make use of this equivalence in order to extend existing automaton induction methods (Giles et al., 1989) to transducers.

## 3 Method

### 3.1 Motivation

Early work on recurrent neural networks (RNNs) recognized an intuitive correspondence between RNNs and finite-state automata (Cleeremans et al., 1989; Giles et al., 1989; Pollack, 1991). Both models process input sequences symbol-by-symbol, moving through a state space (continuous for RNNs and discrete for FSAs) according to a transition function. These similarities motivated attempts to extract automata from RNNs using a method known as *state clustering*. The method, first proposed by Giles et al. (1991, 1992), clusters hidden-state values to form the states of the automaton (Figure 1). These attempts were highly successful at recovering automata from RNNs trained on synthetic data generated by regular languages such as the Tomita grammars (Tomita, 1982). Furthermore, Casey (1996) proved that an RNN which robustly learns to model a regular language *must* organize its state space to correspond with the target FSA.

### 3.2 Overview

We propose an algorithm for FST inference inspired by these earlier findings. Prior work focused on interpretability, using state clustering to understand the internal dynamics of standard RNN models. Because we are more interested in creating high-accuracy transducers, we modify the architecture and training process of the RNN to facilitate transducer extraction. Our method is as follows:

1. Perform **alignment** of input and output strings (§3.3.1).
2. Train an RNN on **transduction** (§3.3.2).

3. Collect **hidden-state values** from the RNN for input strings from the training set and additional synthetically-generated input strings.
4. Perform **clustering** on hidden-state values to form the states of the FST (§3.3.4).
5. For each cluster, select the most common transition per input symbol, and use a **state splitting algorithm** to resolve non-deterministic transitions (§3.3.5).

### 3.3 Steps

#### 3.3.1 Alignment

A unique challenge in transducer induction is that an *alignment* between pairs of input and output characters must be learned in addition to learning the structure of the automaton. In our method, alignment learning is performed prior to structure learning. We perform alignment of all the input and output strings in the **training dataset only** using the *Chinese Restaurant Process Alignment (CRPAlign)*, a Bayesian Monte Carlo alignment algorithm first described in Cotterell et al. (2016). We provide this algorithm in Appendix A.

For examples where the input and output string are the same length, the alignment algorithm simply aligns each pair of characters. For example, given (run, ran), the alignment would be:

```
r u n
r a n
```

However, if the input and output strings are different lengths, the method will insert *epsilons* ( $\epsilon$ ), which represent an empty string. For (run, runs), the alignment might be:

```
r u n  $\epsilon$ 
r u n s
```

Epsilons in the output string are left alone. Epsilons in the input string, though, are problematic as they can introduce non-determinism when we later convert the aligned pairs into transducer transitions.<sup>3</sup> To solve this, we remove epsilons in the input string and merge the corresponding outputs. We use two merging strategies. For G2P and normalization, we always merge outputs to the right:

```
a  $\epsilon$  c  $\rightarrow$  a c
x y z  $\rightarrow$  x yz
```

<sup>3</sup>Because an epsilon-input transition is optional, potentially resulting in multiple possible paths through the automaton if there is also a valid non-epsilon transition from a given state.

For inflection, we use a greedy merging strategy based on global probabilities. We count all pairs of aligned characters involving epsilon inputs, merge the most common pair, and repeat until no more epsilon inputs remain. In the preceding example, if  $(a, x)(\epsilon, y)$  is the most common pair, the merged example would be:

```
a  $\epsilon$  c  $\rightarrow$  a c
x y z  $\rightarrow$  xy z
```

These aligned pairs form the transitions of the FST. For example, an aligned pair  $(a, xy)$  would correspond to a transition with input  $a$  and output  $xy$ . Next, we will use an RNN to induce the structure of the transducer, given these alignments.

#### 3.3.2 RNN Training

We use one-layer simple RNNs (also known as *Elman RNNs*). RNNs can be trained on different tasks/training objectives, with implications for the types of hidden states that are learned. Virtually all prior work trained RNNs on binary classification, where the RNN should predict whether or not a string is a valid member of the target formal language (Giles et al., 1991, 1992; Watrous and Kuhn, 1992; Omlin and Giles, 1996; Weiss et al., 2018a). This approach works well on simple datasets, but we hypothesized, and then empirically verified, that it does not induce well-separated hidden states for challenging, real-world datasets.

Instead, we propose a novel **transduction training objective** that directly matches the structure of a finite-state transducer, providing explicit supervision of every hidden state in a training example. We formulate the transduction task as follows. For each position in the aligned sequence, we compute the hidden state using the standard update rule.

$$h_t = \sigma(W_h h_{t-1} + W_x x_t) \quad (1)$$

where  $h_t$  is the hidden state at time  $t$ ,  $x_t$  is the input symbol embedding, and  $W_h, W_x$  are trainable weight matrices. At each timestep, we train the model to predict the next output symbol given some input symbol. We use a linear layer that takes a concatenated hidden state and the embedding for the **next input symbol**  $x_{t+1}$  and predicts the corresponding output symbol  $y_{t+1}$ :

$$p(y) = \text{softmax}(W_y \cdot \text{concat}(h_t, x_{t+1})) \quad (2)$$

Our objective function is standard cross-entropy loss over the output symbol logits,  $CEL(p(y), y^*)$  where  $y^*$  is the true output symbol. This training

objective closely matches key features of an FST: intermediate states are solely determined by the sequence of input symbols up to that point, and states are differentiated by their outgoing transitions. In addition, we use a **spectral norm** penalty  $\mathcal{L}_{SN}$ , estimated using the power iteration on the weights of each linear layer for updating the hidden state (Miyato et al., 2018). We hypothesized that by encouraging a small spectral norm, and thus a small Lipschitz constant for the hidden state update function, the RNN will tend to have finite-state-like dynamics (i.e. fixed point attractors), and we found empirically that this enabled more robust FST extraction. We weight the spectral norm penalty using a hyperparameter  $\lambda_{SN}$ , giving our full objective as:

$$\mathcal{L} = \mathcal{L}_{CE}(p(y), y^*) + \lambda_{SN} \mathcal{L}_{SN} \quad (3)$$

During training, our token vocabulary includes any merged symbols produced during alignment, unmerged single-character symbols, and a token for each discrete morphological tag. We do not perform any additional tokenization or merging.

We choose to use Elman RNNs because their simple update rule imposes a Markov assumption, where the next hidden state is dependent only on the current state. Gated RNNs (LSTMs and GRUs) violate this assumption, since by design, the memory cells retain information about the full sequence of prior states. Furthermore, Wang et al. (2018) find that these complex hidden-state dynamics make extracting finite-state automata more difficult. We find empirically that increasing the number of layers does not improve the extracted FST. Details for the training process in [Appendix B](#).

### 3.3.3 Collecting Hidden States

Next, we collect hidden state activation values obtained by inputting strings to the RNN. We use two sources of input-side strings: 1) the examples in the training dataset, and 2) synthetically-generated strings that are plausible for the task.

The latter is necessary because of a difference in how neural models and FSTs encode their *domain*, i.e. the strings that can be inputted to the model. An RNN can produce an output for any input string, allowing it to generalize to strings outside of its training set. Meanwhile, an FST must explicitly encode the set of accepted strings through the input side of the transition labels. If a string does not have a valid sequence of transitions, the FST cannot produce an output. Thus, if we only collect activations for the strings in the training set, we

are very unlikely to observe every possible transition. To solve this, we create additional strings that provide greater coverage over the domain.

The particular strategy we use for synthetic string generation depends on the task. For inflection, the domain is finite, consisting of a fixed number of lemma and feature tag combinations, so we can generate synthetic strings by swapping feature tags for lemmas in the training set. For G2P and normalization, the domain is theoretically infinite. We approximate the domain by creating an n-gram language model from the training set strings. Then, we generate all possible strings from these n-grams up to a max length (we use 6 for efficiency).

Finally, for every training and synthetic example, we feed the string to the RNN and collect the hidden state activation values and predicted output string (ignoring predictions for the training set, where we have a gold output already). Every symbol in the input string is assigned a hidden state, which is computed by recursively applying the update rule for all of the preceding symbols.

### 3.3.4 State Clustering

Thanks to the training objective, we expect hidden states that are close together to be more likely to correspond to the same state of the transducer. We standardize activations and perform k-means clustering (McQueen, 1967). We also tested other clustering algorithms (OPTICS, DBSCAN) but found no clear performance benefit.

The clusters of activations form the states of the FST. Next, we create transitions between states by aggregating the transitions between activations assigned to each cluster ([Figure 3](#), part 1). For the activations collected from the training data, we use the ground-truth input and output labels. For the activations predicted for the full domain, we use the output labels as predicted by the RNN transducer.

A key desideratum for our extracted transducer is **input determinism**. In some cases, activations within a single cluster may violate input determinism and have multiple transitions for the same input symbol with different outputs or different destination states. In this case, we try to recover using the state splitting algorithm described in the following section.

### 3.3.5 State Splitting

Our state splitting algorithm (diagram in [Figure 3](#), pseudocode in [Appendix F](#)) is closely inspired by the refinement approach of Weiss et al. (2018a).




Figure 3: State splitting algorithm, where the minimum threshold  $\lambda_{trans} = 2$ . If a cluster has multiple possible transitions each over the threshold (1), the points are split using an SVM or logistic regression (2). If there are still conflicting transitions, but one of them is under the threshold, it is removed (3). Finally, clusters are materialized into FST states (4).

Starting at the initial state, we check whether the aggregated transitions violate input determinism. Specifically, we look for multiple possible transitions for the same input symbol that occur more than a minimum threshold,  $\lambda_{trans}$ , a tuned hyperparameter. If there is no violation, we continue to downstream states. If there is a violation, we identify the conflicting activations and fit a classifier (either SVM or logistic regression; we try both during the hyperparameter sweep) to separate the cluster into two or more new states (including the non-conflicting activations). Since upstream states may now be input non-deterministic, we requeue these states. After splitting concludes, we remove inaccessible states and minimize the transducer using FSA minimization that treats input/output labels as single symbols (Hopcroft, 1971).

## 4 Experiments

### 4.1 Datasets

We evaluate our method on benchmarks for three tasks that can be represented as regular transductions. We evaluate the constructed transducers based on exact-match accuracy for output strings.

We evaluate on held-out test datasets, seeking transducers that generalize to unseen examples.<sup>4</sup>

**Morphological inflection** We use the 2020 SIGMORPHON Shared Task datasets for morphological inflection (Vylomova et al., 2020). The shared task included datasets for 90 languages; we use the 24 languages (full breakdown in Table 7) used in Beemer et al. (2020), allowing us to compare to their hand-crafted transducers. Morphological inflection is formulated as predicting an inflected word-form given a lemma and morphological tags, such as the following English verb:

[3P][Sg][Prs] run  $\rightarrow$  runs

**Grapheme-to-phoneme** We also use the 2020 SIGMORPHON Shared Task data on grapheme-to-phoneme prediction (Gorman et al., 2020) for 15 languages (Table 8). The task consists of predicting the phonemic (using the International Phonetic Alphabet) string for a word written in the language’s orthography, as in the following French example:

hébergement  $\rightarrow$  eβeʁʒemã

**Historical normalization** We use datasets from Bollmann (2019) for historical normalization (which involves mapping spelling variants of the same word to a single canonical form) for seven languages (Table 9). An example in historical English is given below:

thaire  $\rightarrow$  their

### 4.2 Baselines

We compare against the main algorithms for domain-agnostic transducer induction from unaligned data. We do not compare to domain-specific methods requiring expert task knowledge.

**OSTIA** OSTIA (Oncina et al., 1993) is a seminal algorithm for transducer induction from unaligned data. OSTIA involves building a prefix tree for the labeled examples, pushing output substrings as early as possible, and merging states when it does not violate the input determinism condition. Transducers produced by OSTIA perfectly model the training dataset, but often struggle to generalize to other examples. Furthermore, the algorithm has polynomial complexity with the number of examples. We provide a highly-optimized Cython

<sup>4</sup>In other grammatical inference setups, the goal is to learn a minimal transducer that perfectly models the training dataset. We can trivially construct a prefix tree transducer with perfect accuracy on seen data, but that performs no generalization.

implementation,<sup>5</sup> but even still the algorithm may take nearly 100 days to run on large datasets. For practicality, we set a time limit of one day and stop merging after reaching the limit; these runs are marked with an asterisk (\*) in the results table.

**Data-Driven OSTIA** DD-OSTIA is a modification to OSTIA which can converge with fewer examples, by considering states in breadth-first order (Oncina, 1998). We use the variation described in de la Higuera (2010), which omits the costly equivalence measure and performs merges greedily.

**No Change** For historical normalization, many examples do not have any changes. We compute a simple baseline that outputs the input unchanged.

**Human Expert** For inflection, we compare against the transducers from Beemer et al. (2020), which were handcrafted by linguistics students after studying inflection datasets. These transducers are highly accurate thanks to the students’ knowledge of common morphological processes, but required roughly forty hours for an expert to create.

### 4.3 Hyperparameter Tuning

Since our system cannot be optimized end-to-end, any error from the intermediate model is propagated to the final transducer. Instead, we perform extensive hyperparameter sweeps to minimize error for each component, using a held-out evaluation set (distinct from the test set) to select the best model. Because the models involved are very small, this can be completed rapidly, with the largest languages requiring roughly a day using an NVIDIA A100 GPU. We describe the tuned hyperparameters for the RNN (Appendix B) and extraction algorithm (Appendix C).

## 5 Results

We present our main experimental comparison for inflection in Table 1, G2P in Table 2, and historical normalization in Table 3.

**Performance across Tasks** Across tasks, our method is nearly always superior to OSTIA-based methods and simple no-change baselines. However, our method is typically more effective for inflection than the other two tasks. One possible explanation is that our RNNs are unidirectional. For inflection, the RNN can separate states based on the initial feature tags, avoiding conflicting transitions.

<sup>5</sup><https://github.com/mhulden/pyfoma>

	Expert	OSTIA	DD-OSTIA	Ours
aka	1.000	0.533	0.615	<b>0.975</b>
ceb	0.865	0.135	0.108	<b>0.865</b>
crh	0.964	0.020	0.005*	<b>0.888</b>
czn	0.725	0.013	0.036	<b>0.666</b>
dje	1.000	0.125	0.438	<b>0.750</b>
gaa	1.000	0.391	0.331	<b>0.959</b>
izh	0.929	0.009	0.000	<b>0.438</b>
kon	0.987	0.141	0.179	<b>0.846</b>
lin	1.000	0.043	0.065	<b>0.978</b>
mao	0.667	0.000	0.048	<b>0.452</b>
mlg	1.000	0.189	0.307	<b>0.969</b>
nya	1.000	0.618	0.533	<b>0.993</b>
ood	0.710	0.010	0.013	<b>0.570</b>
orm	0.990	0.037	0.012	<b>0.798</b>
sot	1.000	0.424	0.424	<b>0.727</b>
swa	1.000	<b>0.569</b>	0.541	0.526
syc	0.883	0.031	0.007	<b>0.870</b>
tgk	0.938	0.000	0.000	<b>0.938</b>
tgl	0.778	0.019	0.013	<b>0.383</b>
xty	0.817	0.122	0.000	<b>0.808</b>
zpv	0.789	0.175	0.145	<b>0.706</b>
zul	0.833	0.000	0.000	<b>0.692</b>

Table 1: Accuracy of learned transducers for morphological inflection datasets on held-out test sets. \* marks OSTIA runs that reached the time limit before ending.

Meanwhile, G2P often has transitions dependent on the right-side context, which the RNN is unable to distinguish. For example:

cat → kæ?  
cent → sɛn t

The choice of  $(c, k)$  or  $(c, s)$  is impossible to predict from only the left context. There are two clear paths to address this. First, an alternate alignment algorithm could avoid this issue by delaying outputs until they can be distinguished—in the prior example, the symbols could be aligned as:

$(c, \epsilon)$  (a, kæ) (t, ?)  
 $(c, \epsilon)$  (e, sɛ) (n, n) (t, t)

The other option is to replace the unidirectional RNN with a bidirectional RNN (or transformer). Then, one could extract transducers by first learning a *bimachine* (Reutenauer and Schutzenberger, 1991), which can be thought of as a transducer containing two components and which emits its output depending on the joint states of both left-to-right and right-to-left reading components. A bimachine can then be converted to a regular FST.

Another issue arises when the input and output string have the same length, but should not be aligned with a one-to-one mapping. This occurs when there is an equal number of insertions and

	OSTIA	DD-OSTIA	Ours
fre	0.047	0.098	<b>0.200</b>
dut	0.009	0.011*	<b>0.149</b>
arm	0.067	0.091	<b>0.589</b>
geo	0.013	0.056*	<b>0.596</b>
bul	0.038	0.049	<b>0.236</b>
gre	0.087	0.080	<b>0.307</b>
ady	0.013	0.049	<b>0.184</b>
kor	0.029	0.033	<b>0.057</b>
ice	0.029	0.062*	<b>0.278</b>
hin	0.029	0.100	<b>0.411</b>
lit	0.144	0.167	<b>0.171</b>
jpn	0.029	0.071	<b>0.369</b>
rum	0.298	0.271	<b>0.598</b>
vie	0.022	0.038	<b>0.233</b>
hun	0.053	0.140	<b>0.522</b>

Table 2: Accuracy of learned transducers for grapheme-to-phoneme datasets on held-out test sets. \* marks OSTIA runs that reached the time limit.

	No Change	OSTIA	DD-OSTIA	Ours
deu	0.087	<b>0.251*</b>	0.223*	0.214
hun	0.055	0.160	0.147*	<b>0.316</b>
swe	0.431	0.348	0.338*	<b>0.579</b>
por	0.395	<b>0.559</b>	0.546*	0.503
slv	0.754	0.644	0.628*	<b>0.801</b>
isl	0.338	<b>0.581</b>	0.563*	0.507
spa	0.550	0.568	0.566	<b>0.646</b>

Table 3: Accuracy of learned transducers for historical normalization datasets on held-out test sets. \* marks OSTIA runs that reached the time limit.

deletions, such as in the following G2P example ("rite"):

$$(r, r) (i, a) (t, t) (e, ?)$$

Here, the preferred alignment would be:

$$(r, r) (i, ar) (t, ?) (e, e)$$

but our alignment algorithm is unable to produce this.

**Comparison to OSTIA** While OSTIA attempts to learn generalizable transducers via greedy state merging, it does this by simply checking whether merging two states would violate an example in the training dataset. This means that two dissimilar states could be incorrectly merged so long as there is no counter-example. Meanwhile, our method leverages the hidden-state space geometry learned by the RNN, which should naturally organize so that states with similar outgoing transitions and prefixes are close. In addition, OSTIA learns to align input and output symbols by greedily pushing

output symbols as early as possible in the prefix tree. However, this can result in poor generalization. Our method achieves better generalization by using a statistical algorithm that optimizes alignments globally across the dataset.

An alternative method could use a standard neural inflection model to predict output strings for the full domain, and then use OSTIA on the expanded dataset. This is a coarser version of our approach, treating the neural network as a black box and disregarding any information from its hidden state geometry. Accuracy would match that of the best possible neural model (since OSTIA retains perfect recall on its training data), but OSTIA’s reliance on naive state merging would likely produce very large and/or uninterpretable transducers.

**Comparison to Human Expert** In many cases, the automatically extracted transducers are competitive to human expert-created ones (and achieve perfect performance in some cases). For practical use, an expert could then correct erroneous outputs with small edits, saving significant effort compared to building the transducer manually. The constructed FSTs tend to be far larger than expert-crafted transducers; for example, the inflection ceb constructed FST has 280 states while the expert-crafted FST has 30 states (see [Appendix G](#)).

## 6 Analysis

We randomly selected four datasets to perform ablations: the czn and kon inflection datasets, the geo G2P dataset, and the swe normalization dataset. We use the optimal RNN hyperparameters (except for the third ablation, where we run a sweep), and we run a full sweep for FST extraction.

### 6.1 Ablation: CRPAlign

The CRPAlign algorithm makes use of global probabilities to predict alignments. We ablate this alignment algorithm and use a simple local minimum-edit distance (MED) for a small subset of datasets in [Figure 4](#). For the datasets where the input and output alphabets are overlapping (inflection and normalization), MED achieves very similar scores to CRPAlign. However, for G2P where the output alphabet is disjoint from the input alphabet, CRPAlign is far more effective, as MED will always need to replace the entire string.




Figure 4: Scores with CRPAlign algorithm and simple minimum-edit distance (MED).

## 6.2 Ablation: Extraction with Synthetic Data

We omit the synthetic data described in 3.3.3 and only collect activations for the training examples. For the inflection and G2P datasets, the ablated performance is worse by up to 8 points. This supports our hypothesis that collecting additional activations to attain better coverage of the input domain is beneficial to generalization.




Figure 5: Scores with and without synthetic data.

## 6.3 Ablation: Transduction Objective

We initially considered three different training objectives for the RNN: binary classification, next-aligned-pair prediction, and transduction (which we selected). For binary classification, we generated negative samples by randomly replacing output characters in training examples, and trained the RNN to distinguish correct and incorrect samples. For next-pair prediction, we trained a language modeling head using sequences of aligned pairs of input and output symbols. We provide results for some datasets in Figure 6, with no clear trends. While we had hypothesized that transduction would be necessary for well-separated states, these results would indicate that the process was robust to training objective.

## 7 Related Work

This work is closely inspired by early work on RNNs that identified similarities to finite-state automata, providing empirical and theoretical results




Figure 6: Scores for transducers extracted from RNNs trained for transduction, language modeling, and binary classification.

that suggested that RNNs trained on regular languages would learn to replicate finite-state dynamics (Cleeremans et al., 1989; Giles et al., 1989; Pollack, 1991; Casey, 1996). The basic procedure for state clustering was first proposed by Giles et al. (1991, 1992), though we have made heavy modifications to adapt it to transducers. Various works have proposed alternate clustering techniques such as k-means (Frasconi et al., 1996; Gori et al., 1998) and hierarchical clustering (Alquezar and Sanfeliu, 1994). Tiño and Šajda (1995) is one of the few works that examined transducer inference, though they only studied small synthetic languages and assumed a provided alignment between inputs and outputs. Schellhammer et al. (1998) represents the only (to the best of our knowledge) study that extracted automata from naturally-occurring datasets; here, sequences of part-of-speech tags.

More recent work has studied gated RNNs such as LSTMs, finding that it is far more difficult to extract accurate automata (Wang et al., 2018; Weiss et al., 2018b; Hou and Zhou, 2020). Weiss et al. (2018a) proposed an effective method for active learning that repeatedly refines the states of the automaton based on a validation set, inspiring our splitting procedure. Hong et al. (2022) proposed a similar refinement algorithm that performs merging instead of splitting. Other work has attempted to infer automata from transformers with mixed success (Adriaensen and Maene, 2024; Zhang et al., 2024), but it is unclear whether transformers are even capable of learning regular languages robustly (Hahn, 2020; Bhattamishra et al., 2020; Liu et al., 2023), though hard attention transformers can be constructed to act as transducers (Strobl et al., 2025).

Other work has studied weighted transducer inference using neural networks on real-world tasks such as speech recognition (Lecorvé and Motlicek, 2012) and morphological inflection (Rastogi et al., 2016). Our approach could be adapted to produce

weighted transducers by creating probability distributions for transitions.

## 8 Conclusion

We proposed a method for automatically constructing finite-state transducers from unaligned data, using RNNs as intermediate representations that enable transducer learning via gradient descent. We test our system on morphological inflection, grapheme-to-phoneme prediction, and historical normalization. Our system produces transducers that generalize well to unseen data on two of three tasks, far outperforming existing algorithms and approaching the accuracy of transducers created by human experts. These results provide a tractable way to create highly-performant symbolic systems, and they validate the theoretical claim that recurrent neural networks organize their hidden states with dynamics similar to finite-state automata.

## Limitations

Our method is intended for data which can be efficiently represented as an unweighted finite-state transducer. For datasets which violate Markov assumptions or require probabilistic outputs, this method will not work without modification. As noted, our method struggles on right-side dependencies. We run a wide variety of experiments on several tasks using appropriate baselines, but we do not test methods such as LLM in-context learning or reinforcement learning approaches to transducer inference. We did not have the budget to run full ablations across every dataset, which may reveal other trends.

## Ethical Considerations

Our method is intended to enable transducer inference for many languages and tasks, enabling highly performant language technology for under-resourced languages. However, when using data from minority languages, there is always risk of misuse. We strive to use data as intended and do not claim ownership over this data. Furthermore, our work used significant computing resources, which carry an environmental cost.

## Acknowledgments

Foundation under Grant No. 2149404, "CAREER: From One Language to Another." Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and

do not necessarily reflect the views of the National Science Foundation.

## References

- Rik Adriaensen and Jaron Maene. 2024. [Extracting Finite State Machines from Transformers](#). *arXiv preprint*. ArXiv:2410.06045 [cs].
- R. Alquezar and A. Sanfeliu. 1994. [A hybrid connectionist-symbolic approach to regular grammatical inference based on neural learning and hierarchical clustering](#). In Jaime G. Carbonell, Jörg Siekmann, G. Goos, J. Hartmanis, J. Leeuwen, Rafael C. Carasco, and Jose Oncina, editors, *Grammatical Inference and Applications*, volume 862, pages 203–211. Springer Berlin Heidelberg, Berlin, Heidelberg. Series Title: Lecture Notes in Computer Science.
- Sarah Beemer, Zak Boston, April Bukoski, Daniel Chen, Princess Dickens, Andrew Gerlach, Torin Hopkins, Parth Anand Jawale, Chris Koski, Akanksha Malhotra, Piyush Mishra, Saliha Muradoglu, Lan Sang, Tyler Short, Sagarika Shreevastava, Elizabeth Spaulding, Testumichi Umada, Beilei Xiang, Changbing Yang, and Mans Hulden. 2020. [Linguist vs. machine: Rapid development of finite-state morphological grammars](#). In *Proceedings of the 17th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 162–170, Online. Association for Computational Linguistics.
- Satwik Bhattamishra, Kabir Ahuja, and Navin Goyal. 2020. [On the Ability and Limitations of Transformers to Recognize Formal Languages](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7096–7116, Online. Association for Computational Linguistics.
- Marcel Bollmann. 2019. [A large-scale comparison of historical text normalization systems](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3885–3898, Minneapolis, Minnesota. Association for Computational Linguistics.
- Mike Casey. 1996. [The Dynamics of Discrete-Time Computation, with Application to Recurrent Neural Networks and Finite State Machine Extraction](#). *Neural Computation*, 8(6):1135–1178.
- HongSeok Choi, Dongha Choi, and Hyunju Lee. 2022. [Early stopping based on unlabeled samples in text classification](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 708–718, Dublin, Ireland. Association for Computational Linguistics.
- Axel Cleeremans, David Servan-Schreiber, and James L. McClelland. 1989. [Finite state automata and simple recurrent networks](#). *Neural Computation*, 1(3):372–381.

- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, David Yarowsky, Jason Eisner, and Mans Hulden. 2016. [The SIGMORPHON 2016 shared Task—Morphological reinflection](#). In *Proceedings of the 14th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 10–22, Berlin, Germany. Association for Computational Linguistics.
- Colin de la Higuera. 2010. *Grammatical Inference: Learning Automata and Grammars*. Cambridge University Press, USA.
- Paolo Frasconi, Marco Gori, Marco Maggini, and Giovanni Soda. 1996. [Representation of finite state automata in Recurrent Radial Basis Function networks](#). *Machine Learning*, 23(1):5–32.
- Daniel Gildea and Daniel Jurafsky. 1995. [Automatic induction of finite state transducers for simple phonological rules](#). In *33rd Annual Meeting of the Association for Computational Linguistics*, pages 9–15, Cambridge, Massachusetts, USA. Association for Computational Linguistics.
- C. Giles, Guo-Zheng Sun, Hsing-Hen Chen, Yee-Chun Lee, and Dong Chen. 1989. [Higher Order Recurrent Networks and Grammatical Inference](#). In *Advances in Neural Information Processing Systems*, volume 2. Morgan-Kaufmann.
- C. L. Giles, C. B. Miller, D. Chen, H. H. Chen, G. Z. Sun, and Y. C. Lee. 1992. [Learning and Extracting Finite State Automata with Second-Order Recurrent Neural Networks](#). *Neural Computation*, 4(3):393–405.
- C. L. Giles, C. B. Miller, D. Chen, G. Z. Sun, H. H. Chen, and Y. C. Lee. 1991. [Extracting and Learning an Unknown Grammar with Recurrent Neural Networks](#). In *Advances in Neural Information Processing Systems*, volume 4. Morgan-Kaufmann.
- M. Gori, M. Maggini, E. Martinelli, and G. Soda. 1998. [Inductive inference from noisy examples using the hybrid finite state filter](#). *IEEE Transactions on Neural Networks*, 9(3):571–575. Conference Name: IEEE Transactions on Neural Networks.
- 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.
- Michael Hahn. 2020. [Theoretical limitations of self-attention in neural sequence models](#). *Transactions of the Association for Computational Linguistics*, 8:156–171.
- Lars Hellsten, Brian Roark, Praseon Goyal, Cyril Allauzen, Françoise Beaufays, Tom Ouyang, Michael Riley, and David Rybach. 2017. [Transliterated Mobile Keyboard Input via Weighted Finite-State Transducers](#). In *Proceedings of the 13th International Conference on Finite State Methods and Natural Language Processing (FSMNL 2017)*, pages 10–19, Umeå, Sweden. Association for Computational Linguistics.
- Dat Hong, Alberto Maria Segre, and Tong Wang. 2022. [AdaAX: Explaining Recurrent Neural Networks by Learning Automata with Adaptive States](#). In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 574–584, Washington DC USA. ACM.
- John Hopcroft. 1971. An  $n \log n$  algorithm for minimizing states in a finite automaton. In *Theory of machines and computations*, pages 189–196. Elsevier.
- Bo-Jian Hou and Zhi-Hua Zhou. 2020. [Learning with Interpretable Structure from Gated RNN](#). *arXiv preprint*. ArXiv:1810.10708 [cs].
- Katharina Kann, Kyunghyun Cho, and Samuel R. Bowman. 2019. [Towards realistic practices in low-resource natural language processing: The development set](#). 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 3342–3349, Hong Kong, China. Association for Computational Linguistics.
- Katharina Kann and Hinrich Schütze. 2016. [MED: The LMU system for the SIGMORPHON 2016 shared task on morphological reinflection](#). In *Proceedings of the 14th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 62–70, Berlin, Germany. Association for Computational Linguistics.
- Gwénoél Lecorvé and Petr Motlicek. 2012. Conversion of recurrent neural network language models to weighted finite state transducers for automatic speech recognition. In *Interspeech*, pages 1668–1671.
- Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Roshtamzadeh, and Ameet Talwalkar. 2018. [Hyperband: A novel bandit-based approach to hyperparameter optimization](#). *Journal of Machine Learning Research*, 18(185):1–52.
- Bingbin Liu, Jordan T. Ash, Surbhi Goel, Akshay Krishnamurthy, and Cyril Zhang. 2023. [Transformers Learn Shortcuts to Automata](#). *arXiv preprint*. ArXiv:2210.10749 [cs].
- Kavya Manohar, A. R. Jayan, and Rajeev Rajan. 2022. [Mlphon: A multifunctional grapheme-phoneme conversion tool using finite state transducers](#). *IEEE Access*, 10:97555–97575.
- James B McQueen. 1967. Some methods of classification and analysis of multivariate observations. In *Proc. of 5th Berkeley Symposium on Math. Stat. and Prob.*, pages 281–297.

- Takeru Miyato, Toshiaki Kataoka, Masanori Koyama, and Yuichi Yoshida. 2018. [Spectral normalization for generative adversarial networks](#). In *International Conference on Learning Representations*.
- Christian W. Omlin and C. Lee Giles. 1996. [Extraction of rules from discrete-time recurrent neural networks](#). *Neural Networks*, 9(1):41–52.
- J. Oncina, P. Garcia, and E. Vidal. 1993. [Learning subsequential transducers for pattern recognition interpretation tasks](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(5):448–458. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- Jose Oncina. 1998. The data driven approach applied to the ostia algorithm. In *International Colloquium on Grammatical Inference*, pages 50–56. Springer.
- Tom Ouyang, David Rybach, Françoise Beaufays, and Michael Riley. 2017. [Mobile Keyboard Input Decoding with Finite-State Transducers](#). *arXiv preprint*. ArXiv:1704.03987 [cs].
- Jordan B. Pollack. 1991. [The induction of dynamical recognizers](#). *Machine Learning*, 7(2-3):227–252.
- Pushpendre Rastogi, Ryan Cotterell, and Jason Eisner. 2016. [Weighting finite-state transductions with neural context](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 623–633, San Diego, California. Association for Computational Linguistics.
- Christophe Reutenauer and Marcel-Paul Schutzenberger. 1991. [Minimization of rational word functions](#). *SIAM Journal on Computing*, 20(4):669–685.
- Ingo Schellhammer, Joachim Diederich, Michael Towsey, and Claudia Brugman. 1998. [Knowledge Extraction and Recurrent Neural Networks: An Analysis of an Elman Network trained on a Natural Language Learning Task](#). In *New Methods in Language Processing and Computational Natural Language Learning*.
- Jasper Snoek, Hugo Larochelle, and Ryan P Adams. 2012. [Practical bayesian optimization of machine learning algorithms](#). In *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc.
- Lena Strobl, Dana Angluin, David Chiang, Jonathan Rawski, and Ashish Sabharwal. 2025. [Transformers as transducers](#). *Transactions of the Association for Computational Linguistics*, 13:200–219.
- Peter Tiño and Jozef Šajda. 1995. [Learning and Extracting Initial Mealy Automata with a Modular Neural Network Model](#). *Neural Computation*, 7(4):822–844.
- Masaru Tomita. 1982. [Learning of Construction of Finite Automata from Examples Using Hill-Climbing](#). RR: Regular Set Recognizer. Technical report, Defense Technical Information Center, Fort Belvoir, VA.
- Juan Miguel Vilar. 1996. Query learning of subsequential transducers. In *Grammatical Interference: Learning Syntax from Sentences*, pages 72–83, Berlin, Heidelberg. Springer Berlin Heidelberg.
- 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 Krizhanovsky, Paula Czarnowska, Irene Nikkarinen, Andrew Krizhanovsky, Tiago Pimentel, Lucas Torroba Henning, and 9 others. 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.
- Jhing-Fa Wang, Po-Chun Lin, and Bo-Wei Chen. 2012. Design and applications of embedded systems for speech processing. In *Embedded Systems-High Performance Systems, Applications and Projects*. IntechOpen.
- Qinglong Wang, Kaixuan Zhang, Alexander G. Ororbia II, Xinyu Xing, Xue Liu, and C. Lee Giles. 2018. [A Comparative Study of Rule Extraction for Recurrent Neural Networks](#). *arXiv preprint*. ArXiv:1801.05420 [cs].
- Raymond L. Watrous and Gary M. Kuhn. 1992. [Induction of Finite-State Languages Using Second-Order Recurrent Networks](#). *Neural Computation*, 4(3):406–414.
- Gail Weiss, Yoav Goldberg, and Eran Yahav. 2018a. [Extracting Automata from Recurrent Neural Networks Using Queries and Counterexamples](#). In *Proceedings of the 35th International Conference on Machine Learning*, pages 5247–5256. PMLR. ISSN: 2640-3498.
- Gail Weiss, Yoav Goldberg, and Eran Yahav. 2018b. [On the Practical Computational Power of Finite Precision RNNs for Language Recognition](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 740–745, Melbourne, Australia. Association for Computational Linguistics.
- Lawrence Wolf-Sonkin, Vlad Schogol, Brian Roark, and Michael Riley. 2019. [Latin script keyboards for South Asian languages with finite-state normalization](#). In *Proceedings of the 14th International Conference on Finite-State Methods and Natural Language Processing*, pages 108–117, Dresden, Germany. Association for Computational Linguistics.

Yihao Zhang, Zeming Wei, and Meng Sun. 2024. [Automata Extraction from Transformers](#). *arXiv preprint*. ArXiv:2406.05564 [cs].

Žiga Golob, Jerneja Žganec Gros, Mario Žganec, Boštjan Vesnicer, and Simon Dobrišek. 2012. [Fst-based pronunciation lexicon compression for speech engines](#). *International Journal of Advanced Robotic Systems*, 9(5):211.

## A Chinese Restaurant Process Alignment

We model alignment between two symbol sequences as a *monotone 1-1* path through an alignment lattice, where each step aligns a symbol pair  $(a, b)$  and either side may be  $\epsilon$  (insertion/deletion). Symbol-pair types that are frequent in the corpus are preferred: during training we maintain corpus-wide counts of all observed  $(a, b)$  pairs and convert these counts into smoothed *CRP-style* (Pólya-urn) probabilities for scoring alignments.

Rather than searching directly for a single globally optimal set of alignments, we draw samples from the posterior over alignments using a Gibbs sampler that updates one word pair at a time:

1. **Initialize.** Assign each training pair an initial monotone alignment and accumulate the resulting pair-type counts over the whole corpus.
2. **Gibbs sampling.** For a fixed number of iterations, sweep through the training pairs. For each pair:
  - (a) **Remove:** subtract the counts contributed by the pair’s current alignment from the global count table.
  - (b) **Resample:** sample a new monotone alignment from its conditional distribution given the current counts. This is implemented with dynamic programming over the lattice: a forward pass computes the total probability mass of all paths (in log space, using log-sum-exp), and a backward pass samples a single path by repeatedly choosing among diagonal (match), horizontal (insertion), and vertical (deletion) moves proportional to their path weights.
  - (c) **Add:** add the counts from the newly sampled alignment back into the global table.
3. **Final alignment.** After burn-in, we average the count tables across sampled states

to obtain more stable pair probabilities, and then produce a final deterministic alignment for each word pair by Viterbi decoding (minimum-cost alignment) under these averaged probabilities.<sup>6</sup>

## B RNN Training

We train Elman RNNs with a single layer with the following hyperparameters, where parameters with multiple values were tuned via a hyperparameter sweep.

Parameter	Value(s)
Optimizer	AdamW (default params)
Activation	tanh
Spec-norm weight	0.1
Model dim.	{16, 32, 64, 128}
Dropout	{0, 0.1, 0.3}
Learning rate	{ $2E-4$ , $1E-3$ , $2E-3$ , $1E-2$ }
Label smoothing	0.1
Batch size	{ $2 \times 10^k : 1 \leq k \leq 12$ }
Epochs	{200, 600, 1000}

Table 4: Hyperparameters for RNN training.

Batch size is limited to the four largest values that are less than one-fifth of the size of the training set. No early stopping is used, since this can result in stopping before the target metric has reached its peak (Choi et al., 2022). The hyperparameter sweep uses Bayesian Optimization, which optimizes a surrogate model to predict which combination of hyperparameters should be selected (Snoek et al., 2012). We optimize for the validation set loss and use Hyperband with min and max epochs 75 and 1000 respectively, which allows for bad-performing runs to be stopped early (Li et al., 2018; Kann et al., 2019). We limit to 50 runs for each dataset, except for datasets which have more than 5,000 examples, for which we limit to 25 runs.

## C FST Extraction Sweep

We run a Bayesian sweep over the hyperparameters for FST extraction with 100 runs for most datasets and 25 runs for datasets larger than 5,000 examples. We optimize the following parameters to maximize F1 on the evaluation set:

## D Ablation Results

Results for the ablations are given in Table 6.

Parameter	Value(s)
Clustering algorithm	k-means
Num. clusters	$\{n : 50 \leq n \leq n_{max}\}$
Splitting Classifier	{SVM, Logistic Regression}
$\lambda_{trans}$ (Split Threshold)	{None, 2, 3, 4, 5, 10, 15, 20, 25, 30, 40, 50}

Table 5: Hyperparameters for FST extraction algorithm.  $n_{max}$  is the maximum number of unique hidden-state values for a given dataset. If the split threshold is None, then the splitting algorithm is not used.

Setting	czn	kon	geo	swe
Base	0.666	0.846	0.596	0.539
MED (Abl 1)	0.646	0.878	0.407	0.539
Train only (Abl 2)	0.643	0.814	0.516	0.546
LM (Abl 3)	0.636	0.840	0.511	0.579
Class. (Abl 3)	0.580	0.859	0.560	0.549

Table 6: Numerical results for ablation figures.

Code	Language	# Train	# Dev	# Test
aka	Akan	2,793	380	763
ceb	Cebuano	420	58	111
crh	Crimean Tatar	5,215	745	1,490
czn	Zapotecan	1,088	154	305
dje	Zarma	56	9	16
gaa	Ga	607	79	169
izh	Ingrian	763	112	224
kon	Kongo	568	76	156
lin	Lingala	159	23	46
mao	Maori	145	21	42
mlg	Malagasy	447	62	127
nya	Chewa	3,031	429	853
ood	O’odham	1,123	160	314
orm	Oromo	1,424	203	405
ote	Mezquital Otomi	22,962	3,231	6,437
san	Sanskrit	22,968	3,188	6,272
sot	Sotho	345	50	99
swa	Swahili	3,374	469	910
syc	Syriac	1,917	275	548
tgk	Tajik	53	8	16
tgl	Tagalog	1,870	236	478
xty	Yoloxóchitl Mixtec	2,110	299	600
zpv	Chichicapan Zapotec	805	113	228
zul	Zulu	322	42	78

Table 7: Inflection datasets from Vylomova et al. (2020).

## E Datasets

We report statistics for each of the datasets used in Table 7, Table 8 and Table 9. For historical normalization, if there were multiple distinct sources (often based on different time periods), we selected on source.

<sup>6</sup>The implementation is available as a stand-alone aligner at <https://github.com/mhulden/crpalgn>.

Code	Language	# Train	# Dev	# Test
ady	Adyghe	3,600	450	450
arm	Armenian	3,600	450	450
bul	Bulgarian	3,600	450	450
dut	Dutch	3,600	450	450
fre	French	3,600	450	450
geo	Georgian	3,600	450	450
gre	Modern Greek	3,600	450	450
hin	Hindi	3,600	450	450
hun	Hungarian	3,600	450	450
ice	Icelandic	3,600	450	450
jpn	Japanese	3,600	450	450
kor	Korean	3,600	450	450
lit	Lithuanian	3,600	450	450
rum	Romanian	3,600	450	450
vie	Vietnamese	3,600	450	450

Table 8: G2P datasets (Gorman et al., 2020).

Code	Language	# Train	# Dev	# Test
deu	German - Anselm	8,268	8,242	10,000*
hun	Hungarian	46,907	8,987	8,946
isl	Icelandic	10,864	2,490	2,470
por	Portuguese - 19th c.	15,699	3,544	3,670
slv	Slovenian - Gaj	35,436	8,330	8,614
spa	Spanish - 19th c.	5,577	1,086	1,204
swe	Swedish	8,465	1,308	9,576

Table 9: Normalization datasets (Bollmann, 2019)

## F Clustering and Splitting Algorithm

The pseudocode for the splitting algorithm is given in Appendix F.

## G Example Transducers

We visualize transducers for the inflection ceb dataset, comparing the expert-crafted (Figure 7) and automatically constructed (Figure 8) transducers.

---

**Algorithm 1** State clustering and splitting. Assume  $\text{TRANS}(q, \sigma)$  gives the transitions for state  $q$  and input symbol  $\sigma$ , returned as a set of tuples  $(\pi, q_{dest})$  with the output symbol and destination state. Likewise,  $\text{TRANSOVERTHRESHOLD}$  returns only transitions that occur more than the threshold. SVM splits a state to resolve multiple possible transitions for a given input.

---

```

1: procedure RESOLVETRANSITIONS( $q_0, \lambda_{trans}$ )
2:    $Q \leftarrow [q_0]$ 
3:    $visited \leftarrow \{\}$ 
4:   while length of  $Q > 0$  do
5:      $q \leftarrow Q_0$ 
6:      $Q \leftarrow Q_{1..|Q|}$ 
7:     for  $\sigma \in \Sigma$  do
8:        $T_q \leftarrow \text{TRANSOVERTHRESHOLD}(q, \sigma, \lambda_{trans})$ 
9:       if  $|T_q| = 0$  then
10:        continue
11:       else if  $|T_q| = 1$  then
12:          $\text{TRANS}(q, \sigma) \leftarrow T_q$ 
13:          $(\pi, q_{dest}) \leftarrow T_{q_0}$ 
14:          $Q \leftarrow Q \cdot [q_{dest}]$ 
15:       else
16:          $Q_{upstream} \leftarrow \text{SPLIT}(q)$ 
17:          $Q \leftarrow Q \cdot Q_{upstream}$ 
18: procedure SPLIT( $q, \lambda_{trans}$ )
19:    $\sigma_{worst} = \text{argmax}_{\sigma} |\text{TRANSOVERTHRESHOLD}(q, \sigma, \lambda_{trans})|$ 
20:    $Q_{new} \leftarrow \text{SVM}(q, \sigma_{worst})$ 
21:   return  $\{q_{source} : [\exists q_{new} \in Q_{new} : (*, q_{new}) \in \text{TRANS}(q_{source}, *)]\}$ 

```

---

