

LEARNING TO RECOMBINE AND RESAMPLE DATA FOR COMPOSITIONAL GENERALIZATION

Ekin Akyürek
MIT CSAIL
akyurek@mit.edu

Afra Feyza Akyürek
Boston University
akyurek@bu.edu

Jacob Andreas
MIT CSAIL
jda@mit.edu

ABSTRACT

Flexible neural sequence models outperform grammar- and automaton-based counterparts on a variety of tasks. However, neural models perform poorly in settings requiring compositional generalization beyond the training data—particularly to rare or unseen subsequences. Past work has found symbolic scaffolding (e.g. grammars or automata) essential in these settings. We present a family of *learned data augmentation* schemes that support a large category of compositional generalizations without appeal to latent symbolic structure. Our approach to data augmentation has two components: *recombination* of original training examples via a prototype-based generative model and *resampling* of generated examples to encourage extrapolation. Training an ordinary neural sequence model on a dataset augmented with recombined and resampled examples significantly improves generalization in two language processing problems—instruction following (SCAN) and morphological analysis (SIGMORPHON 2018)—where our approach enables learning of new constructions and tenses from as few as eight initial examples.

1 INTRODUCTION

How can we build machine learning models with the ability to learn new concepts in context from little data? Human language learners acquire new word meanings from a single exposure (Carey & Bartlett, 1978), and immediately incorporate words and the concepts they represent productively and compositionally into a larger linguistic or conceptual system (Berko, 1958; Piantadosi & Aslin, 2016). Despite the remarkable success of neural network models on many learning problems in recent years—including one-shot learning of classifiers and policies (Santoro et al., 2016; Wang et al., 2016)—this kind of few-shot learning of composable concepts remains beyond the reach of standard neural models in both diagnostic and naturalistic settings (Lake & Baroni, 2018; Bahdanau et al., 2019).

Consider the *few-shot morphology learning* problem shown in Fig. 1, in which a learner must predict various linguistic features (e.g. 3rd person, Singular, PREsent tense) from word forms, with only a small number of examples of the PAST tense in the training set. Neural sequence-to-sequence models (e.g. Bahdanau et al., 2015) trained on this kind of imbalanced data fail to predict past-tense tags on held-out inputs of any kind (Section 5). Previous attempts to address this and related shortcomings in neural models have focused on explicitly encouraging rule-like behavior, e.g. by modeling data with a symbolic system (Jia & Liang, 2016) or applying rule-based data augmentation (Andreas, 2020). These procedures are highly task-specific and do not generalize well to less structured problems involving a mix of rule-like and

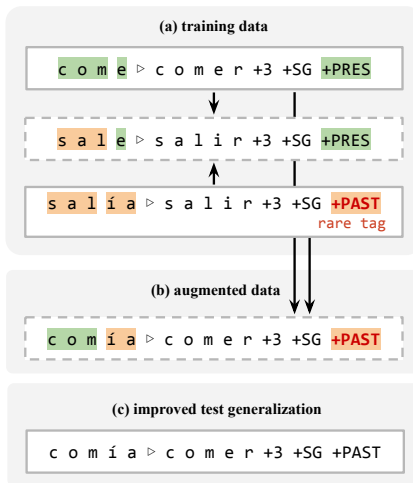


Figure 1: We first train a generative model to reconstruct training pairs $(x \triangleright y)$ by constructing them from other training pairs (dashed box: target)(a). We then perform data augmentation by *sampling* from this model, preferentially generating samples in which y contains rare tokens or substructures (dashed box: sample) (b). Conditional models trained on the augmented dataset accurately predict outputs y from new inputs x requiring compositional generalization (c).

exceptional behavior. More fundamentally, they fail to answer the question of whether explicit rules are necessary for compositional inductive bias, and whether it is possible to achieve human-like compositional inductive generalization *without* appeal to an underlying symbolic generative process.

This paper investigates a data augmentation procedure based on two weaker principles, **recombination** and **resampling**, and finds that it is sufficient to induce many of the compositional generalizations studied in previous work. Concretely: **(1)** We introduce a new multi-prototype generative model that synthesizes new input–output pairs by learning to copy fragments of reference training examples (recombination). **(2)** We show that this model, sampled in low-density regions of the training data (resampling; Chawla et al., 2002), can be used to bootstrap an ordinary sequence-to-sequence model that achieves good performance in-distribution and generalizes compositionally beyond it (Fig. 1a). This procedure is effective in challenging synthetic sequence modeling tasks (Lake & Baroni, 2018) and morphological analysis in multiple natural languages (Cotterell et al., 2018).

Neither recombination nor resampling require explicit representation of symbolic structure. Together, they provide performance comparable to or better than state-of-the-art neuro-symbolic approaches in both domains studied. Our results suggest that some failures of systematicity in neural models can be explained by simpler structural constraints on data distributions and corrected with weaker inductive bias than previously described. We implemented our experiments in Knet (Yuret, 2016) using Julia (Bezanson et al., 2017) and the code for all experiments is released under a public repository ¹

2 BACKGROUND AND RELATED WORK

Compositional generalization Systematic compositionality—the capacity to identify rule-like regularities from limited data and generalize these rules to novel situations—is an essential feature of human reasoning (Fodor et al., 1988). While details vary, a common feature of existing attempts to formalize systematicity (e.g. Gordon et al., 2020) is the intuition that learners should make accurate predictions in situations featuring novel combinations of input or output substructures. For example, learners should generalize from actions seen in isolation to more complex commands involving those actions (Lake et al., 2019), and from relations of the form $r(a, b)$ to $r(b, a)$ (Keysers et al., 2019; Bahdanau et al., 2018). In machine learning, previous studies have found that standard neural architectures fail to generalize systematically even when achieving high in-distribution accuracy in a variety of settings (Lake & Baroni, 2018; Bastings et al., 2018; Johnson et al., 2017).

Data resampling and augmentation Learning to predict sequential outputs with rare or novel substructures is related to the widely studied problem of *class imbalance* in classification problems. There, undersampling of the majority class or oversampling of the minority class has been found to improve the quality of predictions for rare phenomena (Japkowicz et al., 2000). This can be combined with targeted *data augmentation* with synthetic examples of the minority class (Chawla et al., 2002). Generically, given a training dataset \mathcal{D} , learning with class resampling and data augmentation involves defining an augmentation distribution $\tilde{p}(x, y \mid \mathcal{D})$ and sample weighting function $w(x, y)$ and maximizing a training objective of the form:

$$\mathcal{L}(\theta) = \underbrace{\frac{1}{|\mathcal{D}|} \sum_{x \in \mathcal{D}} \log p_{\theta}(y \mid x)}_{\text{Original training data}} + \underbrace{\mathbb{E}_{(x, y) \sim \tilde{p}} w(x, y) \log p_{\theta}(y \mid x)}_{\text{Augmented data}} . \quad (1)$$

In addition to task-specific model architectures (Andreas et al., 2016; Russin et al., 2019), recent years have seen a renewed interest in data augmentation as a flexible and model-agnostic tool for encouraging controlled generalization (Ratner et al., 2017). Existing proposals for sequence models are mainly rule-based—in sequence modeling problems, specifying a synchronous context-free grammar (Jia & Liang, 2016) or regular string rewriting systems to generate new examples (Andreas, 2020). While rule-based data augmentation is highly effective in structured problems featuring crisp correspondences between inputs and outputs, the effectiveness of such approaches involving more complicated, context-dependent relationships between inputs and outputs has not been well-studied.

Learned data augmentation What might compositional data augmentation look like without rules as a source of inductive bias? As Fig. 1 suggests, an ideal data augmentation procedure (\tilde{p} in Eq. 1) should automatically identify which ways of transforming and combining examples are valid. A

¹Public repository for experiments: <https://github.com/ekinakyurek/compgen>

promising starting point is provided by **prototype-based** models, a number of which (Gu et al., 2018; Guu et al., 2018; Khandelwal et al., 2019) have been recently proposed in the sequence modeling setting. Such models generate a distribution $p(d)$ via:

$$d \sim p_{\text{rewrite}}(\cdot | d'; \theta) \quad \text{where} \quad d' \sim \text{Unif}(\mathcal{D}); \quad (2)$$

for a dataset \mathcal{D} and a learned sequence rewriting model $p_{\text{rewrite}}(d | d'; \theta)$. (To avoid confusion, we will use the symbol d to denote a *datum*. Because a data augmentation procedure must produce complete input–output examples, each d is an (x, y) pair for the conditional tasks evaluated in this paper.) While recent variants implement p_{rewrite} with neural networks, these models are closely related to classical kernel density estimators (Rosenblatt, 1956).

Existing work uses prototype-based models as replacements for standard sequence models. We argue here that they are better suited to use as data augmentation procedures—they can produce high-precision examples in the neighborhood of existing training data, then be used to bootstrap simpler predictors that extrapolate more effectively. However, as experiments in Section 5 will show, existing prototype-based models give mixed results on challenging generalizations of the kind depicted in Fig. 1 when used for either direct prediction or data augmentation—performing well in some settings but barely above baseline in others.

Accordingly, we next describe two improvements that transform prototype-based language models into an effective learned data augmentation scheme. Section 3 describes an implementation of p_{rewrite} that encourages greater sample diversity and well-formedness via a multi-prototype copying mechanism. Section 4 describes heuristics for sampling prototypes d' and model outputs d to focus data augmentation on the most informative examples. Section 5 investigates the empirical performance of both components of the approach, finding that they together provide they a simple but surprisingly effective tool for enabling compositional generalization.

3 PROTOTYPE-BASED SEQUENCE MODELS FOR DATA RECOMBINATION

We begin with a brief review of existing prototype-based sequence models. Our presentation mostly follows the *retrieve-and-edit* approach of Guu et al. (2018), but versions of the approach in this paper could also be built on retrieval-based models implemented with memory networks (Miller et al., 2016; Gu et al., 2018) or transformers (Khandelwal et al., 2019; Guu et al., 2020). The generative process described in Eq. 2 implies a marginal sequence probability:

$$p(d) = \frac{1}{|\mathcal{D}|} \sum_{d' \in \mathcal{D}} p_{\text{rewrite}}(d | d'; \theta) \quad (3)$$

Maximizing this quantity over the training set with respect to θ will encourage p_{rewrite} to act as a model of *valid data transformations*: To be assigned high probability, every training example must be explained by at least one other example and a parametric rewriting operation. (The trivial solution where p_θ is the identity function, with $p_\theta(d | d' = d) = 1$, can be ruled out manually in the design of p_θ .) When \mathcal{D} is large, the sum in Eq. 3 is too large to enumerate exhaustively when computing the marginal likelihood. Instead, we can optimize a lower bound by restricting the sum to a **neighborhood** $\mathcal{N}(d) \subset \mathcal{D}$ of training examples around each d :

$$p(d) \geq \frac{1}{|\mathcal{D}|} \sum_{d' \in \mathcal{N}(d)} p_{\text{rewrite}}(d | d'; \theta). \quad (4)$$

The choice of \mathcal{N} is discussed in more detail in Section 4. Now observe that:

$$\log p(d) \geq \log \left(|\mathcal{N}(d)| \sum_{d' \in \mathcal{N}(d)} \frac{1}{|\mathcal{N}(d)|} p_{\text{rewrite}}(d | d'; \theta) \right) - \log |\mathcal{D}| \quad (5)$$

$$\geq \frac{1}{|\mathcal{N}(d)|} \sum_{d' \in \mathcal{N}(d)} \log p_{\text{rewrite}}(d | d'; \theta) + \log \left(\frac{|\mathcal{N}(d)|}{|\mathcal{D}|} \right) \quad (6)$$

where the second step uses Jensen’s inequality. If all $|\mathcal{N}(d)|$ are the same size, maximizing this lower bound on log-likelihood is equivalent to simply maximizing

$$\sum_{d' \in \mathcal{N}(d)} \log p_{\text{rewrite}}(d | d'; \theta) \quad (7)$$

over \mathcal{D} —this is the ordinary conditional likelihood for a string transducer (Ristad & Yianilos, 1998) or sequence-to-sequence model (Sutskever et al., 2014) with examples $(d, d' \in \mathcal{N}(d))$. Past work also includes a continuous latent variable z , defining:

$$p_{\text{rewrite}}(d \mid d') = \mathbb{E}_{z \sim p(z)} [p_{\text{rewrite}}(d \mid d', z; \theta)] \tag{8}$$

As discussed in Section 5, the use of a continuous latent variable appears to make no difference for the tasks in this paper. The remainder of our presentation focuses on the simpler model in Eq. 7.

We have motivated prototype-based models by arguing that p_{rewrite} learns a model of transformations licensed by the training data. However, when generalization involves complex compositions, we will show that neither a basic RNN implementation of p_{rewrite} or a single prototype is enough; we must provide the learned rewriting model with a larger inventory of parts and encourage reuse of those parts as faithfully as possible. This motivates the two improvements on the prototype-based modeling framework described in the remainder of this section: generalization to multiple prototypes (Section 3.1) and a new rewriting model (Section 3.2).

3.1 n -PROTOTYPE MODELS

To improve *compositionality* in prototype-based models, we equip them with the ability to condition on multiple examples simultaneously. We extend the basic prototype-based language model to n prototypes, which we now refer to as a **recombination** model p_{recomb} :

$$d \sim p_{\text{recomb}}(\cdot \mid d'_{1:n}; \theta) \quad \text{where} \quad d'_{1:n} \stackrel{\text{def}}{=} (d'_1, d'_2, \dots, d'_n) \sim p_{\Omega}(\cdot) \tag{9}$$

Rather than making $p_{\Omega}(d'_{1:n})$ uniform on \mathcal{D}^n , we assume access to a set of *compatible prototypes* $\Omega \subset \mathcal{D}^n$ (Section 4) and let $p_{\Omega} \stackrel{\text{def}}{=} \text{Unif}(\Omega)$. We update Eq. 6 to feature a corresponding multi-prototype neighborhood $\mathcal{N} : \mathcal{D} \rightarrow \Omega$. The only terms that have changed are the conditioning variable and the constant term, and it is again sufficient to choose θ to optimize $\sum_{d'_{1:n} \in \mathcal{N}(d)} \log p_{\text{recomb}}(d \mid d'_{1:n})$ over \mathcal{D} . However, as described next, effective implementation of p_{recomb} for multiple prototypes requires improvements to the model architecture used in previous work.

3.2 RECOMBINATION NETWORKS

Past work has found that latent-variable neural sequence models often ignore the latent variable and attempt to directly model sequence marginals (Bowman et al., 2015). When an ordinary sequence-to-sequence model with attention is used to implement p_{recomb} , even in the one-prototype case, generated sentences often have little overlap with their prototypes (Weston et al., 2018). We describe a specific model architecture for p_{recomb} that enables it to explicitly reuse fragments of multiple prototypes, rather than functioning as a generic noise model, by facilitating *copying* from independent streams containing prototypes and previously generated input tokens.




Figure 2: (a) RNN encoders produce contextual embeddings for prototype tokens. (b) In the decoder, a gated copy mechanism reuses prototypes and generated output tokens via an attention mechanism (dashed lines).

We take $p_{\text{recomb}}(d \mid d'_{1:n}; \theta)$ to be a neural (multi-)sequence-to-sequence model (c.f. Sutskever et al., 2014) which decomposes probability autoregressively: $p_{\text{recomb}}(d \mid d'_{1:n}; \theta) = \prod_t p(d^t \mid d^{<t}, d'_{1:n}; \theta)$. As shown in Fig. 2, three LSTM encoders—two for the prototypes and one for the input prefix—compute sequences of token representations h_{proto} and h_{out} respectively. Given the current decoder hidden state h_{out}^t , the model first attends to both prototype and output tokens:

$$\alpha_{\text{out}}^i \propto \exp(h_{\text{out}}^t \top W_o h_{\text{out}}^i) \quad i < t \tag{10}$$

$$\alpha_{\text{proto}}^{kj} \propto \exp(h_{\text{out}}^t \top W_p h_{\text{proto}}^{kj}) \quad k \leq n, j \leq |d'_k| \tag{11}$$

To enable copying from each sequence, we project attention weights α_{out} and α_{proto}^k onto the output vocabulary to produce a sparse vector of probabilities:

$$p_{\text{copy,out}}^t(d^t = w) = \sum_{i < t} \mathbb{1}[d_i = w] \cdot \alpha_{\text{out}}^i \quad (12)$$

$$p_{\text{copy,proto-k}}^t(d^t = w) = \sum_{j \leq |d_k^t|} \mathbb{1}[d_{k,j}^t = w] \cdot \alpha_{\text{proto}}^{kj} \quad (13)$$

Unlike rule-based data recombination procedures, however, p_{recomb} is not required to copy from the prototypes, and can predict output tokens directly using *values* retrieved by the attention mechanism:

$$h_{\text{pre}}^t = [h_{\text{out}}^t, \sum_i \alpha_{\text{out}}^i h_{\text{out}}^i, \sum_{k,j} \alpha_{\text{proto}}^{kj} h_{\text{proto}}^{kj}] \quad (14)$$

$$p_{\text{write}}^t \propto \exp(W_{\text{write}} h_{\text{pre}}^t) \quad (15)$$

To produce a final distribution over output tokens at time t , we combine predictions from each stream:

$$\beta_{\text{gate}} = \text{softmax}(W_{\text{gate}} h_{\text{out}}^t) \quad (16)$$

$$p(d^t = w \mid d^{<t}, d_{1:n}^t; \theta) = \beta_{\text{gate}} \cdot [p_{\text{write}}^t(w), p_{\text{copy,out}}^t(w), p_{\text{copy,proto-1}}^t(w), \dots, p_{\text{copy,proto-n}}^t(w)] \quad (17)$$

This copy mechanism is similar to the one proposed by Merity et al. (2016) and See et al. (2017). We compare 1- and 2-prototype models to an ordinary sequence model and baselines in Section 5.

4 SAMPLING SCHEMES

The models above provide generic procedures for generating well-formed combinations of training data, but do nothing to ensure that the generated samples are of a kind useful for compositional generalization. While the training objective in Eq. 7 encourages the learned $p(d)$ to lie close to the training data, an effective data augmentation procedure should intuitively provide *novel* examples of *rare* phenomena. To generate augmented training data, we combine the generative models of Section 3 with a simple sampling procedure that upweights useful examples.

4.1 RESAMPLING AUGMENTED DATA

In classification problems with imbalanced classes, a common strategy for improving accuracy on the rare class is to *resample* so that the rare class is better represented in training data (Japkowicz et al., 2000). When constructing an augmented dataset using the models described above, we apply a simple rejection sampling scheme. In Eq. 1, we set:

$$w(d) = \prod_t \mathbb{1}[p(d^t) < \epsilon]. \quad (18)$$

Here $p(d^t)$ is the marginal probability that the token d^t appears in any example and ϵ is a hyperparameter. The final model is then trained using Eq. 1, retaining those augmented samples for which $w(d) = 1$. For extremely imbalanced problems, like the ones considered in Section 5, this weighting scheme effectively functions as a **rare tag constraint**: only examples containing rare words or tags are used to augment the original training data.

4.2 NEIGHBORHOODS, PROTOTYPE PRIORS

How can we ensure that the data augmentation procedure generates any samples with positive weight in Eq. 18? The prototype-based models described in Section 3 offer an additional means of control over the generated data. Aside from the implementation of p_{recomb} , the main factors governing the behavior of the model are the choice of neighborhood function $\mathcal{N}(d)$ and, for $n \geq 2$, the set of prior compatible prototypes Ω . Defining these so that rare tags also preferentially appear in prototypes helps ensure that the generated samples contribute to generalization. Let d_1 and d_2 be prototypes. As a notational convenience, given two sequences d_1, d_2 , let $d_1 \setminus d_2$ the *set* of tokens in d_1 but not d_2 , and $d_1 \Delta d_2$ denote the set of tokens not common to d_1 and d_2 .

1-prototype neighborhoods Guu et al. (2018) define a one-prototype \mathcal{N} based on a Jaccard distance threshold Jaccard (1901). For experiments with one-prototype models we employ a similar strategy, choosing an initial neighborhood of candidates such that

$$\mathcal{N}(d) \stackrel{\text{def}}{=} \{d_1 \in \mathcal{D} : (\alpha \cdot |d \Delta d_1| + \beta \cdot \text{lev}(d, d_1)) < \delta\} \quad (19)$$

where lev is string edit distance (Levenshtein, 1966) and α , β and δ are hyperparameters (discussed in Appendix B).

2-prototype neighborhoods The $n \geq 2$ prototype case requires a more complex neighborhood function—intuitively, for an input d , we want each (d_1, d_2, \dots) in the neighborhood to collectively contain enough information to reconstruct d . Future work might treat the neighborhood function itself as latent, allowing the model to identify groups of prototypes that make d probable; here, as in existing one-prototype models, we provide heuristic implementations for the $n = 2$ case.

Long-short recombination: For each $(d_1, d_2) \in \mathcal{N}(d)$, d_1 is chosen to be similar to d , and d_2 is chosen to be similar to the *difference* between d and d_1 . (The neighborhood is so named because one of the prototypes will generally have fewer tokens than the other one.)

$$\mathcal{N}(d) \stackrel{\text{def}}{=} \{(d_1, d_2) \in \Omega : \text{lev}(d, d_1) < \delta, \text{lev}([d \setminus d_1], d_2) < \delta, |d \setminus d_1| > 0, |d \setminus d_1 \setminus d_2| = 0\} \quad (20)$$

Here $[d \setminus d_1]$ is the sequence obtained by removing all tokens in d_1 from d . Recall that we have defined $p_\Omega(d_{1:n}) \stackrel{\text{def}}{=} \text{Unif}(\Omega)$ for a set Ω of “compatible” prototypes. For experiments using long-short combination, all prototypes are treated as compatible; that is, $\Omega = \mathcal{D} \times \mathcal{D}$.

Long-long recombination: $\mathcal{N}(d)$ contains pairs of prototypes that are individually similar to d and collectively contain all the tokens needed to reconstruct d :

$$\mathcal{N}(d) \stackrel{\text{def}}{=} \{(d_1, d_2) \in \Omega : \text{lev}(d, d_1) < \delta, \text{lev}(d, d_2) < \delta, |d \Delta d_1| = 1, |d \setminus d_1 \setminus d_2| = 0\} \quad (21)$$

For experiments using long-long recombination, we take $\Omega = \{(d_1, d_2) \in \mathcal{D} \times \mathcal{D} : |d_1 \Delta d_2| = 1\}$.

5 DATASETS & EXPERIMENTS

We evaluate our approach on two tests of compositional generalization: the SCAN instruction following task (Lake & Baroni, 2018) and a few-shot morphology learning task derived from the SIGMORPHON 2018 dataset (Cotterell et al., 2018). Our experiments are designed to explore the effectiveness of learned data recombination procedures in controlled and natural settings. Both tasks involve conditional sequence prediction: while preceding sections have discussed augmentation procedures that produce data points $d = (x, y)$, learners are evaluated on their ability to predict an output y from an input x : actions y given instructions x , or morphological analyses y given words x .

For each task, we compare a baseline with no data augmentation, the rule-based GECA data augmentation procedure, and a sequence of learned data augmentation models that measure the importance of resampling and recombination. The basic **Learned Aug** model trains an RNN to generate (x, y) pairs, then trains a conditional model on the original data and samples from the generative model. **Resampling** filters these samples as described in Section 4. **Recomb-n** models replace the RNN with a prototype-based model as described in Section 3. Additional experiments (Table 1b) compare data augmentation to prediction of y via **direct inference** (Appendix E) in the prototype-based model and several other model variants.

5.1 SCAN

SCAN (Lake & Baroni, 2018) is a synthetic dataset featuring simple English commands paired with sequences of actions. Our experiments aim to show that learned recombination models perform well at one-shot concept learning and zero-shot generalization on controlled tasks where rule-based models succeed. We experiment with two splits of the dataset, *jump* and *around right*. In the *jump* split, which tests one-shot learning, the word *jump* appears in a single command in the training set but in more complex commands in the test set (e.g. *look and jump twice*). The *around right* split (Loula et al., 2018) tests zero-shot generalization by presenting learners with constructions like *walk around left* and *walk right* in the training set, but *walk around right* only in the test set.

Despite the apparent simplicity of the task, ordinary neural sequence-to-sequence models completely fail to make correct predictions on SCAN test set (Table 1). As such it has been a major focus of research on compositional generalization in sequence-to-sequence models, and a number of heuristic procedures and specialized model architectures and training procedures have been developed to solve it (Russin et al., 2019; Gordon et al., 2020; Lake, 2019; Andreas, 2020). Here we show that the

Table 1: Accuracy results on the SCAN dataset. (a) Comparison of augmentation with n -prototype models to previous work. Connecting lines indicate that model components are inherited from the parent (e.g. the row labeled *recomb-2* also includes resampling). Data augmentation with *recomb-2* + resampling achieves comparable performance to GECA on the *jump* split and nontrivial performance on the *around right* split; data augmentation with a *recomb-1* + resampling model or ordinary RNN do not generalize robustly to either split. All differences except between GECA and *recomb-2* + resampling in *jump* are significant (paired t -test, $p \ll 0.001$). (Dashes indicate that all samples were rejected by the resampling procedure when decoding with temperature $T = 1$.) (b) Ablation experiments on the *jump* split. Introducing the latent variable used in previous work (Guu et al., 2018) does not change performance, while removing the copy mechanism results in a complete failure of generalization. While it is possible to perform conditional inference of $p(y | x)$ given the generative model in Eq. 3 (*direct inference*), this gives significantly worse results than data augmentation (see Section 5.3).

	(a)		(b)	
	<i>around right</i>	<i>jump</i>		<i>jump</i>
baseline	0.00 \pm 0.00	0.00 \pm 0.00	<i>recomb-2</i> (no resampling)	0.88 \pm 0.08
GECA	0.82 \pm 0.11	0.87 \pm 0.05	+ VAE	0.88 \pm 0.07
└ resampling	0.73 \pm 0.12	0.83 \pm 0.07	+ resampling	0.86 \pm 0.09
learned aug (basic)	0.00 \pm 0.00	0.00 \pm 0.00	- copying	0.00 \pm 0.00
└ resampling	-	-	direct inference	0.57 \pm 0.05
└└ <i>recomb-1</i>	0.17 \pm 0.07	-		
└└ <i>recomb-2</i>	0.65 \pm 0.11	0.86 \pm 0.09		

generic prototype recombination procedure described above does so as well. We use long–short recombination for the *jump* split and long–long recombination for the *around right split*. We use a recombination network to generate 400 samples $d = (x, y)$ and then train an ordinary LSTM with attention (Bahdanau et al., 2018) on the original and augmented data to predict y from x . Training hyperparameters are provided in Appendix D.

Table 1 shows the results of training these models on the SCAN dataset. 2-prototype recombination is essential for successful generalization on both splits. Additional ablations (Table 1b) show that the continuous latent variable used by Guu et al. (2018) does not affect performance, but that the copy mechanism described in Section 3.2 and the use of the *recomb-2* model for data augmentation rather than direct inference are necessary for accurate prediction.

5.2 SIGMORPHON 2018

The SIGMORPHON 2018 dataset consists of words paired with morphological analyses (*lemmas*, or base forms, and tags for linguistic features like tense and case, as depicted in Fig. 1). We use the data to construct a morphological *analysis* task (Akyürek et al., 2019) (predicting analyses from surface forms) to test models’ few-shot learning of new morphological paradigms. In three languages of varying morphological complexity (Spanish, Swahili, and Turkish) we construct splits of the data featuring a training set of 1000 examples and three test sets of 100 examples. One test set consists exclusively of words in the past tense, one in the future tense and one with other word forms (present tense verbs, nouns and adjectives). The training set contains exactly eight past-tense and eight future-tense examples; all the rest are other word forms. Our experiments evaluate our approach’s ability to efficiently learn noisy morphological rules, long viewed a key challenge for connectionist approaches to language learning (Rumelhart & McClelland, 1986). As approaches may be sensitive to the choice of the eight examples from which the model must generalize, we construct five different splits per language and use the Spanish past-tense data as a development set. As above, we use *long–long* recombination with similarity criteria applied to y only. We augment the training data with 180 samples from p_{recomb} and again train an ordinary LSTM with attention for final predictions. Details are provided in Appendix B.

Table 2 shows aggregate results across languages. We report the model’s F_1 score for predicting morphological analyses of words in the *few-shot* training condition (past and future) and the standard training condition (other word forms). Here, learned data augmentation with both one- and two-prototype models consistently matches or outperforms GECA. The improvement is sometimes dramatic: for few-shot prediction in Swahili, *recomb-1* augmentation reduces the error rate by 40% relative to the baseline and 21% relative to GECA. An additional **baseline + resampling** experiment upweights the existing rare samples rather than synthesizing new ones; results demonstrate that recombination, and not simply reweighting, is important for generalization. Table 2 also includes a

Table 2: F_1 score for morphological analysis on rare (FUT+PST) and frequent (OTHER) word forms. Learned 1- and 2-prototype recombination with resampling consistently match or outperform both a no-augmentation baseline and GECA; *recomb-1* + resampling is best overall. Bold numbers are not significantly different from the best result in each column under a paired t-test ($p < 0.05$ after Bonferroni correction; nothing is bold if all differences are insignificant). The NOVEL portion of the table shows model accuracy on examples whose exact tag set never appeared in the training data. (There were no such words in the test set for the Spanish OTHER.) Differences between GECA and the best learned data augmentation procedure (*recomb-1* + resampling) are larger than in the full evaluation set. *The Spanish past tense was used as a development set.

		Spanish		Swahili		Turkish	
		FUT+PST*	OTHER	FUT+PST	OTHER	FUT+PST	OTHER
ALL	baseline	0.66 ± 0.01	0.88 ± 0.01	0.75 ± 0.02	0.90 ± 0.01	0.69 ± 0.04	0.85 ± 0.03
	└ resampling	0.65 ± 0.01	0.88 ± 0.01	0.77 ± 0.01	0.90 ± 0.02	0.69 ± 0.04	0.84 ± 0.04
	GECA	0.66 ± 0.01	0.88 ± 0.01	0.76 ± 0.02	0.90 ± 0.02	0.69 ± 0.02	0.87 ± 0.01
	└ resampling	0.72 ± 0.02	0.88 ± 0.01	0.81 ± 0.02	0.89 ± 0.01	0.75 ± 0.03	0.85 ± 0.02
	learned aug (basic)	0.66 ± 0.02	0.88 ± 0.01	0.77 ± 0.02	0.90 ± 0.01	0.70 ± 0.02	0.87 ± 0.01
	└ resampling	0.70 ± 0.02	0.86 ± 0.01	0.84 ± 0.01	0.90 ± 0.01	0.73 ± 0.04	0.85 ± 0.03
	└┐ <i>recomb-1</i>	0.72 ± 0.02	0.87 ± 0.01	0.85 ± 0.01	0.90 ± 0.02	0.77 ± 0.02	0.87 ± 0.02
└┐ <i>recomb-2</i>	0.71 ± 0.01	0.87 ± 0.02	0.82 ± 0.02	0.90 ± 0.01	0.75 ± 0.03	0.86 ± 0.03	
NOVEL	baseline	0.63 ± 0.03	-	0.72 ± 0.02	0.42 ± 0.12	0.68 ± 0.04	0.66 ± 0.15
	GECA + resampling	0.67 ± 0.03	-	0.79 ± 0.02	0.26 ± 0.20	0.73 ± 0.04	0.71 ± 0.10
	<i>recomb-1</i> + resampling	0.69 ± 0.02	-	0.83 ± 0.02	0.42 ± 0.12	0.75 ± 0.03	0.82 ± 0.04

finer-grained analysis of *novel word forms*: words in the evaluation set whose exact morphological analysis never appeared in the training set. Resampling and learned recombination again significantly outperform both the baseline and GECA-based data augmentation in the few-shot FUT+PAST condition and the ordinary OTHER condition, underscoring the effectiveness of this approach for “in-distribution” compositional generalization.

5.3 ANALYSIS

Why is learned data augmentation effective? Samples from the best learned data augmentation models for SCAN and SIGMORPHON may be found in the Appendix G.3 . We programatically analyzed 400 samples from *recomb-2* models in SCAN and found that 40% of novel samples are exactly correct in the *around right* split and 74% in the *jump* split. A manual analysis of 50 Turkish samples indicated that only 14% of the novel samples were exactly correct. The augmentation procedure has a high error rate! However, our analysis found that malformed samples either (1) feature malformed *xs* that will never appear in a test set (a phenomenon also observed in GECA), or (2) are mostly correct at the token level (high F_1 score). Data augmentation thus contributes a mixture of *label noise*—which may exert a positive regularizing effect (Bishop, 1995)—and well-formed examples, a small number of which are sufficient to induce generalization (Bastings et al., 2018). Without resampling, SIGMORPHON models generate almost no examples of rare tags.

Why does data augmentation outperform direct inference? A partial explanation is provided by the preceding analysis, which notes that the accuracy of the data augmentation procedure as a generative model is comparatively low. Additionally, the data augmentation procedure selects only the highest-confidence samples from the model, so the quality of predicted *ys* conditioned on random *xs* will in general be even lower. A conditional model trained on augmented data is able to compensate for errors in augmentation or direct inference (Table 12).

Why is resampling without recombination effective? One surprising feature of Table 2 is performance of the *learned aug (basic) + resampling* model. While less effective than the recombination-based models, augmentation with samples from an ordinary RNN trained on (x, y) pairs improves performance for some test splits. One possible explanation is that resampling effectively acts as a *posterior constraint* on the final model’s predictive distribution, guiding it toward solutions in which rare tags are more probable than observed in the original training data. Future work might model this constraint explicitly, e.g. via posterior regularization (as in Li & Rush, 2020).

6 CONCLUSIONS

We have described a method for improving compositional generalization in sequence-to-sequence models via data augmentation with learned prototype recombination models. These are the first results we are aware of demonstrating that generative models of data are effective as data augmentation schemes in sequence-to-sequence learning problems, even when the generative models are themselves unreliable as base predictors. Our experiments demonstrate that it is possible to achieve compositional generalization on-par with complex symbolic models in clean, highly structured domains, and outperform them in natural ones, using basic neural modeling tools and no symbolic scaffolding.

REFERENCES

- Ekin Akyürek, Erenay Dayanık, and Deniz Yuret. Morphological analysis using a sequence decoder. *Transactions of the Association for Computational Linguistics*, 7:567–579, 2019.
- Jacob Andreas. Good-enough compositional data augmentation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 2020.
- Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. Neural module networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 39–48, 2016.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1409.0473>.
- Dzmitry Bahdanau, Shikhar Murty, Michael Noukhovitch, Thien Huu Nguyen, Harm de Vries, and Aaron Courville. Systematic generalization: what is required and can it be learned? *arXiv preprint arXiv:1811.12889*, 2018.
- Dzmitry Bahdanau, Harm de Vries, Timothy J O’Donnell, Shikhar Murty, Philippe Beaudoin, Yoshua Bengio, and Aaron Courville. Closure: Assessing systematic generalization of clevr models. *arXiv preprint arXiv:1912.05783*, 2019.
- Jasmijn Bastings, Marco Baroni, Jason Weston, Kyunghyun Cho, and Douwe Kiela. Jump to better conclusions: SCAN both left and right. In *Proceedings of the EMNLP BlackboxNLP Workshop*, pp. 47–55, Brussels, Belgium, 2018.
- Jean Berko. The child’s learning of english morphology. *Word*, 14(2-3):150–177, 1958.
- Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017. URL <https://doi.org/10.1137/141000671>.
- Chris M Bishop. Training with noise is equivalent to tikhonov regularization. *Neural computation*, 7(1):108–116, 1995.
- Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew M Dai, Rafal Jozefowicz, and Samy Bengio. Generating sentences from a continuous space. *arXiv preprint arXiv:1511.06349*, 2015.
- Susan Carey and Elsa Bartlett. Acquiring a single new word. 1978.
- Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.
- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Arya D. McCarthy, Katharina Kann, Sabrina J. Mielke, Garrett Nicolai, Miikka Silfverberg, David Yarowsky, Jason Eisner, and Mans Hulden. The CoNLL–SIGMORPHON 2018 shared task: Universal morphological reinflection. In *Proceedings of the CoNLL–SIGMORPHON 2018 Shared Task: Universal Morphological Reinflection*, pp. 1–27, Brussels, October 2018. Association for Computational Linguistics. doi: 10.18653/v1/K18-3001. URL <https://www.aclweb.org/anthology/K18-3001>.

- Jerry A Fodor, Zenon W Pylyshyn, et al. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1-2):3–71, 1988.
- Jonathan Gordon, David Lopez-Paz, Marco Baroni, and Diane Bouchacourt. Permutation equivariant models for compositional generalization in language. In *International Conference on Learning Representations*, 2020.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O.K. Li. Incorporating copying mechanism in sequence-to-sequence learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1631–1640, Berlin, Germany, August 2016. Association for Computational Linguistics. doi: 10.18653/v1/P16-1154. URL <https://www.aclweb.org/anthology/P16-1154>.
- Jiatao Gu, Yong Wang, Kyunghyun Cho, and Victor OK Li. Search engine guided neural machine translation. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Kelvin Guu, Tatsunori B. Hashimoto, Yonatan Oren, and Percy Liang. Generating sentences by editing prototypes. *Transactions of the Association for Computational Linguistics*, 6:437–450, 2018. doi: 10.1162/tacl_a_00030. URL <https://www.aclweb.org/anthology/Q18-1031>.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. Realm: Retrieval-augmented language model pre-training. *arXiv preprint arXiv:2002.08909*, 2020.
- Paul Jaccard. Comparative study of the distribution of flora in a portion of the Alps and Jura. *Bull Soc Vaudoise Sci Nat*, 37:547–579, 1901.
- Nathalie Japkowicz et al. Learning from imbalanced data sets: a comparison of various strategies. In *AAAI workshop on learning from imbalanced data sets*, volume 68, pp. 10–15. Menlo Park, CA, 2000.
- Robin Jia and Percy Liang. Data recombination for neural semantic parsing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 12–22, Berlin, Germany, August 2016. Association for Computational Linguistics. doi: 10.18653/v1/P16-1002. URL <https://www.aclweb.org/anthology/P16-1002>.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2901–2910, 2017.
- Daniel Keysers, Nathanael Schärli, Nathan Scales, Hylke Buisman, Daniel Furrer, Sergii Kashubin, Nikola Momchev, Danila Sinopalnikov, Lukasz Stafiniak, Tibor Tihon, et al. Measuring compositional generalization: A comprehensive method on realistic data. *arXiv preprint arXiv:1912.09713*, 2019.
- Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. Generalization through memorization: Nearest neighbor language models. *arXiv preprint arXiv:1911.00172*, 2019.
- Brenden Lake and Marco Baroni. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2873–2882, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR. URL <http://proceedings.mlr.press/v80/lake18a.html>.
- Brenden M Lake. Compositional generalization through meta sequence-to-sequence learning. In *Advances in Neural Information Processing Systems*, pp. 9788–9798, 2019.
- Brenden M Lake, Tal Linzen, and Marco Baroni. Human few-shot learning of compositional instructions. *arXiv preprint arXiv:1901.04587*, 2019.
- Vladimir I Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, volume 10, pp. 707–710, 1966.

- Xiang Lisa Li and Alexander M Rush. Posterior control of blackbox generation. *arXiv preprint arXiv:2005.04560*, 2020.
- Joao Loula, Marco Baroni, and Brenden M Lake. Rearranging the familiar: Testing compositional generalization in recurrent networks. *arXiv preprint arXiv:1807.07545*, 2018.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*, 2016.
- Alexander Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. Key-value memory networks for directly reading documents. *arXiv preprint arXiv:1606.03126*, 2016.
- Steven Piantadosi and Richard Aslin. Compositional reasoning in early childhood. *PloS one*, 11(9), 2016.
- Alexander J Ratner, Henry Ehrenberg, Zeshan Hussain, Jared Dunnmon, and Christopher Ré. Learning to compose domain-specific transformations for data augmentation. In *Advances in neural information processing systems*, pp. 3236–3246, 2017.
- Eric Sven Ristad and Peter N Yianilos. Learning string-edit distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(5):522–532, 1998.
- Murray Rosenblatt. Remarks on some nonparametric estimates of a density function. *The Annals of Mathematical Statistics*, pp. 832–837, 1956.
- David E Rumelhart and James L McClelland. On learning the past tenses of english verbs. 1986.
- Jake Russin, Jason Jo, Randall C O’Reilly, and Yoshua Bengio. Compositional generalization in a deep seq2seq model by separating syntax and semantics. *arXiv preprint arXiv:1904.09708*, 2019.
- Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *International conference on machine learning*, pp. 1842–1850, 2016.
- Abigail See, Peter J Liu, and Christopher D Manning. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*, 2017.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pp. 3104–3112, 2014.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2017.
- Jane X Wang, Zeb Kurth-Nelson, Dhruva Tirumala, Hubert Soyer, Joel Z Leibo, Remi Munos, Charles Blundell, Dharshan Kumaran, and Matt Botvinick. Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*, 2016.
- Jason Weston, Emily Dinan, and Alexander H Miller. Retrieve and refine: Improved sequence generation models for dialogue. *arXiv preprint arXiv:1808.04776*, 2018.
- Deniz Yuret. Knet: beginning deep learning with 100 lines of julia. In *Machine Learning Systems Workshop at NIPS*, volume 2016, pp. 5, 2016.

A DETAILS OF MODEL ARCHITECTURE

A.1 PROTOTYPE ENCODER

We use a single layer BiLSTM network to encode h_{proto}^{kj} as follows:

$$h_{\text{proto}}^k = \text{proj}(\text{BiLSTM}(W_e d_k^l)) \quad (22)$$

Morphology The hidden and embedding sizes are 1024. No dropout is applied. We project bi-directional embeddings to the hidden size with a linear projection. We concatenate the backward and the forward hidden states.

SCAN We choose the hidden size as 512, and embedding size as 64. We apply 0.5 dropout to the input. We project hidden vectors in the attention mechanism.

A.2 DECODER

The decoder is implemented by a single layer. In addition to the hidden state and memory cell, we also carry out a *feed* vector through time:

$$h_{\text{pre}}^t = [h_{\text{out}}^t, \sum_i \alpha_{\text{out}}^i h_{\text{out}}^i, \sum_{j \leq |d_1|} \alpha_{\text{proto}}^{1j} h_{\text{proto}}^{1j}, \dots, \sum_{j \leq |d_n|} \alpha_{\text{proto}}^{nj} h_{\text{proto}}^{nj}] \quad (23)$$

$$\text{feed}^t = \text{Linear}_{\text{feed}}(h_{\text{pre}}^t) \quad (24)$$

The input to the LSTM decoder at time step t is the concatenation of the previous token’s representation, previous feed vector, and a latent z vector (if VAE model).

$$\text{input}^t = [W_d d^{t-1}, \text{feed}^{t-1}, z] \quad (25)$$

Morphology We use an LSTM network with a hidden size of 1024, and an embedding size of 1024. *feed* is the same size as the hidden state:

No dropout is applied in the decoder. Output calculations are provided in the original paper in equation Eq. 16.

The query vector for the attentions is just the hidden state. Further details of the attention are provided in Appendix A.3.

$$\text{query}^t = h_{\text{out}}^t \quad (26)$$

We initialize the decoder hidden states with the final hidden states of the BiLSTM encoder.

SCAN The decoder is implemented by a single layer LSTM network with hidden size of 512, and embedding size of 64. The embedding parameters are shared with the encoder. Here the size of the feed vector is equal to embedding size, 64.

We have no self-attention for this decoder in the feed vector. There is an attention projection with dimension is 128. The details of the attention mechanism given in Appendix A.3. Finally, we use transpose of the embedding matrix to project *feed* to the output space.

$$\text{output}^t = W_e^\top \text{feed}^t \quad (27)$$

output ^{t} contains unnormalized scores before the final softmax layer. We apply 0.7 dropout to h_{out}^t during both training and test. The copy mechanism will be further described in Appendix A.3.

The input to the LSTM decoder is the same as Eq. 25 except the decoder embedding matrix, W_d , shares parameters with encoder embedding matrix W_e . We applied 0.5 dropout to the embeddings d_{t-1} .

The *query* vector for the attention is calculated by:

$$\text{query}^t = [h_{\text{out}}^t, \text{input}^t] \quad (28)$$

A.3 ATTENTION AND COPYING

We use the attention mechanism described in Vaswani et al. (2017) with slight modifications.

Morphology We use linear transformation for *key* while retaining the size of 1024, and leave *query* and *value* transformations as the identity normalized by the square root of the attention dimension. The query vector is described in the decoder Appendix A.2. The copy mechanism for the morphology task is explained in the paper in detail.

SCAN We use the nonlinear tanh transformation for *key*, *query* and *value*. That the attention scores are calculated separately for each prototype using different parameters as well as the normalization i.e. obtaining α 's is performed separately for each prototype.

The copy mechanism for this task is slightly different and follows Gu et al. (2016) where we normalize prototype attention scores and output scores together. Let $\bar{\alpha}_i$ represent attention weights for each prototype sequence *before* normalization. Then, we concatenate them to the output vector in Eq. 27.

$$\text{final}^t = [\text{output}^t, \bar{\alpha}_1, \dots, \bar{\alpha}_n]$$

Then we obtain a probability vector via a final softmax layer:

$$\text{prob}^t = \text{softmax}(\text{final}^t)$$

That size of this probability vector is vocabulary size plus the total length all prototypes. We then project this into the output space by:

$$p^t(w) = \text{prob}^t(\text{indices}(w))$$

where *indices* finds all corresponding scores in prob^t for token w where there might be more than one element for a given w . This is because one score can come from the *output* ^{t} region, and others from the prototype regions of prob^t . During training we applied 0.5 dropout to the indices from *output* ^{t} . Thus, the model is encouraged to copy more.

B DETAILS OF NEIGHBORHOODS AND SAMPLING

In the Eq. 20 and Eq. 21 we expressed the generic form of neighborhood sets. Here we provide the implementation details.

SCAN In the *jump* split, we use long-short recombination with $\delta = 0.5$. In *around right* we use long-long recombination with $\delta = 0.5$, and Ω requires first and second prototypes to differ by only single token. We randomly pick $k < 10 * 3$ (10 different first prototypes, and 3 different second prototypes for each of them) prototype pairs that satisfy these conditions. For the *recomb-1* experiment, we use the same neighborhood setup except we only consider the $k < 10$ first prototypes.

Sampling In the *jump* split, we used beam search with beam size 4 in the decoder. We calculate the mean and standard deviation over the lengths of both among the first d'_1 and the second d'_2 prototypes in the train set. Then, during the sampling, we expect the first and second prototypes whose length is shorter than their respective mean plus standard deviation. This decision is based on the fact that the part of the Ω that the model is exposed to is determined by the empirical distribution, $\hat{\Omega}$, that arises from training neighborhoods. So, in the sampling time we try to pick prototypes from a distribution that are close to properties of that empirical distribution. In *around right*, we use temperature sampling with $T = 0.4$. If a model cannot sample the expected number of both novel and unique samples within a reasonable time, we increase temperature T .

Morphology We use long-long recombination, as explained in the paper, with slight modifications which leverage the structure of the task. For the *recomb-1* model $\mathcal{N}(d)$ utilizes tag similarity, lemma similarity and is constructed using a score function:

$$\text{score}_1(d, d'_1) = (|d_{\text{tags}} \Delta d'_{1\text{tags}}|, \text{jaccard}(d_{\text{lemma}}, d'_{1\text{lemma}})) \quad (29)$$

Given d , we sort training examples by using score_1 as the comparison key and pick the first four neighbors to form $\mathcal{N}(d)$. Note that the sort function, first sorts by the first element in the sort tuple then the second.

For the *recomb-2* model $\mathcal{N}(d)$ uses the same score function for the first prototype as in the *recomb-1* case. Regarding the second prototype, it uses another score function:

$$\text{score}_2(d, d'_1, d'_2) = (d_{\text{tags}} \neq d'_{2\text{tags}}, \quad (30)$$

$$d'_{1\text{tags}} \neq d'_{2\text{tags}}, \quad (31)$$

$$(d_{\text{tags}} \setminus d'_{1\text{tags}} \setminus d'_{2\text{tags}}) = 0, \quad (32)$$

$$|d'_{1\text{tags}} \Delta d'_{2\text{tags}}| \quad (33)$$

Given x , and a scored first prototype, we do one more sort over training examples by using score_2 as the comparison key. Then we pick first four neighbors for $\mathcal{N}(d)$.

Sampling We use a mix strategy of temperature sampling with $T = 0.5$ and greedy sampling in which we use the former for d_{input} and the latter for d_{output} . We sample 180 unique and novel examples. Note that Ω resorts to conditions in Eq. 31 and Eq. 32.

C DETAILS OF TRAINING OF GENERATIVE MODEL

Morphology All of the hyper parameters mentioned here are optimized by a grid search on the Spanish validation set. We train our models for 25 epochs². We use Adam optimizer with learning rate 0.0001. The generative model is trained on morphological reinlection order ($d_{\text{lemma}}d_{\text{tags}} \triangleright d_{\text{inflection}}$) from left to right, then the samples from the model are reordered for morphological analysis task ($d_{\text{inflection}} \triangleright d_{\text{lemma}}d_{\text{tags}}$).

SCAN We use different number of epochs for *jump* and *around right* splits where all models are trained for 8 epochs in the former and 3 epochs in the latter. We use Adam optimizer with learning rate 0.002, and gradient norm clip with 1.0.

D DETAILS OF SEQ2SEQ BASELINE MODEL

After generating novel samples, we either concatenate them to the training data (in morphology), or we use an augmentation ratio (in SCAN), p_{aug} , to sample training batches from a mixture of the original training data and the augmented data. Our conditional model is indeed the same with the generative model (non-VAE version) used in morphology experiments, which we explained in detail both in the paper and in here—the only change needed is to replace d_{proto} with x , and d with y . So, the conditional model also has the copy mechanism through which it can copy from the input. We didn't observe any gain from enabling self attention here.

Every conditional model's size is the same with its corresponding generative model which was used for augmentation. This is to ensure that the conditional model and the generative model have the same capacity. We train conditional models for 150 epochs for SCAN and we used augmentation ratios of $p_{\text{aug}} = 0.01$ and $p_{\text{aug}} = 0.2$ in *jump* and *around right*, respectively. For morphology, we train the conditional models for 100 epochs, and we use all generated examples for augmentation.

E DETAILS OF DIRECT INFERENCE

Our procedure for the inference is that we define neighborhood functions where the d is replaced by the x . Then, we produced candidate ys given the x for each different prototype combinations. Then, we re-rank these candidates via Eq. 7. Finally, we pick the highest scoring candidate and take it's output as the prediction.

²When training 2-proto and 1-proto models, we increment epoch counter when the entire neighborhood for every d is processed. For 0-proto, one epoch is defined canonically i.e. the entire train set.

F DETAILS OF VAE

Prior $p(z)$: We use the same prior as Guu et al. (2018) given in Eq. 34, primarily due to two reasons: (i) being consistent with the previous work, (ii) with this choice of prior, along with a suitable choice of proposal network, the KL divergence term in ELBO reduces to a constant. In this prior, z is defined by a norm and direction vector. The norm is sampled from the uniform distribution between zero and a maximum possible norm, $\mu_{\max} = 10.0$, and the direction is sampled from the uniform distribution on the unit circle which corresponds to von-Mises Fisher distribution with the concentration parameter zero.

$$z = z_{\text{norm}} \cdot z_{\text{dir}} \quad \text{where} \quad z_{\text{norm}} \sim U(0, \mu_{\max}), \quad z_{\text{dir}} \sim \text{vmF}(\vec{u}, 0) \tag{34}$$

For SCAN, size of z is 32, and for morphology size of z is 2.

Proposal Network $q(z|d, d_{1:n})$: Similar to the prior, the posterior network decomposes z into its norm and direction vectors. The norm vector is sampled from a uniform distribution at $(|\mu|, \min(|\mu| + \epsilon, \mu_{\max}))$, and the direction is sampled from the von-Mises Fisher distribution $\text{vmF}(\mu, \kappa)$ where $\kappa = 25, \epsilon = 1.0$.

$$\begin{aligned}
 h_{\text{final}} &= (\bar{h}_{\text{proto}})_{\text{start}} + (\vec{h}_{\text{proto}})_{\text{end}} \\
 \mu &= \tanh(W_z [h_{d\text{final}}, h_{q\text{final}}]) \\
 z_{\text{norm}} &\sim U(|\mu|, \min(|\mu| + \epsilon, \mu_{\max})) \\
 z_{\text{dir}} &\sim \text{vmF}(\mu, \kappa) \\
 z &= z_{\text{norm}} \cdot z_{\text{dir}}
 \end{aligned}
 \tag{35}$$

G ADDITIONAL RESULTS

G.1 MORPHOLOGY RESULTS

In the paper, Fig. Table 2 shows morphology results for non-vae models and 8 *hints* (past- and future-tense examples in the training set). Here, we provide additional results for other configurations of the number of hints and use of VAE models.

G.1.1 HINTS=4

Table 3: Exact Match Accuracy (non-vae)

	Spanish		Swahili		Turkish	
	FUT+PST*	OTHER	FUT+PST	OTHER	FUT+PST	OTHER
baseline	0.078 ±0.029	0.63 ±0.09	0.107 ±0.034	0.532 ±0.029	0.067 ±0.020	0.57 ±0.04
geca	0.072 ±0.019	0.63 ±0.05	0.039 ±0.011	0.496 ±0.027	0.052 ±0.014	0.54 ±0.08
geca + resampling	0.16 ±0.04	0.65 ±0.05	0.27 ±0.08	0.52 ±0.04	0.12 ±0.04	0.554 ±0.029
learned aug	0.063 ±0.012	0.65 ±0.04	0.066 ±0.034	0.52 ±0.04	0.074 ±0.021	0.57 ±0.04
learned aug + resampling	0.098 ±0.021	0.65 ±0.05	0.29 ±0.06	0.480 ±0.035	0.092 ±0.029	0.54 ±0.06
recomb-1	0.063 ±0.017	0.674 ±0.021	0.061 ±0.017	0.520 ±0.028	0.055 ±0.021	0.554 ±0.030
recomb-1 + resampling	0.13 ±0.04	0.64 ±0.04	0.29 ±0.04	0.48 ±0.04	0.15 ±0.04	0.52 ±0.06
recomb-2	0.061 ±0.010	0.656 ±0.030	0.08 ±0.06	0.524 ±0.026	0.073 ±0.019	0.58 ±0.05
recomb-2 + resampling	0.108 ±0.021	0.64 ±0.05	0.18 ±0.04	0.542 ±0.035	0.067 ±0.026	0.55 ±0.06

Table 4: F1 Accuracy (non-vae)

	Spanish		Swahili		Turkish	
	FUT+PST*	OTHER	FUT+PST	OTHER	FUT+PST	OTHER
baseline	0.609 ±0.025	0.873 ±0.034	0.746 ±0.013	0.897 ±0.005	0.561 ±0.032	0.867 ±0.015
geca	0.606 ±0.019	0.871 ±0.017	0.722 ±0.018	0.884 ±0.007	0.565 ±0.034	0.856 ±0.028
geca + resampling	0.675 ±0.017	0.870 ±0.022	0.802 ±0.023	0.892 ±0.010	0.65 ±0.05	0.850 ±0.019
learned aug	0.597 ±0.021	0.871 ±0.024	0.737 ±0.010	0.897 ±0.011	0.58 ±0.04	0.853 ±0.026
learned aug + resampling	0.646 ±0.007	0.872 ±0.028	0.826 ±0.013	0.887 ±0.010	0.637 ±0.032	0.835 ±0.034
recomb-1	0.596 ±0.020	0.884 ±0.010	0.727 ±0.012	0.893 ±0.010	0.557 ±0.032	0.868 ±0.010
recomb-1 + resampling	0.663 ±0.029	0.874 ±0.014	0.812 ±0.017	0.886 ±0.011	0.67 ±0.04	0.84 ±0.04
recomb-2	0.598 ±0.019	0.874 ±0.007	0.730 ±0.023	0.894 ±0.008	0.581 ±0.035	0.865 ±0.024
recomb-2 + resampling	0.658 ±0.012	0.872 ±0.015	0.778 ±0.011	0.897 ±0.007	0.609 ±0.032	0.850 ±0.029

G.1.2 HINTS=8

Note that Table 2 have the F1 results, we only give the exact match and extra VAE model results here.

Table 5: Exact Match Accuracy (non-vae)

	Spanish		Swahili		Turkish	
	FUT+PST*	OTHER	FUT+PST	OTHER	FUT+PST	OTHER
baseline	0.151 \pm 0.017	0.65 \pm 0.04	0.15 \pm 0.04	0.554 \pm 0.034	0.23 \pm 0.06	0.55 \pm 0.04
geca	0.136 \pm 0.030	0.638 \pm 0.026	0.15 \pm 0.05	0.55 \pm 0.06	0.21 \pm 0.05	0.550 \pm 0.032
geca + resampling	0.249 \pm 0.034	0.64 \pm 0.04	0.25 \pm 0.05	0.532 \pm 0.033	0.27 \pm 0.07	0.524 \pm 0.026
learned aug	0.163 \pm 0.030	0.652 \pm 0.033	0.18 \pm 0.05	0.560 \pm 0.026	0.23 \pm 0.04	0.548 \pm 0.019
learned aug + resampling	0.181 \pm 0.026	0.590 \pm 0.032	0.34 \pm 0.06	0.552 \pm 0.029	0.24 \pm 0.04	0.53 \pm 0.05
recomb-1	0.155 \pm 0.018	0.628 \pm 0.020	0.161 \pm 0.017	0.560 \pm 0.025	0.22 \pm 0.04	0.538 \pm 0.025
recomb-1 + resampling	0.218 \pm 0.032	0.616 \pm 0.034	0.35 \pm 0.04	0.53 \pm 0.04	0.30 \pm 0.04	0.52 \pm 0.04
recomb-2	0.131 \pm 0.028	0.634 \pm 0.027	0.19 \pm 0.11	0.56 \pm 0.04	0.24 \pm 0.05	0.528 \pm 0.032
recomb-2 + resampling	0.203 \pm 0.035	0.63 \pm 0.05	0.27 \pm 0.07	0.552 \pm 0.031	0.25 \pm 0.05	0.54 \pm 0.06

Table 6: F1 Accuracy (vae)

	Spanish		Swahili		Turkish	
	FUT+PST*	OTHER	FUT+PST	OTHER	FUT+PST	OTHER
learned aug + resampling +vae	0.689 \pm 0.018	0.859 \pm 0.010	0.845 \pm 0.014	0.896 \pm 0.011	0.730 \pm 0.032	0.850 \pm 0.015
recomb-1 + resampling +vae	0.717 \pm 0.014	0.870 \pm 0.007	0.843 \pm 0.014	0.898 \pm 0.010	0.736 \pm 0.030	0.859 \pm 0.031
recomb-2 + resampling +vae	0.710 \pm 0.008	0.865 \pm 0.012	0.824 \pm 0.015	0.896 \pm 0.011	0.751 \pm 0.027	0.848 \pm 0.027

G.1.3 HINTS=16

Table 7: Exact Match Accuracy (non-vae)

	Spanish		Swahili		Turkish	
	FUT+PST*	OTHER	FUT+PST	OTHER	FUT+PST	OTHER
baseline	0.27 \pm 0.05	0.65 \pm 0.04	0.28 \pm 0.06	0.544 \pm 0.029	0.40 \pm 0.04	0.614 \pm 0.032
geca	0.26 \pm 0.06	0.65 \pm 0.06	0.26 \pm 0.05	0.530 \pm 0.028	0.37 \pm 0.05	0.570 \pm 0.035
geca + resampling	0.34 \pm 0.05	0.63 \pm 0.04	0.32 \pm 0.07	0.506 \pm 0.034	0.42 \pm 0.05	0.590 \pm 0.035
learned aug	0.25 \pm 0.04	0.65 \pm 0.04	0.32 \pm 0.06	0.538 \pm 0.028	0.39 \pm 0.04	0.58 \pm 0.05
learned aug + resampling	0.230 \pm 0.035	0.61 \pm 0.04	0.42 \pm 0.06	0.54 \pm 0.04	0.42 \pm 0.05	0.578 \pm 0.027
recomb-1	0.27 \pm 0.05	0.63 \pm 0.06	0.32 \pm 0.05	0.55 \pm 0.04	0.35 \pm 0.06	0.60 \pm 0.05
recomb-1 + resampling	0.28 \pm 0.04	0.61 \pm 0.07	0.418 \pm 0.035	0.548 \pm 0.023	0.35 \pm 0.06	0.56 \pm 0.04
recomb-2	0.22 \pm 0.06	0.62 \pm 0.07	0.28 \pm 0.04	0.56 \pm 0.04	0.40 \pm 0.06	0.596 \pm 0.024
recomb-2 + resampling	0.262 \pm 0.025	0.61 \pm 0.07	0.405 \pm 0.028	0.53 \pm 0.04	0.43 \pm 0.06	0.61 \pm 0.04

Table 8: F1 Accuracy (non-vae)

	Spanish		Swahili		Turkish	
	FUT+PST*	OTHER	FUT+PST	OTHER	FUT+PST	OTHER
baseline	0.733 \pm 0.014	0.881 \pm 0.012	0.811 \pm 0.018	0.893 \pm 0.011	0.750 \pm 0.026	0.875 \pm 0.021
geca	0.736 \pm 0.019	0.884 \pm 0.018	0.800 \pm 0.024	0.889 \pm 0.012	0.74 \pm 0.04	0.863 \pm 0.019
geca + resampling	0.782 \pm 0.024	0.867 \pm 0.012	0.830 \pm 0.021	0.885 \pm 0.013	0.794 \pm 0.032	0.865 \pm 0.018
learned aug	0.738 \pm 0.020	0.877 \pm 0.008	0.816 \pm 0.024	0.893 \pm 0.012	0.752 \pm 0.024	0.868 \pm 0.020
learned aug + resampling	0.745 \pm 0.019	0.870 \pm 0.012	0.866 \pm 0.016	0.894 \pm 0.013	0.787 \pm 0.031	0.863 \pm 0.021
recomb-1	0.738 \pm 0.021	0.877 \pm 0.019	0.820 \pm 0.018	0.896 \pm 0.014	0.735 \pm 0.033	0.874 \pm 0.026
recomb-1 + resampling	0.770 \pm 0.020	0.867 \pm 0.023	0.872 \pm 0.005	0.892 \pm 0.010	0.778 \pm 0.024	0.861 \pm 0.022
recomb-2	0.716 \pm 0.019	0.876 \pm 0.022	0.815 \pm 0.017	0.897 \pm 0.016	0.752 \pm 0.034	0.873 \pm 0.017
recomb-2 + resampling	0.765 \pm 0.023	0.868 \pm 0.021	0.856 \pm 0.015	0.888 \pm 0.016	0.808 \pm 0.018	0.868 \pm 0.027

G.2 SIGNIFICANCE TESTS

Note that in Tables 9, 10 and 11 the p -values for pairwise differences between the baseline and prototype-based models

Table 9: Turkish language p -values for paired t -test in PST+FUT tenses for the average F_1 (micro) scores over several runs without Bonferroni correction.

	baseline	geca	learned aug	recomb-1	recomb-2	geca + resampling	learned aug + resampling	recomb-1 + resampling	recomb-2 + resampling
baseline									
geca	0.259314								
learned aug	0.352506	0.802058							
recomb-1	0.707534	0.129244	0.187597						
recomb-2	0.233578	0.0230554	0.0331794	0.363375					
geca + resampling	1.0125e-16	3.7044e-12	4.07678e-15	6.04788e-17	1.71167e-19				
learned aug + resampling	8.00807e-10	1.37553e-07	6.51548e-08	6.35314e-11	3.76501e-13	0.0167999			
recomb-1 + resampling	3.85877e-26	1.60117e-20	2.76421e-22	6.41228e-26	2.07776e-26	0.0109365	1.948e-06		
recomb-2 + resampling	2.56689e-15	3.4083e-13	2.08177e-14	2.92113e-18	1.79928e-19	0.981886	0.0190462	0.0101878	

Table 10: Spanish language p -values for paired t -test in PST+FUT tenses for the average F_1 (micro) scores over several runs without Bonferroni correction.

	baseline	geca	learned aug	recomb-1	recomb-2	geca + resampling	learned aug + resampling	recomb-1 + resampling	recomb-2 + resampling
baseline									
geca	0.394748								
learned aug	0.761129	0.635337							
recomb-1	0.428606	0.974851	0.620601						
recomb-2	0.199768	0.601998	0.317494	0.625078					
geca + resampling	2.27478e-25	6.11513e-29	3.30904e-24	1.38242e-24	2.19894e-27				
learned aug + resampling	1.09224e-10	9.34816e-13	1.40474e-11	4.70624e-13	4.78418e-14	0.000137083			
recomb-1 + resampling	4.00039e-27	8.88347e-30	4.06546e-25	1.2159e-28	7.2465e-29	0.495734	1.35727e-05		
recomb-2 + resampling	1.17709e-17	1.29429e-21	4.07864e-18	1.12332e-19	1.66638e-21	0.313143	0.00925477	0.103819	

Table 11: Swahili language p -values for paired t -test in PST+FUT tenses for the average F_1 (micro) scores over several runs without Bonferroni correction.

	baseline	geca	learned aug	recomb-1	recomb-2	geca + resampling	learned aug + resampling	recomb-1 + resampling	recomb-2 + resampling
baseline									
geca	0.606002								
learned aug	0.000857131	0.00384601							
recomb-1	6.27581e-05	0.00101351	0.769589						
recomb-2	1.75947e-05	0.000207507	0.263242	0.402064					
geca + resampling	2.58696e-21	8.85433e-19	4.57259e-11	1.33673e-11	1.87968e-08				
learned aug + resampling	7.09377e-53	1.46895e-47	2.3846e-38	6.32242e-38	2.09274e-30	3.26321e-10			
recomb-1 + resampling	1.66361e-58	1.28557e-54	9.46035e-44	2.05703e-45	9.46848e-37	1.60241e-17	0.0749463		
recomb-2 + resampling	2.16531e-31	3.52334e-25	7.80047e-20	1.08762e-19	1.26218e-15	0.0756646	1.52594e-06	2.51776e-11	

G.3 GENERATED SAMPLES

All samples are randomly selected unless otherwise indicated.

G.3.1 SCAN

In Table 12, we present three test samples from the SCAN task along with the predictions by direct inference and the conditional model trained on the augmented data with *recomb-2*. Note that the augmentation procedure was able to create novel samples whose input (x) happens to be in the test set (Examples 1 and 3) while y may or may not be correct (Example 1).

	Example 1 (<i>jump</i>)	Example 2 (<i>jump</i>)	Example 3 (<i>around right</i>)
Input ($x = \hat{x}$)	walk twice after jump twice	run right after jump twice	jump left and jump around right
True label (y)	JUMP JUMP WALK WALK	JUMP JUMP RTURN RUN	TURN LEFT JUMP TURN RIGHT JUMP TURN RIGHT JUMP TURN RIGHT JUMP
\hat{y} in augmented dataset	JUMP JUMP JUMP WALK	(not generated)	TURN LEFT JUMP TURN RIGHT JUMP TURN RIGHT JUMP TURN RIGHT JUMP
Predicted \hat{y}			
└ direct inference	JUMP JUMP JUMP WALK	LOOK LOOK RTURN JUMP	TURN LEFT JUMP TURN LEFT JUMP TURN LEFT JUMP TURN LEFT JUMP TURN LEFT JUMP
└ <i>recomb-2</i>	JUMP JUMP WALK WALK	JUMP JUMP RTURN RUN	TURN LEFT JUMP TURN RIGHT JUMP TURN RIGHT JUMP TURN RIGHT JUMP TURN RIGHT JUMP

Table 12: Comparison of generative and unconditional model predictions with and without data augmentation. The conditional model trained on augmented data (final row) is able to compensate for errors in data augmentation (Example 1) and generalize to examples never generated by the data augmentation procedure (Example 2), and often times if a sample is correctly created by the augmentation procedure, the conditional model also gets it right (Example 3).

Below are a set of samples from the learned aug (basic) model for SCAN dataset’s *jump* and *around right* splits, in order:

IN: run opposite and walk opposite right twice OUT: RUN TURN RIGHT TURN RIGHT RUN TURN RIGHT TURN RIGHT WALK

IN: look around right thrice after run around thrice thrice OUT: TURN RIGHT RUN TURN RIGHT RUN TURN RIGHT RUN TURN RIGHT RUN TURN RIGHT RUN TURN RIGHT RUN TURN RIGHT RUN TURN RIGHT RUN TURN RIGHT RUN

TURN RIGHT RUN TURN RIGHT RUN TURN RIGHT RUN TURN RIGHT LOOK TURN RIGHT LOOK TURN RIGHT LOOK TURN RIGHT LOOK TURN RIGHT LOOK

IN: look opposite right twice and walk around twice OUT: TURN RIGHT TURN RIGHT LOOK TURN RIGHT TURN RIGHT LOOK TURN RIGHT WALK TURN LEFT WALK TURN LEFT WALK TURN LEFT WALK

IN: run opposite and thrice OUT: RUN TURN LEFT RUN RUN

IN: walk opposite right thrice turn turn right thrice OUT: TURN RIGHT TURN RIGHT TURN RIGHT TURN RIGHT TURN RIGHT TURN RIGHT TURN RIGHT TURN RIGHT TURN RIGHT WALK

IN: jump opposite right twice jump look around left OUT: TURN RIGHT TURN RIGHT JUMP TURN LEFT TURN LEFT JUMP TURN LEFT TURN LEFT LOOK TURN LEFT TURN LEFT LOOK

IN: walk around left thrice after jump left left OUT: TURN LEFT WALK TURN LEFT WALK TURN LEFT WALK TURN LEFT WALK TURN LEFT WALK TURN LEFT WALK TURN LEFT WALK TURN LEFT WALK TURN LEFT WALK TURN LEFT WALK TURN LEFT WALK TURN LEFT WALK

IN: run opposite right twice walk run left thrice OUT: TURN RIGHT TURN RIGHT RUN TURN RIGHT TURN RIGHT WALK TURN LEFT TURN LEFT RUN TURN LEFT TURN LEFT RUN

Below are a set of samples from the *recomb-1* model for SCAN dataset's *around right* split. Note that there were no samples with rare tags generated by the model for the *jump* split:

IN: run around right after walk around left OUT: TURN LEFT WALK TURN LEFT WALK TURN LEFT RUN TURN LEFT RUN TURN LEFT WALK TURN LEFT WALK TURN LEFT WALK TURN LEFT WALK

IN: look around right after jump around left OUT: TURN LEFT LOOK TURN LEFT JUMP TURN LEFT LOOK TURN LEFT LOOK TURN LEFT JUMP TURN LEFT JUMP TURN LEFT JUMP TURN LEFT JUMP

IN: look around right and jump around left OUT: TURN RIGHT LOOK TURN RIGHT LOOK TURN RIGHT LOOK TURN LEFT JUMP TURN LEFT JUMP TURN LEFT JUMP TURN LEFT JUMP TURN LEFT JUMP

IN: walk around right and turn right twice OUT: TURN RIGHT WALK TURN RIGHT WALK TURN RIGHT WALK TURN RIGHT WALK TURN RIGHT TURN RIGHT

Below are 4 samples from the *recomb-2* model for each of SCAN dataset's *jump* and *around right* splits, respectively:

IN: jump opposite left thrice after jump opposite left thrice OUT: TURN LEFT TURN LEFT JUMP TURN LEFT TURN LEFT JUMP TURN LEFT TURN LEFT JUMP TURN LEFT TURN LEFT WALK TURN LEFT TURN LEFT WALK TURN LEFT TURN LEFT WALK

IN: jump left thrice and jump left thrice OUT: TURN LEFT LOOK TURN LEFT LOOK TURN LEFT LOOK TURN LEFT JUMP TURN LEFT JUMP TURN LEFT JUMP IN: jump opposite right and turn around left OUT: TURN RIGHT TURN RIGHT JUMP

TURN LEFT TURN LEFT TURN LEFT TURN LEFT

IN: turn around left and jump around left OUT: TURN LEFT TURN LEFT TURN LEFT TURN LEFT TURN LEFT JUMP TURN LEFT JUMP TURN LEFT JUMP TURN LEFT JUMP

IN: look right twice after run around right OUT: TURN RIGHT RUN TURN RIGHT RUN TURN RIGHT RUN TURN RIGHT RUN TURN RIGHT LOOK TURN RIGHT LOOK

IN: turn right twice after look around right OUT: TURN RIGHT LOOK TURN RIGHT LOOK TURN RIGHT LOOK TURN RIGHT LOOK TURN RIGHT LOOK TURN RIGHT TURN RIGHT

IN: look twice and run around right OUT: LOOK LOOK TURN RIGHT RUN TURN RIGHT RUN TURN RIGHT RUN TURN RIGHT RUN

IN: walk opposite right twice and jump around right OUT: TURN RIGHT TURN RIGHT WALK TURN RIGHT TURN RIGHT WALK TURN RIGHT JUMP TURN RIGHT JUMP TURN RIGHT JUMP TURN RIGHT JUMP

G.3.2 MORPHOLOGY

Below are a set of samples from the learned aug (basic) model in SIGMORPHON format.

şahmiçe şahmiçende N;LOC;SG;PSS2S
 karadan havaya füze karadan havaya füzel N;DAT;PL;PSS3P
 ernek erneklerine N;DAT;PL;PSS3P
 kiler kilerime N;DAT;SG;PSS1S
 mahlep mahlebimizi N;ACC;SG;PSS1P
 süzmek süzerler V;IND;3;PL;PRS;POS;DECL

âlap âlaps N;LGSPEC1;3S;SG;PRS
jôle jôleleri N;ACC;PL

envejecerse envejeciéndose V.CVB;PRS
colaxar colaxa V;IND;PRS;3;SG
pergedrer no pergedremos V;NEG;IMP;1;PL
mantear no mantees V;NEG;IMP;2;SG
flaggear no flaggeen V;NEG;IMP;3;PL
malacostar malacostarí V;COND;3;SG
desinstar desinse V;POS;IMP;3;SG
concretizar no concretices V;NEG;IMP;2;SG

Below are a set of samples from the learned aug (basic) + resampling model.

şaşırmaq şaşırılmıyor musun? V;IND;2;PL;PST;PROG;POS;INTR
ayılmak ayılmaya V;IND;1;SG;PST;DECL
pleşmek pleşmiyor muyuz? V;IND;1;PL;PST;PROG;NEG;INTR
imciyetmek imciyetmezdeğiz V;IND;1;PL;FUT;NEG;DECL
kuvaşmak kuvaşmayacağıız V;IND;1;PL;FUT;NEG;DECL
yermek yermeyeceğiz V;IND;1;PL;FUT;NEG;DECL
yarıtmak yarıtılmayacağıız V;IND;1;PL;FUT;NEG;DECL
kelimek kelimeyeceğiz V;IND;1;PL;FUT;NEG;DECL

trasescar trasescáis V;IND;PST;2;PL;IPFV
tronar tronar V;IND;FUT;1;SG
terzcalminar terzcalminan V;IND;PST;3;PL;IPFV
esubronizar esubronizamos V;IND;PST;1;PL;IPFV
urdir urdiremos V;IND;FUT;1;PL
conder conderemos V;IND;FUT;1;PL
florear florearían V;IND;PST;3;PL;LGSPEC1;SG
sabrordar sabrordamos V;IND;PST;1;PL;IPFV

Below are a set of samples from the *recomb-1* + resampling model (the best performing model in Table 2). Here we additionally annotate samples with error categories.

kovulmak kovulmaz mısınız V;IND;2;PL;FUT;NEG;INTR (Inflection and tags don't match.)
düşünmek düşündüler V;IND;3;PL;PST;POS;DECL (Correct and novel.)
sütmek sütmez miyiz? V;IND;2;SG;FUT;NEG;INTR (Inflection and tags don't match.)
bakmak bakmayacak mıyım? V;IND;1;PL;FUT;NEG;INTR (Inflection and tags don't match.)
döndürmek döndürecek misiniz? V;IND;2;PL;FUT;POS;INTR (Correct and novel.)
türkçeleştirtmek türkçeleştiriyor m V;IND;2;PL;PST;PROG;NEG;INTR (Wrong inflection, novel tag.)
çalmak çalmayız V;IND;2;PL;FUT;POS;DECL (Inflection and tags don't match.)
üsürmek üsürmezsin V;IND;2;SG;PST;NEG;DECL (Inflection and tags don't match.)

duplicar duplicaráis V;IND;FUT;2;PL (Correct and novel)
efundar efundan V;SBJV;FUT;3;PL (Inflection and tags don't match)
deshumanizar deshumanicas V;SBJV;PST;2;SG (Inplausible inflection.)
emular emulares V;SBJV;FUT;2;SG (Correct and also in train set.)
languidecer languidecíamos V;IND;PST;1;SG;IPFV (Inflection and tags don't match)
nominar nominamos V;SBJV;FUT;1;PL (Novel tags, incorrect inflection.)
finciar finciare V;SBJV;FUT;1;SG (Correct and novel.)
abastar abasto V;IND;PST;1;SG (Inflection and tags don't match)

G.4 ATTENTION HEATMAP

Here we provide a visualization copy and attention mechanism in *recomb-2* model for SCAN experiments.




Figure 3: Generation of a sample. We plot normalized output scores on the left, and attention weights to the different prototypes on the right. The prototypes are on the y axes. The model is *recomb-2* model trained on SCAN *jump* split.

H COMPUTE

We use a single 32GB NVIDIA V100 Volta GPU for each experiment. For every experiment, the whole pipeline which consists of training of the generative model, sampling and training of the conditional model takes less than an hour.