

Lexicon Learning for Few-Shot Neural Sequence Modeling

Ekin Akyürek Jacob Andreas
Massachusetts Institute of Technology
{akyurek, jda}@mit.edu

Abstract

Sequence-to-sequence transduction is the core problem in language processing applications as diverse as semantic parsing, machine translation, and instruction following. The neural network models that provide the dominant solution to these problems are brittle, especially in low-resource settings: they fail to generalize correctly or systematically from small datasets. Past work has shown that many failures of systematic generalization arise from neural models’ inability to disentangle *lexical* phenomena from *syntactic* ones. To address this, we augment neural decoders with a *lexical translation mechanism* that generalizes existing copy mechanisms to incorporate learned, decontextualized, token-level translation rules. We describe how to initialize this mechanism using a variety of lexicon learning algorithms, and show that it improves systematic generalization on a diverse set of sequence modeling tasks drawn from cognitive science, formal semantics, and machine translation.¹

1 Introduction

Humans exhibit a set of structured and remarkably consistent inductive biases when learning from language data. For example, in both natural language acquisition and toy language-learning problems like the one depicted in Fig. 1, human learners exhibit a preference for systematic and compositional interpretation rules (Guasti 2017, Chapter 4; Lake et al. 2019). These inductive biases in turn support behaviors like one-shot learning of new concepts (Carey and Bartlett, 1978). But in natural language processing, recent work has found that state-of-the-art neural models, while highly effective at in-domain prediction, fail to generalize in human-like ways when faced with rare phenomena

¹Our code is released under <https://github.com/ekinakyurek/lexical>




Figure 1: A fragment of the Colors dataset from Lake et al. (2019), a simple sequence-to-sequence translation task. The output vocabulary is only the colored circles `r`, `g`, `b`, `y`. Humans can reliably fill in the missing test labels on the basis of a small training set, but standard neural models cannot. This paper describes a neural sequence model that obtains improved generalization via a learned *lexicon* of token translation rules.

and small datasets (Lake and Baroni, 2018), posing a fundamental challenge for NLP tools in the low-data regime.

Pause for a moment to fill in the missing labels in Fig. 1. While doing so, which training examples did you pay the most attention to? How many times did you find yourself saying *means* or *maps to*? Explicit representations of lexical items and their meanings play a key role diverse models of syntax and semantics (Joshi and Schabes, 1997; Pollard and Sag, 1994; Bresnan et al., 2015). But one of the main findings in existing work on generalization in *neural* models is that they fail to cleanly separate **lexical** phenomena from **syntactic** ones (Lake and Baroni, 2018). Given a dataset like the one depicted in Fig. 1, models conflate (lexical) information about the correspondence between `zup` and `y` with the (syntactic) fact that `y` appears only in a sequence of length 1 at training time. Longer input sequences containing the word `zup` in new syntactic contexts cause models to output tokens only seen in longer sequences (Section 5).

In this paper, we describe a parameterization for sequence decoders that facilitates (but does not enforce) the learning of context-independent word meanings. Specifically, we augment decoder output layers with a **lexical translation mechanism** which generalizes neural copy mechanisms (e.g. See et al., 2017) and enables models to generate token-level translations purely attentionally. While the lexical translation mechanism is quite general, we focus here on its ability to improve few-shot learning in sequence-to-sequence models. On a suite of challenging tests of few-shot semantic parsing and instruction following, our model exhibits strong generalization, achieving the highest reported results for neural sequence models on datasets as diverse as COGS (Kim and Linzen 2020, with 24155 training examples) and Colors (Lake et al. 2019, with 14). Our approach also generalizes to real-world tests of few-shot learning, improving BLEU scores (Papineni et al., 2002) by 1.2 on a low-resource English–Chinese machine translation task (2.2 on test sentences requiring one-shot word learning).

In an additional set of experiments, we explore effective procedures for initializing the lexical translation mechanism using lexicon learning algorithms derived from information theory, statistical machine translation, and Bayesian cognitive modeling. We find that both mutual-information- and alignment- based lexicon initializers perform well across tasks. Surprisingly, however, we show that both approaches can be matched or outperformed by a *rule-based* initializer that identifies high-precision word-level token translation pairs. We then explore joint learning of the lexicon and decoder, but find (again surprisingly) that this gives only marginal improvements over a fixed initialization of the lexicon.

In summary, this work:

- Introduces a new, lexicon-based output mechanism for neural encoder–decoder models.
- Investigates and improves upon lexicon learning algorithms for initialising this mechanism.
- Uses it to solve challenging tests of generalization in instruction following, semantic parsing and machine translation.

A great deal of past work has suggested that neural models come equipped with an inductive bias that makes them fundamentally ill-suited to

human-like generalization about language data, especially in the low-data regime (e.g. Fodor et al., 1988; Marcus, 2018). Our results suggest that the situation is more complicated: by offloading the easier lexicon learning problem to simpler models, neural sequence models are actually quite effective at modeling (and generalizing about) about syntax in synthetic tests of generalization and real translation tasks.

2 Related Work

Systematic generalization in neural sequence models The desired inductive biases noted above are usually grouped together as “systematicity” but in fact involve a variety of phenomena: one-shot learning of new concepts and composition rules (Lake and Baroni, 2018), zero-shot interpretation of novel words from context cues (Gandhi and Lake, 2020), and interpretation of known concepts in novel syntactic configurations (Keysers et al., 2020; Kim and Linzen, 2020). What they share is a common expectation that learners should associate specific production or transformation rules with specific input tokens (or phrases), and generalize to use of these tokens in new contexts.

Recent years have seen tremendous amount of modeling work aimed at encouraging these generalizations in neural models, primarily by equipping them with symbolic scaffolding in the form of program synthesis engines (Nye et al., 2020), stack machines (Grefenstette et al., 2015; Liu et al., 2020), or symbolic data transformation rules (Gordon et al., 2019; Andreas, 2020). A parallel line of work has investigated the role of continuous representations in systematic generalization, proposing improved methods for pretraining (Furrer et al., 2020) and procedures for removing irrelevant contextual information from word representations (Arthur et al., 2016; Russin et al., 2019; Thrush, 2020). The latter two approaches proceed from similar intuition to ours, aiming to disentangle word meanings from syntax in encoder representations via alternative attention mechanisms and adversarial training. Our approach instead focuses on providing an explicit lexicon to the decoder; as discussed below, this appears to be considerably more effective.

Copying and lexicon learning In neural encoder–decoder models, the clearest example of benefits from special treatment of word-level production rules is the *copy mechanism*. A great deal of past work has found that neural models

Inputs	Outputs	Lexicon Entries
<i>A crocodile blessed William . William needed to walk .</i>	crocodile(x_1) AND bless.agent (x_2, x_1) AND bless.theme (x_2, William) need.agent (x_1 , William) AND need.xcomp(x_1, x_3) AND walk.agent (x_3, William)	<i>blessed</i> \mapsto bless <i>needed</i> \mapsto need <i>William</i> \mapsto William
<i>Many moons orbit around Saturn Earth is a planet .</i>	許多 衛星 繞著 土星 運行. 地球 是一個 行星.	<i>saturn</i> \mapsto 土星 <i>earth</i> \mapsto 地球 <i>moon</i> \mapsto 衛星
<i>walk around left turn right turn left jump jump opposite right after look left</i>	LTURN IWALK LTURN IWALK LTURN IWALK LTURN IWALK RTURN LTURN IJUMP LTURN ILOOK RTRUN IJUMP RTURN IJUMP	<i>walk</i> \mapsto IWALK <i>jump</i> \mapsto IJUMP <i>right</i> \mapsto RTURN <i>left</i> \mapsto LTURN <i>look</i> \mapsto ILOOK

Table 1: We present example (input,output) pairs from COGS, English-to-Chinese machine translation and SCAN datasets. We also present some of the lexicon entries which can be learned by proposed lexicon learning methods and that are helpful to make generalizations required in each of the datasets.

benefit from learning a structural copy operation that selects output tokens directly from the input sequence without requiring token identity to be carried through all neural computation in the encoder and the decoder. These mechanisms are described in detail in Section 3, and are widely used in models for language generation, summarization and semantic parsing. Our work generalizes these models to structural operations on the input that replace copying with general context-independent token-level translation.

As will be discussed, the core of our approach is a (non-contextual) lexicon that maps individual input tokens to individual output tokens. Learning lexicons like this is of interest in a number of communities in NLP and language science more broadly. A pair of representative approaches (Brown et al., 1993; Frank et al., 2007) will be discussed in detail below; other work on lexicon learning for semantics and translation includes Liang et al. (2009); Goldwater (2007); Haghghi et al. (2008) among numerous others.

Finally, and closest to the modeling contribution in this work, several previous papers have proposed alternative generalized copy mechanisms for tasks other than semantic lexicon learning. Concurrent work by Prabhu and Kann (2020) introduces a similar approach for grapheme-to-phoneme translation (with a fixed functional lexicon rather than a trainable parameter matrix), and Nguyen and Chiang (2018) and Gū et al. (2019) describe less expressive mechanisms that cannot smoothly interpolate between lexical translation and ordinary decoding at the token level. Pham et al. (2018) incorporate lexicon entries by rewriting *input* sequences prior to ordinary sequence-to-sequence translation. Akyürek et al. (2021) describe a model in which a copy mechanism is combined with a retrieval-

based generative model; like the present work, that model effectively disentangles syntactic and lexical information by using training examples as implicit representations of lexical correspondences.

We generalize and extend this previous work in a number of ways, providing a new parameterization of attentive token-level translation and a detailed study of initialization and learning. But perhaps the most important contribution of this work is the observation that many of the hard problems studied as “compositional generalization” have direct analogues in more conventional NLP problems, especially machine translation. Research on systematicity and generalization would benefit from closer attention to the ingredients of effective translation at scale.

3 Sequence-to-Sequence Models With Lexical Translation Mechanisms

This paper focuses on sequence-to-sequence **language understanding** problems like the ones depicted in Table 1, in which the goal is to map from a natural language **input** $x = [x_1, x_2, \dots, x_n]$ to a structured **output** $y = [y_1, y_2, \dots, y_m]$ —a logical form, action sequence, or translation. We assume input tokens x_i are drawn from a **input vocabulary** \mathcal{V}_x , and output tokens from a corresponding **output vocabulary** \mathcal{V}_y .

Neural encoder–decoders Our approach builds on the standard neural encoder–decoder model with attention (Bahdanau et al., 2014). In this model, an **encoder** represents the input sequence $[x_1, \dots, x_n]$ as a sequence of representations $[e_1, \dots, e_n]$

$$e = \text{encoder}(x) \quad (1)$$

Next, a **decoder** generates a distribution over output sequences y according to the sequentially:

$$\log p(y | x) = \sum_{i=1}^y \log p(y_i | y_{<i}, e, x) \quad (2)$$

Here we specifically consider decoders with **attention**.² When predicting each output token y_i , we assign each input token an **attention weight** α_i^j as in Eq. (3). Then, we construct a context representation c_i as the weighted sum of encoder representations e_i :

$$\alpha_i^j \propto \exp(h_i^\top W_{att} e_j) \quad (3)$$

$$c_i = \sum_{j=1}^{|x|} \alpha_i^j e_j \quad (4)$$

The output distribution over \mathcal{V}_y , which we denote $p_{write,i}$, is calculated by a final projection layer:

$$p(y_i = w | x) = p_{write,i}(w) \propto \exp(W_{write}[c_i, h_i]) \quad (5)$$

Copying A popular extension of the model described above is the **copy mechanism**, in which output tokens can be copied from the input sequence in addition to being generated directly by the decoder (Jia and Liang, 2016; See et al., 2017). Using the decoder hidden state h_i from above, the model first computes a **gate probability**:

$$p_{gate} = \sigma(w_{gate}^\top h_i) \quad (6)$$

and then uses this probability to interpolate between the distribution in Eq. (5) and a **copy distribution** that assigns to each word in the output vocabulary a probability proportional to that word’s weight in the attention vector over the input:

$$p_{copy}(y_i = w | x) = \sum_{j=1}^{|x|} \mathbb{1}[x_j = w] \cdot \alpha_i^j \quad (7)$$

$$p(y_i = w | x) = p_{gate} \cdot p_{write}(y_i = w | x) + (1 - p_{gate}) \cdot p_{copy}(y_i = w | x) \quad (8)$$

(note that this implies $\mathcal{V}_y \supseteq \mathcal{V}_x$).

Content-independent copying is particularly useful in tasks like summarization and machine translation where rare words (like names) are often reused between the input and output.

²All experiments in this paper use LSTM encoders and decoders, but it could be easily integrated with CNNs or transformers (Gehring et al. 2017; Vaswani et al. 2017). We only assume access to a final layer h_i , and final attention weights α_i ; their implementation does not matter.




Figure 2: An encoder-decoder model with a lexical translation mechanism applied to English-to-Chinese translation. At decoder step $t = 4$, attention is focused on the English token *Saturn*. The lexical translation mechanism is activated by p_{gate} , and the model outputs the token **土星** directly from the lexicon. **地球** means *Earth* and appears much more frequently than *Saturn* in the training set.

Our model: Lexical translation When the input and output vocabularies are significantly different, copy mechanisms cannot provide further improvements on a sequence-to-sequence model. However, even for disjoint vocabularies as in Fig. 1, there may be strict correspondences between individual words on input and output vocabularies, e.g. $zup \mapsto \mathbf{y}$ in Fig. 1. Following this intuition, the **lexical translation mechanism** we introduce in this work extends the copy mechanism by introducing an additional layer of indirection between the input sequence x and the output prediction y_i as shown in Fig. 2. Specifically, after selecting an input token $x_j \in \mathcal{V}_x$, the decoder can “translate” it to a context-independent output token $\in \mathcal{V}_y$ prior to the final prediction. We equip the model with an additional **lexicon parameter** L , a $|\mathcal{V}_x| \times |\mathcal{V}_y|$ matrix in which $\sum_w L_{vw} = 1$, and finally define

$$p_{lex}(y_i = w | x) = \sum_{j=1}^{|x|} L_{x_j w} \cdot \alpha_i^j \quad (9)$$

$$p(y_i = w | x) = p_{gate} \cdot p_{write}(y_i = w | x) + (1 - p_{gate}) \cdot p_{lex}(y_i = w | x) \quad (10)$$

The model is visualized in Fig. 2. Note that when $\mathcal{V}_x = \mathcal{V}_y$ and $L = I$ is diagonal, this is identical to the original copy mechanism. However, this approach can in general be used to produce a larger set of tokens. As shown in Table 1, coher-

ent token-level translation rules can be identified for many tasks; the lexical translation mechanism allows them to be stored explicitly, using parameters of the base sequence-to-sequence model to record general structural behavior and more complex, context-dependent translation rules.

4 Initializing the Lexicon

The lexicon parameter L in the preceding section can be viewed as an ordinary fully-connected layer inside the copy mechanism, and trained end-to-end with the rest of the network. As with other neural network parameters, however, our experiments will show that the initialization of the parameter L significantly impacts downstream model performance, and specifically benefits from initialization with a set of input–output mappings learned with an offline lexicon learning step. Indeed, while not widely used in neural sequence models (though c.f. Section 2), lexicon-based initialization was a standard feature of many complex non-neural sequence transduction models, including semantic parsers (Kwiatkowski et al., 2011) and phrase-based machine translation systems (Koehn et al., 2003).

But an important distinction between our approach and these others is the fact that we can handle outputs that are not (transparently) compositional. Not every fragment of an input will correspond to a fragment of an output: for example, *thrice* in SCAN has no corresponding output token and instead describes a structural transformation. Moreover, the lexicon is not the only way to generate: complex mappings can also be learned by p_{write} without going through the lexicon at all.

Thus, while most existing work on lexicon learning aims for complete coverage of all word meanings, the model described in Section 3 benefits from a lexicon with *high-precision* coverage of *rare phenomena* that will be hard to learn in a normal neural model. Lexicon learning is widely studied in language processing and cognitive modeling, and several approaches with very different inductive biases exist. To determine how to best initialize L , we begin by reviewing three algorithms in Section 4.1, and identify ways in which each of them fail to satisfy the high precision criterion above. In Section 4.2, we introduce a simple new lexicon learning rule that addresses this shortcoming.

4.1 Existing Approaches to Lexicon Learning

Statistical alignment In the natural language processing literature, the IBM translation models (Brown et al., 1993) have served as some of the most popular procedures for learning token-level input–output mappings. While originally developed for machine translation, they have also been used to initialize semantic lexicons for semantic parsing (Kwiatkowski et al., 2011) and grapheme-to-phoneme conversion (Rama et al., 2009). We initialize the lexicon parameter L using Model 2.

Model 2 defines a generative process in which source words y_i are generated from target words x_j via latent alignments a_i . Specifically, given a (source, target) pair with n source words and m target words, the probability that the target word i is aligned to the source word j is:

$$p(a_i = j) \propto \exp\left(-\left|\frac{i}{m} - \frac{j}{n}\right|\right) \quad (11)$$

Finally, each target word is generated by its aligned source word via a parameter θ : $p(y_i = w) = \theta(v, x_{a_i})$. Alignments a_i and lexical parameters θ can be jointly estimated using the expectation–maximization algorithm (Dempster et al., 1977).

In neural models, rather than initializing lexical parameters L directly with corresponding IBM model parameters θ , we run Model 2 in both the forward and reverse directions, then extract counts by *intersecting* these alignments and applying a softmax with temperature τ :

$$L_{vw} \propto \exp\left(\tau^{-1} \sum_{(x,y)} \sum_{i=1}^{|y|} \mathbb{1}[x_{a_i} = v] \mathbb{1}[y_i = w]\right) \quad (12)$$

For all lexicon methods discussed in this paper, if an input v is not aligned to any output w , we map it to itself if $\mathcal{V}_x \subseteq \mathcal{V}_y$. Otherwise we align it uniformly to any unmapped output words (a *mutual exclusivity bias*, Gandhi and Lake 2020).

Mutual information Another, even simpler procedure for building a lexicon is based on identifying pairs that have high *pointwise mutual information*. We estimate this quantity directly from co-occurrence statistics in the training corpus:

$$\text{pmi}(v; w) = \log \frac{\#(v, w)}{\#(v)\#(w)} + \log |D_{\text{train}}| \quad (13)$$

where $\#(w)$ is the number of times the word w appears in the training corpus and $\#(w, v)$ is the

number of times that w appears in the input and v appears in the output. Finally, we populate the parameter L via a softmax transformation: $L_{vw} \propto \exp((1/\tau) \text{pmi}(v; w))$.

Bayesian lexicon learning Last, we explore the Bayesian cognitive model of lexicon learning described by Frank et al. (2007). Like IBM model 2, this model is defined by a generative process; here, however, the lexicon itself is part of the generative model. A lexicon ℓ is an (unweighted, many-to-many) map defined by a collection of pairs (x, y) with a *description length* prior: $p(\ell) \propto e^{-|\ell|}$ (where $|\ell|$ is the number of (input, output) pairs in the lexicon). As in Model 2, given a meaning y and a natural-language description x , each x_i is generated independently. We define the probability of a word being used *non-referentially* as $p_{\text{NR}}(x_i | \ell) \propto 1$ if $x_i \notin \ell$ and κ otherwise. The probability of being used referentially is: $p_{\text{R}}(x_j | y_i, \ell) \propto \mathbb{1}_{(x_j, y_i) \in \ell}$. Finally,

$$p(x_j | y_i, \ell) = (1 - \gamma)p_{\text{NR}}(x_j | \ell) + \gamma \sum_{i=1}^{|y|} p_{\text{R}}(x_j | y_i, \ell) \quad (14)$$

To produce a final lexical translation matrix L for use in our experiments, we set $L_{vw} \propto \exp((1/\tau) p((v, w) \in \ell))$: each entry in L is the posterior probability that the given entry appears in a lexicon under the generative model above. Parameters are estimated using the Metropolis–Hastings algorithm, with details described in Appendix C.

4.2 A Simpler Lexicon Learning Rule

Example lexicons learned by the three models above are depicted in Fig. 3 for the SCAN task shown in Table 1. Lexicons learned for remaining tasks can be found in Appendix B. It can be seen that all three models produce errors: the PMI and Bayesian lexicons contain too many entries (in both cases, numbers are associated with the turn right action and prepositions are associated with the turn left action). For the IBM model, one of the alignments is confident but wrong, because the *around* preposition is associated with turn left action. In order to understand these errors, and to better characterize the difference between the demands of lexical translation model initializers and past lexicon learning schemes, we explore a simple logical procedure for extracting lexicon entries




Figure 3: Learned lexicons for the *around right* split in SCAN ($\tau = 0.1$). The rule-based lexicon learning procedure (*Simple*) produces correct alignments, while other methods fail due to the correlation between *around* and *left* in training data.

that, surprisingly, matches or outperforms all three baseline methods in most of our experiments.

What makes an effective, precise lexicon learning rule? As a first step, consider a maximally restrictive criterion (which we’ll call C_1) that extracts only pairs (v, w) for which the presence of v in the input is a *necessary and sufficient condition* for the presence of w in the output.

$$\text{ nec.}(v, w) = \forall xy. (w \in y) \rightarrow (v \in x) \quad (15)$$

$$\text{ suff.}(v, w) = \forall xy. (v \in x) \rightarrow (w \in y) \quad (16)$$

$$C_1(v, w) = \text{ nec.}(v, w) \wedge \text{ suff.}(v, w) \quad (17)$$

C_1 is too restrictive: in many language understanding problems, the mapping from *surface forms* to meanings is many-to-one (in Table 1, both *blessed* and *bless* are associated with the logical form *bless*). Such mappings cannot be learned by the algorithm described above. We can relax the necessity condition slightly, requiring *either* that v is a necessary condition for w , or is part of a group that collectively explains all occurrences of w :

$$\text{ no-winner}(w) = \nexists v'. C_1(v', w) \quad (18)$$

$$C_2(v, w) = \text{ suff.}(v, w) \wedge (\text{ nec.}(v, w) \vee \text{ no-win.}(w)) \quad (19)$$

As a final refinement, we note that C_2 is likely to capture function words that are present in most sentences, and exclude these by restricting the lexicon to words below a certain frequency threshold:

$$C_3 = C_2 \wedge |\{v' : \text{ suff.}(v', w)\}| \leq \epsilon \quad (20)$$

The lexicon matrix L is computed by taking the word co-occurrence matrix, zeroing out all entries where C_3 does not hold, then computing a softmax: $L_{vw} \propto C_3(v, w) \exp((1/\tau) \#(v, w))$. Surprisingly, as shown in Fig. 3 and evaluated below, this rule (which we label Simple) produces the most effective lexicon initializer for three of the four tasks we study. The simplicity (and extreme conservativity) of this rule highlight the different demands on L made by our model and more conventional (e.g. machine translation) approaches: the lexical translation mechanism benefits from a small number of precise mappings rather than a large number of noisy ones.

5 Experiments

We investigate the effectiveness of the lexical translation mechanism on sequence-to-sequence models for four tasks, three focused on compositional generalization and one on low-resource machine translation. In all experiments, we use an LSTM encoder-decoder with attention as the base predictor. We compare our approach (and variants) with two other baselines: GECA (Andreas 2020; a data augmentation scheme) and SynAtt (Russin et al. 2019; an alternative seq2seq model parameterization). Hyper-parameter selection details are given in the Appendix C. Unless otherwise stated, we use $\tau = 0$ and do not fine-tune L after initialization.

5.1 Colors

Task The Colors sequence translation task (see Appendix A for full dataset) was developed to measure *human* inductive biases in sequence-to-sequence learning problems. It poses an extreme test of low-resource learning for neural sequence models: it has only 14 training examples that combine four named colors and three composition operations that perform concatenation, repetition and wrapping. Liu et al. (2020) solve this dataset with a symbolic stack machine; to the best of our knowledge, our approach is the first “pure” neural sequence model to obtain non-trivial accuracy.

Results Both the Simple and IBMM2 initializers produce a lexicon that maps only color words to colors. Both, combined with the lexical translation mechanism, obtain an average test accuracy of 79% across 16 runs, nearly matching the human accuracy of 81% reported by Lake et al. (2019). The

two test examples most frequently predicted incorrectly require generalization to longer sequences than seen during training. More details (including example-level model and human accuracies) are presented in the appendix Appendix A). These results show that LSTMs are quite effective at learning systematic sequence transformation rules from ≈ 3 examples per function word when equipped with lexical translations. Generalization to longer sequences remains as an important challenge for future work.

5.2 SCAN

Task SCAN (Lake and Baroni, 2018) is a larger collection of tests of systematic generalization that pair synthetic English commands (e.g. *turn left twice and jump*) to action sequences (e.g. LTURN LTURN IJUMP) as shown in Table 1. Following previous work, we focus on the *jump* and *around right* splits, each of which features roughly 15,000 training examples, and evaluate models’ ability to perform 1-shot learning of new primitives (*jump*) and zero-shot interpretation of composition rules (*around right*). While these tasks are now solved by a number of specialized approaches, they remain a challenge for conventional neural sequence models, and an important benchmark for new models.

Results In the *jump* split, all initializers improve significantly over the base LSTM when combined with lexical translation. Most methods achieve 99% accuracy at least once across seeds. These results are slightly behind GECA (in which all runs succeed) but ahead of SynAtt.³ Again, they show that lexicon learning is effective for systematic generalization, and that simple initializers (PMI and Simple) outperform complex ones.

5.3 COGS

Task COGS (Compositional Generalization for Semantic Parsing; Kim and Linzen 2020) is an automatically generated English-language semantic parsing dataset that tests systematic generalization in learning language-to-logical-form mappings. It includes 24155 training examples. Compared to the Colors and SCAN datasets, it has a larger vocabulary (876 tokens) and finer-grained inventory of syntactic generalization tests (Table 3).

Results Notably, because some tokens appear in both inputs and logical forms in the COGS

³SynAtt results here are lower than reported in the original paper, which discarded runs with a test accuracy of 0%.

	Colors	jump (SCAN)	around right (SCAN)	COGS
LSTM	0.00 \pm 0.00	0.00 \pm 0.00	0.09 \pm 0.05	0.51 \pm 0.05
GECA	0.41 \pm 0.11	1.00 \pm 0.00	0.98 \pm 0.02	0.48 \pm 0.05
SyntAtt	0.57 \pm 0.26	0.57 \pm 0.38	0.28 \pm 0.26	0.15 \pm 0.14
LSTM + copy	-	-	-	0.66 \pm 0.03
LSTM + Lex.: Simple	0.79 \pm 0.02	0.92 \pm 0.17	0.95 \pm 0.01	0.82 \pm 0.01
LSTM + Lex.: PMI	0.41 \pm 0.19	0.95 \pm 0.08	0.02 \pm 0.04	0.82 \pm 0.00
LSTM + Lex.: IBMM2	0.79 \pm 0.02	0.79 \pm 0.27	0.00 \pm 0.00	0.82 \pm 0.00
LSTM + Lex.: Bayesian	0.51 \pm 0.21	0.82 \pm 0.21	0.02 \pm 0.04	0.70 \pm 0.04

Table 2: Exact match accuracy results for baselines and lexicon learning models on 4 different compositional generalization splits. Errors are standard deviation among 16 different seeds for Colors, 10 seeds for COGS and SCAN. Unbolded numbers are significantly ($p < 0.01$) worse than the best result in the column. Models with lexical translation mechanisms and Simple initialization consistently improve over ordinary LSTMs.

Categories	LSTM	+ copy	+ simple
primitive \rightarrow {subj, obj, inf}			
active \rightarrow passive			
obj PP \rightarrow subj PP			
passive \rightarrow active			
recursion			
unacc \rightarrow transitive			
obj \rightarrow subj proper			
subj \rightarrow obj common			
PP dative \leftrightarrow obj dative			
all			

Table 3: COGS accuracy breakdown according to syntactic generalization types for word usages. The label $a \rightarrow b$ indicates that syntactic context a appears in the training set and b in the test set.

task, even a standard sequence-to-sequence model with copying significantly outperforms the baseline models in the original work of Kim and Linzen (2020), solving most tests of generalization over syntactic roles for nouns (but performing worse at generalizations over verbs, including passive and dative alternations). As above, the lexical translation mechanism (with any of the proposed initializers) provides further improvements, mostly for verbs that baselines model incorrectly (Table 3).

5.4 Machine Translation

Task To demonstrate that this approach is useful beyond synthetic tests of generalization, we evaluate it on a low-resource English–Chinese translation task (the Tatoeba⁴ dataset processed by Kelly 2021). For our experiments, we split the data randomly into 19222 training and 2402 test pairs.

Results Results are shown in Table 4. Models with a lexical translation mechanism obtain modest improvements (up to 1.5 BLEU) over the baseline. Notably, if we restrict evaluation to test sentences

⁴<https://tatoeba.org/>

	ENG-CHN	
	full	1-shot
LSTM	24.18 \pm 0.37	17.47 \pm 0.64
LSTM + GECA	23.90 \pm 0.55	17.94 \pm 0.43
LSTM + Lex.: PMI	24.36 \pm 0.09	18.46 \pm 0.13
LSTM + Lex.: Simple	24.35 \pm 0.09	18.46 \pm 0.19
LSTM + Lex.: IBMM2	25.49 \pm 0.42	19.62 \pm 0.64

Table 4: BLEU scores for English–Chinese translation. *full* shows results on the full test set, and *1-shot* shows results for text examples in which the English text contains a token seen only once during training.

featuring English words that appeared only once in the training set, BLEU improves by more than 2 points, demonstrating that this approach is particularly effective at one-shot word learning (or *fast mapping*; Carey and Bartlett 1978). Fig. 2 shows an example from this dataset, in which the model learns to reliably translate *Saturn* from a single training example. GECA, which makes specific generative assumptions about data distributions, does not generalize to a more realistic low resource MT problem. However, the lexical translation mechanism remains effective in natural tasks with large vocabularies and complex grammars.

5.5 Fine-Tuning the Lexicon

In all the experiments above, the lexicon was discretized ($\tau = 0$) and frozen prior to training. In this final section, we revisit that decision, evaluating whether the parameter L can be learned from scratch, or effectively fine-tuned along with decoder parameters. Experiments in this section focus on the COGS dataset.

Offline initialization of the lexicon is crucial. Rather than initializing L using any of the algorithms described in Section 3, we initialized L to a uniform distribution for each word and optimized

	COGS
LSTM	0.51 \pm 0.06
└ Lex.: Uniform	0.56 \pm 0.07
└ Lex.: Simple	0.82 \pm 0.01
└ Soft	0.83 \pm 0.00
└ Learned	0.83 \pm 0.01

Table 5: Ablation experiments on the COGS dataset. *Uniform* shows results for a lexicon initialized to a uniform distribution. *Soft* sets $\tau = 0.1$ with the Simple lexicon learning rule (rather than 0 in previous experiments). *Learned* shows results for a soft lexicon fine-tuned during training. Soft lexicons with or without learning improve significantly ($p < 0.01$) but very slightly over fixed initialization.

it during training. This improves over the base LSTM (*Uniform* in Table 5), but performs significantly worse than pre-learned lexicons.

Benefits from fine-tuning are minimal. We first increased the temperature parameter τ to 0.1 (providing a “soft” lexicon); this gave a 1% improvement on COGS (Table 5. *Soft*). Finally, we updated this soft initialization via gradient descent; this provided no further improvement (Table 5, *Learned*). One important feature of COGS (and other tests of compositional generalization) is perfect *training* accuracy is easily achieved; thus, there is little pressure on models to learn generalizable lexicons. This pressure must instead come from inductive bias in the initializer.

6 Conclusion

We have described a *lexical translation mechanism* for representing token-level translation rules in neural sequence models. We have additionally described a simple initialization scheme for this lexicon that outperforms a variety of existing algorithms. Together, lexical translation and proper initialization enable neural sequence models to solve a diverse set of tasks—including semantic parsing and machine translation—that require 1-shot word learning and 0-shot compositional generalization. Future work might focus on generalization to longer sequences, learning of atomic but non-concatenative translation rules, and online lexicon learning in situated contexts.

Acknowledgements

This work was supported by the Machine-LearningApplications initiative at MIT CSAIL and the MIT-IBM Watson AI lab. Computing re-

sources were provided by a gift from NVIDIA through the NVAIL program and by the Lincoln Laboratory Supercloud.

References

- Ekin Akyürek, Afra Feyza Akyürek, and Jacob Andreas. 2021. Learning to recombine and resample data for compositional generalization. In *International Conference on Learning Representations*.
- Jacob Andreas. 2020. Good-enough compositional data augmentation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7556–7566.
- Philip Arthur, Graham Neubig, and Satoshi Nakamura. 2016. Incorporating discrete translation lexicons into neural machine translation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1557–1567.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Joan Bresnan, Ash Asudeh, Ida Toivonen, and Stephen Wechsler. 2015. *Lexical-functional syntax*. John Wiley & Sons.
- Peter F Brown, Stephen A Della Pietra, Vincent J Della Pietra, and Robert L Mercer. 1993. The mathematics of statistical machine translation: Parameter estimation. *Computational linguistics*, 19(2):263–311.
- Susan Carey and Elsa Bartlett. 1978. Acquiring a single new word. *Papers and Reports on Child Language Development*, 2.
- Arthur P Dempster, Nan M Laird, and Donald B Rubin. 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22.
- Chris Dyer, Victor Chahuneau, and Noah A Smith. 2013. A simple, fast, and effective reparameterization of IBM Model 2. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–648.
- Jerry A Fodor, Zenon W Pylyshyn, et al. 1988. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1-2):3–71.
- Michael C. Frank, Noah D. Goodman, and J. Tenenbaum. 2007. A Bayesian framework for cross-situational word-learning. In *NeurIPS*.
- Daniel Furrer, Marc van Zee, Nathan Scales, and Nathanael Schärli. 2020. Compositional generalization in semantic parsing: Pre-training vs. specialized architectures. *arXiv preprint arXiv:2007.08970*.

- Kanishk Gandhi and Brenden M Lake. 2020. Mutual exclusivity as a challenge for deep neural networks. *Advances in Neural Information Processing Systems*, 33.
- Jonas Gehring, M. Auli, David Grangier, Denis Yarats, and Yann Dauphin. 2017. Convolutional sequence to sequence learning. In *ICML*.
- Sharon J Goldwater. 2007. *Nonparametric Bayesian Models of Lexical Acquisition*. Citeseer.
- Jonathan Gordon, David Lopez-Paz, Marco Baroni, and Diane Bouchacourt. 2019. Permutation equivariant models for compositional generalization in language. In *International Conference on Learning Representations*.
- Edward Grefenstette, Karl Moritz Hermann, Mustafa Suleyman, and Phil Blunsom. 2015. Learning to transduce with unbounded memory. In *NIPS*.
- Jetic Gū, Hassan S Shavarani, and Anoop Sarkar. 2019. Pointer-based fusion of bilingual lexicons into neural machine translation. *arXiv preprint arXiv:1909.07907*.
- Maria Teresa Guasti. 2017. *Language acquisition: The growth of grammar*. MIT press.
- A. Haghighi, Percy Liang, Taylor Berg-Kirkpatrick, and D. Klein. 2008. Learning bilingual lexicons from monolingual corpora. In *ACL*.
- Robin Jia and Percy Liang. 2016. Data recombination for neural semantic parsing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12–22.
- Aravind K Joshi and Yves Schabes. 1997. Tree-adjointing grammars. In *Handbook of formal languages*, pages 69–123. Springer.
- Charles Kelly. 2021. [\[link\]](#).
- Daniel Keysers, Nathanael Schärli, Nathan Scales, Hylke Buisman, Daniel Furrer, Sergii Kashubin, Nikola Momchev, Danila Sinopalnikov, Lukasz Stafiniak, Tibor Tihon, Dmitry Tsarkov, Xiao Wang, Marc van Zee, and Olivier Bousquet. 2020. Measuring compositional generalization: A comprehensive method on realistic data. In *ICLR*.
- Najoung Kim and Tal Linzen. 2020. COGS: A compositional generalization challenge based on semantic interpretation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9087–9105.
- Philipp Koehn, F. Och, and D. Marcu. 2003. Statistical phrase-based translation. In *HLT-NAACL*.
- T. Kwiatkowski, Luke Zettlemoyer, S. Goldwater, and Mark Steedman. 2011. Lexical generalization in CCG grammar induction for semantic parsing. In *EMNLP*.
- B. Lake, Tal Linzen, and M. Baroni. 2019. Human few-shot learning of compositional instructions. In *CogSci*.
- Brenden Lake and Marco Baroni. 2018. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In *International Conference on Machine Learning*, pages 2873–2882. PMLR.
- Percy Liang, Michael I Jordan, and Dan Klein. 2009. Learning semantic correspondences with less supervision. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 91–99.
- Qian Liu, Shengnan An, Jian-Guang Lou, Bei Chen, Zeqi Lin, Yan Gao, Bin Zhou, Nanning Zheng, and Dongmei Zhang. 2020. Compositional generalization by learning analytical expressions. *Advances in Neural Information Processing Systems*, 33.
- Gary Marcus. 2018. Deep learning: A critical appraisal. *arXiv preprint arXiv:1801.00631*.
- Toan Q. Nguyen and David Chiang. 2018. Improving lexical choice in neural machine translation. *ArXiv*, abs/1710.01329.
- Maxwell Nye, Armando Solar-Lezama, Josh Tenenbaum, and Brenden M Lake. 2020. Learning compositional rules via neural program synthesis. In *Advances in Neural Information Processing Systems*, volume 33, pages 10832–10842. Curran Associates, Inc.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Ngoc-Quan Pham, Jan Niehues, and Alex Waibel. 2018. Towards one-shot learning for rare-word translation with external experts. In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 100–109.
- Carl Pollard and Ivan A Sag. 1994. *Head-driven phrase structure grammar*. University of Chicago Press.
- Nikhil Prabhu and K. Kann. 2020. Making a point: Pointer-generator transformers for disjoint vocabularies. In *AAACL*.
- Taraka Rama, Anil Kumar Singh, and Sudheer Kolachina. 2009. Modeling letter-to-phoneme conversion as a phrase based statistical machine translation problem with minimum error rate training. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Companion Volume: Student Research Workshop and Doctoral Consortium*, pages 90–95.

Jake Russin, Jason Jo, Randall C O'Reilly, and Yoshua Bengio. 2019. Compositional generalization in a deep seq2seq model by separating syntax and semantics. *arXiv preprint arXiv:1904.09708*.

Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 1073–1083.

Tristan Thrush. 2020. Compositional neural machine translation by removing the lexicon from syntax. *arXiv preprint arXiv:2002.08899*.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 6000–6010.

A Colors Dataset & Detailed Results

Here we present the full dataset in Table 6 from Lake et al. (2019), and detailed comparisons of each model with human results in Table 7.

TRAIN		TEST	
INPUT	OUTPUT	INPUT	OUTPUT
dax		zup fep	
lug		zup kiki dax	
wif		wif kiki zup	
zup		zup blicket lug	
lug fep		dax blicket zup	
dax fep		wif kiki zup fep	
lug blicket wif		zup fep kiki lug	
wif blicket dax		lug kiki wif blicket zup	
lug kiki wif		zup blicket wif kiki dax fep	
dax kiki lug		zup blicket zup kiki zup fep	
lug fep kiki wif			
wif kiki dax blicket lug			
lug kiki wif fep			
wif blicket dax kiki lug			

Table 6: Full Colors dataset with Train and Test examples (Lake et al., 2019)

Test Examples	Simple/IBM-M2	Bayesian	GECA	SyntAtt	Human
zup fep	1.0±0.00	0.88±0.33	1.0±0.00	0.7±0.5	0.88
zup kiki dax	1.0±0.00	0.88±0.33	1.0±0.00	0.7±0.5	0.86
wif kiki zup	1.0±0.00	0.8±0.4	1.0±0.00	0.8±0.4	0.86
dax blicket zup	1.0±0.00	0.88±0.33	1.0±0.00	0.8±0.4	0.88
zup blicket lug	0.94±0.24	0.8±0.4	1.0±0.00	0.8±0.4	0.79
wif kiki zup fep	1.0±0.00	0.3±0.5	0.0±0.00	0.4±0.005	0.85
zup fep kiki lug	1.0±0.00	0.2±0.4	0.0±0.00	0.8±0.4	0.85
lug kiki wif blicket zup	1.0±0.00	0.4±0.5	0.0±0.00	0.4±0.5	0.65
zup blicket wif kiki dax fep	0.0±0.00	0.0±0.00	0.0±0.00	0.0±0	0.70
zup blicket zup kiki zup fep	0.0±0.00	0.0±0.00	0.0±0.00	0.0±0.00	0.75

Table 7: Colors dataset exact match breakdown for each individual test example. Human results are taken from (Lake et al., 2019)Fig2.

B Learned Lexicons

Here we provide lexicons for each model and dataset (see Fig. 2 and Fig. 3 for remaining datasets). For COGS, we show a representative subset of words.




Figure 4: Learned lexicons from SCAN dataset *jump* split with $\tau = 0.1$




Figure 5: Learned lexicons from Colors dataset with $\tau = 0.1$




Figure 6: Learned lexicons from COGS dataset with $\tau = 0.1$. We only show important rare words responsible for our model’s improvements over the baseline.

C Hyper-parameter Settings

C.1 Neural Seq2Seq

Most of the datasets we evaluate do not come with a out-of-distribution validation set, making principled hyperparameter tuning difficult. We were unable to reproduce the results of Kim and Linzen (2020) with the hyperparameter settings reported there with our base LSTM setup, and so adjusted them until training was stabilized. Like the original paper, we used a unidirectional 2-layer LSTM with 512 hidden units, an embedding size of 512, gradient clipping of 5.0, a Noam learning rate scheduler with 4000 warm-up steps, and a batch size of 512. Unlike the original paper, we found it necessary to reduce learning rate to 1.0, increase dropout value to 0.4, and the reduce maximum step size timeout

to 8000.

We use same parameters for all COGS, SCAN, and machine translation experiments. For SCAN and Colors, we applied additional dropout ($p=0.5$) in the last layer of p_{write} .

Since Colors has 14 training examples, we need a different batch size, set to 1/3 of the training set size ($= 5$). Qualitative evaluation of gradients in training time revealed that stricter gradient clipping was also needed ($= 0.5$). Similarly, we decreased warm-up steps to 32 epochs. All other hyper-parameters remain the same.

C.2 Lexicon Learning

Simple Lexicon The only parameter in the simple lexicon is ϵ , set to 3 in all experiments.

Bayesian The original work of Frank et al. (2007) did not report hyperparameter settings or sampler details. We found $\alpha = 2$, $\gamma = 0.95$ and $\kappa = 0.1$ to be effective. The M-H proposal distribution inserts or removes a word from the lexicon with 50% probability. For deletions, an entry is removed uniformly at random. For insertions, an entry is added with probability proportional to the empirical joint co-occurrence probability of the input and output tokens. Results were averaged across 5 runs, with a burn-in period of 1000 and a sample drawn every 10 steps.

IBM Model 2 We used the FastAlign implementation (Dyer et al., 2013) and experimented with a variety of hyperparameters in the alignment algorithm itself (favoring diagonal alignment, optimizing tension, using dirichlet priors) and diagonalization heuristics (grow-diag, grow-diag-final, grow-diag-final-and, union). We found that optimizing tension and using the “intersect” diagonalization heuristic works the best overall.

D Baseline Results

D.1 GECA

We reported best results for SCAN dataset from reproduced results in (Akyürek et al., 2021). For other datasets (COGS and Colors), we performed a hyperparameter search over augmentation ratios of 0.1 and 0.3 and hidden sizes of {128, 256, 512}. We report the best results for each dataset.

D.2 SyntAtt

We used the public GitHub repository of SyntAtt⁵ and reproduced reported results for the SCAN dataset. For other datasets, we also explored “syntax action” option, in which both contextualized context (syntax) and un-contextualized embeddings (semantics) used in final layer Russin et al. (2019). We additionally performed a search over hidden layer sizes {128,256,512} and depths {1,2}. We report the best results for each dataset.

E Datasets & Evaluation & Tokenization

E.1 Datasets and Sizes

	around_right	jump	COGS	Colors	ENG-CHN
train	15225	14670	24155	14	19222
validation	-	-	3000	-	2402
test	4476	7706	21000	10	2402

E.2 Evaluation

We report exact match accuracies and BLEU scores. In both evaluations we include punctuation. For BLEU we use NLTK⁶ library’s default implementation.

E.3 Tokenization

We use *Moses* library⁷ for English tokenization, and *jieba*⁸ library for Chinese tokenization. In other datasets, we use default space tokenization.

F Computing Infrastructure

Experiments were performed on a DGX-2 with NVIDIA 32GB VOLTA-V100 GPUs. Experiments take at most 2.5 hours on a single GPU.

⁵(https://github.com/jlruissin/syntactic_attention)

⁶<https://www.nltk.org/>

⁷<https://pypi.org/project/mosestokenizer/>

⁸<https://github.com/fxsjy/jieba>