

Compositional Generalization without Trees using Multiset Tagging and Latent Permutations

Matthias Lindemann¹ and Alexander Koller² and Ivan Titov^{1,3}

¹ ILCC, University of Edinburgh, ² LST, Saarland University, ³ ILLC, University of Amsterdam

m.m.lindemann@sms.ed.ac.uk, koller@coli.uni-saarland.de, ititov@inf.ed.ac.uk

Abstract

Seq2seq models have been shown to struggle with compositional generalization in semantic parsing, i.e. generalizing to unseen compositions of phenomena that the model handles correctly in isolation. We phrase semantic parsing as a two-step process: we first tag each input token with a multiset of output tokens. Then we arrange the tokens into an output sequence using a new way of parameterizing and predicting permutations. We formulate predicting a permutation as solving a regularized linear program and we backpropagate through the solver. In contrast to prior work, our approach does not place a priori restrictions on possible permutations, making it very expressive. Our model outperforms pretrained seq2seq models and prior work on realistic semantic parsing tasks that require generalization to longer examples. We also outperform non-tree-based models on structural generalization on the COGS benchmark. For the first time, we show that a model without an inductive bias provided by trees achieves high accuracy on generalization to deeper recursion depth.¹

1 Introduction

Sequence-to-sequence models have been very successfully applied to many structural tasks in NLP such as semantic parsing. However, they have also been shown to struggle with compositional generalization (Lake and Baroni, 2018; Finegan-Dollak et al., 2018; Kim and Linzen, 2020; Hupkes et al., 2020), i.e. the model fails on examples that contain unseen compositions or deeper recursion of phenomena that it handles correctly in isolation. For example, a model which correctly parses ‘Mary knew that Jim slept’ should also be able to parse sentences with deeper recursion than it has seen during training such as ‘Paul said that Mary knew that Jim slept’. This sort of generalization is easy for humans but challenging for neural models.

¹<https://github.com/namednil/multiset-perm>

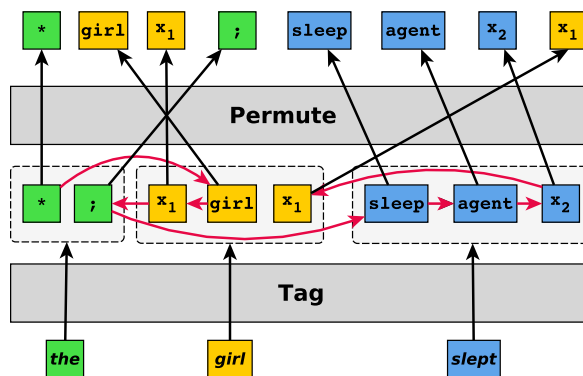


Figure 1: We model seq2seq tasks by first predicting a multiset (dashed boxes) for every input token, and then permuting the tokens to put them into order. We realize the permutation using the red edges: if there is an edge from token i to j , then i is the predecessor of j in the output. Every token is visited exactly once.

In order for a model to generalize compositionally in semantic parsing, it has to identify reusable ‘fragments’ and be able to systematically combine them in novel ways. One way to make a model sensitive to fragments is to make it rely on a tree that makes the compositional structure explicit. However, this complicates the training because these trees need to be obtained or induced. This is computationally demanding or at least requires structural preprocessing informed by domain knowledge.

In this paper, we propose a two-step sequence-based approach with a structural inductive bias that does not rely on trees: the ‘fragments’ are multisets of output tokens that we predict for every input token in a first step. A second step then arranges the tokens we predicted in the previous step into a single sequence using a permutation model. In contrast to prior work (Wang et al., 2021; Lindemann et al., 2023), there are no hard constraints on the permutations that our model can predict. We show that this enables structural generalization on tasks that go beyond the class of synchronous context-free languages.

We overcome two key technical challenges in this work: Firstly, we do not have supervision for the correspondence between input tokens and output tokens. Therefore, we induce the correspondence during training. Secondly, predicting permutations without restrictions is computationally challenging. For this, we develop a differentiable GPU-friendly algorithm.

On realistic semantic parsing tasks our approach outperforms previous work on generalization to longer examples than seen at training. We also outperform all other non-tree models on the structural generalization tasks in semantic parsing on COGS (Kim and Linzen, 2020). For the first time, we also show that a model without an inductive bias towards trees can obtain high accuracy on generalization to deeper recursion on COGS.

To summarize, our main contributions are:

- a flexible seq2seq model that performs well on structural generalization in semantic parsing without assuming that input and output are related to each other via a tree structure.
- a differentiable algorithm to parameterize and predict permutations without a priori restrictions on what permutations are possible.

2 Overview and Motivation

Our approach consists of two stages. In the **first stage** (multiset tagging), we predict a multiset \mathbf{z}_i of tokens for every input token \mathbf{x}_i from the contextualized representation of \mathbf{x}_i . This is motivated by the observation that input tokens often systematically correspond to a fragment of the output (like *slept* corresponding to *sleep* and *agent* and a variable in Fig. 1). Importantly, we expect this systematic relationship to be largely invariant to phrases being used in new contexts or deeper recursion. We refer to the elements of the multisets as *multiset tokens*.

In the **second stage** (permutation), we order the multiset tokens to arrive at a sequential output. Conceptually, we do this by going from left to right over the *output* and determining which multiset token to put in every position. Consider the example in Fig. 1. For the first output position, we simply select a multiset token ($*$ in the example). All subsequent tokens are put into position by ‘jumping’ from the token that was last placed into the output to a new multiset token. In Fig. 1, we jump from $*$ to *girl* (shown by the outgoing red edge from $*$). This indicates that *girl* is the successor of $*$ in the output and hence the *second* output token. From

girl we jump to one of the x_1 tokens to determine what the *third* output token is and so on. Since we predict a permutation, we must visit every multiset token exactly once in this process.

The jumps are inspired by reordering in phrase-based machine translation (Koehn et al., 2007) and methods from dependency parsing, where directly modeling bi-lexical relationships on hidden states has proven successful (Kiperwasser and Goldberg, 2016). Note also that *any* permutation can be represented with jumps. In contrast, prior work (Wang et al., 2021; Lindemann et al., 2023) has put strong restrictions on the possible permutations that can be predicted. Our approach is more flexible and empirically it also scales better to longer inputs, which opens up new applications and datasets.

Setup. We assume we are given a dataset $\mathcal{D} = \{(\mathbf{x}^1, \mathbf{y}^1), \dots\}$ of input utterances \mathbf{x} and target sequences \mathbf{y} . If we had gold alignments, it would be straightforward to train our model. Since we do not have this supervision, we have to discover during training which tokens of the output \mathbf{y} belong into which multiset \mathbf{z}_i . We describe the model and the training objective of the multiset tagging in Section 3. After the model is trained, we can annotate our training set with the most likely \mathbf{z} , and then train the permutation model.

For predicting a permutation, we associate a score with each possible jump and search for the highest-scoring sequence of jumps. We ensure that the jumps correspond to a permutation by means of constraints, which results in a combinatorial optimization problem. The flexibility of our model and its parametrization come with the challenge that this problem is NP-hard. We approximate it with a regularized linear program which also ensures differentiability. Our permutation model and its training are described in Section 4. In Section 5, we discuss how to solve the regularized linear program and how to backpropagate through the solution.

3 Learning Multisets

For the multiset tagging, we need to train a model that predicts the multisets $\mathbf{z}_1, \dots, \mathbf{z}_n$ of tokens by conditioning on the input. We represent a multiset \mathbf{z}_i as an integer-valued vector that contains for every vocabulary item v the multiplicity of v in \mathbf{z}_i , i.e. $\mathbf{z}_{i,v} = k$ means that input token i contributes k occurrences of output type v . If v is not present in multiset \mathbf{z}_i , then $\mathbf{z}_{i,v} = 0$. For example, in Fig. 1, $\mathbf{z}_{3,\text{sleep}} = 1$ and $\mathbf{z}_{2,x_1} = 2$. As discussed in Sec-

tion 2, we do not have supervision for $\mathbf{z}_1, \dots, \mathbf{z}_n$ and treat them as latent variables. To allow for efficient exact training, we assume that all $\mathbf{z}_{i,v}$ are independent of each other conditioned on the input:

$$P(\mathbf{z} | \mathbf{x}) = \prod_{i,v} P(\mathbf{z}_{i,v} | \mathbf{x}) \quad (1)$$

where v ranges over the *entire* vocabulary.

Parameterization. We parameterize $P(\mathbf{z}_{i,v} | \mathbf{x})$ as follows. We first pass the input \mathbf{x} through a pretrained RoBERTa encoder (Liu et al., 2019), where $\text{ENCODER}(\mathbf{x})$ is the output of the final layer. We then add RoBERTa’s static word embeddings from the first, non-contextualized, layer to that:

$$\mathbf{h}_i = \text{ENCODER}(\mathbf{x})_i + \text{EMBED}(\mathbf{x}_i) \quad (2)$$

We then pass \mathbf{h}_i through a feedforward network obtaining $\tilde{\mathbf{h}}_i = \text{FF}(\mathbf{h}_i)$ and define a distribution over the multiplicity of v in the multiset \mathbf{z}_i :

$$P(\mathbf{z}_{i,v} = k | \mathbf{x}_i) = \frac{\exp\left(\tilde{\mathbf{h}}_i^T \mathbf{w}^{v,k} + b^{v,k}\right)}{\sum_l \exp\left(\tilde{\mathbf{h}}_i^T \mathbf{w}^{v,l} + b^{v,l}\right)} \quad (3)$$

where the weights \mathbf{w} and biases b are specific to v and the multiplicity k . In contrast to standard sequence-to-sequence models, this softmax is not normalized over the vocabulary but over the multiplicity, and we have distinct distributions for every vocabulary item v . Despite the independence assumptions in Eq. (1), the model can still be strong because \mathbf{h}_i takes the entire input \mathbf{x} into account.

Training. The probability of generating the overall multiset \mathbf{m} as the union of all \mathbf{z}_i is the probability that for every vocabulary item v , the total number of occurrences of v across all input positions i sums to \mathbf{m}_v :

$$P(\mathbf{m} | \mathbf{x}) = \prod_v P(\mathbf{z}_{1,v} + \dots + \mathbf{z}_{n,v} = \mathbf{m}_v | \mathbf{x})$$

This can be computed recursively:

$$P(\mathbf{z}_{1,v} + \dots + \mathbf{z}_{n,v} = \mathbf{m}_v | \mathbf{x}) = \sum_k P(\mathbf{z}_{1,v} = k | \mathbf{x}) P(\mathbf{z}_{2,v} + \dots + \mathbf{z}_{n,v} = \mathbf{m}_v - k | \mathbf{x})$$

Let $\mathbf{m}(\mathbf{y})$ be the multiset of tokens in the gold sequence \mathbf{y} . We train our model with gradient ascent to maximize the marginal log-likelihood of $\mathbf{m}(\mathbf{y})$:

$$\sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}} \log P(\mathbf{m}(\mathbf{y}) | \mathbf{x}) \quad (4)$$

Like Lindemann et al. (2023), we found it helpful to initialize the training of our model with high-confidence alignments from an IBM-1 model (Brown et al., 1993) (see Appendix C.3 for details).

Preparation for permutation. The scoring function of the permutation model expects a sequence as input. There is no a priori obvious order for the elements within the individual multisets \mathbf{z}_i . We handle this by imposing a canonical order $\text{ORDER}(\mathbf{z}_i)$ on the elements of \mathbf{z}_i by sorting the multiset tokens by their id in the vocabulary. They are then concatenated to form the input $\mathbf{z}' = \text{ORDER}(\mathbf{z}_1) \dots \text{ORDER}(\mathbf{z}_n)$ to the permutation model.

4 Relaxed Permutations

After predicting a multiset for every input token and arranging the elements within each multiset to form a sequence \mathbf{z}' , we predict a permutation of \mathbf{z}' .

We represent a permutation as a matrix \mathbf{V} that contains exactly one 1 in every row and column and zeros otherwise. We write $\mathbf{V}_{i,j} = 1$ if position i in \mathbf{z}' is mapped to position j in the output \mathbf{y} . Let \mathcal{P} be the set of all permutation matrices.

Now we formalize the parametrization of permutations as discussed in Section 2. We associate a score predicted by our neural network with each permutation \mathbf{V} and search for the permutation with the highest score. The score of a permutation decomposes into a sum of scores for binary ‘features’ of \mathbf{V} . We use two types of features.

The first type of feature is active if the permutation \mathbf{V} maps input position i to output position j (i.e. $\mathbf{V}_{ij} = 1$). We associate this feature with the score $s_{i \rightarrow j}$ and use these scores only to model what the first and the last token in the output should be. That is, $s_{i \rightarrow 1}$ models the preference to map position i in the input to the first position in the output, and analogously $s_{i \rightarrow n}$ models the preference to put i into the last position in the output. For all output positions j that are neither the first nor the last position, we simply set $s_{i \rightarrow j} = 0$.

The second type of feature models the jumps we introduced in Section 2. We introduce a feature that is 1 iff \mathbf{V} contains a jump from k to i , and associate this with a score $s_{k \curvearrowright i}$. In order for there to be a jump from k to i , the permutation \mathbf{V} must map input i to some output position j ($\mathbf{V}_{ij} = 1$) and it must also map input position k to output position $j - 1$ ($\mathbf{V}_{k,j-1} = 1$). Hence, the product $\mathbf{V}_{k,j-1} \mathbf{V}_{ij}$ is 1 iff there is a jump from k to i at output position

j . Based on this, the sum $\sum_j \mathbf{V}_{k,j-1} \mathbf{V}_{ij}$ equals 1 if there is *any* output position j at which there is a jump from k to i and 0 otherwise. This constitutes the second type of feature.

Multiplying the features with their respective scores, we want to find the highest-scoring permutation under the following overall scoring function:

$$\arg \max_{\mathbf{V} \in \mathcal{P}} \sum_{i,j} \mathbf{V}_{ij} s_{i \rightarrow j} + \sum_{i,k} s_{k \curvearrowright i} \left(\sum_j \mathbf{V}_{k,j-1} \mathbf{V}_{ij} \right) \quad (5)$$

Let $\mathbf{V}^*(s)$ be the solution to Eq. (5) as a function of the scores. Unfortunately, $\mathbf{V}^*(s)$ does not have sensible derivatives because \mathcal{P} is discrete. This makes $\mathbf{V}^*(s)$ unsuitable as a neural network layer. In addition, Eq. (5) is NP-hard (see Appendix A.1).

We now formulate an optimization problem that approximates Eq. (5) and which has useful derivatives. Firstly, we relax the permutation matrix \mathbf{V} to a bistochastic matrix \mathbf{U} , i.e. \mathbf{U} has non-negative elements and every row and every column each sum to 1. Secondly, note that Eq. (5) contains quadratic terms. As we will discuss in the next section, our solver assumes a linear objective, so we replace $\mathbf{V}_{k,j-1} \mathbf{V}_{ij}$ with an auxiliary variable \mathbf{W}_{ijk} . The variable \mathbf{W}_{ijk} is designed to take the value 1 if and only if $\mathbf{U}_{i,j} = 1$ and $\mathbf{U}_{k,j-1} = 1$. We achieve this by coupling \mathbf{W} and \mathbf{U} using constraints. Then, the optimization problem becomes:

$$\arg \max_{\mathbf{U}, \mathbf{W}} \sum_{i,j} \mathbf{U}_{ij} s_{i \rightarrow j} + \sum_{i,j,k} \mathbf{W}_{ijk} s_{k \curvearrowright i} \quad (6)$$

$$\text{subject to } \sum_i \mathbf{U}_{ij} = 1 \quad \forall j \quad (6a)$$

$$\sum_j \mathbf{U}_{ij} = 1 \quad \forall i \quad (6b)$$

$$\sum_k \mathbf{W}_{ijk} = \mathbf{U}_{ij} \quad \forall j > 1, i \quad (6c)$$

$$\sum_i \mathbf{W}_{ijk} = \mathbf{U}_{k(j-1)} \quad \forall j > 1, k \quad (6d)$$

$$\mathbf{U}, \mathbf{W} \geq 0 \quad (6e)$$

Finally, in combination with the linear objective, the argmax operation still causes the solution $\mathbf{U}^*(s)$ of Eq. (6) as a function of s to have no useful derivatives. This is because an infinitesimal change in s has no effect on the solution $\mathbf{U}^*(s)$ for almost all s .

To address this, we add an entropy regularization term $\tau (H(\mathbf{U}) + H(\mathbf{W}))$ to the objective Eq. (6), where $H(\mathbf{U}) = -\sum_{ij} \mathbf{U}_{ij} (\log \mathbf{U}_{ij} - 1)$, and $\tau > 0$ determines the strength of the regularization. The entropy regularization ‘smooths’ the solution $\mathbf{U}^*(s)$ in an analogous way to softmax being a smoothed version of argmax. The parameter τ behaves analogously to the softmax temperature: the smaller τ , the sharper $\mathbf{U}^*(s)$ will be. We discuss how to solve the regularized linear program in Section 5.

Predicting permutations. At test time, we want to find the highest scoring permutation, i.e. we want to solve Eq. (5). We approximate this by using $\mathbf{U}^*(s)$ instead, the solution to the entropy regularized version of Eq. (6). Despite using a low temperature τ , there is no guarantee that $\mathbf{U}^*(s)$ can be trivially rounded to a permutation matrix. Therefore, we solve the linear assignment problem with $\mathbf{U}^*(s)$ as scores using the Hungarian Algorithm (Kuhn, 1955). The linear assignment problem asks for the permutation matrix \mathbf{V} that maximizes $\sum_{ij} \mathbf{V}_{ij} \mathbf{U}^*(s)_{ij}$.

4.1 Parametrization

We now describe how we parameterize the scores s to permute the tokens into the right order. We first encode the original input utterance x like in Eq. (2) to obtain a hidden representation \mathbf{h}_i for input token \mathbf{x}_i . Let a be the function that maps $a(i) \mapsto j$ if the token in position i in z' came from the multiset that was generated by token \mathbf{x}_j . For example, in Fig. 1, $a(6) = 3$ since sleep was predicted from input token *sleep*. We then define the hidden representation \mathbf{h}'_i as the concatenation of $\mathbf{h}_{a(i)}$ and an embedding of z'_i :

$$\mathbf{h}'_i = [\mathbf{h}_{a(i)}; \text{EMBED}(z'_i)] \quad (7)$$

We parameterize the scores for starting the output with token i as

$$s_{i \rightarrow 1} = \mathbf{w}_{\text{start}}^T \text{FF}_{\text{start}}(\mathbf{h}'_i)$$

and analogously for ending it with token i :

$$s_{i \rightarrow n} = \mathbf{w}_{\text{end}}^T \text{FF}_{\text{end}}(\mathbf{h}'_i)$$

We set $s_{i \rightarrow j} = 0$ for all other i, j .

We parameterize the jump scores $s_{k \curvearrowright i}$ using Geometric Attention (Csordás et al., 2022) from \mathbf{h}'_k to \mathbf{h}'_i . Intuitively, Geometric Attention favours selecting the ‘matching’ element \mathbf{h}'_i that is closest to \mathbf{h}'_k in terms of distance $|i - k|$ in the string. We refer to Csordás et al. (2022) for details.




Figure 2: Two possible permutations (represented with red arrows) of z' that yield the same result y . Only the permutation marked with solid arrows corresponds to the right linguistic relation that will generalize well.

4.2 Learning Permutations

We now turn to training the permutation model. At training time, we have access to the gold output y and a sequence z' from the output of the multiset tagging (see the end of Section 3). We note that whenever y (or z') contains one vocabulary item at least twice, there are multiple permutations that can be applied to z' to yield y . Many of these permutations will give the right result for the wrong reasons and the permutation that is desirable for generalization is latent. For example, consider Fig. 2. The token *agent* is followed by the entity who performs the action, whereas *theme* is followed by the one affected by the action. The permutation indicated by dashed arrows generalizes poorly to a sentence like *Emma knew that James slid* since *slid* will introduce a theme role rather than an agent role (as the sliding is happening to James). Thus, this permutation would then lead to the incorrect output *know theme Emma ... slide agent James*, in which Emma is affected by the knowing event and James is the one who slides something.

In order to train the permutation model in this setting, we use a method that is similar to EM for structured prediction.² During training, the model output $U^*(s)$ and $W^*(s)$ often represents a soft permutation that does *not* permute z' into y . Our goal is to push the model output into the space of (soft) permutations that lead to y . More formally, let $Q \in Q(y, z')$ be the set of bistochastic matrices such that $Q_{i,j} = 0$ iff $z'_i \neq y_j$. That is, any permutation included in $Q(y, z')$ leads to the gold output y when applied to z' .

First, we project the current prediction $U^*(s)$ and $W^*(s)$ into $Q(y, z')$ using the KL divergence

²It can also be derived as a lower bound to the marginal log-likelihood $\log P_\theta(y|x, z')$ (see Appendix A.4), and be seen as a form of posterior regularization (Ganchev et al., 2010)

as a measure of distance (E-step):

$$\hat{U}, \hat{W} = \arg \min_{U \in Q(y, z'), W} \text{KL}(U \parallel U^*(s)) + \text{KL}(W \parallel W^*(s)) \quad (8)$$

subject to Eq. (6a) to Eq. (6e).

Similar to the M-step of EM, we then treat \hat{U} and \hat{W} as soft gold labels and train our model to minimize the KL divergence between labels and model:

$$\text{KL}(\hat{U} \parallel U^*(s)) + \text{KL}(\hat{W} \parallel W^*(s))$$

Eq. (8) can be solved in the same way as the entropy regularized version of Eq. (6) because expanding the definition of KL-divergence leads to a regularized linear program with a similar feasible region (see Appendix A.3 for details).

5 Inference for Relaxed Permutations

Now we describe how to solve the entropy regularized form of Eq. (6) and how to backpropagate through it. This section may be skipped on the first reading as it is not required to understand the experiments; we note that the resulting algorithm (Algorithm 1) is conceptually relatively simple. Before describing our method, we explain the general principle.

5.1 Bregman's Method

Bregman's method (Bregman, 1967) is a method for constrained convex optimization. In particular, it can be used to solve problems of the form

$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in C_0 \cap C_1 \dots \cap C_n, \mathbf{x} \geq 0} \mathbf{s}^T \mathbf{x} + \underbrace{\tau H(\mathbf{x})}_{\text{regularizer}} \quad (9)$$

where C_0, \dots, C_n are linear equality constraints, $H(\mathbf{x}) = -\sum_i \mathbf{x}_i (\log \mathbf{x}_i - 1)$ is a form of entropy regularization, and τ determines the strength of the regularization. Note that our parameterization of permutations (Eq. (6)) has this form.

Bregman's method is a simple iterative process. We start with the scores \mathbf{s} and then cyclically iterate over the constraints and project the current estimate \mathbf{x}^i onto the chosen constraint until convergence:

$$\begin{aligned} \mathbf{x}^0 &= \exp \frac{\mathbf{s}}{\tau} \\ \mathbf{x}^{i+1} &= \arg \min_{\mathbf{x} \in C_i \text{ mod } (n-1)} \text{KL}(\mathbf{x} \parallel \mathbf{x}^i) \end{aligned} \quad (10)$$

where $\text{KL}(\mathbf{x} \parallel \mathbf{y}) = \sum_i \mathbf{x}_i \log \frac{\mathbf{x}_i}{\mathbf{y}_i} - \mathbf{x}_i + \mathbf{y}_i$ is the generalized KL divergence. We call

Algorithm 1 Bregman’s method for Eq. (6) with entropic regularization

```

function BREGMAN( $s, \tau$ )
   $\mathbf{U}_{ij} = \exp(\tau^{-1} s_{i \rightarrow j})$ 
   $\mathbf{W}_{ijk} = \exp(\tau^{-1} s_{k \rightsquigarrow i})$ 
  while within budget and not converged do
     $\mathbf{U}, \mathbf{W} = \text{KL project}(\mathbf{U}, \mathbf{W}; \text{6a,6c})$  with Prop. 2
     $\mathbf{U}, \mathbf{W} = \text{KL project}(\mathbf{U}, \mathbf{W}; \text{6a,6d})$  with Prop. 2
     $\mathbf{U} = \text{KL project}(\mathbf{U}, \text{6b})$  with Prop. 1
  end while
  return  $\mathbf{U}, \mathbf{W}$ 
end function

```

$\arg \min_{\mathbf{x} \in C} \text{KL}(\mathbf{x} \mid \mathbf{x}^i)$ a KL projection. In order to apply Bregman’s method, we need to be able to compute the KL projection $\arg \min_{\mathbf{x} \in C} \text{KL}(\mathbf{x} \mid \mathbf{x}^i)$ for all C_0, \dots, C_n in closed-form. We discuss how to do this for Eq. (6) in the next section.

As an example, consider a problem of the form Eq. (9) with a single linear constraint $C^\Delta = \{\mathbf{x} \mid \sum_i \mathbf{x}_i = 1\}$. In this case, Bregman’s method coincides with the softmax function. This is because the KL projection $\mathbf{x}^* = \arg \min_{\mathbf{x} \in C^\Delta} \text{KL}(\mathbf{x} \parallel \mathbf{y})$ for $\mathbf{y} > \mathbf{0}$ has the closed-form solution $\mathbf{x}_i^* = \frac{y_i}{\sum_i y_i}$.

If we have a closed-form expression for a KL projection (such as normalizing a vector), we can use automatic differentiation to backpropagate through it. To backpropagate through the entire solver, we apply automatic differentiation to the composition of all projection steps.

5.2 Bregman’s Method for Eq. (6)

In order to apply Bregman’s method to solve the entropy regularized version of Eq. (6), we need to decompose the constraints into sets which we can efficiently project onto. We choose the following three sets here: (i), containing Eq. (6a) and Eq. (6c), and (ii), containing Eq. (6a) and Eq. (6d), and finally (iii), containing only Eq. (6b). We now need to establish what the KL projections are for our chosen sets. For (iii), the projection is simple:

Proposition 1. (Benamou et al. (2015), Prop. 1) For $\mathbf{A}, \mathbf{m} > \mathbf{0}$, the KL projection $\arg \min_{\mathbf{U}} \text{KL}(\mathbf{U} \parallel \mathbf{A})$ subject to $\sum_j \mathbf{U}_{ij} = \mathbf{m}_i$ is given by $\mathbf{U}_{ij}^* = \mathbf{m}_i \frac{\mathbf{A}_{ij}}{\sum_{j'} \mathbf{A}_{ij'}}$.

Let us now turn to (i) and (ii). The constraints Eq. (6d) and Eq. (6c) are structurally essentially the same, meaning that we can project onto (ii) in basically the same manner as onto (i). We project onto (i), with the following proposition:

Proposition 2. For $\mathbf{A}, \mathbf{B} > \mathbf{0}$, the KL projection

$$\begin{aligned} & \arg \min_{\mathbf{U}, \mathbf{W}} \text{KL}(\mathbf{U} \parallel \mathbf{A}) + \text{KL}(\mathbf{W} \parallel \mathbf{B}) \\ & \text{subject to} \quad \sum_i \mathbf{U}_{ij} = 1 \quad \forall j \quad (11) \\ & \quad \quad \quad \sum_k \mathbf{W}_{ijk} = \mathbf{U}_{ij} \quad \forall j, i \end{aligned}$$

is given by:

$$\begin{aligned} \mathbf{U}_{ij}^* &= \frac{\mathbf{T}_{ij}}{\sum_{i'} \mathbf{T}_{i'j}} \\ \mathbf{W}_{ijk}^* &= \mathbf{U}_{ij}^* \frac{\mathbf{B}_{ijk}}{\sum_{k'} \mathbf{B}_{ijk'}} \end{aligned}$$

where $\mathbf{T}_{ij} = \sqrt{\mathbf{A}_{ij} \cdot \sum_k \mathbf{B}_{ijk}}$.

The proof can be found in Appendix A.2. We present the overall algorithm in Algorithm 1, and note that it is easy to implement for GPUs. In practice, we implement all projections in log space for numerical stability.

6 Evaluation

6.1 Doubling Task

Our permutation model is very expressive and is not limited to synchronous context-free languages. This is in contrast to the formalisms that other approaches rely on (Wang et al., 2021; Lindemann et al., 2023). To evaluate if our model can structurally generalize beyond the synchronous context-free languages in practice, we consider the function $F = \{(w, ww) \mid w \in \Sigma^*\}$. This function is related to processing challenging natural language phenomena such as reduplication and cross-serial dependencies. We compare our model with an LSTM-based seq2seq model with attention and a Transformer in the style of Csordás et al. (2021) that uses a relative positional encoding. Since the input is a sequence of symbols rather than English, we replace RoBERTa with a bidirectional LSTM and use randomly initialized embeddings. The models are trained on inputs of lengths 5 to 10 and evaluated on longer examples. The results can be found in Fig. 3. All models get perfect or close to perfect accuracy on inputs of length 11 but accuracy quickly deteriorates for the LSTM and the Transformer. In contrast, our model extrapolates very well to longer sequences.

6.2 COGS

COGS is a synthetic benchmark for compositional generalization introduced by Kim and Linzen




Figure 3: Accuracy by input length for doubling task.

(2020). Models are tested for 21 different cases of generalization, 18 of which focus on using a lexical item in new contexts (**Lex**). There are 1000 instances per generalization case. Seq2seq models struggle in particular with the *structural* generalization tasks (Yao and Koller, 2022), and we focus on those: (i) generalization to deeper **PP** recursion than seen during training ("Emma saw a hedgehog on a chair in the garden beside the tree..."), (ii) deeper **CP** recursion ("Olivia mentioned that James saw that Emma admired that the dog slept"), and (iii) PPs modifying subjects when PPs modified only objects in the training data (**OS**).

We follow previous work and use a lexicon (Akyurek and Andreas, 2021) to map some input tokens to output tokens (see Appendix C.2 for details). We also use this mechanism to handle the variable symbols in COGS.

We report the means and standard deviations for 10 random seeds in Table 1. Our approach obtains high accuracy on CP and PP recursion but exact match accuracy is low for OS. This is in part because our model sometimes predicts semantic representations for OS that are equivalent to the gold standard but use a different order for the conjuncts. Therefore, we report accuracy that accounts for this in Table 2. In both tables, we also report the impact of using a simple copy mechanism instead of the more complex lexicon induction mechanism (-Lex). Our model outperforms all other non-tree-based models by a considerable margin.

Structural generalization without trees. All previous methods that obtain high accuracy on recursion generalization on COGS use trees. Some approaches directly predict a tree over the input (Liu et al., 2021; Weißenhorn et al., 2022), while others use derivations from a grammar for data augmentation (Qiu et al., 2022) or decompose the input along a task-specific parse tree (Drozdov et al., 2022). Our results show that trees are not as important for compositional generalization as their success in the literature may suggest, and that weaker

Model	OS	CP	PP	Lex	Total
Q'22♣	-	-	-	-	99
D'22♣	-	-	-	-	99
LexLSTM	0	0	1	96	82
Dangle	0	14	14	98	85
Dangle [◇]	0	25	35	99	88
T5-base [◇]	0	13	18	99	86
Ours [◇]	9 _{±13}	79 _{±11}	85 _{±14}	97 _{±2}	92 _{±2}
Ours [◇] - Lex	5 _{±6}	71 _{±26}	82 _{±20}	79 _{±1}	75 _{±2}

Table 1: Exact match accuracy on **COGS** by generalization type. [◇]refers to models using pretrained transformers, and [♣]refers to models implicitly or explicitly using trees. Q'22 is Qiu et al. (2022), D'22 is Drozdov et al. (2022). LexLSTM is the model by Akyurek and Andreas (2021) and Dangle is the model by Zheng and Lapata (2021). We report their results for T5.

Model	OS	CP	PP	Lex	Total
Lear [♣]	93	100	99	99	99
AM parser ^{◇♣}	72	100	97	76	78
Dangle	8	14	14	99	87
Ours [◇]	33 _{±24}	82 _{±11}	91 _{±5}	97 _{±2}	93 _{±2}
Ours [◇] - Lex	35 _{±22}	73 _{±27}	92 _{±5}	79 _{±1}	77 _{±2}

Table 2: Accuracy on **COGS** when the order of conjuncts is disregarded or established in post-processing. For Lear (Liu et al., 2021) and the AM parser (Groschwitz et al., 2018; Weißenhorn et al., 2022), we report numbers of Weißenhorn et al. (2022).

structural assumptions already reap some of the benefits.

Logical forms with variables. COGS uses logical forms with variables, which were removed in conversion to variable-free formats for evaluation of some approaches (Zheng and Lapata, 2021; Qiu et al., 2022; Drozdov et al., 2022). Recently, Wu et al. (2023) have argued for keeping the variable symbols because they are important for some semantic distinctions; we keep the variables.

6.3 ATIS

While COGS is a good benchmark for compositional generalization, the data is synthetic and does not contain many phenomena that are frequent in semantic parsing on real data, such as paraphrases that map to the same logical form. ATIS (Dahl et al., 1994) is a realistic English semantic parsing dataset with executable logical forms. We follow the setup of our previous work (Lindemann et al., 2023) (L'23) and use the variable-free FunQL rep-

Model	iid	Length
LSTM seq2seq	75.98 \pm 1.30	4.95 \pm 2.16
Transformer	75.76 \pm 1.43	1.15 \pm 1.41
BART-base \diamond	86.96 \pm 1.26	19.03 \pm 4.57
L'23 \clubsuit	68.26 \pm 1.53	29.91 \pm 2.91
L'23 $\dagger \clubsuit$	74.15 \pm 1.35	35.41 \pm 4.09
Ours \diamond	76.65 \pm 1.67	41.39 \pm 13.47
Ours	73.93 \pm 1.43	38.79 \pm 7.11

Table 3: Accuracy on ATIS. \dagger indicates grammar-based decoding. L'23 is the model of Lindemann et al. (2023).

resentation (Guo et al., 2020). Apart from the usual iid split, we evaluate on a length split, where a model is trained on examples with few conjuncts and has to generalize to longer logical forms with more conjuncts. For a fair comparison with previous work, we do not use the lexicon/copying. We also evaluate a version of our model without RoBERTa that uses a bidirectional LSTM and GloVe embeddings instead. This mirrors the model of L'23.

Table 3 shows mean accuracy and standard deviations over 5 runs. Our model is competitive with non-pretrained models in-distribution, and outperforms all other models on the length generalization. The high standard deviation on the length split stems from an outlier run with 18% accuracy – the second worst-performing run achieved an accuracy of 44%. Even without pretraining, our model performs very well. In particular, without grammar-based decoding our model performs on par or outperforms L'23 *with* grammar-based decoding.

The runtime of the model in L'23 is dominated by the permutation model and it takes up to 12 hours to train on ATIS. Training the model presented here only takes around 2 hours for both stages.

Performance breakdown. In order for our approach to be accurate, both the multiset tagging model and the permutation model have to be accurate. Table 4 explores which model acts as the bottleneck in terms of accuracy on ATIS and COGS. The answer depends on the dataset: for the synthetic COGS dataset, predicting the multisets correctly is easy except for OS, and the model struggles more with getting the permutation right. In contrast, for ATIS, the vast majority of errors can be attributed to the first stage.

		Freq	Seq	Seq/Freq
COGS	OS	50 \pm 34	9 \pm 13	11 \pm 15
	CP	97 \pm 5	79 \pm 11	82 \pm 13
	PP	99 \pm 0	85 \pm 14	85 \pm 15
ATIS	iid	77.6 \pm 1.4	76.7 \pm 1.7	98.7 \pm 0.5
	Length	42.2 \pm 13.6	41.4 \pm 13.5	97.8 \pm 0.8

Table 4: Performance breakdown of the first and second stage. ‘Freq’ refers to accuracy measured on the predicted multiset; it measures performance of the first stage. ‘Seq’ measures the accuracy of both stages. ‘Seq/Freq’ is the percentage of correct predictions given that the multiset is predicted correctly.

Model	Calendar	Doc	Email
BART-base \diamond	36.7 \pm 3.0	0.6 \pm 0.3	20.5 \pm 9.8
L'23 \clubsuit	57.2 \pm 19.9	36.1 \pm 5.6	43.9 \pm 3.8
L'23 $\dagger \clubsuit$	69.5 \pm 13.9	42.4 \pm 5.7	55.6 \pm 2.7
Ours \diamond	74.3 \pm 3.5	57.8 \pm 5.5	60.6 \pm 4.8
Ours	65.6 \pm 2.8	41.4 \pm 4.9	47.6 \pm 4.5

Table 5: Accuracy on length splits by domain on Okapi.

6.4 Okapi

Finally, we consider the recent Okapi (Hosseini et al., 2021) semantic parsing dataset, in which an English utterance from one of three domains (Calendar, Document, Email) has to be mapped to an API request. We again follow the setup of L'23 and evaluate on their length split, where a model has to generalize to longer logical forms. In contrast to all other datasets we consider, Okapi is quite noisy because it was collected with crowd workers. This presents a realistic additional challenge on top of the challenge of structural generalization.

The results of 5 runs can be found in Table 5. Our model outperforms both BART (Lewis et al., 2020) and the model of L'23. In the comparison without pretraining, our model also consistently achieves higher accuracy than the comparable model of L'23 without grammar-based decoding.

7 Related Work

Predicting permutations. Mena et al. (2018) and Lyu and Titov (2018) use variational autoencoders based on the Sinkhorn algorithm to learn latent permutations. The Sinkhorn algorithm (Sinkhorn, 1964) is also an instance of Bregman’s method and solves the entropy regularized version of Eq. (6) without the \mathbf{W} -term. This parameterization is con-

siderably weaker than ours since it cannot capture our notion of ‘jumps’.

Wang et al. (2021) compute soft permutations as an expected value by marginalizing over the permutations representable by ITGs (Wu, 1997). This approach is exact but excludes some permutations. In particular, it excludes permutations needed for COGS.³ In addition, the algorithm they describe takes a lot of resources as it is both $O(n^5)$ in memory and compute. Devatine et al. (2022) investigate sentence reordering methods. They use bigram scores, which results in a similar computational problem to ours. However, they deal with it by restricting what permutations are possible to enable tractable dynamic programs. Eisner and Tromble (2006) propose local search methods for decoding permutations for machine translation.

Outside of NLP, Kushinsky et al. (2019) have applied Bregman’s method to the quadratic assignment problem, which Eq. (5) is a special case of. Since they solve a more general problem, using their approach for Eq. (6) would require $O(n^4)$ rather than $O(n^3)$ variables in the linear program.

Compositional generalization. Much research on compositional generalization has focused on lexical generalization with notable success (Andreas, 2020; Akyurek and Andreas, 2021; Conklin et al., 2021; Csordás et al., 2021). Structural generalization remains more challenging for seq2seq models (Yao and Koller, 2022).

Zheng and Lapata (2022) modify the transformer architecture and re-encode the input and partially generated output for every decoding step to disentangle the information in the representations. Structure has also been introduced in models by means of grammars: Qiu et al. (2022) heuristically induce a quasi-synchronous grammar (QCFG, Smith and Eisner (2006)) and use it for data augmentation for a seq2seq model. Kim (2021) introduces neural QCFGs which perform well on compositional generalization tasks but are very compute-intensive. Other works directly parse into trees or graphs inspired by methods from syntactic parsing (Liu et al., 2021; Herzig and Berant, 2021; Weißenhorn et al., 2022; Jambor and Bahdanau, 2022; Petit and Corro, 2023).

Several approaches, including ours, have decoupled the presence or absence of output tokens from their order: Wang et al. (2021) train a model end-

to-to-end to permute the input (as discussed above) and then monotonically translate it into an output sequence. Lindemann et al. (2023) also present an end-to-end differentiable model that first applies a ‘fertility step’ which predicts for every word how many copies to make of its representation, and then uses the permutation method of Wang et al. (2021) to reorder the representation before translating them. Cazzaro et al. (2023) first translate the input monotonically and feed it into a second model. They use alignments from an external aligner to train the first model. The second model is a tagger or a pretrained seq2seq model and predicts the output as a permutation of its input. We compare against such a baseline for permutations in Appendix B, finding that it does not work as well as ours in the compositional generalization setups we consider.

8 Conclusion

In this paper, we have presented a flexible new seq2seq model for semantic parsing. Our approach consists of two steps: We first tag each input token with a multiset of output tokens. Then we arrange those tokens into a sequence using a permutation model. We introduce a new method to predict and learn permutations based on a regularized linear program that does not restrict what permutations can be learned. The model we present has a strong ability to generalize compositionally on synthetic and natural semantic parsing datasets. Our results also show that trees are not necessarily required to generalize well to deeper recursion than seen at training time.

Limitations

The conditional independence assumptions are a limitation for the applicability of our multiset tagging model. For example, the independence assumptions are too strong to apply it to natural language generation tasks such as summarization. From a technical point of view, the independence assumptions are important to be able to induce the latent assignment of output tokens to multisets efficiently. Future work may design multiset tagging methods that make fewer independence assumptions.

While our method for predicting permutations is comparatively fast and only has a memory requirement of $O(n^3)$, inference on long sequences, e.g. with more than 100 tokens, remains somewhat

³This is mostly due to the treatment of definite descriptions, which appear clustered together at the start of the logical form.

slow. In future work, we plan to investigate other approximate inference techniques like local search and dual decomposition.

Regarding the importance of trees for compositional generalization, our model has no explicit structural inductive bias towards trees. However, we do not exclude that the pretrained RoBERTa model that we use as a component *implicitly* captures trees or tree-like structures to a certain degree.

Acknowledgements

We thank Bailin Wang and Jonas Groschwitz for technical discussions; we thank Hao Zheng for discussions and for providing system outputs for further analysis. We also say thank you to Christine Schäfer and Agostina Calabrese for their comments on this paper.

ML is supported by the UKRI Centre for Doctoral Training in Natural Language Processing, funded by the UKRI (grant EP/S022481/1), the University of Edinburgh, School of Informatics and School of Philosophy, Psychology & Language Sciences, and a grant from Huawei Technologies. IT is supported by the Dutch National Science Foundation (NWO Vici VI.C.212.053).

References

- Ekin Akyurek and Jacob Andreas. 2021. [Lexicon learning for few shot sequence modeling](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4934–4946, Online. Association for Computational Linguistics.
- Jacob Andreas. 2020. [Good-enough compositional data augmentation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7556–7566, Online. Association for Computational Linguistics.
- Jean-David Benamou, Guillaume Carlier, Marco Cuturi, Luca Nenna, and Gabriel Peyré. 2015. [Iterative bregman projections for regularized transportation problems](#). *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138.
- Lev M Bregman. 1967. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR computational mathematics and mathematical physics*, 7(3):200–217.
- 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.
- Francesco Cazzaro, Davide Locatelli, Ariadna Quattoni, and Xavier Carreras. 2023. [Translate first reorder later: Leveraging monotonicity in semantic parsing](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 227–238, Dubrovnik, Croatia. Association for Computational Linguistics.
- Henry Conklin, Bailin Wang, Kenny Smith, and Ivan Titov. 2021. [Meta-learning to compositionally generalize](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3322–3335, Online. Association for Computational Linguistics.
- Róbert Csordás, Kazuki Irie, and Juergen Schmidhuber. 2021. [The devil is in the detail: Simple tricks improve systematic generalization of transformers](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 619–634, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Róbert Csordás, Kazuki Irie, and Jürgen Schmidhuber. 2022. [The neural data router: Adaptive control flow in transformers improves systematic generalization](#). In *International Conference on Learning Representations*.
- Deborah A. Dahl, Madeleine Bates, Michael Brown, William Fisher, Kate Hunicke-Smith, David Pallett, Christine Pao, Alexander Rudnicky, and Elizabeth Shriberg. 1994. [Expanding the scope of the ATIS task: The ATIS-3 corpus](#). In *Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994*.
- John M. Danskin. 1967. *The Theory of Max-Min and its Application to Weapons Allocation Problems*. Springer Berlin Heidelberg.
- Nicolas Devatine, Caio Corro, and François Yvon. 2022. [Ré-ordonnement via programmation dynamique pour l’adaptation cross-lingue d’un analyseur en dépendances \(sentence reordering via dynamic programming for cross-lingual dependency parsing\)](#). In *Actes de la 29e Conférence sur le Traitement Automatique des Langues Naturelles. Volume 1 : conférence principale*, pages 183–197, Avignon, France. ATALA.
- Andrew Drozdov, Nathanael Schärli, Ekin Akyürek, Nathan Scales, Xinying Song, Xinyun Chen, Olivier Bousquet, and Denny Zhou. 2022. [Compositional semantic parsing with large language models](#). *arXiv preprint arXiv:2209.15003*.
- Jason Eisner and Roy W Tromble. 2006. [Local search with very large-scale neighborhoods for optimal permutations in machine translation](#). In *Proceedings of the HLT-NAACL Workshop on Computationally*

- Hard Problems and Joint Inference in Speech and Language Processing*, pages 57–75.
- Catherine Finegan-Dollak, Jonathan K. Kummerfeld, Li Zhang, Karthik Ramanathan, Sesh Sadasivam, Rui Zhang, and Dragomir Radev. 2018. [Improving text-to-SQL evaluation methodology](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 351–360, Melbourne, Australia. Association for Computational Linguistics.
- Kuzman Ganchev, Joao Graça, Jennifer Gillenwater, and Ben Taskar. 2010. [Posterior regularization for structured latent variable models](#). *The Journal of Machine Learning Research*, 11:2001–2049.
- Jonas Groschwitz, Matthias Lindemann, Meaghan Fowlie, Mark Johnson, and Alexander Koller. 2018. [AMR dependency parsing with a typed semantic algebra](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1831–1841, Melbourne, Australia. Association for Computational Linguistics.
- Jiaqi Guo, Qian Liu, Jian-Guang Lou, Zhenwen Li, Xueqing Liu, Tao Xie, and Ting Liu. 2020. [Benchmarking meaning representations in neural semantic parsing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1520–1540, Online. Association for Computational Linguistics.
- Jonathan Herzig and Jonathan Berant. 2021. [Span-based semantic parsing for compositional generalization](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 908–921, Online. Association for Computational Linguistics.
- Saghar Hosseini, Ahmed Hassan Awadallah, and Yu Su. 2021. [Compositional generalization for natural language interfaces to web apis](#). *arXiv preprint arXiv:2112.05209*.
- Dieuwke Hupkes, Verna Dankers, Mathijs Mul, and Elia Bruni. 2020. [Compositionality decomposed: how do neural networks generalise?](#) *Journal of Artificial Intelligence Research*, 67:757–795.
- Dora Jambor and Dzmitry Bahdanau. 2022. [LAGr: Label aligned graphs for better systematic generalization in semantic parsing](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3295–3308, Dublin, Ireland. Association for Computational Linguistics.
- 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, Online. Association for Computational Linguistics.
- Yoon Kim. 2021. [Sequence-to-sequence learning with latent neural grammars](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 26302–26317. Curran Associates, Inc.
- Eliyahu Kiperwasser and Yoav Goldberg. 2016. [Simple and accurate dependency parsing using bidirectional LSTM feature representations](#). *Transactions of the Association for Computational Linguistics*, 4:313–327.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open source toolkit for statistical machine translation](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- Harold W Kuhn. 1955. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97.
- Yam Kushinsky, Haggai Maron, Nadav Dym, and Yaron Lipman. 2019. [Sinkhorn algorithm for lifted assignment problems](#). *SIAM Journal on Imaging Sciences*, 12(2):716–735.
- 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.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Matthias Lindemann, Alexander Koller, and Ivan Titov. 2023. [Compositional generalisation with structured reordering and fertility layers](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2172–2186, Dubrovnik, Croatia. Association for Computational Linguistics.
- Chenyao Liu, Shengnan An, Zeqi Lin, Qian Liu, Bei Chen, Jian-Guang Lou, Lijie Wen, Nanning Zheng, and Dongmei Zhang. 2021. [Learning algebraic recombination for compositional generalization](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1129–1144, Online. Association for Computational Linguistics.

- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *ArXiv*, abs/1907.11692.
- Chunchuan Lyu and Ivan Titov. 2018. **AMR parsing as graph prediction with latent alignment**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 397–407, Melbourne, Australia. Association for Computational Linguistics.
- Gonzalo Mena, David Belanger, Scott Linderman, and Jasper Snoek. 2018. Learning latent permutations with gumbel-sinkhorn networks. In *International Conference on Learning Representations*.
- Alban Petit and Caio Corro. 2023. **On graph-based reentrancy-free semantic parsing**. *Transactions of the Association for Computational Linguistics*.
- Linlu Qiu, Peter Shaw, Panupong Pasupat, Pawel Nowak, Tal Linzen, Fei Sha, and Kristina Toutanova. 2022. **Improving compositional generalization with latent structure and data augmentation**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4341–4362, Seattle, United States. Association for Computational Linguistics.
- Richard Sinkhorn. 1964. **A relationship between arbitrary positive matrices and doubly stochastic matrices**. *The Annals of Mathematical Statistics*, 35(2):876–879.
- David Smith and Jason Eisner. 2006. **Quasi-synchronous grammars: Alignment by soft projection of syntactic dependencies**. In *Proceedings on the Workshop on Statistical Machine Translation*, pages 23–30, New York City. Association for Computational Linguistics.
- Bailin Wang, Mirella Lapata, and Ivan Titov. 2021. **Structured reordering for modeling latent alignments in sequence transduction**. In *Thirty-Fifth Conference on Neural Information Processing Systems*.
- Pia Weißenhorn, Lucia Donatelli, and Alexander Koller. 2022. **Compositional generalization with a broad-coverage semantic parser**. In *Proceedings of the 11th Joint Conference on Lexical and Computational Semantics*, pages 44–54, Seattle, Washington. Association for Computational Linguistics.
- Dekai Wu. 1997. **Stochastic inversion transduction grammars and bilingual parsing of parallel corpora**. *Computational Linguistics*, 23(3):377–403.
- Zhengxuan Wu, Christopher D Manning, and Christopher Potts. 2023. **Recogs: How incidental details of a logical form overshadow an evaluation of semantic interpretation**. *arXiv preprint arXiv:2303.13716*.
- Yuekun Yao and Alexander Koller. 2022. **Structural generalization is hard for sequence-to-sequence models**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Hao Zheng and Mirella Lapata. 2021. **Compositional generalization via semantic tagging**. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1022–1032, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Hao Zheng and Mirella Lapata. 2022. **Disentangled sequence to sequence learning for compositional generalization**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4256–4268, Dublin, Ireland. Association for Computational Linguistics.

A Math Details

A.1 NP-hardness

We show that Eq. (5) can be used to decide the Hamiltonian Path problem. Let $G = (V, E)$ be a graph with nodes $V = \{1, 2, \dots, n\}$. A Hamiltonian path $P = v_1, v_2, \dots, v_n$ is a path in G (i.e. $(v_i, v_{i+1}) \in E$ for all i) such that each node of G appears exactly once. Deciding if a graph has a Hamiltonian path is NP-complete.

Reduction of Hamiltonian path to Eq. (5).

Note that a necessary but not sufficient condition for P to be a Hamiltonian path is that P is a permutation of V . This will be ensured by the constraints on the solution in Eq. (5).

We construct a score function

$$s_{k \rightarrow i} = \begin{cases} 1 & \text{if } (k, i) \in E \\ 0 & \text{else} \end{cases} \quad (12)$$

and let $s_{i \rightarrow j} = 0$ for all i, j . If we find the solution of Eq. (5) for the score function Eq. (12), we obtain a permutation P of V , which may or may not be a path in G . In a path of n nodes, there are $n - 1$ edges that are crossed. If the score of the solution is $n - 1$, then all node pairs (v_i, v_{i+1}) that are adjacent in P must have had a score of 1, indicating an edge $(v_i, v_{i+1}) \in E$. Therefore, P must be a Hamiltonian path. If the score of the solution is less than $n - 1$, then there is no permutation of V that is also a path, and hence G has no Hamiltonian path.

A.2 Proof of Proposition 2

We now prove Prop. 2 using a very similar technique as Kushinsky et al. (2019). As the constraints

$$\begin{aligned}
& \text{KL}(\mathbf{x} \parallel \mathbf{z}) + \text{KL}(\mathbf{Y}^*(\mathbf{x}) \parallel \mathbf{W}) \\
&= \text{KL}(\mathbf{x} \parallel \mathbf{z}) + \sum_{i,j} \frac{\mathbf{x}_i \mathbf{W}_{i,j}}{\sum_{j'} \mathbf{W}_{i,j'}} (\log \frac{\mathbf{x}_i \mathbf{W}_{i,j}}{\sum_{j'} \mathbf{W}_{i,j'}} - 1) \\
&= \text{KL}(\mathbf{x} \parallel \mathbf{z}) + \sum_i \frac{\mathbf{x}_i \sum_j \mathbf{W}_{i,j}}{\sum_{j'} \mathbf{W}_{i,j'}} (\log \frac{\mathbf{x}_i}{\sum_{j'} \mathbf{W}_{i,j'}} - 1) \\
&= \sum_i \mathbf{x}_i (\log(\frac{\mathbf{x}_i}{\mathbf{z}_i}) - 1 + \log \frac{\mathbf{x}_i}{\sum_{j'} \mathbf{W}_{i,j'}} - 1) \\
&= \sum_i \mathbf{x}_i (2 \log \mathbf{x}_i - \log \mathbf{z}_i - \log \left(\sum_{j'} \mathbf{W}_{i,j'} \right) - 2) \\
&= 2 \sum_i \mathbf{x}_i (\log \mathbf{x}_i - 1 - \underbrace{\frac{1}{2} (\log \mathbf{z}_i + \log \left[\sum_{j'} \mathbf{W}_{i,j'} \right])}_{\log \mathbf{q}_i}) \\
&= 2 \cdot \text{KL}(\mathbf{x} \parallel \mathbf{q})
\end{aligned}$$

Figure 4: Rewriting the objective function of Eq. (14). We note that the generalized KL divergence $\text{KL}(\mathbf{x} \parallel \mathbf{y}) = \sum_i \mathbf{x}_i \log \frac{\mathbf{x}_i}{\mathbf{y}_i} - \mathbf{x}_i + \mathbf{y}_i$ simplifies to $\text{KL}(\mathbf{x} \parallel \mathbf{y}) = \sum_i \mathbf{x}_i (\log \frac{\mathbf{x}_i}{\mathbf{y}_i} - 1)$ because we want to find the argmin wrt to \mathbf{x} .

in Prop. 2 for any value of j do not interact with constraints for other values of j , we can assume w.l.o.g. that j takes a single value only and drop it in the notation. We want to solve:

$$\begin{aligned}
& \arg \min_{\mathbf{x}, \mathbf{Y}} \quad \text{KL}(\mathbf{x} \parallel \mathbf{z}) + \text{KL}(\mathbf{Y} \parallel \mathbf{W}) \\
& \text{subject to} \quad \sum_i \mathbf{x}_i = 1 \\
& \quad \quad \quad \sum_k \mathbf{Y}_{ik} = \mathbf{x}_i \quad \forall i
\end{aligned} \tag{13}$$

We can find $\mathbf{Y}^* = \arg \min_{\mathbf{Y}} \text{KL}(\mathbf{Y} \parallel \mathbf{W})$ subject to $\sum_j \mathbf{Y}_{i,j} = \mathbf{x}_i$ based on Prop. 1: $\mathbf{Y}_{i,j}^* = \frac{\mathbf{x}_i \mathbf{W}_{i,j}}{\sum_j \mathbf{W}_{i,j}}$. That is, we can express \mathbf{Y}^* as a function of \mathbf{x} (which we write $\mathbf{Y}^*(\mathbf{x})$), and therefore our overall problem is now a problem in one variable (\mathbf{x}):

$$\begin{aligned}
& \arg \min_{\mathbf{x}} \quad \text{KL}(\mathbf{x} \parallel \mathbf{z}) + \text{KL}(\mathbf{Y}^*(\mathbf{x}) \parallel \mathbf{W}) \\
& \text{subject to} \quad \sum_i \mathbf{x}_i = 1
\end{aligned} \tag{14}$$

We can now rewrite the objective function as

$$\begin{aligned}
& \arg \min_{\mathbf{x}} \quad \text{KL}(\mathbf{x} \parallel \mathbf{q}) \\
& \text{subject to} \quad \sum_i \mathbf{x}_i = 1
\end{aligned} \tag{15}$$

where $\mathbf{q}_i = \sqrt{\mathbf{z}_i \cdot \sum_{j'} \mathbf{W}_{i,j'}}$. This step is justified in detail in Fig. 4.

The rewritten optimization problem has the right form to apply Prop. 1 a second time. We obtain:

$$\mathbf{x}_i^* = \frac{\mathbf{q}_i}{\sum_{i'} \mathbf{q}_{i'}}$$

By plugging this into $\mathbf{Y}^*(\mathbf{x})$, we obtain the solution to the overall optimization problem.

A.3 Reduction of Eq. (8) to Eq. (6)

In this section, we show how Eq. (8) can be reduced to a problem of the entropy regularized version of Eq. (6). This is useful because it means we can use Algorithm 1 to solve Eq. (8).

First, we show that computing a KL projection is equivalent to solving an entropy-regularized linear program. Let C be the feasible region of the linear constraints.

$$\begin{aligned}
& \arg \max_{\mathbf{x} \in C} \mathbf{s}^T \mathbf{x} - \tau \sum_i \mathbf{x}_i (\log \mathbf{x}_i - 1) \\
&= \arg \min_{\mathbf{x} \in C} \tau \sum_i \mathbf{x}_i (\log \mathbf{x}_i - \frac{\mathbf{s}_i}{\tau} - 1) \\
&= \arg \min_{\mathbf{x} \in C} \text{KL}(\mathbf{x} \parallel \exp(\frac{\mathbf{s}}{\tau}))
\end{aligned}$$

Due to this, Eq. (8) is equivalent to a linear program that has the same feasible region as Eq. (6) except for the additional constraint $\mathbf{U} \in \mathcal{Q}(\mathbf{y}, \mathbf{z}')$. Note that $\mathbf{U} \in \mathcal{Q}(\mathbf{y}, \mathbf{z}')$ essentially rules out certain correspondences. Therefore we can approximately enforce $\mathbf{U} \in \mathcal{Q}(\mathbf{y}, \mathbf{z}')$ by masking $\mathbf{U}^*(s)$ such that any forbidden correspondence receives a very low score.

A.4 Derivation of loss function as ELBO

We now show how the training procedure we use to train our permutation model can be derived from a form of evidence lower bound (ELBO).

Ideally, our permutation model would be a distribution $P_\theta(\mathbf{R}|\mathbf{x}, \mathbf{z}')$ over permutation matrices \mathbf{R} and we would maximize the marginal likelihood, i.e. marginalizing over all permutations:

$$P(\mathbf{y}|\mathbf{x}, \mathbf{z}') = \sum_{\mathbf{R} \in \mathcal{P}} P_\theta(\mathbf{R}|\mathbf{x}, \mathbf{z}') P(\mathbf{y}|\mathbf{z}', \mathbf{R}) \quad (16)$$

where $P(\mathbf{y}|\mathbf{z}', \mathbf{R}) = \prod_j \sum_i R_{ij} \cdot \mathbb{1}(y_j = z_i)$ with $\mathbb{1}$ being the indicator function. $P(\mathbf{y}|\mathbf{z}', \mathbf{R})$ returns 1 iff applying the permutation \mathbf{R} to \mathbf{z}' results in \mathbf{y} . Unfortunately, computing Eq. (16) exactly is intractable in general due to the sum over permutation matrices. We instead use techniques from variational inference and consider the following evidence lower bound (ELBO):

$$\begin{aligned} \log P(\mathbf{y}|\mathbf{x}, \mathbf{z}') &\geq \max_Q \mathbb{E}_{\mathbf{R} \sim Q(\mathbf{R}|\mathbf{x}, \mathbf{z}', \mathbf{y})} \log P(\mathbf{y}|\mathbf{z}', \mathbf{R}) \\ &\quad - \text{KL}(Q(\mathbf{R}|\mathbf{x}, \mathbf{z}', \mathbf{y}) \parallel P_\theta(\mathbf{R}|\mathbf{x}, \mathbf{z}')) \end{aligned} \quad (17)$$

where $Q(\mathbf{R}|\mathbf{x}, \mathbf{z}', \mathbf{y})$ is an approximate variational posterior. We now relax the restriction that $P(\mathbf{R}|\mathbf{x}, \mathbf{z}')$ places non-zero mass only on permutation matrices and use the following definition of $P_\theta(\mathbf{R}|\mathbf{x}, \mathbf{z}')$:

$$P_\theta(\mathbf{R}_{ij} = 1|\mathbf{x}, \mathbf{z}') = \mathbf{U}^*(s)_{ij}$$

where $\mathbf{U}^*(s)$ is the solution to Eq. (6) with added entropy regularization.

It turns out, in our case, we can easily construct a variational posterior Q that has zero reconstruction loss (the first term on the right side in Eq. (17)): we can choose any $Q(\mathbf{R}|\mathbf{x}, \mathbf{z}', \mathbf{y}) \in \mathcal{Q}(\mathbf{y}, \mathbf{z}')$ where $\mathcal{Q}(\mathbf{y}, \mathbf{z}')$ is the set of bistochastic matrices such that $Q(\mathbf{R}|\mathbf{x}, \mathbf{z}', \mathbf{y})_{i,j} = 0$ iff $\mathbf{z}'_i \neq \mathbf{y}_j$. To see that this gives zero reconstruction error, consider position j in the output: The probability mass is distributed across precisely those positions i in \mathbf{z}' where the




Figure 5: Accuracy on the document domain of Okapi by number of conjuncts in the gold logical form.

		BART perm	Ours
COGS	OS	16 \pm 8	33 \pm 24
	CP	17 \pm 3	82 \pm 11
	PP	16 \pm 5	91 \pm 5
ATIS	iid	77.10 \pm 1.61	76.65 \pm 1.67
	Length	23.81 \pm 9.45	41.39 \pm 13.47

Table 6: Performance of our permutation model in comparison to using BART for predicting the permutation. We ignore the order of conjuncts for evaluation.

right kind of token lives. In other words, any alignment with non-zero probability will reconstruct the output token at position j .

Therefore we can use the following lower bound to the log-likelihood:

$$\begin{aligned} \log P(\mathbf{y}|\mathbf{x}, \mathbf{z}') &\geq \\ &\quad - \min_{Q \in \mathcal{Q}(\mathbf{y}, \mathbf{z}')} \text{KL}(Q(\mathbf{R}|\mathbf{x}, \mathbf{z}', \mathbf{y}) \parallel P_\theta(\mathbf{R}|\mathbf{x}, \mathbf{z}')) \end{aligned} \quad (18)$$

During training, we need to compute the gradient of Eq. (18). By Danskin’s theorem (Danskin, 1967), this is:

$$-\nabla_\theta \text{KL}(Q^* \parallel P_\theta(\mathbf{R}|\mathbf{x}, \mathbf{z}')) \quad (19)$$

where $Q^* \in \mathcal{Q}(\mathbf{y}, \mathbf{z}')$ is the minimizer of Eq. (18). Note that Q^* can equivalently be characterised as $\hat{\mathbf{U}}$ (Eq. (8)).

In practice, we also add $-\text{KL}(\hat{\mathbf{W}}|\mathbf{W}^*(s))$ to our objective in Eq. (18) to speed up convergence; this does not change the fact that we use a lower bound.

B Additional results and analysis

Okapi. In Fig. 5 we show the accuracy of our model on the document domain in comparison with previous work by number of conjuncts in the logical form.

Permutation baseline. A simpler approach for predicting a permutation of the output \mathbf{z}' from

the multiset tagging is to use a seq2seq model. In order to compare our approach to such a baseline, we concatenate the original input x with a separator token and z' . We then feed this as input to a BART-base model which is trained to predict the output sequence y . At inference time, we use beam search and enforce the output to be a permutation of the input. As detailed in Table 6, this approach works well in-distribution and it also shows a small improvement over finetuning BART directly on the length split of ATIS. However, it does not perform as well as our approach. On COGS, our model outperforms the permutation baseline by an even bigger margin. Unseen variable symbols could be a challenge for BART on COGS which might explain part of the gap in performance.

This approach towards predicting permutations is similar to that of Cazzaro et al. (2023) except that they do not constrain the beam search to permutations. We found that not constraining the output to be a permutation worked worse in the compositional generalization setups.

C Further model details

C.1 Parametrization of permutation model

We do not share parameters between the multiset tagging model and the permutation model.

Tokens that appear more than once in the same multiset have the same representation \mathbf{h}'_i in Eq. (7). In order to distinguish them, we concatenate another embedding to \mathbf{h}'_i : if the token z'_i is the k -th instance of its type in its multiset, we concatenate an embedding for k to \mathbf{h}'_i . For example, in Fig. 1, $z'_5 = 'x_1'$ and it is the second instance of $'x_1'$ in its multiset, so we use the embedding for 2.

We found it helpful to make the temperature τ of the scores for Algorithm 1 dependent on the number of elements in the permutation, setting $\tau = (\log n)^{-1}$, so that longer sequences have slightly sharper distributions.

Since the permutation model is designed to model exactly permutations, during training, z and y must have the same elements. This is not guaranteed because z is the prediction of the multiset model which may not have perfect accuracy on the training data. For simplicity, we disregard instances where z and y do not have the same elements. In practice, this leads to a very small loss in training data for the permutation model.

C.2 Lexicon mechanism

The lexicon L is a lookup table that deterministically maps an input token x_i to an output token $L(x_i)$, and we modify the distribution for multiset tagging as follows:

$$P'(\mathbf{z}_{i,v} = k|\mathbf{x}) = \begin{cases} P(\mathbf{z}_{i,\mathcal{L}} = k|\mathbf{x}_i) & \text{if } v = L(\mathbf{x}_i) \\ P(\mathbf{z}_{i,v} = k|\mathbf{x}) & \text{else} \end{cases}$$

where $P(\mathbf{z}_{i,v} = k|\mathbf{x})$ is as defined in Eq. (3) and \mathcal{L} is a special lexicon symbol in the vocabulary. $P(\mathbf{z}_{i,\mathcal{L}}|\mathbf{x}_i)$ is a distribution over the multiplicity of $L(\mathbf{x}_i)$, independent of the identity of $L(x_i)$. We use the ‘simple’ lexicon induction method by Akyurek and Andreas (2021). Unless otherwise specified during learning, $L(\mathbf{x}_i) = \mathbf{x}_i$ like in a copy mechanism.

Handling of variables in COGS. For the COGS dataset, a model has to predict variable symbols. The variables are numbered (0-based) by the input token that introduced it (e.g. in Fig. 1, slept, the third token, introduces a variable symbol x_2). In order to robustly predict variable symbols for sentences with unseen length, we use a similar mechanism as the lexicon look up table: we introduce another special symbol in the vocabulary, Var. If Var is predicted with a multiplicity of k at i -th input token, it adds the token x_{i-1} to its multiset k times.

C.3 Initialization of Multiset Tagging model

If there are l alignments with a posterior probability of at least χ that an input token i produces token v , we add the term $\lambda \log P(\mathbf{z}_{i,v} \geq l | \mathbf{x})$ to Eq. (4). λ is the hyperparameter determining the strength. This additional loss is only used during the first g epochs.

D Datasets and Preprocessing

We show basic statistics about the data we use in Table 7. Except for the doubling task, all our datasets are in English. COGS uses a small fragment of English generated by a grammar, see Kim and Linzen (2020) for details.

Doubling task. For the doubling task, we use an alphabet of size $|\Sigma| = 11$. To generate inputs with a specific range of lengths, we first draw a length from the range uniformly at random. The symbols in the input are also drawn uniformly at random and then concatenated into a sequence. Examples of lengths 5 - 10 are used as training, examples of

Dataset	Split/Version	Train	Dev	Test
Doubling		4,000	500	1,000
COGS		24,155	3,000	21,000
ATIS	iid	4,465	497	448
	length	4,017	942	331
Okapi	Calendar	1,145	200	1061
	Document	2,328	412	514
	Email	2,343	200	991

Table 7: Number of examples per dataset/split.

length 11 are used as development data (e.g. for hyperparameter selection), and examples of length 11 - 20 are used as test data.

D.1 Preprocessing

COGS. Unlike [Zheng and Lapata \(2022\)](#); [Qiu et al. \(2022\)](#); [Drozdov et al. \(2022\)](#) we do not apply structural preprocessing to the original COGS meaning representation and keep the variable symbols: all our preprocessing is local and aimed at reducing the length of the logical form (to keep runtimes low). We delete any token in {"", " ", "_", "(", ")", "x", ".", ";", "AND"} as these do not contribute to the semantics and can be reconstructed easily in post-processing. The tokens {"agent", "theme", "recipient", "ccomp", "xcomp", "nmod", "in", "on", "beside"} are always preceded by a "." and we merge "." and any of those tokens into a single token.

Example:

```
* cookie ( x _ 3 ) ; * table ( x _ 6 )
; lend . agent ( x _ 1 , Dylan ) AND
lend . theme ( x _ 1 , x _ 3 ) AND lend .
recipient ( x _ 1 , x _ 9 ) AND cookie .
nmod . beside ( x _ 3 , x _ 6 ) AND girl
( x _ 9 )
```

Becomes

```
* cookie 3 * table 6 lend .agent 1 Dylan
lend .theme 1 3 lend .recipient 1 9 cookie
.nmod .beside 3 6 girl 9
```

ATIS. We follow the pre-processing by [Lindemann et al. \(2023\)](#) and use the variable-free FunQL representation as annotated by [Guo et al. \(2020\)](#). We use spacy 3.0.5 (model en_core_web_sm) to tokenize the input.

Okapi. Again, we follow the preprocessing of [Lindemann et al. \(2023\)](#). We use spacy 3.0.5 (model en_core_web_sm) to tokenize both the input utterances and the output logical forms.

E Details on evaluation metrics

We provide code for all evaluation metrics in our repository.

Doubling. We use exact match accuracy on the string.

COGS. For COGS we use exact match accuracy on the sequence in one evaluation setup. The other evaluation setup disregards the order of conjuncts: we first remove the ‘preamble’ (which contains all the definite descriptions) from the conjunctions. We count a prediction as correct if the set of definite descriptions in the preamble matches the set of definite descriptions in the gold logical form *and* the set of clauses in the prediction match the set of clauses in the gold logical form.

ATIS. We allow for different order of conjuncts between system output and gold parse in computing accuracy. We do this by sorting conjuncts before comparing two trees node by node. This is the same evaluation metric as used by [Lindemann et al. \(2023\)](#).

Okapi. We follow [Hosseini et al. \(2021\)](#); [Lindemann et al. \(2023\)](#) and disregard the order of the parameters for computing accuracy. We use a case-insensitive string comparison.

F Hyperparameters

We use the same hyperparameters for all splits of a dataset. For our model, we only tune the hyperparameters of the multiset tagging model; the permutation model is fixed, and we use the same configuration for all tasks where we use RoBERTa. For model ablations where we use an LSTM instead of RoBERTa, we use the same hyperparameters for Okapi and ATIS, and a smaller model for the doubling task. These configurations were determined by hand without tuning. For BART, we use the same hyperparameter as [Lindemann et al. \(2023\)](#).

We follow the random hyperparameter search procedure of [Lindemann et al. \(2023\)](#) for the multiset tagging models and the LSTM/transformer we train from scratch: we sample 20 configurations and evaluate them on the development set. We run the two best-performing configurations again with a different random seed and pick the one with the highest accuracy (comparing the union of the predicted multisets with the gold multiset). We then train and evaluate our model with entirely different random seeds.

Dataset	Model	First stage	Second stage
Doubling	LSTM	5	-
	Transformer	3	-
	Ours	3	14*
COGS	Ours	20	80
ATIS	Ours	30	50
	Ours/LSTM	30	110
Okapi/Calendar	Ours	9	35
	Ours/LSTM	5	30
Okapi/Email	Ours	12	40
	Ours/LSTM	9	40
Okapi/Document	Ours	12	70
	Ours/LSTM	10	55

Table 8: Average runtime for train/evaluate on dev and test in **minutes**. For the doubling task, we note that our model has converged usually after $\frac{1}{4}$ of the time in the table.

Dataset	Model	First stage	Second stage
Doubling	LSTM		2.462
	Transformer		10.424
	Ours	0.273	0.463
COGS	Ours	125.091	127.16
ATIS	Ours	125.506	127.119
	Ours/LSTM	3.345	1.816
Okapi/Calendar	Ours	124.927	127.03
	Ours/LSTM	1.493	1.876

Table 9: Number of parameters in millions in the models we train. This includes the 124.646m params of RoBERTa when we finetune it.

The chosen hyperparameters along with the search space are provided in the github repository.

G Number of parameters, computing infrastructure and runtime

We show the number of parameters in the models we train in Table 9.

All experiments were run on GeForce GTX 1080 Ti or GeForce GTX 2080 Ti with 12GB RAM and Intel Xeon Silver or Xeon E5 CPUs.

The runtime of one run contains the time for training, evaluation on the devset after each epoch and running the model on the test set. We show runtimes of the model we train in Table 8. Since we evaluate on 5 random seeds (10 for COGS due to high variance of results), our experiments overall took around 64 hours of compute time on our computing infrastructure.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Limitations, at the end of the paper
- A2. Did you discuss any potential risks of your work?
No apparent societal risks.
- A3. Do the abstract and introduction summarize the paper’s main claims?
Abstract and 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

6, Appendix D

- B1. Did you cite the creators of artifacts you used?
6
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
The data we create programmatically is likely too simple to be protected by copyright.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
6
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
6, Appendix D
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
6, Appendix D

C Did you run computational experiments?

6

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Appendix H

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

6, Appendix G, code submission.

C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

6, Appendix B

C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Appendix D

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.