Latent Structure Learning Through Stochastic Invariant

Anonymous Author(s)
Affiliation
Address
email

Abstract

Structured latent variables allow incorporating meaningful prior knowledge into deep learning models. However, learning with such variables remains challenging because of their discrete nature. Nowadays, the standard learning approach is to define a latent variable as a perturbed algorithm output and to use a differentiable surrogate for training. In general, the surrogate puts additional constraints on the model and inevitably leads to biased gradients. As an alternative, we develop score function estimators for a variety of structured latent variables. In particular, we highlight a family of recursive algorithms with a common feature we call stochastic invariant. The feature allows us to construct reliable gradient estimates and control variates without additional constraints on the model. In our experiments, we consider various structured latent variables models and achieve results competitive with relaxation-based counterparts.

1 Introduction

To this day, the majority of deep learning architectures consists of differentiable computation blocks and relies on gradient estimates for learning. At the same time, architectures with discrete intermediate components are a good fit for incorporating inductive biases [2, 45] or dynamic control flow [23, 11]. One of the approaches to learning such architectures is to replace the discrete component with a stochastic latent variable and use the expected objective as a differentiable surrogate for learning.

In practice, the expectation has high computational cost, thus one typically resorts to stochastic estimates for the expectation and its gradient. Particularly, the two prevalent approaches to estimate the gradient of the objective are the score function estimator [42] and the reparameterization trick [15, 37] for relaxed discrete variables [25, 14]. The former puts mild assumptions on the distribution and the objective, requiring the gradient of log-probability with respect to the distribution parameters, and provides unbiased estimates for the objective gradient. However, the naive estimate suffers from high variance and is less intuitive in implementation. In comparison, the reparameterized gradient estimates seamlessly integrate within the backpropagation algorithm and exhibit low variance out of the box. At the same time, relaxation requires an architecture defined on the extended domain of the relaxed variable and introduces bias to the gradient estimate.

In the recent years, the attention of the community shifted towards models with structured latent variables. Informally, a structured variable models a distribution over structured objects such as graphs [3, 31], sequences [8] or matchings [27]. A latent variable may alter the computation graph or represent a generative process of data. Often, a structured variable is represented as a sequence of categorical random variables with a joint distribution incorporating the structure constraints (e.g., the fixed number of edges in an adjacency matrix of a tree). Recent works on structured latent variables address learning largely through the reparameterization trick using relaxed variables. In fact, the Gumbel-softmax trick naturally translates to structured variables when $\arg\max$ operator is applied

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

over a structured domain rather than component-wise [31]. In contrast, score function estimators are now less common in structured domain, with a few exceptions such as [47, 13]. The primary difficulty is the sample score function: neither Gibbs distributions, nor distribution defined through a generative process have a general shortcut to compute it.

In our work, we develop a framework to define structured variables along with a low-variance score function estimator. Our goal is to allow learning models that do not admit relaxed variables and to improve optimization by alleviating the bias of the relaxed estimators. To achieve the goal we define the structured variable as an output of an algorithm with a perturbed input. Then, we outline a family of algorithms with a common property we call *stochastic invariant*. The property was inspired by the observation in [31, Appendix, Sec. B], where the authors showed that the Kruskal's algorithm [21] and the CLE algorithm [6] are recursively applying the Gumbel-Max trick. We construct new algorithms with the same property and show how to use the property to learn structured latent variables. In the experimental section, we report performance on par with relaxation-based methods and apply the framework in a setup that does not allow relaxations.

2 Gradient Estimation for Discrete Variables

We study the estimates of the gradients of an expected objective $\nabla_\lambda \mathbb{E}_X \mathcal{L}(X)$, where a discrete random variable X has parametric distribution $\mathbb{P}_X(\cdot;\lambda)$ with a finite support. As a convention, we use capital letters to denote random variables and small letters to denote the variable realization. In this section, we briefly introduce the two approaches to gradient estimation crucial for our exposition.

The basic *score function estimator* [42], also known as REINFORCE, defines an unbiased estimate for the gradient using sample x as $\mathcal{L}(x)\nabla_\lambda \log \mathbb{P}_X(X = x;\lambda)$. The estimator does not make any assumptions about $\mathcal{L}(\cdot)$, but requires an efficient sampling algorithm for X and the score function $\nabla_\lambda \log \mathbb{P}_X(X = x;\lambda)$. For a categorical random variable X with outcome probabilities $\mathbb{P}_X(X = k;\lambda) \propto \lambda_k$ computation of $\nabla_\lambda \log \mathbb{P}_X(X = x;\lambda)$ is linear in the number of outcomes d. Therefore the gradient estimation is fast when d is small, whereas the combinatorial number of outcomes may complicate gradient estimation for structured variables.

The Gumbel-Softmax estimator introduced in [14, 25] is an alternative estimator based on

Lemma 1. *(Exponential-min trick) If $E_i \sim \mathrm{Exp}(\lambda_i), i \in \{1,\dots,d\}$ are independent, then for $X := \mathrm{argmin}_i\, E_i$*

1. *the outcome probability is $\mathbb{P}_X(X = x;\lambda) \propto \lambda_x$;*
2. *random variables $E'_i := E_i - E_X, i \in \{1,\dots,d\}$ are mutually independent given X with $E'_i \mid X \sim \mathrm{Exp}(\lambda_i)$ when $i \neq X$ and $E'_i = 0$ otherwise.*[1]

On the forward pass, the Gumbel-Softmax estimator replaces the categorical variable X with a differentiable surrogate $\tilde{x} = \mathrm{soft\,max}(\frac{g}{\tau})$, where the input $g := (-\log e_1, \dots, -\log e_d)$ is a component-wise transformation of exponential samples. Due to Lemma 1, the surrogate converges the one-hot encoding of a categorical sample x as $\mathrm{soft\,max}$ converges to $\arg\max$ when $\tau \to 0$. On the backward pass, the estimator uses the chain rule to construct the gradient $\nabla_\lambda \mathcal{L}(\tilde{x})$ using the reparameterization trick [15, 37] to define the partial derivative of a sample as $\frac{\partial e_i}{\partial \lambda_i} = -\frac{e_i}{\lambda_i}$. The components of g are distributed according to the Gumbel distribution with location parameters $\theta_i := \log \lambda_i$, hence the name. Although the formulation in term of the Gumbel distribution parameterized with θ is more common, for convenience of the exposition in Section 3 we will mainly use exponential random variables E with parameters λ.

As opposed to the score function estimator, the Gumbel-Softmax estimator requires a differentiable loss $\mathcal{L}(\cdot)$. Such requirement imposes an additional restriction on a model architecture. The architecture must accept relaxed samples as well as hard samples, a non-trivial requirement for the models where discrete variables define branching [23] or distribution factorization [17, 12]. Also, the two estimators perform differently in practice. Compared to the Gumbel-Softmax estimator, the vanilla score function estimator has notoriously higher variance and requires a control variate to improve the gradient descent convergence. We will use two strategies to construct the control variates. The first strategy is another application of Lemma 1. It defines a family of sample-dependent

[1] As a convention, we will assume that $0 \stackrel{d}{=} \mathrm{Exp}(\infty)$.

control variates for the score function estimator [40, 10]. With X and E defined as in Lemma 1, the estimator generates samples e, x, adds a control variate $c(e)$ and uses an independent sample $\tilde{e}$ from the conditional distribution $E \mid X = x$ to eliminate the introduced bias

$$(\mathcal{L}(x) - c(\tilde{e}))\, \nabla_\lambda \log \mathbb{P}_X(x; \lambda) - \nabla_\lambda c(\tilde{e}) + \nabla_\lambda c(e). \tag{1}$$

The latter two terms are the reparameterized gradients of the control variate. In general, the above estimate extends to any pair of random variables (X, E) that admits the conditional reparameterization trick for $E \mid X = x$. In [40], the control variate used the relaxed loss $\mathcal{L}(\cdot)$, whereas [10] proposed to learn the control variate to improve the training dynamic. In our experiments, we will use the estimator of [10] and refer to it as RELAX.

The second family of control variates we consider uses $K > 1$ samples $x_1, \ldots, x_K$ to reduce the variance. Besides averaging the independent estimates, it uses the objective sample mean $\bar{\mathcal{L}} := \frac{\sum_{i=1}^{K} \mathcal{L}(x_i)}{K}$ to reduce the variance even further:

$$\frac{1}{K-1} \sum_{i=1}^{K} \left(\mathcal{L}(x_i) - \overline{\mathcal{L}}\right) \nabla_\lambda \log \mathbb{P}_X(X = x_i; \lambda). \tag{2}$$

Despite being quite simple, the above leave-one-out estimator [18] proved to be competitive with multiple recent works [5, 38]. In our experiments, we refer to such estimator as REINFORCE+. To facilitate fair comparison, in a batch training setup we reduce the batch size proportionally to K.

3 Algorithms With Stochastic Invariants

In our work, we consider perturbation models for structured latent variables. Specifically, we consider a vector E of i.i.d. exponential random variables with the rates $\lambda \in \mathbb{R}_+^d$ and define the structured variable X as the output of an algorithm with the perturbed input E. Assuming the algorithm has a specific recursive form, we introduce an additional variable T corresponding to the execution trace of the algorithm and present score function estimators based on the variable. We highlight the central points of the section in Table 1.

Table 1: Similarly to the perturbed input E, the trace T has tractable score function, yet the score function estimator yields lower variance. Such estimator is a practical alternative to an estimator using X, since the score of X involves computationally demanding marginalization. Additionally, E admits conditional reparameterization given T.

	Input E	Execution Trace $T = T(E)$	Output $X = X(T)$
Score $\nabla_\lambda \log \mathbb{P}(\cdot)$	Tractable	Tractable	Intractable
REINFORCE variance	High	Moderate	Low
Conditional $p_E(E \mid \cdot)$	$\delta(E - e)$	$\sum \mathrm{Exp}(\sum \lambda_k)$	Intractable

3.1 Stochastic Invariant in $\arg\operatorname{top} k$

We start with an algorithm for finding top-k elements to illustrate our framework. We present the recursive form of $\arg\operatorname{top} k$ in Algorithm 1. For each recursion level, the algorithm finds the minimum element, decrements k and calls itself to find the subset excluding the minimum variable. For reasons explained below, the algorithm subtracts the minimum from the sequence $E'_j = E_j - E_T$ before the recursion. This step does not change the output and may seem redundant.

Assuming the input of the algorithm is a vector E of i.i.d. exponential variables with rate parameters λ, the first argument in the recursive call E' is again a vector of i.i.d. exponential variables (given T) due to Lemma 1. In other words, the input distribution class is *invariant* throughout the recursion. Besides that, for each recursion level variable T has categorical distribution (conditioned on T found in the above calls) with output probabilities proportional to $\lambda_k, k \in K$.

We will use upper indices to denote the recursion depth and, with a slight abuse of notation, denote the concatenation of variables T for each recursion depth as $T := (T^1, \ldots, T^k)$. The output X is a set and does not take into account the order in T. Intuitively, T acts as *the execution trace* of the

Algorithm 1 $F_{\text{top-k}}(E, K, k)$ - finds k smallest elements in a sequence E, where K is the set of indices *(keys)* of E

Input: E, K, k
Output: X
  if $k = 0$ **then**
    return
  end if
  {Find the smallest element}
  $T \Leftarrow \arg\min_{j \in K} E_j$
  for $j \in K$ **do**
    $E'_j \Leftarrow E_j - E_T$
  end for
  {Exclude $\arg\min$ index T, decrement k}
  $K', k' \Leftarrow K \setminus \{T\}, k - 1$
  $E' \Leftarrow \{E'_k \mid k \in K'\}$
  {Solve for $k' = k - 1$}
  $X' \Leftarrow F_{\text{top-k}}(E', K', k')$
  return $\{T\} \cup X'$

Algorithm 2 $F_{\text{struct}}(E, K, R)$ - returns structured variable X based on utilities E indexed by K and an auxiliary variable R

Input: E, K, R
Output: X
  if $f_{\text{stop}}(K, R)$ **then**
    return
  end if
  $P_1, \ldots, P_m \Leftarrow f_{\text{split}}(K, R)$ $\{\sqcup_{i=1}^m P_i = K\}$
  for $i = 1$ to m **do**
    $T_i \Leftarrow \arg\min_{j \in P_i} E_j$
    for $j \in P_i$ **do**
      $E'_j \Leftarrow E_j - E_{T_i}$
    end for
  end for
  $K', R' \Leftarrow f_{\text{map}}(K, R, \{T_i\}_{i=1}^m)$ $\{K' \subsetneq K\}$
  $E' \Leftarrow \{E'_k \mid k \in K'\}$
  $X' \Leftarrow F_{\text{struct}}(E', K', R')$ {Recursive call}
  return $f_{\text{combine}}(X', K, R, \{T_i\}_{i=1}^m)$

Figure 1: The recursive algorithm for $\arg\operatorname{top} k$ and the general algorithm with the stochastic invariant put side-by-side. Both algorithm perform the exponential-min trick and proceed with recursion using a subset of variables. The output X combines the current trace T and the recursion output X'.

algorithm, whereas X contains only partial information about T. The marginal probability of x is the sum $\mathbb{P}_X(X = x; \lambda) = \sum_{t \in X^{-1}(x)} \mathbb{P}_T(T = t; \lambda)$ over all possible orderings of $x = \{x_1, \ldots, x_k\}$ denoted as $X^{-1}(x)$. The direct computation of such sum is prohibitive even for moderate k.

The $\arg\operatorname{top} k$ illustration is a well-known extension of the exponential-min trick. In particular, the distribution of T is known as the Plackett-Luce distribution [34] and coincides with k categorical samples without replacement. Following the recursion, the observation probability factorizes according to the chain rule with i-th factor governed by equation $\mathbb{P}_{T_i}(T_i = t_i \mid t_1, \ldots, t_{i-1}; \lambda) = \frac{\lambda_{t_i}}{\sum_{j=1}^d \lambda_j - \sum_{j=1}^{i-1} \lambda_{t_j}}$

3.2 General Algorithm With the Stochastic Invariant

We put the general Algorithm 2 next to Algorithm 1 for comparison. Again, our generalization aims to preserve the stochastic invariant of the algorithm: the first argument E of the algorithm always has exponential distribution (conditioned on the trace T from the above recursion levels). As before, we pass the indices of E as a separate variable K. In general, K may be an arbitrary finite set. We also introduce an auxiliary variable R to accumulate the information from the above recursion calls.

The core feature of the generalization is that Algorithm 2 applies the exponential-min trick simultaneously over m disjoint sets rather than the whole index set K. Specifically, we construct *a partition* of indices $P_1, \ldots, P_m$ and find the $\arg\min$ across the partition sets. To generate the partition, we introduce a new sub-routine f_{split} taking the index set K and the auxiliary variable R as inputs. The partition size m may also be variable. After the m simultaneous exponential-min tricks, the algorithm constructs arguments K' and R' passed in the recursive call. Whereas the top-k algorithm excluded a single variable, the generalized algorithm calls f_{map} to select a subset of variables K' and to accumulate the necessary information for the next call in variable R'. After the recursive call the algorithm calls f_{combine} to construct the structured variable X using the recursive call output X' and the other variables. In the top-k example, f_{combine} appended the minimum variable index T to the set X'. Finally, we introduce a sub-routine f_{stop} to define the stop condition of the recursion.

Importantly, the algorithm only accesses E values though the exponential-min trick. As a result, the intermediate variables K' and R' as well as the output X depend on E only through T. In other words, the execution trace is a function of perturbation $T = T(E)$ and the structured variable $X = X(T)$ is a function of trace. Additionally, due to Lemma 1, the trace components $T_1, \ldots, T_m$

have categorical distributions, whereas $E'_k, k \in K$ are exponential random variables. We prove these properties by induction w.r.t. the recursion depth in Appendix B.

Given the above, we arrive to two modifications of Algorithm 2 generalizing Lemma 2 and the Plackett-Luce distribution from the illustration. Algorithm 3 computes the log-probability $\log \mathbb{P}_T(t; \lambda)$ of a trace realization t. We will use the algorithm output to construct score function estimators. Again, the pseudo-code introduces index j to denote the recursion depth and assumes the input $t = \{t_i^j\}_{i,j}$ is the concatenation of trace variables for all recursion depths $j = 1, \ldots, k$. Similarly, in Appendix C we present an algorithm returning a sample from $E \mid T = t$ given trace realization t. We prove the correctness of the two modifications by induction in Section B.

3.3 Further Examples

This subsection contains an overview of algorithms with stochastic invariants along with the corresponding structured variables. We present the details and the pseudo-code in Appendix F.

Analogous to the $\arg\operatorname{top} k$ and the subset variable, the insertion sorting algorithm is an algorithm with the stochastic invariant. In the case of sorting, we do not omit the order of the trace variable T and return the permutation $X = T$. The resulting variable X has the Plackett-Luce distribution. We use the variable as a latent variable for insertion-based non-monotonic generation [12]. As an alternative to the Plackett-Luce distribution, we consider a square parameter matrix and find a matching between rows and columns. We perturb the matrix and iteratively find the minimum element in the matrix. We exclude the row and the column containing the element and proceed to the next recursion step. Notably, in contrast to this algorithm, the Hungarian algorithm [22] for the minimum weight matching does not have the stochastic invariant.

As [31] observe, Kruskal's algorithm [21] and Chu-Liu-Edmonds [6] algorithm recursively apply the exponential-min trick, thus have the stochastic invariant. The former constructs the minimum spanning (MST) tree edge-by-edge. The corresponding trace variable T is a sequence of edges, whereas X is an unordered set of edges. Interestingly, we could not represent Prim's algorithm [35, 4] for the MST as an algorithm with the stochastic invariant. The Chu-Liu-Edmonds algorithm is an analog of Kruskal's algorithm for directed graphs. It returns the minimum tree X with a fixed root.

Additionally, we construct a latent binary tree variable specifically for the non-monotonic generation model [41]. In this model, each token in the sentence corresponds to a node in the tree. We assign weights to nodes and perturb the weights to obtain a sample. During the recursion, we set the minimum weight node to be the parent. We put the nodes on the left-hand side to the left subtree and the nodes on the right-hand side to the right subtree.

4 Gradient Estimation Using the Stochastic Invariant

We have introduced a probabilistic model involving an exponential variable E and the structured variable X defined as an output of an algorithm with input E. Additionally, we have defined an intermediate trace variable $T = T(E)$ such that X is a function $X = X(T)$. In this section, we will apply T to estimate grardients of $\mathbb{E}_X \mathcal{L}(X)$.

In our setup, the score function $\nabla_\lambda \log \mathbb{P}_E(E = e; \lambda)$ is available out of the box. However, the score function estimator

$$g_E := \mathcal{L}(X(e)) \nabla_\lambda \log \mathbb{P}_E(E = e; \lambda), \tag{3}$$

which we will refer to as E-REINFORCE, is rarely used in practice. In fact, the variance of the score function estimator using E as an action space exceeds the variance of the estimator using X. On the other hand, the score function for the structured variable $\nabla_\lambda \log \mathbb{P}_X(X = x; \lambda)$ involves marginalization w.r.t. E and may require significant computation resources to estimate.

To mitigate the variance of g_E, we follow the observation of [40, Appendix, Sec. B] and define another estimator as

$$g_T := \mathcal{L}(X(t)) \nabla_\lambda \log \mathbb{P}_T(T = t; \lambda), \tag{4}$$

which is essentially the score function estimator that uses T rather than E. Below we will refer to it as T-REINFORCE. Such estimate can be seen as the score function estimator g_E marginalized over E given T (Appendix B contains the detailed derivation)

$$\mathcal{L}(X) \nabla_\lambda \log \mathbb{P}_T(T; \lambda) = \mathbb{E}_{E|T} \left[\mathcal{L}(X) \nabla_\lambda \log \mathbb{P}_E(E; \lambda) \mid T \right]. \tag{5}$$

Algorithm 3 $F_{\text{log-prob}}(t, \lambda, K, R)$ - returns $\log \mathbb{P}_T(t; \lambda)$ for trace t, rates λ, K and R as in Alg. 2

Input: t, λ, K, R
Output: $\log \mathbb{P}_T(t; \lambda)$
  if $f_{\text{stop}}(K, R)$ **then**
    return
  end if
  $P_1, \ldots, P_m \Leftarrow f_{\text{split}}(K, R)$
  for $i = 1$ to m **do**
    $\log \mathbb{P}_T(t_i^1; \lambda) \Leftarrow \log \lambda_{t_i^1} - \log \left(\sum_{k \in P_i} \lambda_k\right)$ {Index j in T_i^j denotes the recursion level}
    for $k \in P_i \setminus \{t_i^1\}$ **do**
      $\lambda'_k \Leftarrow \lambda_k$
    end for
    $\lambda'_{t_i^1} \Leftarrow +\infty$ {Because $E'(t_i^1) = 0$}
  end for
  $K', R' \Leftarrow f_{\text{map}}(K, R, \{t_i^1\}_{i=1}^m)$
  $\lambda' \Leftarrow \{\lambda'_k \mid k \in K'\}$
  $\log \mathbb{P}_T(t^{>1} \mid T^1 = t^1; \lambda) \Leftarrow F_{\text{log-prob}}(t^{>1}, \lambda', K', R')$ {Compute log-prob of $t^{>1} := \{t_i^j\}_{j>1}$}
  return $\sum_{i=1}^m \log \mathbb{P}_T(t_i^1; \lambda) + \log \mathbb{P}_T(t^{>1} \mid T^1 = t^1; \lambda)$

As a result, the proposed gradient estimate g_T is unbiased

$$\mathbb{E}_E \mathcal{L}(X) \nabla_\lambda \log \mathbb{P}_E(E; \lambda) = \mathbb{E}_T \mathbb{E}_{E|T} \left[\mathcal{L}(X) \nabla_\lambda \log \mathbb{P}_E(E; \lambda) \mid T\right], \tag{6}$$

whereas the variance of the estimate does not exceed the variance of g_E

$$\text{var}_E \left[g_E\right] = \text{var}_T \left[\mathbb{E}_{E|T} g_E\right] + \mathbb{E}_T \, \text{var}_{E|T} \left[g_E\right] = \text{var}_T \left[g_T\right] + \mathbb{E}_T \, \text{var}_{E|T} \left[g_E\right] \geq \text{var}_T \left[g_T\right]. \tag{7}$$

In fact, in our experiments, we observed a significant difference in optimization due to the reduced variance of the estimator.

As we've argued in the previous section, we can compute the score function for trace variable and apply the estimator g_T in practice. Similarly, marginalization with respect to $T \mid X$ leads to the score function estimator $g_X := \mathcal{L}(x) \nabla_\lambda \log \mathbb{P}_X(X = x; \lambda)$ and reduces the variance even further $g_T : \text{var}_T \left[g_T\right] \geq \text{var}_X \left[g_X\right]$. Therefore, the standard score function estimator is preferable when $\nabla_\lambda \log \mathbb{P}_X(X = x; \lambda)$ is available. In other cases, g_T is a practical alternative.

Besides that, the above arguments also apply to estimators with control variates such as NVIL [28] or REINFORCE+.[2] Finally, the algorithm for conditional reparameterization of $E \mid T$ (Appendix C, Algorithm 6) allows to generalize the action-dependent control variates of [40, 10] to algorithms with the stochastic invariant

$$g_{\text{RELAX}} := \left(\mathcal{L}(x) - c(\tilde{e})\right) \nabla_\lambda \log \mathbb{P}_T(T = t; \lambda) - \nabla_\lambda c(\tilde{e}) + \nabla_\lambda c(e), \tag{8}$$

where e, t, x are realizations of E, T, X respectively and $\tilde{e}$ is a sample from $E \mid T = t$.

5 Related Work

Early models with structured latent variables include HMMs [36], PCFGs [33], make strong assumptions about the model structure, and typically use EM-algorithm variations for training. This paper continues the line of work on perturbation models [30] for distributions over combinatorial sets. Initially, perturbation models approximated Gibbs distributions with an efficient sampling procedure using the MAP oracle for the Gibbs distribution. Later, [14, 25] proposed to relax the component-wise optimization used in the Gumbel-Max trick to facilitate gradient-based learning for Gibbs distributions. The combination of the two approaches, namely a perturbed model together with a bespoke relaxation of the MAP oracle, allows designing learning algorithms for latent subsets [43], permutations [27], trees [3] and sequences [8]. Recently, [31] developed a systematic approach to design relaxations for perturbed models with linear MAP oracle. Unlike the previous works, we mainly focus on the score function estimators [42] for learning.

[2] In Eq.2, we denote multi-sample counterparts of g_T and g_E as T-REINFORCE+ and E-REINFORCE+

Table 2: Results of k-subset selection on Aroma aspect data. MSE ($\times 10^{-2}$) and subset precision (%) is shown for best models selected on validation averaged across different random seeds.

Model	Estimator	$k = 5$		$k = 10$		$k = 15$	
		mean ± std	Prec.	mean ± std	Prec.	mean ± std	Prec.
Simple	*SST (Our Impl.)*	$\mathbf{3.6 \pm 0.1}$	28 ± 1.4	3.21 ± 0.12	29.5 ± 1.7	2.77 ± 0.09	$\mathbf{28.1 \pm 1.2}$
	E-REINFORCE+	3.89 ± 0.2	25 ± 1.4	3.77 ± 0.23	26.7 ± 3.4	3.16 ± 0.16	25.3 ± 1.3
	T-REINFORCE+	3.79 ± 0.13	$\mathbf{30.5 \pm 2.2}$	$\mathbf{3.14 \pm 0.16}$	$\mathbf{31 \pm 2.9}$	$\mathbf{2.69 \pm 0.11}$	27.6 ± 0.9
	RELAX	3.76 ± 0.11	24 ± 1.9	3.5 ± 0.13	28.9 ± 1.9	2.95 ± 0.15	26.1 ± 1.9
Complex	*SST (Our Impl.)*	2.93 ± 0.09	$\mathbf{56 \pm 2.1}$	2.55 ± 0.08	$\mathbf{49.4 \pm 3.1}$	2.51 ± 0.05	40.3 ± 0.9
	E-REINFORCE+	3.03 ± 0.2	49.4 ± 3.3	2.92 ± 0.12	45 ± 2.3	2.76 ± 0.22	42 ± 2.2
	T-REINFORCE+	$\mathbf{2.75 \pm 0.08}$	55.8 ± 3.4	$\mathbf{2.48 \pm 0.05}$	48.6 ± 2.4	$\mathbf{2.4 \pm 0.03}$	$\mathbf{44.2 \pm 1.3}$
	RELAX	2.8 ± 0.08	54.4 ± 2.1	2.58 ± 0.09	47.6 ± 1.6	2.46 ± 0.08	42.1 ± 1.5

We illustrate the framework with the well-known Gumbel-Top-k trick [46]. Among the various applications, [44] used the trick to define a differentiable relaxation for the subset model considered in our paper; meanwhile, [9, 39] used the trick to define score function estimators for latent permutations. Besides that, [19, 20] leveraged the trick for sampling without replacement for a certain family of graphical models and design a gradient estimator using the sampler [18]. Importantly, [31, Appendix, Sec. B] showed that the Kruskal's algorithm and the Chu-Liu-Edmonds algorithm extend the Gumbel-Top-k trick. They used the observation to argue in favor of exponential perturbations. In turn, we generalize the observation and propose a learning algorithm based on the generalization.

The conditional reparameterization scheme, proposed in our work, allows action-dependent control variates [40, 10] for learning the structured variables. However, similarly to [38, 5] we observed that often a simpler leave-one-out baseline [18] has better performance. Besides the control variates, [32] recently adopted conditional reparameterization to construct an improved Gumbel Straight-Through estimator [14] for the categorical latent variables.

6 Applications

In the section below, we study various instances of algorithms with stochastic invariants and compare them against relaxation-based approaches. In practice, the recursive form of Algorithm 2 does not facilitate batch parallelization. Specifically, the recursion depth and the decrease in size of E may differ within a batch. Therefore, Algorithms 2,3 may require further optimization. We discuss our implementations in Appendix, Section D.

6.1 Learning to Explain by Finding a Fixed Size Subset

We evaluated our method on the experimental setup from L2X [1] on the BeerAdvocate [26] dataset. The dataset consists of textual beer reviews and numerical ratings for four beer aspects (*Aroma*, *Taste*, *Appearance*, *Palate*). The model utilizes encoder-decoder architecture. The encoder outputs parameters of top-k distribution over subsets of a review given the entire review. The decoder predicts a rating given the review subset of size k. Following setup of [31] we use $k = \{5, 10, 15\}$ and two CNN architectures: *Simple* — a one-layer CNN and *Complex* — a three-layer CNN and train our models for each aspect using MSE loss. We train models using several score function based estimators and compare them with our implementation of SST. We report means and standard deviations for loss and precision averaged across different random model initializations. Table 2 shows the obtained results for *Aroma* aspect. The best metrics with respect to means are highlighted in bold. Detailed experimental setup and description of models can be found in Appendix E.1

6.2 Learning Latent Spanning Trees with Kruskal's Algorithm

Given a system of interacting particles, the dependencies of their states can be formally described as an undirected graph. We use Neural Relational Inference [16], initially representing a relaxation-based approach, and examine its performance as a generative model and ability to reconstruct system interconnections, when applying score function techniques instead. We build an experiment in

Table 3: Graph Layout experiment results for T=10 iterations. Metrics are obtained by choosing models with best validation ELBO and averaging results across different random seeds on the test set.

ESTIMATOR	$T = 10$					
	ELBO		EDGE PREC.		EDGE REC.	
	MEAN $\pm$ STD	MAX	MEAN $\pm$ STD	MAX	MEAN $\pm$ STD	MAX
SST (Our Impl.)	-1860.79 ± 1116.83	-1374.81	$\mathbf{87 \pm 21}$	**95**	$\mathbf{93 \pm 3}$	**95**
T-REINFORCE	$\mathbf{-1582.72 \pm 571.18}$	$\mathbf{-1192.04}$	70 ± 31	91	86 ± 8	91
RELAX	-2079.18 ± 569.49	-1205.87	43 ± 31	90	81 ± 6	90

Table 4: Unsupervised Parsing on ListOps. We report the average test-performance of the model with the best validation accuracy across different random initializations.

ESTIMATOR	ACCURACY		PRECISION		RECALL	
	MEAN $\pm$ STD	MAX	MEAN $\pm$ STD	MAX	MEAN $\pm$ STD	MAX
SST (Our Impl.)	78.42 ± 8.14	**93.78**	56.84 ± 20.08	**82.40**	30.18 ± 19.10	73.11
E-REINFORCE+	60.25 ± 2.29	64.47	40.87 ± 6.90	45.74	40.74 ± 6.93	45.46
T-REINFORCE+	$\mathbf{87.34 \pm 3.00}$	91.97	$\mathbf{77.93 \pm 7.36}$	79.65	$\mathbf{61.10 \pm 14.11}$	**79.65**
RELAX	79.60 ± 9.36	88.64	54.73 ± 17.48	75.27	53.61 ± 17.14	75.27

line with [31], generating data by translating a ground truth latent spanning tree (corresponding to the connections in the system) into a sequence of positions of points on a real plane, representing dynamics of the system over time. These points are obtained executing force-directed algorithm [7] for $T = 10$ or $T = 20$ iterations and fed into the model.

The architecture of the model consists of a graph neural network (GNN) encoder, producing distribution over spanning trees, and a GNN decoder, producing distribution over time series of points positions. Model is trained in a manner of variational autoencoders (VAEs), optimizing ELBO, a lower bound on the joint log-likelihood of the observed data points at all timesteps.

We measure precision and recall with respect to the encoder samples and the ground truth dependency spanning tree. Table 3 shows the results for T=10 iterations. Overall, score function methods performed better than their relaxation-based counterpart, achieving higher values of ELBO on the test set, but slightly worse performance in terms of structure recovery metrics. The results for T=20 and the detailed experimental setup are described in the Appendix E.2.

6.3 Unsupervised Parsing with Rooted Trees with CLE Algorithm

We study the ability of the proposed score function estimators to recover the latent structure of the data in a setting, where it can be quite accurately described with an arborescence. Following details about data and models outlined by [31], we use a simplified version of the ListOps [29] dataset. It consists of mathematical expressions (e.g. `min[3 med[3 5 4] 2]`), written in prefix form along with results of their evaluation, which are integers in $[0, 9]$. Given a prefix expression, one can algorithmically recover its structure as a parse tree. We bound maximal length of expressions and maximal depth of their parses along with removing the examples with `summod` operator. These limitations sufficiently decrease the amount of memory a model should have to calculate the result and facilitates the usage of GNNs which now become capable of evaluating expressions by a bounded number of message passing steps.

Our model consists of two parts: an encoder and a classifier. The encoder is a pair of LSTMs that generate parameters of the distribution over rooted arborescence on token nodes. The classifier is a GNN, which passes a fixed number of messages over the sampled arborescence and feeds the resulting embedding of the first token into the final MLP. Models are trained simultaneously to minimize cross-entropy. We examine the performance of the models by measuring classification accuracy along with precision and recall with respect to the edges of ground truth parse trees. Table 4 shows score function based estimators, particularly T-REINFORCE+, show more stable performance

in comparison to relaxation-based estimator. Detailed description of the experiment can be found in the Appendix E.3.

6.4 Non-monotonic Generation of Balanced Parentheses with Binary Trees

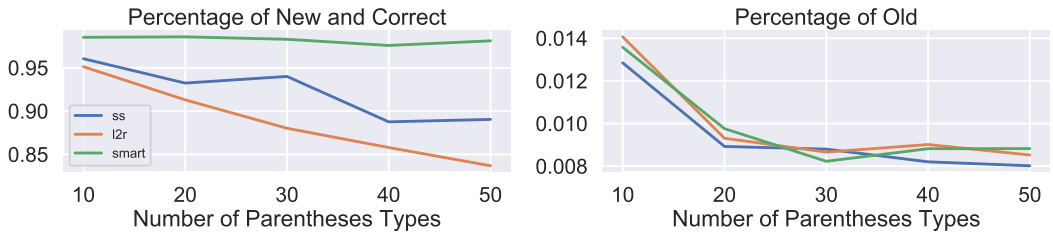

Figure 2: Generative statistics of non-monotonic language model with different orders

We apply our methods in a setting where the application of continuous relaxations is seemingly not possible. We construct a simple dataset by sampling a fixed number of balanced parentheses of various types (from 10 to 50 types) and model them with non-monotonic architecture, defined in [41]. Here, a language model generates sequences by consecutively adding new tokens in between previously generated ones. This process can be seen as modeling joint distribution over binary trees (treated as orders of generation) and sequences (which are generated along these orders).

We refer to this probabilistic model as the decoder, fix its architecture to be a single-layer LSTM and use teacher-forcing for training. More precisely, we compare two different modes of training decoder: fixed order and semi-supervised trainable order. In the first case, the order is obtained by applying a deterministic function to the input sequence. We study two fixed orders: *left-to-right*, corresponding to a degenerate tree with only right child nodes, and a more natural *smart* order, described in Appendix E.4. In the semi-supervised (*ss*) case, 10% of sequences are paired with the smart order, and the overall model is trained as a VAE. We train the decoder by directly backpropagating the output signal and obtain the gradients for encoder using RELAX estimator.

We choose models with the best validation perplexity and evaluate them by generating 200000 unique sequences and measuring portion of those which are present in the train dataset (Old) and those which are balanced and not seen during training (New and Correct). Results of this toy experiment show that it is possible to improve generative metrics by considering non-trivial orders of generation. Experiment is described in details in the Appendix E.4.

7 Conclusion

In this work, we proposed a framework aimed at the design and inference of structured latent variables. We defined the variables as an output of an algorithm with the perturbed input. Then we derived a marginalized score function estimator for the model assuming a specific form of the underlying algorithm. Besides that, we derived a sample-dependent control variate for the estimator. Unlike the relaxation-based gradient estimators, our approach does not require the objective to be a differentiable function of the discrete variable. In the experiments, we have shown that the estimator can achieve the results competitively or better than the relaxation-based counterparts. Notably, the multi-sample estimator with a smaller batch is more stable than the single-sample estimator with a sample-dependent control variate. As a direction for future improvements, we suggest applying the conditional reparameterization scheme used in the control variate to improve the Gumbel straight-through estimators.

Broader Impact

Our work is attempt to broaden the applicability of latent variable models. Due to the fundamental nature of the research, we do not expect any direct harm. Although we cannot rule out indirect harm altogether, in the future latent variable models may improve the interpretability and overall transparency of machine learning applications.

References

[1] Jianbo Chen, Le Song, Martin Wainwright, and Michael Jordan. Learning to explain: An information-theoretic perspective on model interpretation. In *International Conference on Machine Learning*, pages 883–892. PMLR, 2018.

[2] Jean-Baptiste Cordonnier, Aravindh Mahendran, Alexey Dosovitskiy, Dirk Weissenborn, Jakob Uszkoreit, and Thomas Unterthiner. Differentiable patch selection for image recognition. *arXiv preprint arXiv:2104.03059*, 2021.

[3] Caio Corro and Ivan Titov. Differentiable perturb-and-parse: Semi-supervised parsing with a structured variational auto-encoder. In *International Conference on Learning Representations*, 2019.

[4] Sanjoy Dasgupta, Christos H Papadimitriou, and Umesh Virkumar Vazirani. *Algorithms*. McGraw-Hill Higher Education New York, 2008.

[5] Zhe Dong, Andriy Mnih, and George Tucker. Coupled gradient estimators for discrete latent variables.

[6] Jack Edmonds. Optimum branchings. *Journal of Research of the national Bureau of Standards B*, 71(4):233–240, 1967.

[7] Thomas MJ Fruchterman and Edward M Reingold. Graph drawing by force-directed placement. *Software: Practice and experience*, 21(11):1129–1164, 1991.

[8] Yao Fu, Chuanqi Tan, Bin Bi, Mosha Chen, Yansong Feng, and Alexander M Rush. Latent template induction with gumbel-crfs. *arXiv preprint arXiv:2011.14244*, 2020.

[9] Artyom Gadetsky, Kirill Struminsky, Christopher Robinson, Novi Quadrianto, and Dmitry Vetrov. Low-variance black-box gradient estimates for the plackett-luce distribution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 10126–10135, 2020.

[10] Will Grathwohl, Dami Choi, Yuhuai Wu, Geoffrey Roeder, and David Duvenaud. Backpropagation through the void: Optimizing control variates for black-box gradient estimation. *arXiv preprint arXiv:1711.00123*, 2017.

[11] Alex Graves. Adaptive computation time for recurrent neural networks. *arXiv preprint arXiv:1603.08983*, 2016.

[12] Jiatao Gu, Qi Liu, and Kyunghyun Cho. Insertion-based decoding with automatically inferred generation order. *Transactions of the Association for Computational Linguistics*, 7:661–676, 2019.

[13] Serhii Havrylov, Germán Kruszewski, and Armand Joulin. Cooperative learning of disjoint syntax and semantics. *arXiv preprint arXiv:1902.09393*, 2019.

[14] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.

[15] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[16] Thomas Kipf, Ethan Fetaya, Kuan-Chieh Wang, Max Welling, and Richard Zemel. Neural relational inference for interacting systems. In *International Conference on Machine Learning*, pages 2688–2697. PMLR, 2018.

[17] Wouter Kool, Herke van Hoof, and Max Welling. Attention, learn to solve routing problems! In *International Conference on Learning Representations*, 2018.

[18] Wouter Kool, Herke van Hoof, and Max Welling. Buy 4 reinforce samples, get a baseline for free! 2019.

[19] Wouter Kool, Herke Van Hoof, and Max Welling. Stochastic beams and where to find them: The gumbel-top-k trick for sampling sequences without replacement. In *International Conference on Machine Learning*, pages 3499–3508. PMLR, 2019.

[20] Wouter Kool, Herke van Hoof, and Max Welling. Ancestral gumbel-top-k sampling for sampling without replacement. *Journal of Machine Learning Research*, 21(47):1–36, 2020.

[21] Joseph B Kruskal. On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of the American Mathematical society*, 7(1):48–50, 1956.

[22] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.

[23] Tuan Anh Le, Adam R Kosiorek, N Siddharth, Yee Whye Teh, and Frank Wood. Revisiting reweighted wake-sleep for models with stochastic control flow. In *Uncertainty in Artificial Intelligence*, pages 1039–1049. PMLR, 2020.

[24] Tao Lei, Regina Barzilay, and Tommi Jaakkola. Rationalizing neural predictions. *arXiv preprint arXiv:1606.04155*, 2016.

[25] Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.

[26] Julian McAuley, Jure Leskovec, and Dan Jurafsky. Learning attitudes and attributes from multi-aspect reviews. In *2012 IEEE 12th International Conference on Data Mining*, pages 1020–1025. IEEE, 2012.

[27] Gonzalo Mena, David Belanger, Scott Linderman, and Jasper Snoek. Learning latent permutations with gumbel-sinkhorn networks. *arXiv preprint arXiv:1802.08665*, 2018.

[28] Andriy Mnih and Karol Gregor. Neural variational inference and learning in belief networks. In *International Conference on Machine Learning*, pages 1791–1799. PMLR, 2014.

[29] Nikita Nangia and Samuel R Bowman. Listops: A diagnostic dataset for latent tree learning. *arXiv preprint arXiv:1804.06028*, 2018.

[30] George Papandreou and Alan L Yuille. Perturb-and-map random fields: Using discrete optimization to learn and sample from energy models. In *2011 International Conference on Computer Vision*, pages 193–200. IEEE, 2011.

[31] Max B Paulus, Dami Choi, Daniel Tarlow, Andreas Krause, and Chris J Maddison. Gradient estimation with stochastic softmax tricks. *arXiv preprint arXiv:2006.08063*, 2020.

[32] Max B Paulus, Chris J Maddison, and Andreas Krause. Rao-blackwellizing the straight-through gumbel-softmax gradient estimator. *arXiv preprint arXiv:2010.04838*, 2020.

[33] Slav Petrov and Dan Klein. Discriminative log-linear grammars with latent variables. In *Advances in neural information processing systems*, pages 1153–1160. Citeseer, 2008.

[34] Robin L Plackett. The analysis of permutations. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 24(2):193–202, 1975.

[35] Robert Clay Prim. Shortest connection networks and some generalizations. *The Bell System Technical Journal*, 36(6):1389–1401, 1957.

[36] Lawrence R Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.

[37] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pages 1278–1286. PMLR, 2014.

[38] Lorenz Richter, Ayman Boustati, Nikolas Nüsken, Francisco JR Ruiz, and Ömer Deniz Akyildiz. Vargrad: A low-variance gradient estimator for variational inference. *arXiv preprint arXiv:2010.10436*, 2020.

[39] Valentino Santucci, Josu Ceberio, and Marco Baioletti. Gradient search in the space of permutations: an application for the linear ordering problem. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference Companion*, pages 1704–1711, 2020.

[40] George Tucker, Andriy Mnih, Chris J Maddison, John Lawson, and Jascha Sohl-Dickstein. Rebar: Low-variance, unbiased gradient estimates for discrete latent variable models. In *Advances in Neural Information Processing Systems*, pages 2627–2636, 2017.

[41] Sean Welleck, Kianté Brantley, Hal Daumé Iii, and Kyunghyun Cho. Non-monotonic sequential text generation. In *International Conference on Machine Learning*, pages 6716–6726. PMLR, 2019.

[42] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.

[43] Sang Michael Xie and Stefano Ermon. Reparameterizable subset sampling via continuous relaxations. In *IJCAI*, 2019.

[44] Yujia Xie, Hanjun Dai, Minshuo Chen, Bo Dai, Tuo Zhao, Hongyuan Zha, Wei Wei, and Tomas Pfister. Differentiable top-k operator with optimal transport. *arXiv preprint arXiv:2002.06504*, 2020.

[45] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057. PMLR, 2015.

[46] John I Yellott Jr. The relationship between luce's choice axiom, thurstone's theory of comparative judgment, and the double exponential distribution. *Journal of Mathematical Psychology*, 15(2):109–144, 1977.

[47] Dani Yogatama, Phil Blunsom, Chris Dyer, Edward Grefenstette, and Wang Ling. Learning to compose words into sentences with reinforcement learning. *arXiv preprint arXiv:1611.09100*, 2016.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] In Section 3.3, we report that our framework does not cover certain common optimization algorithms. We start Section 6 discussing the issues of integrating recursive algorithms with frameworks aimed at batch-parallel computations.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section7
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Sections 3,4
   (b) Did you include complete proofs of all theoretical results? [Yes] See Appendix B
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] We plan to make code available in the near future.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 6 and the Appendix

(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Section 6. Excluding the parenthesis experiment, where we report a graph containing statistics across multiple models

(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Appendix

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

   (a) If your work uses existing assets, did you cite the creators? [No] We consider synthetic setup and did not find the reference for the BeerAdvocate dataset. We cite the creators of ListOps

   (b) Did you mention the license of the assets? [No]

   (c) Did you include any new assets either in the supplemental material or as a URL? [No]

   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]