

# Adaptive Perturbation-Based Gradient Estimation for Discrete Latent Variable Models

Pasquale Minervini<sup>1, 2</sup>, Luca Franceschi<sup>2</sup>, Mathias Niepert<sup>3</sup>

<sup>1</sup> School of Informatics, University of Edinburgh, Edinburgh, United Kingdom

<sup>2</sup> UCL Centre for Artificial Intelligence, London, United Kingdom

<sup>3</sup> University of Stuttgart, Stuttgart, Germany  
p.minervini@ed.ac.uk

## Abstract

The integration of discrete algorithmic components in deep learning architectures has numerous applications. Recently, Implicit Maximum Likelihood Estimation, a class of gradient estimators for discrete exponential family distributions, was proposed by combining implicit differentiation through perturbation with the path-wise gradient estimator. However, due to the finite difference approximation of the gradients, it is especially sensitive to the choice of the finite difference step size, which needs to be specified by the user. In this work, we present Adaptive IMLE (AIMLE), the first adaptive gradient estimator for complex discrete distributions: it adaptively identifies the target distribution for IMLE by trading off the density of gradient information with the degree of bias in the gradient estimates. We empirically evaluate our estimator on synthetic examples, as well as on Learning to Explain, Discrete Variational Auto-Encoders, and Neural Relational Inference tasks. In our experiments, we show that our adaptive gradient estimator can produce faithful estimates while requiring orders of magnitude fewer samples than other gradient estimators.

## Introduction

There is a growing interest in end-to-end learnable models incorporating discrete algorithms that allow, e.g. to sample from discrete latent distributions (Jang, Gu, and Poole 2017; Paulus et al. 2020) or solve combinatorial optimisation problems (Pogancic et al. 2020; Mandi et al. 2020; Niepert, Minervini, and Franceschi 2021). These discrete components are not continuously differentiable, and an important problem is to efficiently estimate the gradients of their inputs to perform backpropagation. Reinforcement learning, discrete Energy-Based Models (EBMs, LeCun et al. 2006), learning to explain (Chen et al. 2018), discrete Variational Auto-Encoders (VAEs, Kingma and Welling 2014), and discrete world models (Hafner et al. 2020) are additional examples of neural network-based architectures that require the ability to back-propagate through expectations of discrete probability distributions.

The main challenge these approaches have in common is the problem of (approximately) computing gradients of an expectation of a continuously differentiable function  $f$ :

$$\nabla_{\theta} \mathbb{E}_{z \sim p(z; \theta)} [f(z)], \quad (1)$$

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

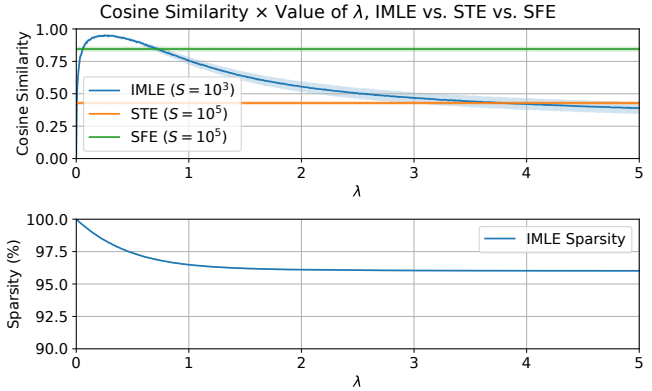


Figure 1: (Top) cosine similarity between the true and estimated gradients  $\nabla_{\theta} \mathbb{E}_{z \sim p(z; \theta)} [\|z - \mathbf{b}\|^2]$ , with  $\mathbf{b} \sim \mathcal{N}(0, \mathbf{I})$  and  $z = \{0, 1\}^{50}$  such that  $\sum_i z_i = 1$ , where estimates are computed using IMLE (Niepert, Minervini, and Franceschi 2021) with  $S = 10^3$  samples, the Straight-Through Estimator (STE, Bengio, Léonard, and Courville 2013) with  $S = 10^5$  samples, and the Score Function Estimator (SFE, Williams 1992) with  $S = 10^5$  samples, and (bottom) sparsity (% of zero elements) of the estimate IMLE gradient – results were averaged across 128 seeds. For  $\lambda \rightarrow 0$ , the IMLE gradient estimate is  $\mathbf{0}$ , while increasing  $\lambda$  leads to increasingly more biased gradient estimates.

where the expectation is taken over a complex discrete probability distribution with intractable marginals and normalisation constant.

In principle, one could use the Score Function Estimator (SFE, Williams 1992). Unfortunately, it suffers from high variance typically exacerbated by the distribution  $p(z; \theta)$  being intractable.

Implicit Maximum Likelihood Estimation (IMLE, Niepert, Minervini, and Franceschi 2021), a recently proposed general-purpose gradient estimation technique, has shown lower variance and outperformed other existing methods, including the score function estimator and problem-specific continuous relaxations, in several settings (Niepert, Minervini, and Franceschi 2021; Betz et al. 2021; Qian et al. 2022). For instance, for the synthetic problem in Fig. 1, the gradient estimate produced by SFE based on  $10^5$  samples is worse

than the estimate based on  $10^3$  — two orders of magnitude fewer — samples using IMLE due to the high variance of the SFE. IMLE combines Perturb-and-MAP sampling with a finite difference method for implicit differentiation originally developed for loss functions defined over marginals (Domke 2010). In IMLE, gradients are approximated as:

$$\nabla_{\theta} \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z}; \theta)} [f(\mathbf{z})] \approx \sum_{i=1}^n \frac{1}{\lambda} \left\{ \text{MAP}(\theta + \epsilon_i) - \text{MAP}(\theta + \epsilon_i - \lambda \nabla_{\mathbf{z}_i} f(\mathbf{z}_i)) \right\}, \quad (2)$$

where  $\text{MAP}(\theta)$  is a maximum-probability state of the distribution  $p(\mathbf{z}; \theta)$ ,  $\epsilon_i \sim \rho(\epsilon)$  is a perturbation drawn from a noise distribution, and  $\mathbf{z}_i = \text{MAP}(\theta + \epsilon_i)$ . Computing MAP the states  $\text{MAP}(\theta)$  instead of sampling from  $p(\mathbf{z}; \theta)$  is especially interesting since, in many cases, it has lower computational complexity than sampling the corresponding distribution (Niepert, Minervini, and Franceschi 2021).

Crucially, the parameter  $\lambda$  determines the step size of the finite difference approximation. When the input to  $f$  is  $p(\mathbf{z}; \theta)$ 's continuously differentiable marginals, we have that smaller values of  $\lambda$  lead to less biased estimates. Hence, in this setting,  $\lambda$  is typically set to a value that depends on the machine precision to void numerical instabilities (Domke 2010). In the setting, we consider, however, that the input to  $f$  is discrete and discontinuous. Setting  $\lambda$  to a very small value, in this case, results in zero gradients. This is illustrated in Fig. 2 (right) for the forward difference method. Hence, the crucial insight is that  $\lambda$  trades off the bias and sparsity of the gradient approximation. In Fig. 1 (top) and (bottom), we plot, respectively, the bias and the sparsity of the gradient estimates for different values of  $\lambda$  on a toy optimisation problem. As we can see, larger values of  $\lambda$  result in a higher bias, and low values of  $\lambda$  result in gradient estimates almost always being zero. With this paper, we propose to make the parameter  $\lambda$  adaptive. We also provide empirical results showing that making  $\lambda$  adaptive reduces the bias and improves the results on several benchmark problems. All source code and datasets are available at <https://github.com/EdinburghNLP/torch-adaptive-imle>.

## Problem Definition

We consider the problem of computing the gradients of an expectation over a discrete probability distribution of a continuously differentiable function  $f$ , that is,

$$\nabla_{\theta} \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z}; \theta)} [f(\mathbf{z})] \quad (3)$$

where  $p(\mathbf{z}; \theta)$  is a discrete probability distribution over binary vectors  $\mathbf{z}$  and with parameters  $\theta$ . Specifically, we are concerned with settings where  $p(\mathbf{z}; \theta)$  is a discrete probability distribution with an intractable normalisation constant. Moreover, we assume that the function  $f$  is a parameterized non-trivial continuously differentiable function which makes existing approaches such as direct loss minimisation and perturbed optimizers (Berthet et al. 2020) not directly applicable.

More formally, let  $\theta \in \Theta \subseteq \mathbb{R}^m$  be a real-valued parameter vector. The probability mass function (PMF) of a discrete

constrained exponential family r.v. is:

$$p(\mathbf{z}; \theta) = \begin{cases} \exp(\langle \mathbf{z}, \theta \rangle - A(\theta)) & \text{if } \mathbf{z} \in \mathcal{C}, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

Here,  $\langle \cdot, \cdot \rangle$  is the inner product.  $A(\theta)$  is the log-partition function, defined as  $A(\theta) = \log(\sum_{\mathbf{z} \in \mathcal{C}} \exp(\langle \mathbf{z}, \theta \rangle))$ , and  $\mathcal{C}$  is an integral polytope of feasible configurations  $\mathbf{z}$ . We call  $\langle \mathbf{z}, \theta \rangle$  the *weight* of the state  $\mathbf{z}$ . The *marginals* (expected value, mean) of the r.v.s  $\mathbf{Z}$  are defined as  $\mu(\theta) := \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z}; \theta)}[\mathbf{z}]$ . Finally, the most probable states also referred to as the *Maximum A-Posteriori* (MAP) states, are defined as  $\text{MAP}(\theta) := \arg \max_{\mathbf{z} \in \mathcal{C}} \langle \mathbf{z}, \theta \rangle$ .

## AIMLE: Adaptive Implicit Maximum-Likelihood Learning

We base our estimator on a finite difference method for implicit differentiation (Domke 2010; Niepert, Minervini, and Franceschi 2021), which is generally applicable to any discrete distribution as defined in Eq. (4). For the initial derivation, we assume that we can compute exact samples from the distribution  $p(\mathbf{z}; \theta)$  using Perturb-and-MAP (Papandreou and Yuille 2011) with noise distribution  $\rho(\epsilon)$ . We write  $\theta + \epsilon$  for a perturbation of the parameters  $\theta$  by a sample  $\epsilon$  from a noise distribution  $\rho(\epsilon)$ . Since, in general, this is not possible for complex distributions, using approximate Perturb-and-MAP samples introduces a bias in the gradient estimates. For now, however, we assume that these perturbations are exact, i.e., that  $\text{MAP}(\theta + \epsilon) \sim p(\mathbf{z}; \theta)$ . Under these assumptions, and by invoking the law of the unconscious statistician (Mohamed et al. 2019), we obtain:

$$\nabla_{\theta} \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z}; \theta)} [f(\mathbf{z})] = \nabla_{\theta} \mathbb{E}_{\epsilon \sim \rho(\epsilon)} [f(\text{MAP}(\theta + \epsilon))].$$

We now approximate  $\text{MAP}(\theta + \epsilon)$  by  $\mu\left(\frac{\theta + \epsilon}{\tau}\right)$  for a small  $\tau > 0$ . That is, we replace a MAP state with a corresponding vector representing the marginal probabilities where we can make the probabilities increasingly spikier by lowering a temperature parameter  $\tau$ . The approximation error between these two terms can be made arbitrarily small since:

$$\mathbb{E}_{\epsilon \sim \rho(\epsilon)} [f(\text{MAP}(\theta + \epsilon))] = \mathbb{E}_{\epsilon \sim \rho(\epsilon)} \left[ f \left( \lim_{\tau \rightarrow 0} \mu \left( \frac{\theta + \epsilon}{\tau} \right) \right) \right].$$

The above equality holds almost everywhere if the noise distribution is such that the probability of two or more components of  $\theta + \epsilon$  being equal is zero. This is the case in the standard setting where  $\rho(\epsilon) \sim \text{Gumbel}(0, 1)$  (Papandreou and Yuille 2011). Therefore, we can write that, for some  $\tau > 0$ ,

$$\nabla_{\theta} \mathbb{E}_{\epsilon \sim \rho(\epsilon)} [\text{MAP}(\theta + \epsilon)] \approx \nabla_{\theta} \mathbb{E}_{\epsilon \sim \rho(\epsilon)} \left[ f \left( \mu \left( \frac{\theta + \epsilon}{\tau} \right) \right) \right].$$

Writing the expectation as an integral, we have:

$$\begin{aligned} & \nabla_{\theta} \mathbb{E}_{\epsilon \sim \rho(\epsilon)} \left[ f \left( \mu \left( \frac{\theta + \epsilon}{\tau} \right) \right) \right] \\ &= \nabla_{\theta} \int_{\mathbb{R}} p(\epsilon) f \left( \mu \left( \frac{\theta + \epsilon}{\tau} \right) \right) d\epsilon. \end{aligned} \quad (5)$$

Now we can exchange differentiation and integration since, for finite  $\tau > 0$ ,  $f$  and  $\mu$  are continuously differentiable:

$$\begin{aligned} \nabla_{\theta} \mathbb{E}_{\epsilon \sim \rho(\epsilon)} \left[ f \left( \mu \left( \frac{\theta + \epsilon}{\tau} \right) \right) \right] & \quad (6) \\ &= \int_{\mathbb{R}} p(\epsilon) \nabla_{\theta} f(\mu) d\epsilon \quad \text{with} \quad \mu := \mu \left( \frac{\theta + \epsilon}{\tau} \right) \\ &= \mathbb{E}_{\epsilon \sim \rho(\epsilon)} [\nabla_{\theta} f(\mu)] \\ &= \mathbb{E}_{\epsilon \sim \rho(\epsilon)} \left[ \lim_{\lambda \rightarrow 0} \frac{1}{\lambda} \left\{ \mu - \mu \left( \frac{\theta + \epsilon}{\tau} - \lambda \nabla_{\mu} f(\mu) \right) \right\} \right]. \end{aligned}$$

The last equality uses implicit differentiation by perturbation (Domke 2010), which is a finite difference method for computing the gradients of a function defined on marginals. Finally, we again approximate the expression  $\mu = \mu \left( \frac{\theta + \epsilon}{\tau} \right)$  with  $z := \text{MAP}(\theta + \epsilon)$  and obtain:

$$\begin{aligned} \nabla_{\theta} \mathbb{E}_{z \sim p(z; \theta)} [f(z)] & \approx \quad (7) \\ \mathbb{E}_{\epsilon \sim \rho(\epsilon)} \left[ \lim_{\lambda \rightarrow 0} \frac{1}{\lambda} \left\{ z - \text{MAP}(\theta + \epsilon - \lambda \nabla_z f(z)) \right\} \right]. \end{aligned}$$

Again, the approximation error of the above expression is arbitrarily small (but not zero) because the derivation shown here is valid for any  $\tau > 0$ . A finite sample approximation of Eq. (7) results in the IMLE gradient estimator of Eq. (2) given in the introduction. While we could, in principle, use the marginals as input to the function  $f$  as in relaxed gradient estimators (Maddison, Mnih, and Teh 2017; Jang, Gu, and Poole 2017; Paulus et al. 2020), computing marginals for the complex distributions we consider here is not tractable in general, and we have to use approximate Perturb-and-MAP samples.

## An Adaptive Optimiser for Finite-Difference based Implicit Differentiation

An important observation that motivates the proposed adaptive version of IMLE is that we need to choose a hyperparameter  $\lambda \in \mathbb{R}_+$  for Eq. (2). Choosing a very small  $\lambda$  leads to most gradients being zero. Consequently, the gradients being back-propagated to the upstream model  $f_v$  are zeros, which prevents the upstream model from being trained. If we choose  $\lambda$  too large, we obtain less sparse gradients, but the gradients are also more biased. Hence, we propose an optimiser that adapts  $\lambda$  during training to trade off non-zero but biased and sparse but unbiased gradients. Similar to adaptive first-order optimizers in deep learning, we replace a single hyperparameter with a set of new ones but show that we obtain consistently better results when using default hyperparameters for the adaptive method.

**Normalisation of the perturbation strength.** Our first observation is that the magnitude of the perturbation  $\lambda$  in the direction of the negative downstream gradient in Eq. (7) highly depends on  $\theta$ , the gradients of the downstream function  $f$ . To mitigate the variations in the downstream gradients norm relative to the parameters  $\theta$ , we propose to set a perturbation magnitude (the norm of the difference between  $\theta$  and the perturbed  $\theta$ ) to be a fraction of the norm of the parameter

vector  $\theta$ . In particular, let  $\alpha \geq 0$  be such a fraction, then we seek  $\lambda$  such that:

$$\begin{aligned} \|\theta - \theta'\|_2 = \alpha \|\theta\|_2 & \Leftrightarrow \lambda \|\nabla_z f(z)\|_2 = \alpha \|\theta\|_2 \\ & \Leftrightarrow \lambda = \alpha \frac{\|\theta\|_2}{\|\nabla_z f(z)\|_2}. \end{aligned} \quad (8)$$

This way, we ensure that a global value for  $\lambda$  roughly translates to the same input-specific magnitude of the perturbation in the direction of the negative gradient.

## Trading off Bias and Sparsity of the Gradient Estimates.

For computing  $\lambda$  as in Eq. (8), we track the sparsity of the gradient estimator with an exponential moving average of the gradient norm. Since the gradients – i.e. the difference between the two MAP states – in Eq. (2) for each  $i$  are always in  $\{-1, 0, 1\}$ , we take the  $L_0$ -norm which is here equivalent to the number of non-zero gradients.

Consider a batch of  $N$  inputs during training. Let  $\hat{\nabla}_{\theta}(j) := \text{MAP}(\theta_j + \epsilon_j) - \text{MAP}(\theta_j + \epsilon_j - \lambda \nabla_{z_j} f(z_j))$  with  $z_j = \text{MAP}(\theta_j + \epsilon_j)$  be a single-sample gradient estimate from Eq. (2) for an input data point  $j \in \{1, \dots, N\}$  with input parameters  $\theta_j$  and without the scaling factor  $1/\lambda$ . We compute, in every training iteration  $t > 0$  and using a discount factor  $0 < \gamma \leq 1$ , the exponential moving average of the number of non-zero gradients per training example:

$$\bar{g}_{t+1} = \gamma \bar{g}_t + (1 - \gamma) \frac{1}{N} \sum_{j=1}^N \|\hat{\nabla}_{\theta}(j)\|_0. \quad (9)$$

Similarly to adaptive optimisation algorithms for neural networks, we introduce an update rule for  $\lambda$ . Let  $c$  be the desired learning rate, expressed as the *number of non-zero gradients per example*. This is the target learning rate, that is, the desired number of gradients we obtain on average per example. A typical value is  $c = 1.0$ , meaning that we aim to adapt the value for  $\lambda$  to obtain, on average, at least one non-zero gradient per example. We now use the following update rule for some fixed  $\eta > 0$ :

$$\alpha_{t+1} = \begin{cases} \alpha_t + \eta & \text{if } \bar{g}_{t+1} \leq c, \\ \alpha_t - \eta & \text{otherwise.} \end{cases} \quad (10)$$

Hence, by increasing or decreasing  $\alpha$  by a constant factor and based on the current exponential moving average of the gradient sparsity, we adapt  $\lambda$  through Eq. (8), which relates  $\lambda$  and  $\alpha$ . Algorithm 1 lists the gradient estimator as a layer in a neural network with a forward and backward pass.

## Forward and centred finite difference approximation.

Gradient estimation in IMLE, as outlined in Eq. (2), is analogous to gradient estimation with forward (one-sided) finite difference approximations, where  $[g(x+h) - g(x)]/h \approx g'(x)$ . A better approximation can be obtained by the *centred* (two-sided) *difference formula*  $[g(x+h) - g(x-h)]/2h$ , which is a second-order approximation to the first derivative (Olver 2013). Following this intuition, we replace  $\frac{1}{\lambda} [\text{MAP}(\theta + \epsilon) - \text{MAP}(\theta + \epsilon - \lambda \nabla_z f(z))]$  in Eq. (2) with  $\frac{1}{2\lambda} [\text{MAP}(\theta + \epsilon + \lambda \nabla_z f(z)) - \text{MAP}(\theta + \epsilon - \lambda \nabla_z f(z))]$ , leading to the update equation in Algorithm 1.

---

**Algorithm 1** Central Difference Perturbation-based Adaptive Implicit Maximum Likelihood Estimation (AIMLE).

---

**function** INIT $\alpha \leftarrow 0$  // Initial value of  $\alpha$   
 $\bar{g} \leftarrow 1$  // Initial gradient norm estimate  
 $\eta \leftarrow 10^{-3}$  // Update step for  $\alpha$ **function** FORWARDPASS( $\theta$ )// Sample from the noise distribution  $\rho(\epsilon)$   
 $\epsilon_1, \dots, \epsilon_N \sim \rho(\epsilon)$   
// MAP states of perturbed  $\theta$   
 $z_i \leftarrow \text{MAP}(\theta + \epsilon_i)$ , for  $i = 1, \dots, N$   
**save**  $\theta, \epsilon_1, \dots, \epsilon_N$ , and  $z_1, \dots, z_N$   
**return**  $z_1, \dots, z_N$ **function** BACKWARDPASS( $\nabla_{z_1} f(z_1), \dots, \nabla_{z_N} f(z_N)$ )**load**  $\theta, \epsilon_1, \dots, \epsilon_N$ , and  $z_1, \dots, z_N$   
 $\lambda = \alpha \frac{1}{N} \sum_{i=1}^N \frac{\|\theta\|_2}{\|\nabla_{z_i} f(z_i)\|_2}$  (see Eq. (8))**for**  $i = 1, \dots, N$  **do** $\theta'_{R_i} \leftarrow \theta - \lambda \nabla_{z_i} f(z_i)$   
 $\theta'_{L_i} \leftarrow \theta + \lambda \nabla_{z_i} f(z_i)$   
 $g_i \leftarrow \frac{1}{2\lambda} [\text{MAP}(\theta'_{L_i} + \epsilon_i) - \text{MAP}(\theta'_{R_i} + \epsilon_i)]$ 

// Moving average of the gradient norm

 $\bar{g} \leftarrow 0.9 \bar{g} + 0.1 \frac{1}{N} \sum_{i=1}^N \|g_i\|_0$   
// Update  $\alpha$  to make  $\bar{g}$  closer to  $c$   
 $\alpha \leftarrow [\alpha + (\eta \text{ if } \bar{g} \leq c \text{ else } -\eta)]_+$ **return**  $\frac{1}{N} \sum_{i=1}^N g_i$ 

---




Figure 2: Finite difference approximation of a downstream function  $f$  on continuous marginals (left) and discrete samples (right). The step size  $\lambda$  trades off bias and sparsity of the gradient approximations for discrete samples and we propose to make the step size  $\lambda$  adaptive.

## Related Work

**Continuous relaxations.** Several works address the gradient estimation problem for discrete random variables, often resorting to continuous relaxations. Maddison, Mnih, and Teh (2017); Jang, Gu, and Poole (2017) propose the Gumbel-Softmax (or concrete) distribution to relax categorical random variables, which was extended by Paulus et al. (2020) to more complex probability distributions. The Gumbel-Softmax distribution only directly applies to categorical variables: for more complex distributions, one has to come up with tailor-made relaxations, or use the STE or SFE – e.g., see Kim, Sabharwal, and Ermon (2016) and Grover et al. (2019). REBAR (Tucker et al. 2017) and RELAX (Grathwohl et al. 2018) use parameterized control variates based on continuous relaxations for the SFE. In this work, we focus explicitly on problems where *only* discrete samples are used during training. Furthermore, REBAR is tailored to categorical distributions, while IMLE and AIMLE are intended for models with complex distributions and multiple constraints. Approaches that do not rely on relaxations are specific to certain distributions (Bengio, Léonard, and Courville 2013; Franceschi et al. 2019; Liu et al. 2019) or assume knowledge of the constraints  $\mathcal{C}$  (Kool, van Hoof, and Welling 2020). AIMLE and IMLE provide a general-purpose framework that does not

require access to the linear constraints and the corresponding integer polytope  $\mathcal{C}$ . SparseMAP (Niculae et al. 2018) is an approach to structured prediction and latent variables, replacing an exponential distribution with a sparser distribution; similarly to our work, it only presupposes the availability of a MAP oracle. LP-SparseMAP (Niculae and Martins 2020) is an extension of SparseMAP that uses a relaxation of the underlying optimisation problem.

**Differentiating through combinatorial solvers.** A series of works about differentiating through combinatorial optimisation problems (Wilder, Dilkina, and Tambe 2019; Elmachtoub and Grigas 2022; Ferber et al. 2020; Mandi and Guns 2020) relax ILPs by adding a regularisation term, and differentiate through the KKT conditions deriving from the application of the cutting plane or the interior-point methods. These approaches are conceptually linked to techniques for differentiating through smooth programs (Amos and Kolter 2017; Donti, Kolter, and Amos 2017; Agrawal et al. 2019; Chen et al. 2020; Domke 2012; Franceschi et al. 2018) that arise in modelling, hyperparameter optimisation, and meta-learning. Black-box Backprop (Pogancic et al. 2020; Rolínek et al. 2020) and DPO (Berthet et al. 2020) are methods that are not tied to a specific ILP solver. Black-box Backprop, originally derived from a continuous interpolation argument, can be interpreted as special instantiations of IMLE and AIMLE. DPO addresses the theory of perturbed optimizers and discusses Perturb-and-MAP in the context of the Fenchel-Young losses. All the combinatorial optimisation-related works assume that either optimal costs or solutions are given as training data, while IMLE and AIMLE can also be applied in the absence of such supervision by making use of implicitly generated target distributions. Other authors focus on devising differentiable relaxations for specific combinatorial problems such as SAT (Evans and Grefenstette 2018) or MaxSAT (Wang et al. 2019). Machine learning intersects with combinatorial optimisation in other contexts, e.g. in learning heuristics to improve the performances of combinatorial solvers — we refer to Bengio, Lodi, and Prouvost (2020) for further details.




Figure 3: Cosine similarity between the estimated gradient and the true gradient ( $y$ -axis) using several estimators (IMLE, AIMLE, STE, and SFE) with  $S$  samples, with  $S \in \{10^0, 10^1, \dots, 10^5\}$  ( $x$ -axis) — the gradient estimates produced by AIMLE, both in its forward and central difference versions, are significantly more similar to the true gradient than the estimates produced by other methods.

Direct Loss Minimisation (DLM, McAllester, Hazan, and Keshet 2010; Song et al. 2016) is also related to our work, but it relies on the assumption that examples of optimal states  $z$  are given. Lorberbom et al. (2019) extend the DLM framework to discrete VAEs using coupled perturbations: their approach is tailored to VAEs, and is not general-purpose. Under a methodological viewpoint, IMLE inherits from classical MLE (Wainwright and Jordan 2008) and Perturb-and-MAP (Papandreou and Yuille 2011). The theory of Perturb-and-MAP was used to derive general-purpose upper bounds for log-partition functions (Hazan and Jaakkola 2012; Shpakova and Bach 2016).

## Experiments

Similarly to Niepert, Minervini, and Franceschi (2021), conducted three different types of experiments. First, we analyse and compare the behaviour of AIMLE with other gradient estimators (STE, SFE, IMLE) in a synthetic setting. Second, we consider a setting where the distribution parameters  $\theta$  are produced by an upstream neural model, denoted by  $f$  in Eq. (3), and the optimal discrete structure is not available during training. Finally, we consider the problem of differentiating through black-box combinatorial solvers, where we use the target distribution derived in Eq. (7). In all our experiments, we fix the AIMLE hyper-parameters and use the target gradient norm  $c$  to  $c = 1$ , and the update step  $\eta$  to  $\eta = 10^{-3}$ , based on the AIMLE implementation described in Algorithm 1. More experimental details and additional experiments on the Warcraft dataset proposed by Pogancic et al. (2020) are available in the appendix.

**Synthetic Experiments.** We conducted a series of experiments with a tractable categorical distribution where  $z \in \{0, 1\}^n$  with  $n \in \{10, 20, 30, 50\}$  and  $\sum_i z_i = 1$ . We set the loss to  $L(\theta) = \mathbb{E}_{z \sim p(z; \theta)} [\|z - \mathbf{b}\|^2]$ , where  $\mathbf{b} \sim \mathcal{N}(0, \mathbf{I})$ .

In Fig. 3, we plot the cosine similarity between the gradient




Figure 4: Cosine similarity between the estimated gradient and the true gradient ( $y$ -axis) using several estimators — namely IMLE with a varying  $\lambda \in [0, 5]$  ( $x$ -axis), AIMLE, STE, and SFE — with  $S = 1000$  samples.

estimates produced by the Straight-Through Estimator (STE, Bengio, Léonard, and Courville 2013) the Score Function Estimator (SFE, Williams 1992), Implicit Maximum Likelihood Estimation (IMLE, Niepert, Minervini, and Franceschi 2021), and AIMLE. For STE and IMLE, we use Perturb-and-MAP with Gumbel noise. For IMLE and AIMLE, we evaluated both their *forward difference* (Forward) and *central difference* (Central) versions. We evaluated all estimators using  $S \in \{10^0, 10^1, \dots, 10^5\}$  samples, and report how  $S$  influences the cosine similarity between the gradient estimate and true gradient. Statistics are over 32 runs. From the results, outlined in Fig. 3, we can see that AIMLE, both in its central and forward difference versions, produces significantly more accurate estimates of the true gradient compared to IMLE, STE, and SFE, with orders of magnitude fewer samples. Furthermore, we report the cosine similarity between the true and the estimated gradient for AIMLE, STE, SFE, and IMLE with a varying value of  $\lambda \in [0, 5]$  — all estimators use  $S = 1000$  samples. Results are outlined in Fig. 4: we can see that AIMLE can produce gradient estimates that are comparable to the best estimates produced by IMLE, without the need of training a  $\lambda$  hyper-parameter.

**Learning to Explain.** The BEERADVOCATE dataset (McAuley, Leskovec, and Jurafsky 2012) consists of free-text reviews and ratings for 4 different aspects of beer: *appearance*, *aroma*, *palate*, and *taste*. Each sentence in the test set has annotations providing the words that best describe the various aspects. Following the experimental setting in Paulus et al. (2020); Niepert, Minervini, and Franceschi (2021), we address the problem of learning a distribution over  $k$ -subsets of words that best explain a given aspect rating, introduced by Chen et al. (2018). The complexity of the MAP problem for the  $k$ -subset distribution is linear in  $k$ .

The training set has 80,000 reviews for the aspect *appearance* and 70,000 reviews for all other aspects. Since the original dataset (McAuley, Leskovec, and Jurafsky 2012) did not provide separate validation and test sets, following Niepert, Minervini, and Franceschi (2021), we compute 10 different evenly sized validation and test splits of the 10,000 held out set and compute mean and standard deviation over 10 models, each trained on one split. Subset precision was computed



Method	Test MSE		Subset Prec.	
	Mean	SD	Mean	SD
Aspect: <i>aroma</i> , $K = 5$				
SoftSub ( $\tau = 1.0$ )	2.515	0.087	55.453	2.338
STE ( $\tau = 0.0$ )	4.660	0.053	44.593	0.523
SST ( $\tau = 0.5$ )	4.788	0.486	56.854	3.752
IMLE (F, $\lambda = 10^3$ , $\tau = 1.0$ )	2.413	0.055	53.744	5.635
IMLE (C, $\lambda = 10^3$ , $\tau = 0.0$ )	<b>2.266</b>	<b>0.050</b>	50.888	5.453
AIMLE (F, $\tau = 1.0$ )	2.499	0.089	44.668	6.936
AIMLE (C, $\tau = 3.0$ )	2.385	0.049	<b>62.056</b>	<b>2.107</b>
Aspect: <i>aroma</i> , $K = 10$				
SoftSub ( $\tau = 2.0$ )	2.543	0.044	44.513	2.958
STE ( $\tau = 0.0$ )	4.310	0.039	39.635	0.281
SST ( $\tau = 0.1$ )	5.213	0.295	24.328	12.463
IMLE (F, $\lambda = 10^3$ , $\tau = 1.0$ )	2.368	0.075	48.215	2.182
IMLE (C, $\lambda = 10^3$ , $\tau = 0.0$ )	<b>2.256</b>	<b>0.043</b>	45.339	3.115
AIMLE (F, $\tau = 2.0$ )	2.402	0.042	48.397	1.967
AIMLE (C, $\tau = 2.0$ )	2.419	0.061	<b>53.260</b>	<b>2.271</b>
Aspect: <i>aroma</i> , $K = 15$				
SoftSub ( $\tau = 2.0$ )	2.711	0.035	37.202	1.374
STE ( $\tau = 0.5$ )	4.062	0.054	36.267	0.161
SST ( $\tau = 0.1$ )	5.787	0.517	24.551	9.827
IMLE (F, $\lambda = 10^3$ , $\tau = 1.0$ )	2.411	0.087	41.850	1.477
IMLE (C, $\lambda = 10^3$ , $\tau = 0.0$ )	2.508	0.396	40.057	7.172
AIMLE (F, $\tau = 1.0$ )	<b>2.408</b>	<b>0.064</b>	41.688	2.246
AIMLE (C, $\tau = 3.0$ )	2.470	0.026	<b>47.109</b>	<b>2.863</b>

Table 1: Detailed results for the aspect *aroma*. Test MSE and subset precision, both  $\times 100$ , for  $k \in \{5, 10, 15\}$ .

using a subset of 993 annotated reviews. We use pre-trained word embeddings from (Lei, Barzilay, and Jaakkola 2016). We extend the implementations provided by Niepert, Minervini, and Franceschi (2021), which use a neural network following the architecture introduced by Paulus et al. (2020) with four convolutional layers and one dense layer. This neural network outputs the parameters  $\theta$  of the distribution  $p(z; \theta)$  over  $k$ -hot binary latent masks with  $k \in \{5, 10, 15\}$ .

We compare AIMLE (both the forward and central difference versions) to relaxation-based baselines L2X (Chen et al. 2018) and SoftSub (Xie and Ermon 2019); to STE with Gumbel perturbations; and to IMLE (Niepert, Minervini, and Franceschi 2021) with Gumbel perturbations. We used the standard hyperparameter settings of Chen et al. (2018) and chose the temperature parameter  $t \in \{0.1, 0.5, 1.0, 2.0\}$  for all methods. For IMLE we choose  $\lambda \in \{10, 100, 1000\}$  based on the validation MSE. We trained separate models for each aspect using MSE as the training loss, using the Adam (Kingma and Ba 2015) optimiser with its default hyperparameters.

Table 1 lists detailed results for the aspect *aroma*. We can see that AIMLE, in its central differences version, systematically produces the highest subset precision values while yielding test MSE values comparable to those produced by IMLE, while not requiring tuning the  $\lambda$  hyper-parameter. In the appendix, we report the results for the other aspects, where we notice that AIMLE produces significantly higher subset precision values in all other aspects.

**Discrete Variational Auto-Encoder.** Following Niepert, Minervini, and Franceschi (2021), we compare IMLE, STE,

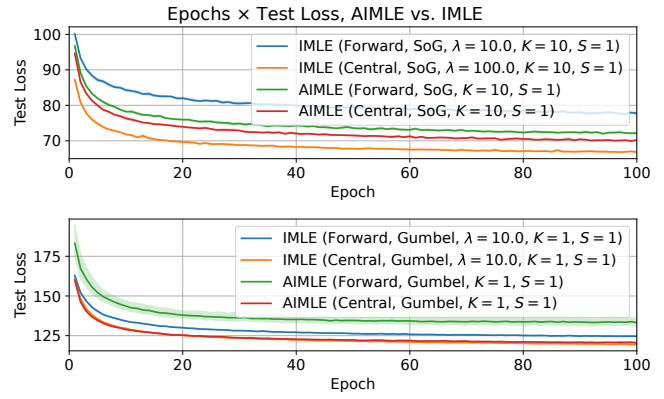


Figure 5: Training dynamics for a DVAE using AIMLE and IMLE for  $K = 10$  (top) and  $K = 1$  (bottom).




Figure 6: Training dynamics for a DVAE using AIMLE and Gumbel-Softmax for  $K = 1$ .

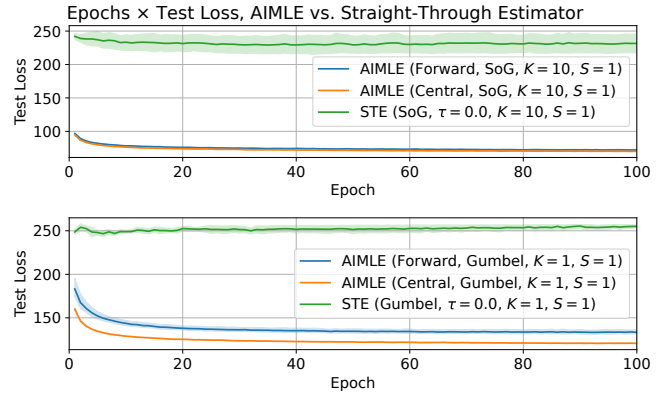


Figure 7: Training dynamics for a DVAE using AIMLE and STE for  $K = 10$  (top) and  $K = 1$  (bottom).

and the Gumbel-Softmax trick (Maddison, Mnih, and Teh 2017; Jang, Gu, and Poole 2017) using a discrete  $K$ -subset Variational Auto-Encoder (VAE). The latent variables model a probability distribution over  $K$ -subsets of — or top- $K$  assignments — binary vectors of length 20; note that, for  $K = 1$ , this is equivalent to a categorical variable with 20 categories. We follow the implementation details in Niepert, Minervini, and Franceschi (2021), where the encoder and the decoder of the VAE consist of three dense layers, where the encoder and decoder activations have sizes 512-256-20 $\times$ 20, and 256-512-784, respectively. The loss is the sum of the reconstruction losses — binary cross-entropy loss on output pixels — and the KL divergence between the marginals of

Edge Distribution	T=10			T=20		
	ELBO	Edge Prec.	Edge Rec.	ELBO	Edge Prec.	Edge Rec.
SST (Hard)	-2301.47 ± 85.86	33.75 ± 9.44	60.40 ± 23.23	-3407.89 ± 221.53	57.40 ± 17.87	70.42 ± 8.22
IMLE (Forward)	-2289.94 ± 4.31	23.94 ± 0.03	95.75 ± 0.14	-3820.68 ± 25.32	20.28 ± 0.12	20.28 ± 0.12
IMLE (Central)	-2341.71 ± 41.68	43.95 ± 7.22	43.95 ± 7.22	-3447.29 ± 550.38	40.25 ± 14.26	40.25 ± 14.26
AIMLE (Forward)	<b>-1877.90 ± 277.53</b>	55.23 ± 11.86	55.23 ± 11.86	<b>-1884.83 ± 124.62</b>	40.48 ± 4.25	40.48 ± 4.25
AIMLE (Central)	-2018.39 ± 357.16	29.32 ± 6.89	41.83 ± 21.51	-1999.57 ± 856.27	70.89 ± 24.77	83.73 ± 1.31

Table 2: Latent Graph Structure Recovery – Stochastic Softmax Tricks (SST,][[paulus2020gradient defining a *spanning tree* over undirected edges (with hard sampling), in comparison with IMLE (Niepert, Minervini, and Franceschi 2021) and AIMLE, where the MAP function is computed by Kruskal’s algorithm (Kruskal 1956). AIMLE yields the lowest test ELBO values in the T=10 (shorter sequences) and T=20 (longer sequences) settings.

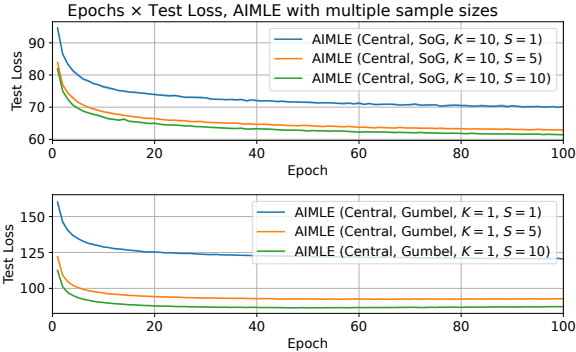


Figure 8: Training dynamics for a DVAE using AIMLE with a different number of samples (i.e.  $S \in \{1, 5, 10\}$ ) for  $K = 10$  (top) and  $K = 1$  (bottom).

the variables and the uniform distribution.

For AIMLE, IMLE, and STE, we use Gumbel(0, 1) perturbations for  $K = 1$ , and Sum-of-Gamma perturbations (SoG, Niepert, Minervini, and Franceschi 2021) for  $K = 10$ , with a temperature of  $\tau = 1$ . For IMLE, we select the hyper-parameter  $\lambda \in \{1, 10, 100\}$  on a held-out validation set based on the validation loss. For IMLE and AIMLE, we report the results for both the forward and central difference versions. We train the DVAE for 100 epochs and report the loss on a held-out test set. In Fig. 5, we show the test losses for IMLE and AIMLE (forward and central difference versions) for  $K = 10$  and  $K = 1$ . We can see that, for  $K = 10$ , AIMLE produces significantly lower test loss values than IMLE, and, for  $K = 1$ , the central difference version of AIMLE produces test loss values comparable to IMLE. We also compare AIMLE to the Gumbel-Softmax trick (see Fig. 6) and STE (see Fig. 7): AIMLE produces significantly lower test losses than STE ( $K \in \{1, 10\}$ ) while producing higher test losses than Gumbel-Softmax ( $K = 1$ ). When increasing the number of samples  $S$  in AIMLE (see Fig. 8), we found that higher values of  $S$  produce significantly lower test loss values, to the point that, for  $K = 1$ , AIMLE with  $S \in \{5, 10\}$  produces lower test loss values than Gumbel-Softmax, but with a higher computational cost. All hyper-parameters and an analysis of the evolution of  $\lambda$  during training are available in the appendix.

### Neural Relational Inference for Recovering Latent Graph Structures.

In this experiment, we investigated the use of

AIMLE for recovering latent graph structures and predicting the evolution of a dynamical system. In Neural Relational Inference (NRI, Kipf et al. 2018), a Graph Neural Network (GNN, Micheli 2009; Scarselli et al. 2009) encoder is used to generate a latent interaction graph, which is then used to produce a distribution over an interacting particle system. NRI is trained as a variational auto-encoder to maximise a lower bound (ELBO) on the marginal log-likelihood of the time series. Based on the implementation provided by Paulus et al. (2020), we compared the Stochastic Softmax Tricks (SST, Paulus et al. 2020) encoder that induces a *spanning tree* over undirected edges, with an encoder producing a maximum spanning tree using Kruskal’s algorithm (Kruskal 1956), and using either IMLE or AIMLE to back-propagate through it. In this setting, Kruskal’s algorithm represents the MAP estimator for a distribution over the latent graph structures. Our dataset consisted of latent prior spanning trees over 10 vertices sampled from the Gumbel(0, 1) prior. Given a tree, we embed the vertices in  $\mathbb{R}^2$  by applying  $T \in \{10, 20\}$  iterations of a force-directed algorithm (Fruchterman and Reingold 1991). The model saw particle locations at each iteration, not the underlying spanning tree. Results are outlined in Table 2: we found that AIMLE performed best, improving on both ELBO and the recovery of latent structure over the structured SST baseline proposed by Paulus et al. (2020).

## Conclusions

We introduced Adaptive Implicit Maximum Likelihood Estimation (AIMLE), an efficient, simple-to-implement, and general-purpose framework for learning hybrid models. AIMLE is an extension of IMLE (Niepert, Minervini, and Franceschi 2021) that, during training, can dynamically select the optimal target distribution by identifying the update step  $\lambda$  that yields the desired gradient norm. Furthermore, AIMLE incorporates insights from finite difference methods, improving its effectiveness in gradient estimation tasks. In our experiments, we show that AIMLE produces better results than relaxation-based approaches for discrete latent variable models and approaches that back-propagate through black-box combinatorial solvers. A limitation is that AIMLE relies on a *warm-up period* for selecting the optimal  $\lambda$ , whose duration varies depending on the update step  $\eta$  — which we fix to  $\eta = 10^{-3}$ . A potential solution to this problem is to use adaptive update steps, such as momentum (Qian 1999).

## Acknowledgments

Pasquale was partially funded by the European Union’s Horizon 2020 research and innovation programme under grant agreement no. 875160, ELIAI (The Edinburgh Laboratory for Integrated Artificial Intelligence) EPSRC (grant no. EP/W002876/1), an industry grant from Cisco, and a donation from Accenture LLP; and is grateful to NVIDIA for the GPU donations. Mathias was partially funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2075.

## References

- Agrawal, A.; Amos, B.; Barratt, S. T.; Boyd, S. P.; Diamond, S.; and Kolter, J. Z. 2019. Differentiable Convex Optimization Layers. In *NeurIPS*, 9558–9570.
- Amos, B.; and Kolter, J. Z. 2017. OptNet: Differentiable Optimization as a Layer in Neural Networks. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, 136–145. PMLR.
- Bengio, Y.; Léonard, N.; and Courville, A. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*.
- Bengio, Y.; Lodi, A.; and Prouvost, A. 2020. Machine learning for combinatorial optimization: a methodological tour d’horizon. *European Journal of Operational Research*.
- Berthet, Q.; Blondel, M.; Teboul, O.; Cuturi, M.; Vert, J.; and Bach, F. R. 2020. Learning with Differentiable Perturbed Optimizers. In *NeurIPS*.
- Betz, P.; Niepert, M.; Minervini, P.; and Stuckenschmidt, H. 2021. Backpropagating through Markov Logic Networks. In *Proceedings of the 15th International Workshop on Neural-Symbolic Learning and Reasoning*, 67–81.
- Chen, J.; Song, L.; Wainwright, M.; and Jordan, M. 2018. Learning to explain: An information-theoretic perspective on model interpretation. In *International Conference on Machine Learning*, 883–892. PMLR.
- Chen, X.; Zhang, Y.; Reisinger, C.; and Song, L. 2020. Understanding Deep Architecture with Reasoning Layer. In *NeurIPS*.
- Domke, J. 2010. Implicit Differentiation by Perturbation. In *NIPS*, 523–531. Curran Associates, Inc.
- Domke, J. 2012. Generic methods for optimization-based modeling. In *Artificial Intelligence and Statistics*, 318–326. PMLR.
- Donti, P. L.; Kolter, J. Z.; and Amos, B. 2017. Task-based End-to-end Model Learning in Stochastic Optimization. In *NIPS*, 5484–5494.
- Elmachtoub, A. N.; and Grigas, P. 2022. Smart "Predict, then Optimize". *Manag. Sci.*, 68(1): 9–26.
- Evans, R.; and Grefenstette, E. 2018. Learning explanatory rules from noisy data. *Journal of Artificial Intelligence Research*, 61: 1–64.
- Ferber, A. M.; Wilder, B.; Dilkina, B.; and Tambe, M. 2020. MIPaaL: Mixed Integer Program as a Layer. In *AAAI*, 1504–1511. AAAI Press.
- Franceschi, L.; Frasconi, P.; Salzo, S.; Grazzi, R.; and Pontil, M. 2018. Bilevel Programming for Hyperparameter Optimization and Meta-Learning. In *ICML*, volume 80 of *Proceedings of Machine Learning Research*, 1563–1572. PMLR.
- Franceschi, L.; Niepert, M.; Pontil, M.; and He, X. 2019. Learning Discrete Structures for Graph Neural Networks. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, 1972–1982. PMLR.
- Fruchterman, T. M. J.; and Reingold, E. M. 1991. Graph Drawing by Force-directed Placement. *Softw. Pract. Exp.*, 21(11): 1129–1164.
- Grathwohl, W.; Choi, D.; Wu, Y.; Roeder, G.; and Duvenaud, D. 2018. Backpropagation through the Void: Optimizing control variates for black-box gradient estimation. In *ICLR (Poster)*. OpenReview.net.
- Grover, A.; Wang, E.; Zweig, A.; and Ermon, S. 2019. Stochastic optimization of sorting networks via continuous relaxations. *arXiv preprint arXiv:1903.08850*.
- Hafner, D.; Lillicrap, T.; Norouzi, M.; and Ba, J. 2020. Mastering atari with discrete world models. *arXiv preprint arXiv:2010.02193*.
- Hazan, T.; and Jaakkola, T. S. 2012. On the Partition Function and Random Maximum A-Posteriori Perturbations. In *ICML*. icml.cc / Omnipress.
- Jang, E.; Gu, S.; and Poole, B. 2017. Categorical reparameterization with gumbel-softmax. *ICLR*.
- Kim, C.; Sabharwal, A.; and Ermon, S. 2016. Exact Sampling with Integer Linear Programs and Random Perturbations. In *AAAI*, 3248–3254. AAAI Press.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *ICLR (Poster)*.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. In *ICLR*.
- Kipf, T. N.; Fetaya, E.; Wang, K.; Welling, M.; and Zemel, R. S. 2018. Neural Relational Inference for Interacting Systems. In *ICML*, volume 80 of *Proceedings of Machine Learning Research*, 2693–2702. PMLR.
- Kool, W.; van Hoof, H.; and Welling, M. 2020. Estimating gradients for discrete random variables by sampling without replacement. *ICLR*.
- Kruskal, J. B. 1956. On the Shortest Spanning Subtree of a Graph and the Traveling Salesman Problem. In *Proceedings of the American Mathematical Society*, 7.
- LeCun, Y.; Chopra, S.; Hadsell, R.; Ranzato, M.; and Huang, F. 2006. A tutorial on energy-based learning. *Predicting structured data*, 1(0).
- Lei, T.; Barzilay, R.; and Jaakkola, T. 2016. Rationalizing Neural Predictions. In *EMNLP*.
- Liu, R.; Regier, J.; Tripuraneni, N.; Jordan, M. I.; and McAuliffe, J. D. 2019. Rao-Blackwellized Stochastic Gradients for Discrete Distributions. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, 4023–4031. PMLR.
- Lorberbom, G.; Jaakkola, T. S.; Gane, A.; and Hazan, T. 2019. Direct Optimization through arg max for Discrete Variational Auto-Encoder. In *NeurIPS*, 6200–6211.



- Maddison, C. J.; Mnih, A.; and Teh, Y. W. 2017. The concrete distribution: A continuous relaxation of discrete random variables. *ICLR*.
- Mandi, J.; Demirovic, E.; Stuckey, P. J.; and Guns, T. 2020. Smart Predict-and-Optimize for Hard Combinatorial Optimization Problems. In *AAAI*, 1603–1610. AAAI Press.
- Mandi, J.; and Guns, T. 2020. Interior Point Solving for LP-based prediction+optimisation. In *NeurIPS*.
- McAllester, D. A.; Hazan, T.; and Keshet, J. 2010. Direct Loss Minimization for Structured Prediction. In *NIPS*, 1594–1602. Curran Associates, Inc.
- McAuley, J. J.; Leskovec, J.; and Jurafsky, D. 2012. Learning Attitudes and Attributes from Multi-aspect Reviews. In *ICDM*, 1020–1025. IEEE Computer Society.
- Micheli, A. 2009. Neural Network for Graphs: A Contextual Constructive Approach. *IEEE Trans. Neural Networks*, 20(3): 498–511.
- Mohamed, S.; Rosca, M.; Figurnov, M.; and Mnih, A. 2019. Monte carlo gradient estimation in machine learning. *arXiv preprint arXiv:1906.10652*.
- Niculae, V.; and Martins, A. F. T. 2020. LP-SparseMAP: Differentiable Relaxed Optimization for Sparse Structured Prediction. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, 7348–7359. PMLR.
- Niculae, V.; Martins, A. F. T.; Blondel, M.; and Cardie, C. 2018. SparseMAP: Differentiable Sparse Structured Inference. In *ICML*, volume 80 of *Proceedings of Machine Learning Research*, 3796–3805. PMLR.
- Niepert, M.; Minervini, P.; and Franceschi, L. 2021. Implicit MLE: Backpropagating Through Discrete Exponential Family Distributions. In *NeurIPS*, Proceedings of Machine Learning Research. PMLR.
- Olver, P. J. 2013. *Introduction to partial differential equations*. Undergraduate Texts in Mathematics. Basel, Switzerland: Springer International Publishing, 2014 edition.
- Papandreou, G.; and Yuille, A. L. 2011. Perturb-and-MAP random fields: Using discrete optimization to learn and sample from energy models. In *ICCV*, 193–200. IEEE Computer Society.
- Paulus, M. B.; Choi, D.; Tarlow, D.; Krause, A.; and Maddison, C. J. 2020. Gradient Estimation with Stochastic Softmax Tricks. In *NeurIPS*.
- Pogancic, M. V.; Paulus, A.; Musil, V.; Martius, G.; and Rolínek, M. 2020. Differentiation of Blackbox Combinatorial Solvers. In *ICLR*. OpenReview.net.
- Qian, C.; Rattan, G.; Geerts, F.; Morris, C.; and Niepert, M. 2022. Ordered Subgraph Aggregation Networks. *arXiv preprint arXiv:2206.11168*.
- Qian, N. 1999. On the momentum term in gradient descent learning algorithms. *Neural Networks*, 12(1): 145–151.
- Rolínek, M.; Swoboda, P.; Zietlow, D.; Paulus, A.; Musil, V.; and Martius, G. 2020. Deep Graph Matching via Blackbox Differentiation of Combinatorial Solvers. In *ECCV*.
- Scarselli, F.; Gori, M.; Tsoi, A. C.; Hagenbuchner, M.; and Monfardini, G. 2009. The Graph Neural Network Model. *IEEE Trans. Neural Networks*, 20(1): 61–80.
- Shpakova, T.; and Bach, F. R. 2016. Parameter Learning for Log-supermodular Distributions. In *NIPS*, 3234–3242.
- Song, Y.; Schwing, A.; Urtasun, R.; et al. 2016. Training deep neural networks via direct loss minimization. In *International Conference on Machine Learning*, 2169–2177.
- Tucker, G.; Mnih, A.; Maddison, C. J.; Lawson, D.; and Sohl-Dickstein, J. 2017. REBAR: Low-variance, unbiased gradient estimates for discrete latent variable models. In *NIPS*, 2627–2636.
- Wainwright, M. J.; and Jordan, M. I. 2008. *Graphical models, exponential families, and variational inference*. Now Publishers Inc.
- Wang, P.; Donti, P. L.; Wilder, B.; and Kolter, J. Z. 2019. SATNet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, 6545–6554. PMLR.
- Wilder, B.; Dilkina, B.; and Tambe, M. 2019. Melding the Data-Decisions Pipeline: Decision-Focused Learning for Combinatorial Optimization. In *AAAI*, 1658–1665. AAAI Press.
- Williams, R. J. 1992. Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning. *Mach. Learn.*, 8: 229–256.
- Xie, S. M.; and Ermon, S. 2019. Reparameterizable Subset Sampling via Continuous Relaxations. In *IJCAI*, 3919–3925. ijcai.org.