



HAL
open science

Discrete Latent Structure in Neural Networks

Vlad Niculae, Caio Corro, Nikita Nangia, Tsvetomila Mihaylova, André Martins

► **To cite this version:**

Vlad Niculae, Caio Corro, Nikita Nangia, Tsvetomila Mihaylova, André Martins. Discrete Latent Structure in Neural Networks. Foundations and Trends in Signal Processing, 2025, Foundations and Trends in Signal Processing, 19 (2), pp.99-211. <10.1561/2000000134>. <hal-05457940>

HAL Id: hal-05457940

<https://hal.science/hal-05457940v1>

Submitted on 14 Jan 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-NC-ND 4.0 - Attribution - Non-commercial use - No Derivative Works - International License

Discrete Latent Structure in Neural Networks

Vlad Niculae¹, Caio Corro², Nikita Nangia³, Tsvetomila Mihaylova^{4,5}
and André F. T. Martins^{4,5,6}

¹*Informatics Institute, University of Amsterdam, Netherlands*

²*INSA Rennes, IRISA, Inria, CNRS, Université de Rennes*

³*New York University, USA*

⁴*Instituto Superior Técnico, Lisbon, Portugal*

⁵*Instituto de Telecomunicações, Lisbon, Portugal*

⁶*Unbabel, Lisbon, Portugal*

Abstract

Many types of data from fields including natural language processing, computer vision, and bioinformatics, are well represented by discrete, compositional structures such as trees, sequences, or matchings. Latent structure models are a powerful tool for learning to extract such representations, offering a way to incorporate structural bias, discover insight about the data, and interpret decisions. However, effective training is challenging, as neural networks are typically designed for continuous computation.

This text explores three broad strategies for learning with discrete latent structure: continuous relaxation, surrogate gradients, and probabilistic estimation. Our presentation relies on consistent notations for a wide range of models. As such, we reveal many new connections between latent structure learning strategies, showing how most consist of the same small set of fundamental building blocks, but use them differently, leading to substantially different applicability and properties.

Contents

1	Introduction	7
1.1	Motivation	7
1.2	Supervised Learning	8
1.3	Latent Representations	12
1.4	Further History and Scope	15
1.5	Roadmap	20
2	Structure Prediction Background	21
2.1	Overview	21
2.2	Incremental Prediction	26
2.3	Global Prediction	34
2.4	Summary	42
3	Continuous Relaxations	44
3.1	Challenges of Deterministic Choices	44
3.2	Regularized Argmax Operators	48
3.3	Categorical Relaxation and Attention	51
3.4	Global Structured Relaxations and Structured Attention	55
3.5	Mean Structure Regularization: Sinkhorn and SparseMAP	58
3.6	Summary	63

4	Surrogate Gradients	64
4.1	Straight-Through Gradients	64
4.2	Straight-Through Variants	65
4.3	Quantization: Straight-Through Friendly Models	68
4.4	Interpretation via Pulled-Back Labels	69
4.5	Summary	73
5	Probabilistic Latent Variables	74
5.1	Formulating the Probabilistic Model.	74
5.2	Explicit Marginalization by Enumeration	75
5.3	Monte Carlo Gradient Estimation	76
5.4	Path Gradient Estimator (The Reparametrization Trick)	77
5.5	Score Function Estimator	85
5.6	Sparsifying the Distribution	90
5.7	Summary	95
6	Conclusions	97
6.1	Overview	97
6.2	Implementations and Libraries	99
	Acknowledgements	102
	References	103

Notation

Vectors, matrices, and indexing.

$u, \mathbf{v}, \mathbf{W}, \mathcal{X}$ a scalar, a vector, a matrix, and a set.

v_i the i th element of vector \mathbf{v} .

\mathbf{w}_j the j th column of matrix \mathbf{W} .

$\|\mathbf{v}\|_p := \left(\sum_{i=1}^d |v_i|^p\right)^{1/p}$, the p -norm of $\mathbf{v} \in \mathbb{R}^d$.

Probabilities and distributions.

Y a random variable with values $y \in \mathcal{Y}$;

$p(Y = y)$ probability that Y take the specific value y ;

$p(y | x)$ short for $p(Y = y | \mathbf{X} = x)$ when unambiguous;

$\mathbb{E}_{p(Y)}[Y] := \sum_{y \in \mathcal{Y}} yp(y)$, the expected value of Y .

Differentiation.

$\partial_i f$ the partial derivative of $f : \mathbb{R}^{d_1} \times \dots \times \mathbb{R}^{d_n} \rightarrow \mathbb{R}^d$ w.r.t. the i th argument. $(\partial_i f)(\mathbf{x}_1, \dots, \mathbf{x}_n)$ is a linear $\mathbb{R}^d \rightarrow \mathbb{R}^{d_i}$ map (the pullback of f), identified with with a $d_i \times d$ matrix: the Jacobian transpose. For single-argument $f : \mathbb{R}^{d_1} \rightarrow \mathbb{R}^d$ we omit the subscript, and if \mathbf{J}_x is the Jacobian of f at \mathbf{x} then $\partial f(\mathbf{x})(\mathbf{v}) = \mathbf{J}_x^\top \mathbf{v}$. This transposed convention is more convenient for backpropagation.

$\partial_\theta(\text{expr.})$ interprets the (possibly complex) expression as a single-argument function of θ and applies ∂ .

Convex sets.

\mathbb{R}_+^d := $\{\boldsymbol{\alpha} \in \mathbb{R}^d; \alpha_i \geq 0 \text{ for all } 1 \leq i \leq d\}$, the non-negative orthant;

Δ_d := $\{\boldsymbol{\alpha} \in \mathbb{R}_+^d; \sum_i \alpha_i = 1\}$, the simplex with d bins, containing all probability distributions over d choices;

$\text{conv}(\mathcal{Z})$ the convex hull of \mathcal{Z} , *i.e.*, the smallest convex set containing \mathcal{Z} .

1

Introduction

1.1 Motivation

Machine learning (ML) is often employed to build predictive models for analyzing rich data, such as images, text, or sound. Most such data is governed by underlying *structured representations*, such as segmentations, hierarchy, or graph structure. For example, natural language sentences can be analyzed in terms of their *dependency structure*, yielding an arborescence of directed grammatical relationships between words (figure 2.1).

It is common for practical ML systems to be structured as **pipelines**, including off-the-shelf components (analyzers) that produce structured representations of the input, used as features in subsequent steps of the pipeline. On the one hand, such architectures require availability of these analyzers (or of the data to train them). Since the analyzer may not be built with the downstream goal in mind, a disadvantage of pipelines is that they are prone to error propagation. On the other hand, they are transparent: the predicted structures can be directly inspected and used to interpret downstream predictions. In contrast, *deep neural networks* rival and even outperform pipelines by learning dense, continuous representations of the data, solely driven by the downstream objective.

However, the popular success of end-to-end deep learning hides some

fundamental challenges. For example, large language models are still based on a pipeline system in which tokenization is an independent pre-processing step. Another known limitation is the so-called structural generalization problem (Yao and Koller, 2022): sequential architectures (both recurrent neural networks and self-attentive networks) have difficulties to generalize to unseen (recursive) combinations of known parts. It is possible to tackle this problem by inducing latent structured representations (Bogin *et al.*, 2021; Liu *et al.*, 2021). Similar limitations are known for length generalization (Anil *et al.*, 2022; Zhou *et al.*, 2024). Another important research direction in the natural language processing community is intermediate plan-based representations for text generation (Narayan *et al.*, 2023; Liu *et al.*, 2023), where latent structures may play an important role, for example when learning with limited information (Xu and Lapata, 2022). Beside natural language processing, latent structure inference is also a key topic in computer vision for unsupervised segmentation and learning object-centric representations (Steenkiste *et al.*, 2018; Greff *et al.*, 2019; Locatello *et al.*, 2020; Elsayed *et al.*, 2022).

This text is about neural network models that induce **discrete latent structure**, combining the strengths of both end-to-end-approaches and pipeline systems. In the following, we do not assume a specific downstream application in natural language processing nor computer vision. Our presentation follows an abstract framework that allows to focus on technical aspects related to end-to-end learning with deep neural network.

1.2 Supervised Learning

We begin by establishing the common setup of predictive machine learning. A prediction function is a function that associates with input $x \in \mathcal{X}$ an output $y \in \mathcal{Y}$. Prediction functions usually rely on a *scoring*

function

$$M(x, y; \boldsymbol{\theta}), \quad (1.1)$$

which returns the score, or preference, for some candidate $y \in \mathcal{Y}$, given an input x . In our setting, M is a parametric function with learnable parameters $\boldsymbol{\theta}$. For simple classification problems, M could be a feed-forward network with x as input, and an $|\mathcal{Y}|$ -dimensional output, such that the y th position of the output is $M(x, y; \boldsymbol{\theta})$. Our notation allows for more involved setting like predicting structured objects (for examples graphs, see Chapter 2). To make predictions, we search for the output of maximum weight

$$\hat{y}(x; \boldsymbol{\theta}) := \arg \max_{y' \in \mathcal{Y}} M(x, y'; \boldsymbol{\theta}). \quad (1.2)$$

In many cases, we are also interested in a distribution over outputs. Assuming \mathcal{Y} is a finite set, a common choice is to rely on a Boltzmann-Gibbs distribution, also called *softmax* (Bridle, 1989), defined as follows:

$$\begin{aligned} p(y | x) &= \frac{\exp M(x, y; \boldsymbol{\theta})}{\sum_{y' \in \mathcal{Y}} \exp M(x, y'; \boldsymbol{\theta})} \quad \text{for } y \in \mathcal{Y}, \\ &\propto \exp M(x, y; \boldsymbol{\theta}). \end{aligned}$$

Note that the most probable output under distribution $p(\cdot | x)$ is equal to $\hat{y}(x; \boldsymbol{\theta})$.

In the supervised learning scenario, we assume access to a dataset \mathcal{D} containing samples of input/output pairs $(x, y) \in \mathcal{D}$. Parameters $\boldsymbol{\theta}$ are fixed to minimize the empirical risk

$$L_{\text{avg}}(\boldsymbol{\theta}) := \frac{1}{|\mathcal{D}|} \sum_{(x, y) \in \mathcal{D}} L(y, x; \boldsymbol{\theta}), \quad (1.3)$$

where L is a loss function (Vapnik, 1991). For practical reasons, the loss function used for classification problems is usually not the targeted evaluation function (for example the 0-1 loss which is equal to 1 if and

only if the model predicts the expected output) but a surrogate loss that is amenable for gradient-based optimization. Statistical consistency of such surrogates has been widely studied (Peter L Bartlett and McAuliffe, 2006; Reid and Williamson, 2010; Williamson *et al.*, 2016). A common choice is the cross-entropy loss,

$$L(x, y; \boldsymbol{\theta}) = -M(x, y; \boldsymbol{\theta}) + \log \sum_{y' \in \mathcal{Y}} \exp M(x, y'; \boldsymbol{\theta}), \quad (1.4)$$

which is simply the model negative log-probability of gold output under a Boltzmann-Gibbs distribution. Then, Eq. 1.3 can be interpreted as *maximum likelihood* estimation of $\boldsymbol{\theta}$. Non-probabilistic losses like the hinge loss or the perceptron loss fit the framework as well.

From a computational point of view, note that both training and use of such a model eventually requires evaluating or optimizing a function of the form

$$g(x, y; \boldsymbol{\theta}),$$

which may refer to either the scoring model M or the loss L . Therefore, we shall use the generic functional notation $g(x, y; \boldsymbol{\theta})$ in the following. Note that in this text, we are interested in computing (or approximating) partial derivatives with respect to all values in $\boldsymbol{\theta}$ via the backpropagation algorithm for automatic differentiation (Linnainmaa, 1970).

Gradient-based learning. The gradient method for minimizing a differentiable function $F : \mathbb{R}^d \rightarrow \mathbb{R}$ iterates

$$\boldsymbol{\theta}^{(t+1)} \leftarrow \boldsymbol{\theta}^{(t)} + \eta^{(t)} (\partial F)(\boldsymbol{\theta}^{(t)}), \quad (1.5)$$

where $\eta^{(t)}$ is a step size schedule, and $\partial F(\cdot)$ is identified with its column-vector Jacobian. This method converges to a stationary point of F under some assumptions on the step size (Bertsekas, 1999, §1.2.2).

Often in machine learning evaluating F is slow and memory-intensive, as it depends on the entire training data; this is the case in Eq. 1.3. In such cases, the stochastic gradient (SG, Robbins and Monro, 1951) method may be preferred. The SG method replaces the gradient with a stochastic direction \mathbf{G} such that

$$\mathbb{E}[\mathbf{G}] = \partial F(\boldsymbol{\theta}^{(t)}), \quad (1.6)$$

followed by updating

$$\boldsymbol{\theta}^{(t+1)} \leftarrow \boldsymbol{\theta}^{(t)} + \eta^{(t)} \mathbf{G}. \quad (1.7)$$

This method also converges to a stationary point under mild assumptions (Bertsekas and Tsitsiklis, 2000) : mainly, requiring smooth F , square-summable decreasing step sizes, and a linear bound on the variance of \mathbf{G} *w.r.t.* the norm of the gradient of F . If F takes the form of an average, *i.e.*, $F(\boldsymbol{\theta}) = \frac{1}{N} \sum_{i=1}^N F_i(\boldsymbol{\theta})$ (for instance Eq. 1.3), then \mathbf{G} may be chosen as a single sample $F_i(\boldsymbol{\theta})$ where i is drawn uniformly from $\{1, \dots, N\}$, or a mini-batch estimator. The gradient and stochastic gradient methods can be extended to a broader family using acceleration, momentum, and adaptivity (Nesterov, 1983; Kingma and Ba, 2015; Loshchilov and Hutter, 2019; d’Aspremont *et al.*, 2021). Algorithms in this family are the *de facto* choice in deep learning at the time of writing. For this reason, our work focuses on compatibility with gradient-based learning.

Backpropagation and the Chain Rule. Given a composition of functions $u : \mathbb{R}^m \rightarrow \mathbb{R}^n$, $v : \mathbb{R}^n \rightarrow \mathbb{R}^p$, $w : \mathbb{R}^p \rightarrow \mathbb{R}^q$, and their composition $(w \circ v \circ u)(\boldsymbol{\theta}) := w(v(u(\boldsymbol{\theta})))$ we have:

$$\partial(w \circ v \circ u)(\boldsymbol{\theta}) = (\partial u)(\boldsymbol{\theta}) \circ (\partial v)(u(\boldsymbol{\theta})) \circ (\partial w)(v(u(\boldsymbol{\theta}))). \quad (1.8)$$

The derivatives are applied in the opposite order compared to the computation. This is known as *backpropagation* or *reverse-mode automatic*

differentiation (Griewank and Walther, 2008) and is popular in deep learning, where models are built using such compositions, with the final layer w having a scalar output (loss). The *forward pass* computes and stores the intermediate values that appear in $w \circ v \circ u$, and the backward pass invokes the ∂ operator to propagate gradients from the output to the input. In the most popular software frameworks today (*e.g.*, PyTorch, 2017), elementary building blocks are provided as composable modules, with implementations providing *forward* calls $f(\boldsymbol{\theta})$ and *backward* calls (vector-Jacobian products) $\partial f(\boldsymbol{\theta})(\mathbf{z})$, and the automatic differentiation engine handles the composition.

1.3 Latent Representations

Our main motivation is to go beyond direct mappings $x \rightarrow y$, toward machine learning models with latent representations. In this text, we take a rather inclusive view of what constitutes a latent representation (Bengio *et al.*, 2013a). We call a latent representation $z \in \mathcal{Z}$ an object designed to capture some relevant property of a data point $x \in \mathcal{X}$, which can be inferred based on x , but is typically unobserved. In particular, we cover but we do not require probabilistic modeling of z (Bishop, 1998). On the other hand, we are explicitly interested in discrete and structured latent representations.

Latent representations are often designed with downstream tasks in mind: we may look for a model of $y \in \mathcal{Y}$ that has access not only to x but also to the representation z :

$$g(x, y, z; \boldsymbol{\theta}_g). \quad (\text{downstream model}) \quad (1.9)$$

Remark. During prediction from a pretrained model, we may think of g as a classifier returning the score of class y . For training the model, however, we may want to think of g as some loss function on top of the same classifier. Mathematically, this distinction is irrelevant for the

purpose of our text, which is the modelling of z , so we henceforth use $g(x, y, z; \boldsymbol{\theta}_g)$ to denote either. Practitioners must however be mindful of this.

A downstream model as in Eq. 1.9 is not directly usable, since z is unknown both at training and at test time. Therefore, the problem we are concerned with in this text is jointly learning to predict z from x using an encoder $f : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$:

$$f(x, z; \boldsymbol{\theta}_f), \quad (\text{encoder model}) \quad (1.10)$$

assigning higher values to better-fitting choices of z to the given x .

The key challenge of learning latent variable models is that we cannot learn $\boldsymbol{\theta}_f$ using standard supervised approaches, since z is not observed. This text is about how to learn a good encoder model f jointly with the downstream model only from pairs (x, y) . During training, the downstream model gets direct supervision, but the encoder model only gets a form of *distant supervision*, its only learning signal is coming in the form of gradients propagated through the downstream model. Joint learning with latent structure in this scenario is the main topic of our text. The next three paragraphs outline the main ways to train end-to-end models in such encoders; the main part of our text (Chapters 3 to 5) later goes into detail.

Pretraining and pipelines. A first strategy is to sidestep the issue altogether and obtain supervision. This poses no challenge mathematically, and is not studied further in this text, but serves as a motivating base case: If in fact some training pairs (x, z) are available, it is promising to first train a model $f(x, z; \boldsymbol{\theta}_f)$ and then deploy a two-step pipeline:

1. predict $\hat{z} = \arg \max_{z' \in \mathcal{Z}} f(x, z'; \boldsymbol{\theta}_f)$,
2. use downstream model $g(x, y, \hat{z}; \boldsymbol{\theta}_g)$.

The parameters of the downstream model θ_g can now be trained in a fully-supervised fashion, since \hat{z} is a known fixed input. This corresponds to the time-tested approach of using off-the-shelf analysis models (parsers, object detection, entity recognizers, etc.) as a pre-processing step. This approach is vulnerable to two main sources of error: *domain shift*, due to the fact that θ_f is likely trained on samples coming from a different distribution than the one \mathcal{D} is drawn from, and *error propagation*, due to the lack of mechanism for improving θ_f if the model makes errors. The latent representation treatment we propose mitigates both these concerns by allowing the fine tuning of θ_f with signal from downstream, see (Peng *et al.*, 2018) for examples.

Deterministic latent representations. A straightforward idea for end-to-end learning would be to characterize the mapping from x to a promising candidate \hat{z} as a function,

$$\hat{z}(x; \theta_f),$$

which implicitly defined by the encoder f . (For example, $\hat{z}(x) = \arg \max_{z \in \mathcal{Z}} f(x, z)$.) Then, an end-to-end model emerges as a composition of functions:

$$g(x, y, \hat{z}(x; \theta_f); \theta_g). \quad (1.11)$$

This resembles the pipeline approach, but now we aim to train θ_f and θ_g jointly using gradient methods. Depending on how \hat{z} is constructed, we may have an end-to-end differentiable relaxed model (Chapter 3) or a discrete model optimized with surrogate gradient heuristics (Chapter 4). Both cases will require further assumptions compared to the pipeline approach with frozen θ_f , but require no supervision on z .

Probabilistic latent variables. Alternatively, we can gain expressivity by modelling latent representations as **random variables** whose distri-

bution is induced by the encoder f . Notationally, we define a random variable Z taking values $z \in \mathcal{Z}$, with distribution $p(Z = z \mid x; \theta_f)$ parametrized in some way using f (e.g., $p(z \mid x; \theta_f) \propto \exp f(x, z; \theta_f)$.) Then, the end-to-end model will consider not a single value of z but the expectation over all possible values $z \in \mathcal{Z}$:

$$\begin{aligned} \bar{g}(x, y; \theta_f, \theta_g) &:= \mathbb{E}_Z [g(x, y, Z; \theta_g)] \\ &= \sum_{z \in \mathcal{Z}} g(x, y, z; \theta_g) p(z \mid x; \theta_f). \end{aligned} \quad (1.12)$$

The expected loss depends on both θ_f and θ_g , and so provides a learning signal to both the encoder and the downstream model. In particular, some choices of g can correspond to a probabilistic treatment of Y as well, making this strategy interesting for generative modelling. We study methods for probabilistic latent variables in Chapter 5. Broadly speaking, these methods tend to require fewer assumptions compared to deterministic ones, but come at a higher computational cost.

Remark. What sets apart a latent representation from an arbitrary “hidden layer” is that the former is designed to capture a specific aspect of x , relevant to the modeler. In this text, we focus on discrete z with structural constraints that can guide it to take a certain form of interest (e.g., alignments, syntax.) This is often (but not necessarily) reflected in the more transparent, informed way in which the way the downstream model g accesses z .

1.4 Further History and Scope

Latent variable models have a long history in ML, specifically for unsupervised learning. In this section, we briefly survey this history and clarify the scope of this work.

Shallow models. Many popular models fall under this umbrella, typically with linear f and g . Factor analysis (FA) is an unsupervised representation learning model ($\mathcal{Y} = \mathbb{R}^d$, $\mathcal{X} = \emptyset$) with continuous latent variables ($\mathcal{Z} = \mathbb{R}^k$) defined by (Barber, 2012, §21.1)

$$f(\mathbf{y}, z; \boldsymbol{\theta}) = -\frac{1}{2}(\mathbf{y} - \mathbf{F}z - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{y} - \mathbf{F}z - \boldsymbol{\mu}), \quad (1.13)$$

where the covariance $\boldsymbol{\Sigma}$ is a diagonal matrix. If $\boldsymbol{\Sigma}$ is further constrained to be isotropic, FA reduces to probabilistic PCA. The discrete counterpart is the Gaussian mixture model (GMM) where $\mathcal{Z} = \{1, 2, \dots, k\}$ is a discrete variable and we have

$$f(\mathbf{y}, z; \boldsymbol{\theta}) = -\frac{1}{2}(\mathbf{y} - \boldsymbol{\mu}_z)^\top \boldsymbol{\Sigma}_z^{-1}(\mathbf{y} - \boldsymbol{\mu}_z). \quad (1.14)$$

For supervised regression of continuous \mathbf{y} given \mathbf{x} , the counterpart of FA is the linear mixed effect model

$$f(\mathbf{y}, z, \mathbf{x}; \boldsymbol{\theta}) = -\frac{1}{2}(\mathbf{y} - \mathbf{F}z - \mathbf{W}\mathbf{x})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{y} - \mathbf{F}z - \mathbf{W}\mathbf{x}). \quad (1.15)$$

and the counterpart of the GMM is the mixture of linear regressions

$$f(\mathbf{y}, z, \mathbf{x}; \boldsymbol{\theta}) = -\frac{1}{2}(\mathbf{y} - \mathbf{W}_z\mathbf{x})^\top \boldsymbol{\Sigma}_z^{-1}(\mathbf{y} - \mathbf{W}_z\mathbf{x}), \quad (1.16)$$

corresponding to learning a separate linear regression model for each cluster component. All of the above can be fit by expectation-maximization algorithms, with the notable exception of probabilistic PCA, for which the exact solution can be found from a single SVD of the design matrix. Extensions to categorical (*i.e.*, classification) models of \mathcal{Y} are mostly studied in the context of mixed effects models within the framework of hierarchical generalized linear models.

Unsupervised linguistic structure discovery. An important line of work in Natural Language Processing has been the use of latent structures for language modeling (*i.e.* learning a distribution over sentences)

in a Bayesian setting, that is by defining a Bayesian network whose observations are sentences and latent variables include structure modeling. Then, parameter inference from raw texts can provide structured representation of texts. Although useful for unsupervised and semi-supervised structured prediction, it is important to bear in mind that part of this line of work is also motivated by the goal of automatically discovering structures that may be useful for linguistic research.

Segmentation models discovers word boundaries (Brent, 1999; Venkataraman, 2001; Goldwater *et al.*, 2006). They are especially useful for discovering word units in languages that do not have explicit boundary markers and speech processing for (non written) low resource languages (Zanon Boito *et al.*, 2022). Unsupervised tagging models learn to group similar words in the same class (Goldwater and Griffiths, 2007). They are mainly based on Hidden Markov Models, possibly with an infinite number of classes (Beal *et al.*, 2001). Syntactic models aim to represent more complex relations between words in a sentence than as a sequence of words. We often differentiate two types of syntactic structures:

- Phrase structures or constituency trees, that model syntax by grouping words in hierarchical spans. Such models are mainly based on latent Probabilistic Context-Free Grammars (Lari and Young, 1990; Clark, 2001; Johnson *et al.*, 2007).
- Dependency trees, that model syntax using bilexical dependencies between words. The main approach is called Dependency Model with Valence (Klein and Manning, 2004).

Beyond the sentence level, previous work considered latent modeling of discourse structures (Chen *et al.*, 2009) and topic segmentation, which aims to model topical changes in a document (Eisenstein and Barzilay, 2008; Du *et al.*, 2013).

Note that these works are not covered in this manuscript. Previous

book of Cohen (2019) covers all basic techniques in purely probabilistic setting (e.g. parameter inference techniques like Markov Chain Monte-Carlo and Variational Inference) including the use of priors to bias models toward linguistically plausible structures. These approaches exploit probability distribution structures and their (simple) parameterization, which is not possible with the neural network setting that we cover in this manuscript. We instead focus on techniques for learning neural models in end-to-end approaches with limited assumptions, including but not limited to techniques described by Kim *et al.* (2018).

Deep models. Sigmoid Belief Networks (SBN, Neal, 1992) and Boltzmann Machines (BM, Ackley *et al.*, 1985) are popular generative neural networks with discrete latent variables that have a long history in machine learning. They are graphical models (Bayesian network in the case of SBN, factor graph in the case of BM) that use implicit parameterization using a small neural network instead of explicit contingency tables. SBNs can naturally describe deep architectures with several layers of latent variables whereas RBMs can be stacked to achieve a similar goal (Hinton *et al.*, 2006; Salakhutdinov and Hinton, 2009). Straightforward approaches to fit these models are based on Markov chain Monte Carlo estimation of the gradient (Neal, 1992; Hinton, 2002; Hinton *et al.*, 2006), which can be slow in practice. Generalization of the Expectation-Maximization (EM, Dempster *et al.*, 1977) algorithm using Mean Field Theory approximation (Parisi and Shankar, 1988) allows fast training of these models (Peterson, 1987; Saul *et al.*, 1996). A downside of EM is that it relies on strong assumption on factors' parameterization (i.e. simple linear projection), and therefore does not extend to complex neural parameterization. This contrasts with methods studied in this manuscript that focus on techniques for learning discrete latent variables that (1) can learn more complex latent structures than binary variables and (2) are compatible with the modern end-to-end learning

framework. Moreover, some of the techniques we described do not have a probabilistic interpretation of latent variables.

Nonlinear models parametrized by neural networks have proven themselves effective for generative modeling. Prominent above them is the *variational auto-encoder* (VAE, Kingma and Welling, 2014; Rezende *et al.*, 2014), which is a Bayesian network where conditional distributions are parameterized by deep neural networks. This means that variational methods used for SBN are not applicable anymore. Key to the success of the VAE is the “evidence lower bound” (ELBO) approximation of the objective:

$$\begin{aligned}
 L(\mathbf{x}; \boldsymbol{\theta}_g) &= -\log \mathbb{E}_{\mathbf{Z}} [g(\mathbf{x}; \mathbf{Z}, \boldsymbol{\theta}_g)] \\
 &= -\log \sum_{z \in \mathcal{Z}} p(z) p(\mathbf{x} | z; \boldsymbol{\theta}_g) \\
 &\leq \underbrace{\text{KL}[p(\mathbf{Z} | \mathbf{x}; \boldsymbol{\theta}_f), p(\mathbf{Z})] - \log \sum_{z \in \mathcal{Z}} p(z | \mathbf{x}; \boldsymbol{\theta}_f) p(z | \mathbf{x}; \boldsymbol{\theta}_g)}_{\text{reconstruction term}} \\
 &= \bar{L}(\mathbf{x}, \boldsymbol{\theta}_g, \boldsymbol{\theta}_f),
 \end{aligned} \tag{1.17}$$

where KL denotes the Kullback–Leibler divergence, $p(\mathbf{Z} | \mathbf{x}; \boldsymbol{\theta}_f)$ the approximate posterior and the the sum is replaced with an integration in case of continuous latent variable \mathbf{Z} . The conditional and approximate posterior distributions are fully specified by the Gibbs distributions

$$p(\mathbf{x} | \mathbf{z}; \boldsymbol{\theta}_g) \propto \exp f(\mathbf{x}; \mathbf{z}, \boldsymbol{\theta}_g), \quad p(\mathbf{z} | \mathbf{x}; \boldsymbol{\theta}_f) \propto \exp g(\mathbf{z}; \mathbf{x}, \boldsymbol{\theta}_f).$$

As such, the reconstruction term of the ELBO is highly similar to the objective Eq. 1.12.

In our framework, we may take $\mathbf{y} = \mathbf{x}$ to represent an autoencoding task, and set, for a Gaussian latent and Gaussian output VAE,

$$\begin{aligned}
 f(\mathbf{z}; \mathbf{x}, \boldsymbol{\theta}_f) &= (\mathbf{z} - \boldsymbol{\mu}_z(\mathbf{x}; \boldsymbol{\theta}_f))^\top \boldsymbol{\Sigma}_z^{-1}(\mathbf{x}; \boldsymbol{\theta}_f) (\mathbf{z} - \boldsymbol{\mu}_z(\mathbf{x}; \boldsymbol{\theta}_f)), \\
 g(\mathbf{x}; \mathbf{z}; \boldsymbol{\theta}_g) &= (\mathbf{x} - \boldsymbol{\mu}_x(\mathbf{z}; \boldsymbol{\theta}_g))^\top \boldsymbol{\Sigma}_x^{-1}(\mathbf{z}; \boldsymbol{\theta}_g) (\mathbf{x} - \boldsymbol{\mu}_x(\mathbf{z}; \boldsymbol{\theta}_g)),
 \end{aligned} \tag{1.18}$$

i. e., a neural network is used to generate the parameters of an observation distribution and of an approximate posterior; this strategy is known as amortization.

In this text, we focus on deep models with discrete, structured latent variables. This differs from works that extend the original VAE with richer priors or structured inference networks (Johnson *et al.*, 2016; Lin *et al.*, 2018; Pearce, 2020; Zhao and Linderman, 2023, amongst others). For a tutorial on latent variable learning with a focus on probabilistic models for language, we refer the reader to the thorough text by Kim *et al.* (2018).

1.5 Roadmap

Before getting into the matter of discrete latent structure, in Chapter 2 we revisit the tools of the trade of (supervised) structure prediction; they will prove essential for the latent case as well.

The central Chapters 3 to 5 form the main part of our text, covering three different directions to take for learning deep networks with discrete latent structure. In Chapter 3 we explore a deterministic approach to learning latent structure, using a fundamental tool in machine learning, namely *relaxation*, at the cost of partially abandoning discreteness. Then, in Chapter 4 we discuss a range of methods that regain discreteness by introducing a gap between the learning objective and the desired model. Finally, in Chapter 5 we study strategies for approximately minimizing the true stochastic objective, allowing for the most flexible latent structure models, at a controllable computational cost. Chapter 6 summarizes the field and provides a table representations of various trade-offs and applicability of the discussed methods, as well as pointers to prominent libraries.

2

Structure Prediction Background

2.1 Overview

Before handling combinatorial structures as latent representations, we must establish a formalism for representing and computationally modelling such structures. In this section, we provide the necessary background, which draws from the field of *structure prediction*. The examples in this section are therefore examples of structured-output models, not of models with latent structure, but many of them will reappear in later sections in latent structure form.

Definition 1. A structured representation $\mathbf{z} \in \mathcal{Z}$ is a discrete object consisting of a combination of interdependent parts $p \in \mathcal{P}$. The set of parts is problem-specific. The set of valid structures $\mathcal{Z} \subseteq \{0, 1\}^{|\mathcal{P}|}$ can encode global structural constraints on which parts can be simultaneously present in a structure, and, like the set of parts, differs based on what is being modeled.

When convenient, we may see \mathbf{z} as a binary vector indexed by \mathcal{P} , such that $z_p = 1$ if part p is present in the structure \mathbf{z} and $z_p = 0$ otherwise. From this perspective, $\mathcal{Z} \subseteq \{0, 1\}^{|\mathcal{P}|}$.¹ This is a powerful

¹Since the parts are a discrete set, indexing by them is equivalent to assigning the parts numeric indices $\{1, 2, \dots, |\mathcal{P}|\}$. We assume an implicit indexing and use either notation interchangeably.

but abstract definition. Figure 2.2 presents some common and useful types of structures, under our established notation. The first two are polar opposite extreme examples, and we describe them in more detail below.

Example 1 (Unconstrained). In order to represent a combination of choices, such as sandwich toppings, we show how we represent a global configuration as a vector of binary variables. The parts are the possible choices, *e.g.*, $\mathcal{P} = \{\text{lettuce}, \text{tomato}, \dots\}$, and each part is encoded with a binary indicator $z_p \in \{0, 1\}$ for each $p \in \mathcal{P}$. Beyond the binary construction, the structures are not constrained:

any subset (including no toppings at all) is acceptable. Thus, $\mathcal{Z} = \{0, 1\}^{|\mathcal{P}|}$. The number of admissible structures is thus exponential in the number of parts. When using such an output space for supervised machine learning, the resulting task is known as multilabel classification (Tsoumakas and Katakis, 2007).

Example 2 (One-of-K). In contrast, the **one-of-K** setup forces us to pick exactly one of the parts from \mathcal{P} . For example, a word in a sentence must have one part-of-speech tag. In this case, $\mathcal{P} = \{\text{NOUN}, \text{VERB}, \dots\}$ and

$$\mathcal{Z} = \left\{ \mathbf{z} \in \{0, 1\}^{|\mathcal{P}|} : \sum_{i=1}^{|\mathcal{P}|} z_i = 1 \right\} = \{\mathbf{e}_1, \dots, \mathbf{e}_{|\mathcal{P}|}\}.$$

The number of admissible configurations is thus exactly the same as the number of parts. In supervised learning, this output space leads to the standard multiclass classification problem (Murphy, 2022, ch. 10.3).

The next three examples illustrate how structure prediction generalizes and explores the space between the two extremes encountered in the first two examples.




Figure 2.1: Some example structures. Left: linear assignment (matching); center: dependency parse tree (directed arborescence); right: binary constituency parse tree (binary tree).

Example 3 (Linear assignment). The linear assignment (one-to-one, or bipartite matching) setup (Hitchcock, 1941) amounts to finding an optimal pairing between two equal-size sets A, B (e.g., horses and jockeys), such that each horse gets one jockey and each jockey gets one horse (figure 2.1, left). Formally,

$$\mathcal{P} = \{i \leftrightarrow j : i \in A, j \in B\},$$

$$\mathcal{Z} = \left\{ z \in \{0, 1\}^{|\mathcal{P}|} : \sum_{i \in A} z_{i \leftrightarrow j} = 1 \text{ for each } j, \sum_{j \in B} z_{i \leftrightarrow j} = 1 \text{ for each } i \right\}$$

Writing $|A| = |B| = n$, notice $|\mathcal{P}| = n^2$. By further noting that any valid assignment is a permutation, we see that $|\mathcal{Z}| = n!$. We will later see how this elementary set already poses challenges, *i.e.*, full probabilistic modeling over this set is intractable.

Example 4 (Non-projective dependency parsing). This example comes from the domain of natural language syntax (Jurafsky and Martin, 2018, Ch. 19) given a sentence, search for the best dependency tree to

describe the dependency (syntactic governance) relationships between its words (figure 2.1, middle): the parent word is interpreted as a *head* word and the child is an *modifier*. A dependency tree is an *arborescence*, that is, a directed rooted subgraph in which every vertex (*i.e.*, word), has at most one incoming arc and is reachable from the root. We next show how to represent such dependency trees in our framework. Given a sentence with words w_1, \dots, w_n , we define the parts of a dependency tree as the possible arcs:

$$\mathcal{P} = \{h \rightarrow m : h \in \{*, w_1, \dots, w_n\}, m \in \{w_1, \dots, w_n\}, h \neq m\},$$

where the asterisk denotes a special token identifying the root. In order for a set of such arcs to correspond to a valid dependency tree, it is required that every node except $*$ have exactly one parent, and that there be no cycles in the resulting graph (Martins *et al.*, 2009a). From a generalization of Kirchoff's matrix-tree theorem (Kirchhoff, 1847) due to Tutte (2001, Thm. VI.27), we obtain $|\mathcal{Z}| = (n + 1)^{n-1}$.

The role of structure prediction is to provide a **prediction rule** for selecting a best structure, and, if possible, a **probability distribution** over all possible structures. In the latter case, we would have a *probabilistic* structure prediction model. This is challenging because the total number of possible structures can grow very quickly (often exponentially) with the size of the problem, see figure 2.2. Given a specified probabilistic structure prediction model $p(\mathbf{z} | x)$, in this book, we consider three challenging problems, which will serve as **computational building blocks** for learning and prediction.

- **Maximization:** find $\arg \max_{\mathbf{z} \in \mathcal{Z}} p(\mathbf{z} | x)$;
- **Expectation:** compute $\mathbb{E}_{\mathbf{Z} \sim p(\mathbf{Z} | x)} [h(\mathbf{Z})]$ for some function h ;
- **Sampling:** draw a stochastic $\mathbf{z} \sim p(\mathbf{Z} | x)$.

structure	description
<p>Unconstrained (multiple choice)</p> <div style="display: flex; flex-direction: column; align-items: flex-start;"> <div style="display: flex; align-items: center; margin-bottom: 5px;"> <input checked="" type="radio"/> lettuce </div> <div style="display: flex; align-items: center; margin-bottom: 5px;"> <input checked="" type="radio"/> tomato </div> <div style="display: flex; align-items: center; margin-bottom: 5px;"> <input type="radio"/> corn </div> <div style="display: flex; align-items: center;"> <input checked="" type="radio"/> red onion </div> </div>	<p>$\mathcal{P} = \{\text{lettuce, tomato, } \dots\}$; Structures can have any combination of parts.</p> <p>$\mathcal{Z} = 2^{ \mathcal{P} }$</p>
<p>One-of-K (XOR)</p> <div style="display: flex; flex-direction: column; align-items: flex-start;"> <div style="display: flex; align-items: center; margin-bottom: 5px;"> <input type="radio"/> noun </div> <div style="display: flex; align-items: center; margin-bottom: 5px;"> <input checked="" type="radio"/> verb </div> <div style="display: flex; align-items: center; margin-bottom: 5px;"> <input type="radio"/> adjective </div> <div style="display: flex; align-items: center;"> <input type="radio"/> preposition </div> </div>	<p>$\mathcal{P} = \{\text{noun, verb, } \dots\}$; Structures have exactly one part ($\mathbf{1}^\top \mathbf{z} = 1$)</p> <p>$\mathcal{Z} = \mathcal{P}$</p>
<p>Non-projective dependency parsing</p>	<p>$\mathcal{P} = \{h \rightarrow m : h \in S \cup \{*\}, m \in S, h \neq m\}$ set of possible arcs. Valid structures are connected and acyclic.</p> <p>$\mathcal{Z} = (S + 1)^{ S -1}$</p>
<p>Linear assignment</p>	<p>$\mathcal{P} = \{i \leftrightarrow j\} = A \times B$, valid structures are exact one-to-one assignments.</p> <p>$\mathcal{Z} = A ! = B !$</p>

Figure 2.2: Illustration of a selection of useful structures, along with their matrix representation.

These building blocks are well studied in the structured output prediction setting, which we review in this chapter. However, in the next chapters we will see that structured latent variable learning also builds upon the same building blocks. Moreover, for latent variable learning, it is usually not enough to just have access to these methods, but we also need to be able to compute or approximate their gradients, with respect to some parameters of the distribution; this is sometimes called “differentiating over the structure” and is the main topic of §3.4.

In this chapter, we review the two main strategies classically employed for structure prediction: incremental prediction and global prediction. Both of them rely on breaking down structures into the constituent parts, but they differ in the trade-offs between representation power and optimality.

2.2 Incremental Prediction

To model a structured variable \mathbf{z} , possibly conditioned on some input x , we may employ a probabilistic model $p(\mathbf{z} | x)$, which must, of course, obey the global constraints (*i.e.*, it must assign zero probability to invalid configurations). The set of possible configurations \mathcal{Z} is discrete and finite. Note that \mathcal{Z} may depend on x — *e.g.*, for part-of-speech tagging, the length of a sentence $|x|$ gives the number of tags in the output sequence.

The probability chain rule gives a way to factor the distribution of \mathbf{z} without any loss of expressiveness:

$$p(\mathbf{z} | x) = \prod_{i=1}^{|\mathcal{P}|} p(z_i | z_1, \dots, z_{i-1}, x) \quad (2.1)$$

This decomposition suggests a method for building up a structure by *incrementally adding parts*.

Connection to multi-label classification. The multi-label setting is a particular instance of unconstrained structure prediction (Example 1). In this context, incremental prediction is known as the *classifier chain* method (Read *et al.*, 2011), which amounts to training N binary classifiers applied in sequence, each classifier having access to the already-predicted labels.

Blockwise prediction. A commonly-occurring pattern involves predicting multiple interdependent *one-of- K* choices (Example 2). If a structure consists of L assignments, each one with K_i choices, for $i = 1, \dots, L$, then the vector \mathbf{z} then can be seen as a concatenation

$$\mathbf{z} = [\mathbf{z}_1, \dots, \mathbf{z}_L]$$

where each $\mathbf{z}_i \in \{\mathbf{e}_1, \dots, \mathbf{e}_{K_i}\}$ is an indicator vector, so $\mathcal{Z} = \{\mathbf{e}_1, \dots, \mathbf{e}_{K_1}\} \times \dots \times \{\mathbf{e}_1, \dots, \mathbf{e}_{K_L}\}$. The dimensionality of \mathbf{z} is thus $|\mathcal{P}| = \sum_{i=1}^L K_i$, and the number of possible configurations is $|\mathcal{Z}| = \prod_{i=1}^L K_i$. We may incrementally predict \mathbf{z} block-wise, using

$$p(\mathbf{z} | x) = \prod_{i=1}^L p(\mathbf{z}_i | \mathbf{z}_1, \dots, \mathbf{z}_{i-1}, x). \quad (2.2)$$

The only difference from Eq. 2.1 is that this time we have a chain of *multi-class* classifiers rather than binary classifiers.

Example 5 (Incremental part-of-speech tagging.) Here, the input x is a sentence of length L , and we seek to generate a sequence \mathbf{z} consisting of L one-of- K assignments: for instance, the sentence “the dog barked” should be tagged as **determiner, noun, verb**. The possible tags are the same for each word, so $K_i = K$, so $|\mathcal{P}| = LK$ and $|\mathcal{Z}| = K^L$.

Despite its simplicity, this paradigm has gained popularity with the rise of neural networks. A particularly powerful extension is to allow

variable-length target structures, $\mathcal{Z} = \bigcup_{L \in \mathbb{N}} \{e_1, \dots, e_K\}^L$. This can be achieved by designating a specific variable assignment as a “stop” control token and imposing that $p(\mathbf{z}_i \mid \mathbf{z}_{i-1} = \text{STOP}) = 0$. The next example illustrates this generalization and underscores many successful recent models in NLP.

Example 6 (Conditional language generation and seq-to-seq). Sentences in natural language can be seen as a sequence of discrete tokens (*e.g.*, words). We may incrementally build up a sentence by appending tokens selected from a *vocabulary* of K words, mapped to $\mathcal{Z} = \{e_1, \dots, e_K\}$, one of which is the special STOP token.

Let $\mathbf{V} \in \mathcal{R}^{d \times K}$ denote a matrix of token embeddings, and a neural network $\mathbf{h}(\mathbf{z}_1, \dots, \mathbf{z}_{i-1}, x, \boldsymbol{\theta}) \in \mathbb{R}^d$. We can parametrize the probability of the next word as

$$p(\mathbf{z}_i \mid \mathbf{z}_1, \dots, \mathbf{z}_{i-1}, x) \propto \exp f(x, \mathbf{z}_1, \dots, \mathbf{z}_{i-1}, \mathbf{z}_i),$$

where

$$f(x, \mathbf{z}_1, \dots, \mathbf{z}_i) := \langle \mathbf{z}_i, \mathbf{V}^\top \mathbf{h}(\mathbf{z}_1, \dots, \mathbf{z}_{i-1}, x) \rangle.$$

This is a neural conditional language generation model, where x represents an input to be converted to a sentence. For image caption generation, x is an image; for speech-to-text transcription, x is a sound, and for translation, x is a sentence in a foreign language. When x is a sequence, such models are often called sequence-to-sequence or seq-to-seq. The parametrization of \mathbf{h} is crucial as it needs to handle variable-length input. Common choices include recurrent networks and transformers (Sutskever *et al.*, 2014; Cho *et al.*, 2014; Vaswani *et al.*, 2017). The parameter sharing between predictors at different positions allows the models to generalize, including to variable-length

sequences. This strategy is sometimes called *autoregressive* generation, drawing parallels with time series analysis.

Transition-based systems, constraints, and automata. The extension of incremental prediction to more complex and constrained structured problems is enabled by transition-based encodings. Here, the assignments z_i are interpreted as *actions* or *transitions*. For instance, the path of a robot through a grid can be seen as L steps, or actions, chosen among {up, down, left, right}. Incremental prediction is popular in parsing, where it is known as transition-based parsing, action-based parsers, or shift-reduce parsing. For parsing, alignments, and other similar more complex structures, the set of admissible actions may be constrained based on the actions already taken.

In general, an arbitrary constraint $h(\mathbf{z}) = 0$ cannot be handled by incremental prediction, because there may be no admissible actions left. Any finer-grained early detection of a violation is only possible for certain forms of constraints, and for well-chosen prediction orderings. Fortunately, a principled way to incorporate constraints into incremental prediction is provided by automata theory. By building an automaton with accepting set equal to the desired feasible set, it is possible to detect wrong paths early.

The next example is a celebrated construction from syntactic parsing of natural language sentences.

Example 7 (Shift-reduce parsing). One flavor of syntactic analysis of language is provided by constituency parsing, where a sentence is recursively divided into smaller phrasal units (figure 2.1). A constituency parse can be represented by a binary tree. We may build up this tree by choosing between two actions (Sagae and Lavie, 2005):

- Shift: add a new leaf node;
- Reduce: merge the two most recently produced nodes.

The process, shown in Table 2.1, corresponds to an encoding of candidate binary parse trees for a sentence of length N as bit vectors of length $d = 2N - 1$, with $z_i = 1$ denoting a shift and $z_i = 0$ a reduce. In order to represent an admissible tree, there are non-trivial constraints. Firstly, reducing is only allowed when there are two nodes available for merging. Equivalently, up to any position k , the shifts must strictly outnumber the reduces, or, $\sum_{i=1}^k z_i > k/2$. In particular, this means the first two actions must always be shifts. Secondly, exactly N shift operations must be performed, since all words must be included in the tree. This can be written as $\sum_{i=1}^d z_i = N$.

The first constraint is naturally phrased in a way that can be checked at every iteration, to decide whether a reduce action is allowed. It is thus already in an FSA form. The second constraint appears to depend on the entire sequence, so we would have to defer its checking until the end. However, if we incrementally satisfy the first constraint, we can never have $\sum_i z_i < N$, so we only have to watch out, incrementally, whether we already “used up” all of the N shifts. If this happens, the algorithm must terminate, filling up the remaining structure with reduce actions.

We note that the choice of transition system for a structured problem is not unique, and moreover, every choice may introduce specific subtle biases. We highlight the bias intrinsic to the shift-reduce model (Drozdov and Bowman, 2017): consider an “uninformed” shift-reduce model that assigns to $p(z_k \mid z_{1:k-1})$ a uniform distribution over the admissible choices at that point. One might mistakenly expect this to correspond to an uniform distribution over all possible binary trees, but this is not the case: the uninformed shift-reduce model prefers highly skewed structures.

Table 2.1: Example of incremental construction of vector \mathbf{z} by a shift-reduce parser. Action S (respectively R) corresponds to a shift (respectively reduce).

Action	\mathbf{z}	Node stack	Input buffer
S	[1]	Sleep	Sleep the clock around
S	[1 1]	Sleep the	the clock around
S	[1 1 1]	Sleep the clock	clock around
R	[1 1 1 0]	Sleep (the clock)	around
S	[1 1 1 0 1]	Sleep (the clock) around	
R	[1 1 1 0 1 0]	Sleep ((the clock) around)	
R	[1 1 1 0 1 0 0]	(Sleep ((the clock) around))	

If instead of uniform probability we introduce a slight preference to shifts, this skew is accentuated, as demonstrated in figure 2.3.

A further complication posed by transition-based systems is that, in some models, multiple sequences of transitions can lead to the same target structure. This phenomenon is known as *spurious ambiguity* (Cohen *et al.*, 2012).

2.2.1 Computation.

We next describe the algorithms available for maximization, sampling, and computing expectations, in incremental structure models such as transition-based models.

Maximization. The reformulation in Eq. 2.1 does not change the fact that there are $\mathcal{O}(2^{|\mathcal{P}|})$ possible paths to consider. As such, exactly finding the highest-probability structure is not feasible in general. A popular heuristic algorithm is the *beam search* algorithm, which keeps track of the k best partial prefixes at each time step and seeks to expand them (see Algorithm 1). The particular case of $k = 1$ is known as *greedy search* and corresponds to directly choosing the most likely




Figure 2.3: Shift-reduce probabilities for binary trees, with shift probability .5 (left) and .8 (right) modulo constraints.




Figure 2.4: The graphical model of incremental prediction is a fully connected acyclic graph. The combinatorial explosion challenge is still present, but delegated to the incremental predictors $p(z_k \mid z_1, \dots, z_{k-1})$.

▷ Approximate maximization using a beam of size k

function BEAMSEARCH(k, g, x)

$B = \{\emptyset\}$

repeat

$B' = \emptyset$

for $[z_1 \dots z_i] \in B$ **do**

if $z_i = \text{STOP}$ **then**

▷ Don't expand sequences whose last token is STOP

$B' \leftarrow B' \cup \{[z_1, \dots, z_i]\}$

else

$w = V^\top h(z_1, \dots, z_i, x)$

for $z_{i+1} \in \{e_1, \dots, e_K\}$ **do**

$\text{score}([z_1, \dots, z_{i+1}]) \leftarrow \text{score}([z_1, \dots, z_i]) + \log \frac{\exp \langle z_{i+1}, w \rangle}{\sum_i \exp w_i}$

$B' \leftarrow B' \cup \{[z_1, \dots, z_{i+1}]\}$

▷ Compute the k -best sequences in B'

$B = \text{top}_{(k)}(\text{score}, B')$

until $\forall [z_1 \dots z_i] \in B : z_i = \text{STOP}$

Algorithm 1: Beam search for conditional language generation and seq-to-seq. Note that, in practice, it is usual to have maximum sequence length to assert that the algorithm will eventually terminate even if not all sequences in the beam reached the STOP token. Moreover, the top-k computation usually renormalizes scores by hypothesis lengths.

transition at each step. The suboptimality of such heuristics is apparent: if the optimal structure starts with an action that seemed unlikely early on, the optimal structure cannot be recovered. The chosen ordering therefore also may impact performance. As such, even if beam search has a long history (Lowerre, 1976), it is still an active research area (Kool *et al.*, 2019; Meister *et al.*, 2020b; Meister *et al.*, 2020a; Meister *et al.*, 2021; Kasai *et al.*, 2024, amongst others). If further factorization assumptions are imposed on the transition model, exact maximization is possible; this is the scenario considered in §2.3.

Sampling. The incremental parametrization of $p(\mathbf{z} \mid x)$ as in Eq. 2.1 can be seen as a dense graphical model, illustrated in figure 2.4. This suggests we may conditionally sample from this distribution by sampling $z_1, \dots, z_{|\mathcal{P}|}$ sequentially: this is known as ancestral sampling. Moreover, samples *without replacement* can be drawn using a *stochastic beam* algorithm (Kool *et al.*, 2020a). Algorithms have been proposed to perform approximate maximization by way of sampling (*e.g.*, Keith *et al.*, 2018).

Expectations. For the same considerations as maximization, it is in general not computationally feasible to compute expectations exactly. However, since exact sampling is available, Monte Carlo methods provide a good approximation. Given samples $\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(s)}$, we have

$$\sum_{\mathbf{z} \in \mathcal{Z}} h(\mathbf{z}) p(\mathbf{z} \mid x) \approx \frac{1}{s} \sum_{i=1}^s h(\mathbf{z}^{(i)}),$$

for any function h .

2.3 Global Prediction

In order to build structured models with global optimality guarantees, and thus open an avenue toward exact maximization and expectations, we will need to sacrifice some expressivity.

Since global prediction models are not always probabilistic, we will consider, for more generality, the *score* of a structure $f(\mathbf{z}, x; \boldsymbol{\theta}_f) : \mathcal{Z} \times \mathcal{X} \rightarrow \mathbb{R}$ as a measure of compatibility of the structure represented by \mathbf{z} with the input (context) x . We may then use the induced Gibbs probability distribution,

$$p(\mathbf{z} \mid x) \propto \exp f(\mathbf{z}, x; \boldsymbol{\theta}_f), \quad (2.3)$$

but other distributions are also possible, and will be discussed later on.

For global prediction, we will make the following assumption.

Assumption 1 (Decomposition into parts). We assume that the score of a structure can additively decompose as:

$$\begin{aligned} f(x, \mathbf{z}, \boldsymbol{\theta}_f) &= \sum_{\substack{p \in \mathcal{P} \\ z_p = 1}} [\mathbf{s}(x, \boldsymbol{\theta}_f)]_p \\ &= \langle \mathbf{z}, \mathbf{s}(x; \boldsymbol{\theta}_f) \rangle, \end{aligned} \tag{2.4}$$

where $\mathbf{s}(\cdot, \boldsymbol{\theta}_f) \in \mathbb{R}^{|\mathcal{P}|}$ is a vector of scores for each possible part, and the score of the structure \mathbf{z} is the sum of scores of the parts that make it up.

The additive decomposition of the score turns into a multiplicative factorization of the probability, as $p(z | x) \propto \prod_{p \in \mathcal{P}} z_p \exp[\mathbf{s}(x, \boldsymbol{\theta}_f)]_p$. For this reason, Assumption 1 is sometimes called a *factorization* assumption. Since the probabilistic interpretation is not necessary, we avoid this term. As we shall see, this assumption removes some but not all computational limitations.

Example 8 (Arc-factored dependency parsing). For the non-projective dependency tree example, the parts \mathcal{P} are the possible arcs $h \rightarrow m$. Following Kiperwasser and Goldberg (2016), consider encoding each word in a sentence of length L with a neural network (e.g., a bidirectional LSTM) with parameters $\tilde{\boldsymbol{\theta}}_f$, giving vectors $\{\mathbf{v}_j : j = 1, \dots, L\}$. We may parametrize an arc-factored dependency parser as

$$[\mathbf{s}(x)]_{h \rightarrow m} = \langle \mathbf{w}, \tanh([\mathbf{v}_h; \mathbf{v}_m]) \rangle. \tag{2.5}$$

The parameters of the arc-factored parser are thus $\boldsymbol{\theta}_f = \{\mathbf{w}\} \cup \tilde{\boldsymbol{\theta}}_f$.

Example 9 (Markov sequence tagging). For sequence tagging (Lafferty *et al.*, 2001), given a sentence of length L and a tag set of size T , we may build a globally-optimal model by imposing a Markov assumption, *i.e.*, scoring *pairs* of consecutive assignments. We may encode this by using *bigram parts*: $\mathcal{P} = \{(i, t, t') : 1 \leq i < L, 1 \leq t, t' \leq T\}$ where t and t' are consecutive tags at position i and $i + 1$ (subject to a specific handling for the last word transition). One way to parameterize this model is

$$[\mathbf{s}(x)]_{i,t,t'} = \langle \mathbf{v}_i, \mathbf{w}_t \rangle + b_{t,t'}, \quad (2.6)$$

where \mathbf{v}_i are vector representations for each word (*e.g.*, from some neural net module), matrix $\mathbf{W} \in \mathbb{R}^{T \times H}$ is an “output layer” containing embeddings of each tag, and $\mathbf{B} \in \mathbb{R}^{T \times T}$ is a pairwise matrix of transition scores, *i.e.*, $b_{\text{det}, \text{noun}}$ is a bonus reward for tagging a determiner right before a noun. The parameters are thus $\boldsymbol{\theta}_f = \{\mathbf{W}, \mathbf{B}\} \cup \widetilde{\boldsymbol{\theta}}_f$.

2.3.1 The Marginal Polytope

The decomposition assumption allows us to recast structure prediction problems as optimization over the convex hull of possible structures

$$\mathcal{M} = \text{conv}(\mathcal{Z}). \quad (2.7)$$

This set is often called the *marginal polytope* (Wainwright and Jordan, 2008). In a sense, the simplest possible marginal polytope is the convex hull of the one-of- K indicator vectors $\mathcal{Z} = \{\mathbf{e}_1, \dots, \mathbf{e}_K\}$, which is the simplex \triangle_K . For any other finite $\mathcal{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_{|\mathcal{Z}|}\}$ the marginal

polytope can be written as a linear transformation of the $|\mathcal{Z}|$ -simplex,²

$$\text{conv}(\mathcal{Z}) = \left\{ \sum_{i=1}^{|\mathcal{Z}|} \alpha_i \mathbf{z}_i : \boldsymbol{\alpha} \in \Delta_K \right\}. \quad (2.8)$$

The (possibly impractically long) vectors $\boldsymbol{\alpha} \in \Delta_{|\mathcal{Z}|}$ can be interpreted as distributions over the set of structures \mathcal{Z} , so any point $\hat{\mathbf{z}} \in \text{conv}(\mathcal{Z})$ can be seen as an *expected structure*:

$$\hat{\mathbf{z}} \in \text{conv}(\mathcal{Z}) \iff \hat{\mathbf{z}} = \mathbb{E}_{\mathbf{Z} \sim \boldsymbol{\alpha}}[\mathbf{Z}], \quad (2.9)$$

for at least one $\boldsymbol{\alpha} \in \Delta_{|\mathcal{Z}|}$.

2.3.2 Computation

We describe here the computational challenges of global structure prediction. A summary of key algorithms for structures of interest is presented in Table 2.2.

Maximization. Finding the highest-scoring structure, *i.e.*, solving $\arg \max_{\mathbf{z} \in \mathcal{Z}} f(\mathbf{z}; x)$, is a difficult combinatorial optimization problem. But, for many specific problems of interest, specialized, case-by-case algorithms are available. For example, for sequence tagging, the Viterbi algorithm retrieves the global optimum (Forney, 1973). Under Assumption 1, finding the highest-scoring structure is equivalent to maximizing a linear function over a polytope,

$$\max_{\mathbf{z} \in \mathcal{Z}} f(\mathbf{z}; x) = \max_{\mathbf{z} \in \text{conv}(\mathcal{Z})} \langle \mathbf{z}, \mathbf{s}(x) \rangle, \quad (2.10)$$

because the maximum of a linear function over a polyhedral domain is always attained at a vertex (a fact that will be revisited in Chapter 3). For this reason, general structured maximization can thus be

²Some authors call every such convex hull a $|\mathcal{Z}|$ -simplex, and $\Delta_{|\mathcal{Z}|}$ the *canonical* or *probability* simplex.

tackled via linear programming, but describing $\text{conv}(\mathcal{Z})$ may require exponentially many constraints. One strategy for approaching complex structured problems is therefore developing compact constrained formulations (Punyakanok *et al.*, 2004; Roth and Yih, 2005; Martins *et al.*, 2009b), either exact or approximate. Problems involving intersections of tractable substructures can be approached via *dual decomposition*, an approach that tackles an outer relaxation of the linear programming formulation (Rush *et al.*, 2010; Komodakis *et al.*, 2007; Martins *et al.*, 2015).

Expectations. Under a parts-factored Gibbs distribution, computing $\hat{z} = \mathbb{E}[Z]$ is known as “marginal inference”, because each coordinate of the result is the marginal probability of a part being selected: $[\hat{z}]_r = p(Z_r = 1 \mid x)$. Like maximization, this is intractable in general, but efficient exact algorithms are available for many structures of interest; summarized in Table 2.2. Even when a marginal inference algorithm is available, computing general expectations $\mathbb{E}[h(Z)]$ is not necessarily possible unless h is affine.

Marginal inference is also the gradient of the log-normalizing constant of the Gibbs distribution with respect to the factorized scores (Wainwright and Jordan, 2008; Eisner, 2016):

$$\partial_s \log \sum_{z \in \mathcal{Z}} \exp \langle z, s \rangle = \mathbb{E}[Z]. \quad (2.11)$$

Marginal inference is therefore necessary in maximum likelihood learning for supervised structure prediction, and, as we shall later see, it also provides ways of learning structured latent variables.

Sampling. In contrast to incremental models, exact sampling from $p(z \mid x)$ is more difficult. For structures tractable with dynamic programming, we describe an algorithm in §2.3.3. Figure 2.5 illustrates samples




Figure 2.5: Samples drawn from a Needleman-Wunsch alignment model using a dynamic programming *Forwards Filtering, Backwards Sampling* strategy.

drawn using the aforementioned method from a Needleman-Wunsch model often used in DNA sequence alignment. For non-projective dependency trees, Zmigrod *et al.* (2021) and Stanojević (2022) propose an algorithm involving repeated calls to marginal inference. A powerful general framework called **perturb-and-MAP** (Papandreou and Yuille, 2011; Hazan and Jaakkola, 2012) attempts to reduce sampling to maximization of $\langle z, s + U \rangle$ where U is a perturbation drawn from a noise distribution; in general, such samples do not come from the exact Gibbs distribution, but have proven useful for learning and uncertainty quantification (Corro and Titov, 2019a; Corro and Titov, 2019b).

2.3.3 Structures Encoded as Paths in a DAG.

A specific case in which global structure prediction is tractable is when we can represent structures as paths from a source s to a target t in a weighted directed acyclic graph (DAG) $\mathcal{G} = (\mathcal{V}, \mathcal{E}, w)$; in this




Figure 2.6: Representing Markov sequence tagging (left) and monotonic sequence alignment (right) as paths in DAGs. Edge weights correspond to part scores.

case, all three computational primitives discussed can be computed in time $\Theta(|\mathcal{V}| + |\mathcal{E}|)$ using straightforward variants of the Viterbi algorithm (Huang, 2008).³ Figure 2.6 shows how we may represent Markov sequence tagging (Lafferty *et al.*, 2001) and monotonic sequence alignments (Vintsyuk, 1968; Needleman and Wunsch, 1970; Wagner and Fischer, 1974) as paths in a DAG. The edges in the DAG correspond to parts of the structure, and the edge weights to the part scores; therefore the total weight of a path is the score of a structure. We give a general form of the algorithms in Algorithm 2. Such algorithms are instances of dynamic programming (Bellman, 1954) on semirings (Mohri, 2002) and were independently (re)discovered a number of times. The maximization algorithm is credited to Viterbi (1967), the algorithm for log-normalization is due to Baum (1972), and the sampling algorithm, known as *forward filtering, backward sampling* (FFBS) is due to Frühwirth-Schnatter (1994).

³In fact, Huang (2008) generalize this even further to *hypergraphs*; with algorithms essentially unchanged. We present the simpler case for clarity.

Table 2.2: Summary of algorithms for maximization and marginal inference, for a selection of structures of interest.

	Maximization	Marginals
Sequence tagging	Viterbi (Viterbi, 1967; Rabiner, 1989)	Forward-Backward (Baum, 1972; Rabiner, 1989)
Constituency trees	CKY (Kasami, 1966; Younger, 1967; Cocke and Schwartz, 1970)	Inside-Outside (Baker, 1979)
Sequence alignments	(Vintsyuk, 1968; Needleman and Wunsch, 1970; Wagner and Fischer, 1974; Sakoe and Chiba, 1978)	Soft-DTW (Cuturi and Blondel, 2017)
Non-projective dependency trees	Maximum Spanning Arbor- borescence (Chu and Liu, 1965; Edmonds, 1967)	Matrix-Tree (Kirchhoff, 1847)
Linear assignment	Kuhn-Munkres (Kuhn, 1955), Jonker and Volgenant (1987), Auction (Bertsekas, 1988)	intractable, #P-complete (Valiant, 1979; Garey and Johnson, 1979; Taskar, 2004)

2.4 Summary

In this chapter, we have provided the formalisms and algorithms needed for computations with structures, as studied for predicting structured outputs. We have covered two very different directions. The first is an incremental one, allowing expressive parametrizations, at the cost of tractability of maximization and marginalization. The second is a global one, limiting the possible parametrization of structure scores, but resulting often in tractable algorithms. In the next chapters, we switch gears to latent structure, and we shall see how the discussed algorithms show up again as the required computational building blocks for learning and prediction with structured latent representations.

function MAXIMIZE(\mathcal{G}) ▷ Viterbi algorithm

for $v \in \text{topologic_order}(\mathcal{V})$ **do**

$d_v \leftarrow \max_{u:uv \in \mathcal{E}} d_u + w(uv)$

$p_v \leftarrow \arg \max_{u:uv \in \mathcal{E}} d_u + w(uv)$

$v \leftarrow t$; path = ()

while $v \neq s$ **do**

path $\leftarrow p_v v + \text{path}$

$v \leftarrow p_v$

return d_t, path ▷ $d_t = \max_{z \in \mathcal{Z}} f(z)$; path = $\arg \max_{z \in \mathcal{Z}} f(z)$

function LOGNORMALIZER(\mathcal{G}) ▷ Forward algorithm

for $v \in \text{topologic_order}(\mathcal{V})$ **do**

$c_v \leftarrow \log \sum_{u:uv \in \mathcal{E}} \exp d_u + w(uv)$

return c_t ▷ $c_t = \log \sum_{z \in \mathcal{Z}} \exp f(z)$

function SAMPLE(\mathcal{G}) ▷ FFBS algorithm

for $v \in \text{topologic_order}(\mathcal{V})$ **do**

$c_v \leftarrow \log \sum_{u:uv \in \mathcal{E}} \exp d_u + w(uv)$

$v \leftarrow t$; path = ()

while $v \neq s$ **do**

sample u w.p. $\begin{cases} \exp(c_u - c_v + w(uv)), & uv \in \mathcal{E} \\ 0, & uv \notin \mathcal{E} \end{cases}$

path $\leftarrow uv + \text{path}$

$v \leftarrow u$

return path ▷ path is a sample from $p(\mathbf{Z} = \mathbf{z}) \propto \exp f(\mathbf{z})$

Algorithm 2: Viterbi algorithms for maximization, log-normalizer, and sampling paths from s to t in a weighted DAG $\mathcal{G} = (\mathcal{V}, \mathcal{E}, w)$ (Huang, 2008). The marginals $\mathbb{E}[\mathbf{Z}]$ can be computed by automatic differentiation of the log-normalizer Eq. 2.11.

3 Continuous Relaxations

3.1 Challenges of Deterministic Choices

Deterministic latent representations are a straightforward way to define the mapping from x to a representation \hat{z} as a function:

$$\hat{z}(x, \theta_f). \quad (3.1)$$

Given such a decision rule, the end-to-end downstream model is obtained through function composition of g with \hat{z} :

$$\tilde{g}(x, y, \theta_f, \theta_g) = g(x, y, \hat{z}(x, \theta_f); \theta_g). \quad (3.2)$$

This strategy resembles a pipeline, but instead of using an off-the-shelf frozen model $\hat{z}(x, \theta_f)$ to predict the latent representation, we will aim to train θ_f and the downstream θ_g together.

Notice that $\tilde{g}(x, y)$ depends on both θ_f and θ_g . The gradient with respect to the decoder parameters poses no problems:

$$\partial_{\theta_g} \tilde{g}(x, y, \theta_f, \theta_g) = \partial_4 g(x, y, \hat{z}(x), \theta_g).$$

In contrast, the dependency on θ_f is indirect, requiring application of the chain rule through \hat{z} (*cf.* Eq. 1.8):

$$\partial_{\theta_f} \tilde{g}(x, y, \theta_f, \theta_g) = \underbrace{\partial_2 \hat{z}(x, \theta_f)}_{\text{grad. through encoder}} \circ \overbrace{\partial_3 g(x, y, \hat{z}(x), \theta_g)}^{\text{downstream grad. w.r.t. } \hat{z}}. \quad (3.3)$$

The breakdown in Eq. 3.3 highlights that deterministic end-to-end learning has the following requirements driven by the two terms. In particular, the following assumption constrains the possible downstream models $g(x, y, \mathbf{z}, \boldsymbol{\theta}_g)$ that can be used with deterministic latents.

Assumption 2 (Continuous decoder). Gradient-based learning of deterministic latent variables requires that:

1. The latent values $\mathbf{z} \in \mathcal{Z}$ live in a vector space, *i. e.*, $\mathcal{Z} \subset \mathbb{R}^d$.
2. The downstream g accepts weighted averages of structures, *i. e.*, it is defined at least on $\text{conv}(\mathcal{Z})$.
3. g is almost everywhere differentiable w.r.t. \mathbf{z} .

Without this assumption, the downstream gradient $\partial_3 g$ cannot be defined: this partial derivative quantifies the response of g for small changes in its third argument (*i. e.*, \mathbf{z}), but if \mathbf{z} is not encoded as a continuous vector (and is instead, *e. g.*, a string label), then small changes cannot be quantified.

Differentiating through discrete mappings. Even if Assumption 2 is satisfied, the first term in Eq. 3.3 may impede learning. Consider the argmax mapping $\hat{\mathbf{z}}(x, \boldsymbol{\theta}_f) := \arg \max_{\mathbf{z} \in \mathcal{Z}} f(x, \mathbf{z}, \boldsymbol{\theta}_f)$. If \mathcal{Z} is a discrete set, this mapping is discontinuous and flat almost everywhere. The discontinuity is due to “jumps” between discrete values, and if the changes are small enough to not warrant jumps then we would get no change at all, *i. e.*, $\partial_{\boldsymbol{\theta}_f} \hat{\mathbf{z}} = \mathbf{0}$. There are two main ways to address this challenge. In this chapter, we consider **continuous, differentiable deterministic mappings** $\hat{\mathbf{z}}$, allowing exact calculation of Eq. 3.3 for end-to-end training. In Chapter 4, we shall see another approach.

To ease into our construction, consider the rule that assigns the **average representation**:

$$\hat{\mathbf{z}}_H(x; \boldsymbol{\theta}_f) := \sum_{\mathbf{z} \in \mathcal{Z}} \alpha(\mathbf{z}, x; \boldsymbol{\theta}_f) \mathbf{z} \quad (3.4)$$

where the weights α are defined as

$$\alpha(\mathbf{z}, x; \boldsymbol{\theta}_f) := \frac{\exp f(\mathbf{z}, x; \boldsymbol{\theta}_f)}{\sum_{\boldsymbol{\zeta} \in \mathcal{Z}} \exp f(\boldsymbol{\zeta}, x; \boldsymbol{\theta}_f)}. \quad (3.5)$$

The weights α are differentiable in $\boldsymbol{\theta}_f$, and therefore so is $\hat{\mathbf{z}}_H$.

This construction is related to the probabilistic latent variable formulation (Eq. 1.3): the weights α can be seen as a distribution of $p(\mathbf{Z} | x)$, thus $\hat{\mathbf{z}}_H = \mathbb{E}_{p(\mathbf{Z}|x)}[\mathbf{Z}]$ is the “expected assignment”. Therefore, this relaxation can be seen as moving the expectation inside the model, replacing $\mathbb{E}[g(\mathbf{Z})]$ by $g(\mathbb{E}[\mathbf{Z}])$. The latter is often much easier to compute.¹

The following examples demonstrate the implication of relaxations.

Example 10. Sigmoid Belief Networks. Consider a Sigmoid Belief Network (Neal, 1992): a probabilistic model defined with observed random variables \mathbf{X} (input) and \mathbf{Y} (output) and latent variables \mathbf{Z} as follows:

$$\begin{aligned} p(\mathbf{Z}_i = 1 | \mathbf{x}) &= \text{sigmoid}(\langle \mathbf{w}_i, \mathbf{x} \rangle), \\ p(\mathbf{Y}_j = 1 | \mathbf{z}) &= \text{sigmoid}(\langle \mathbf{v}_j, \mathbf{z} \rangle). \end{aligned}$$

where \mathbf{Y} and \mathbf{Z} are vectors concatenating the corresponding Bernoulli variables. It is usual to train using the negative log-likelihood loss,

$$L(\mathbf{x}, \mathbf{y}) = -\log \mathbb{E}_{p(\mathbf{Z}_1, \dots, \mathbf{Z}_d | \mathbf{x})} [p(\mathbf{Y} = \mathbf{y} | \mathbf{Z}_1, \dots, \mathbf{Z}_d)]$$

¹An additional insight is that if g is linear in \mathbf{z} , the two are equal, and if g is convex in \mathbf{z} , then the relaxation yields a lower bound, by Jensen’s inequality. However, in deep learning, g is usually non-convex in \mathbf{z} .

However, this objective is intractable as the conditional distribution $p(\mathbf{Y} \mid \mathbf{Z}_1, \dots, \mathbf{Z}_d)$ is parameterized by a shallow neural network. We can approximate this loss function via the following term:

$$\approx -\log p(\mathbf{Y} \mid \hat{z}_1(\mathbf{x}), \dots, \hat{z}_d(\mathbf{x}))$$

where $\hat{z}_i(\mathbf{x})$ is the expected value of latent variable \mathbf{Z}_i :

$$\begin{aligned} \hat{z}_i(\mathbf{x}) &= \mathbb{E}_{\mathbf{Z}_i}[\mathbf{Z}_i \mid \mathbf{x}] = 1 \cdot p(\mathbf{Z}_i = 1 \mid \mathbf{x}) + 0 \cdot p(\mathbf{Z}_i = 0 \mid \mathbf{x}) \\ &= \text{sigmoid}(\langle \mathbf{w}_i, \mathbf{x} \rangle). \end{aligned}$$

In other words, the relaxation yields a standard deterministic feed-forward neural network with one hidden layer with sigmoid activations. Note that this is different from the mean field theory approach that relies on approximate posterior inference in a generalized EM algorithm (Saul *et al.*, 1996), and which is therefore highly dependent on the specific parameterization of the distribution $p(\mathbf{Y} \mid \mathbf{z})$.

The next two examples illustrate cases where relaxations cannot be applied because Assumption 2 does not hold.

Example 11. Coarse module selection. Consider a model with a binary latent variable \mathbf{Z} and a decoder defined as follows:

$$g(x, y, z) = \begin{cases} \text{LSTM}(x, y), & z = 0; \\ \text{transformer}(x, y), & z = 1. \end{cases} \quad (3.6)$$

In other words, \mathbf{Z} determines what architecture to use for prediction. As given, g is undefined for relaxed $z \in (0, 1)$. In order to relax g , we would need to come up with a meaningful interpolation between LSTMs and transformers, to support evaluation at, *e.g.*, $\hat{z} = 0.5$. Due to the heterogeneous nature of the two possible modules, no such

relaxation seems any more useful or efficient than explicitly computing $\mathbb{E}[g(x; y, Z)]$, which requires running both modules, defeating the entire purpose of the relaxation approach.

Example 12. Program induction. Consider the task of inducing latent programs in the form of code, given expected outcomes y (Liang *et al.*, 2010), possibly with helpful natural language descriptions x (*e.g.*, Wong *et al.*, 2021). Let Z be a latent program string, and $g(x, y, Z)$ be computed by compiling and executing Z as code. As compilers cannot interpolate between discrete programs, relaxations of Z are not applicable with this choice of g .

In this section, we present relaxation methods for latent structure. All presented methods share this intuition of returning an average latent choice, including generalizations toward structured and sparse representations.

3.2 Regularized Argmax Operators

If our model were a well-trained pipeline and we had to make an assignment to the latent variable during the forward pass, a sensible choice is the highest-scoring one:

$$\begin{aligned} \hat{z}_0(x; \theta_f) &:= \arg \max_{z \in \mathcal{Z}} f(x, z, \theta_f) \\ &= \arg \max_{z \in \mathcal{Z}} \langle z, \mathbf{s}(x, \theta_f) \rangle. \end{aligned} \tag{3.7}$$

where the second equality comes from Assumption 1. In the sequel, we shall omit the dependency of \mathbf{s} on x and θ_f and just study the mapping $\mathbf{s} \rightarrow \hat{z}$. If we can obtain a continuous, differentiable function $\mathbf{s} \rightarrow \hat{z}$,

then the composition rule (chain rule) gives us gradients *w.r.t.* $\boldsymbol{\theta}_f$ by composing $\boldsymbol{\theta}_f \rightarrow \mathbf{s}(x, \boldsymbol{\theta}_f) \rightarrow \hat{\mathbf{z}}(x, \boldsymbol{\theta}_f)$. We thus zoom in on

$$\hat{\mathbf{z}}_0(\mathbf{s}) := \arg \max_{z \in \mathcal{Z}} \langle \mathbf{z}, \mathbf{s} \rangle.$$

The argmax in the right-hand side is a **discrete** optimization problem, and, as such, the mapping $\boldsymbol{\theta}_f \rightarrow \hat{\mathbf{z}}_0$ is discontinuous. Moreover, this mapping is also almost everywhere flat, because so is the mapping $\mathbf{s} \rightarrow \hat{\mathbf{z}}_0$. The following example illustrates this.

Example 13 (Single Coinflip). Take $\mathcal{Z} = \{0, 1\}$ and let $s \in \mathbb{R}$. The problem $\hat{z}(s) = \arg \max_{z \in \{0,1\}} zs$ can be solved by considering the sign of s . If $s > 0$, then $1s > 0s$ so $z = 1$. If $s < 0$, then $1s < 0s$ so $z = 0$. If $s = 0$, we have $1s = 0s$ so either choice for z is optimal. This means the mapping $\hat{z}(s)$ is not always a function, since at $s = 0$ it is multi-valued. Even ignoring this pathology, if we vary s continuously, the maximizer z almost never changes, except abruptly at $s = 0$. In other words, the derivative $\partial \hat{z}(s)$ is zero almost everywhere, preventing any learning with gradient backpropagation.

Chapter 4 is dedicated to heuristic methods that replace this flat gradient with an informative (but wrong) one. Here, instead, we designed the assignment $\hat{\mathbf{z}}_H$ which circumvents these issues.

Continuous relaxation. The first step in our construction is to apply the result

$$\begin{aligned} \hat{\mathbf{z}}_0(\mathbf{s}) &= \arg \max_{z \in \mathcal{Z}} \langle \mathbf{z}, \mathbf{s} \rangle \\ &\subseteq \arg \max_{z \in \text{conv}(\mathcal{Z})} \langle \mathbf{z}, \mathbf{s} \rangle. \end{aligned} \tag{3.8}$$

This is a consequence of a theorem of Dantzig *et al.* (1955), fundamental to the field of *linear programming*: the maximum of an optimization

problem over a polytope is always achieved at one of the vertices (but, possibly, at more than one.) This result lets us relax a discrete problem to a continuous but constrained one, at the cost of allowing spurious solutions at threshold points.

Example 13 (Single Coinflip, continued). Returning to the one-dimensional case, the convex hull of $\{0, 1\}$ is the closed interval $[0, 1]$. The cases $s > 0$ and $s < 0$ still yield the optimum $z = 1$ and $z = 0$ respectively. However, the case $s = 0$ now admits any $z \in [0, 1]$ as solution. And, promised by Eq. 3.8, this example verifies $\{0, 1\} \subset [0, 1]$.

At first, this might not seem like a step forward, as we made the problem harder by increasing the search space. However, the continuous formulation allows us to flexibly “reconfigure” the maximization problem to obtain desired properties, including continuity and differentiability of $\mathbf{s} \rightarrow \hat{\mathbf{z}}_0$.

Smoothing. The second step relies on the insight that the argmax mapping of an optimization problem with *strongly concave* objective over a convex domain is a smooth function. The objective of Eq. 3.8 is linear, thus concave but not strongly so. We then perform a further relaxation by adding a **strongly concave regularizer** $G(\mathbf{z})$ to the objective, yielding:

$$\hat{\mathbf{z}}_G := \arg \max_{\mathbf{z} \in \text{conv}(\mathcal{Z})} \langle \mathbf{z}, \mathbf{s} \rangle + G(\mathbf{z}). \quad (3.9)$$

Strong concavity, along with non-emptiness of \mathcal{Z} , ensure that the maximizer exists and is unique, so the mapping $\mathbf{s} \rightarrow \hat{\mathbf{z}}_G$ is well-defined, Lipschitz continuous, and thus differentiable almost everywhere (Zălinescu, 2002, Corollary 3.5.11; Kakade *et al.*, 2012, Theorem 3).

We first illustrate a popular choice of H in the one-dimensional example above. Then, in §3.3, we study more relaxations in the categorical

case, followed by a discussion of the structured generalizations in §3.4.

Example 13 (Single Coinflip, continued). Let $H_{FD}(z) := -z \log z - (1-z) \log(1-z)$. This function is sometimes called the Fermi-Dirac entropy (Borwein and Lewis, 2010, Section 3.3). We then have the optimization problem

$$\arg \max_{0 \leq z \leq 1} zs + H_{FD}(z).$$

Ignoring the constraints and taking the gradient of the objective w.r.t. z gives the unconstrained optimality condition $s + \log(1-z) - \log(z) = 0$, with unique solution $z = (1 + \exp(-s))^{-1} = \text{sigmoid}(s)$, the logistic sigmoid function. We can now check that for any s , this solution also verifies the constraints, therefore it is correct. We thus have the relaxation $\hat{z}_{H_{FD}}(s) = \text{sigmoid}(s)$. This result agrees with the theoretical guarantee that $s \rightarrow \text{sigmoid}(s)$ is single-valued, continuous, and differentiable almost everywhere (in this case, it turns out it is differentiable everywhere).

3.3 Categorical Relaxation and Attention

Let us now consider the categorical (one-of- K) case, *i.e.*, $\mathcal{Z} = \{\mathbf{e}_1, \dots, \mathbf{e}_K\}$, as in Example 2, and thus $\text{conv}(\mathcal{Z}) = \Delta^K$.

Points $z \in \Delta^K$ can be interpreted as discrete probability distributions (categoricals), suggesting that a meaningful regularization might be the Shannon entropy:

$$H_1(\mathbf{z}) = - \sum_{i=1}^K z_i \log z_i. \quad (3.10)$$

This function is strongly concave over the simplex, and the resulting

relaxation turns out to correspond to another familiar function:

$$\hat{\mathbf{z}}_{H_1}(\mathbf{s}) := \arg \max_{\mathbf{z} \in \Delta^K} \langle \mathbf{z}, \mathbf{s} \rangle + H_1(\mathbf{z}) =: \text{softmax}(\mathbf{s}). \quad (3.11)$$

We can thus regard softmax as a differentiable relaxation of the one-of- K argmax mapping. Since softmax is defined using the exponential function, the output of softmax is **fully dense**: $\hat{\mathbf{z}}_{H_1}$ is elementwise strictly positive. In other words, there is no possible score vector $\mathbf{s} \in \mathbb{R}^K$ such that $[\hat{\mathbf{z}}_{H_1}(\mathbf{s})]_j = 0$ for some j . And since the vertices are one-hot vectors, this also rules them out: there is no $\mathbf{s} \in \mathbb{R}^K$ such that $\hat{\mathbf{z}}_{H_1}(\mathbf{s}) \in \mathcal{Z}$.

Starting from this motivation, Martins and Astudillo (2016) proposed *sparsemax* as a sparse alternative to softmax. It corresponds to regularizing the argmax with the Gini entropy

$$H_2(\mathbf{z}) = \frac{1}{2} \sum_{i=1}^K z_i(1 - z_i) = -\frac{1}{2} \|\mathbf{z}\|^2 + \frac{1}{2}, \quad (3.12)$$

where the rightmost equality uses $\sum_i z_i = 1$. Thus,

$$\begin{aligned} \hat{\mathbf{z}}_{H_2}(\mathbf{s}) &:= \arg \max_{\mathbf{z} \in \Delta^K} \langle \mathbf{z}, \mathbf{s} \rangle + H_2(\mathbf{z}) \\ &= \arg \min_{\mathbf{z} \in \Delta^K} \frac{1}{2} \|\mathbf{s} - \mathbf{z}\|^2 =: \text{sparsemax}(\mathbf{s}). \end{aligned} \quad (3.13)$$

Contrary to softmax, sparsemax does not rule out sparse points and can even return a vertex, provided the difference between the highest and second-highest score is large enough – this behaviour is analogous to the *margin property* of max-margin losses applied in a latent setting (Blondel *et al.*, 2020). Yet, as a construction of the form of Eq. 3.9, it is continuous and differentiable almost everywhere.

More generally, Blondel *et al.* (2020) identify *generalized entropies* that lead to similar properties as softmax and sparsemax. We highlight

one such useful construction, the α -entmax family of mappings recovers softmax and sparsemax, and arises from regularization with the Tsallis entropy,

$$H_\alpha(\mathbf{z}) = \sum_i \frac{z_i - z_i^\alpha}{\alpha(\alpha - 1)}. \quad (3.14)$$

Solutions can be written in a thresholded form,

$$\hat{\mathbf{z}}_{H_\alpha}(\mathbf{s}) = [(\alpha - 1)\mathbf{s} - \tau\mathbf{1}]_+^{1/\alpha-1}, \quad (3.15)$$

allowing efficient computation via reduction to one-dimensional root finding. Moreover, for a few α values, there exist exact algorithms based on sorting or median-finding algorithms: for $\alpha = 2$, such algorithms are known since Held *et al.* (1974) and Brucker (1984) and recently revisited by Condat (2016) and Duchi *et al.* (2008). For $\alpha = 1.5$, a similar quasilinear algorithm was derived by Peters *et al.* (2019). These algorithms rely on closed-form solutions to systems of linear, and, respectively, quadratic equations and are thus not readily extensible to other α values.

Gradients of relaxed mappings. Recall that our main motivation for the relaxations discussed so far is to obtain a differentiable mapping with nontrivial gradient $\partial\hat{\mathbf{z}}(\mathbf{s})$. For the softmax case, the gradient can be found by routine application of vector calculus:

$$\partial\hat{\mathbf{z}}_{H_1}(\mathbf{s}) = \text{diag}(\text{softmax}(\mathbf{s})) - \text{softmax}(\mathbf{s})\text{softmax}(\mathbf{s})^\top. \quad (3.16)$$

For the general α -entmax family, Peters *et al.* (2019) show that a (generalized) Jacobian of the mapping $\mathbf{s} \rightarrow \hat{\mathbf{z}}_{H_\alpha}$ has the expression

$$\partial\hat{\mathbf{z}}_{H_\alpha}(\mathbf{s}) = \text{diag}(\mathbf{q}) - \frac{\mathbf{q}\mathbf{q}^\top}{\mathbf{1}^\top\mathbf{q}}, \text{ where } q_i = \begin{cases} (\hat{z}_i)^{2-\alpha}, & \hat{z}_i > 0; \\ 0, & \text{otherwise,} \end{cases} \quad (3.17)$$

```

function ENTMAX( $\mathbf{s} \in \mathbb{R}^d, \alpha \in (1, \infty), T \in \mathbb{N}$ )
   $\tau_L \leftarrow \alpha \max(\mathbf{s}) - 1$ 
   $\tau_R \leftarrow \alpha \max(\mathbf{s}) - d^{1-\alpha}$ 
  loop  $T$  times
     $\tau_M \leftarrow (\tau_L + \tau_R)/2$ 
     $\tilde{\mathbf{z}} \leftarrow [(\alpha - 1)\mathbf{s} - \tau_M \mathbf{1}]_+^{1/\alpha-1}$ 
    if  $\langle \mathbf{1}, \tilde{\mathbf{z}} \rangle < 1$  then
       $\tau_L \leftarrow \tau_M$ 
    else
       $\tau_R \leftarrow \tau_M$ 
  return  $\tilde{\mathbf{z}}$ 

```

Algorithm 3: Computing α -entmax by root finding via bisection. The algorithm finds τ approximately by identifying its initial bounds based on the maximum value in \mathbf{s} , then recursively halves the search interval.

and $\hat{\mathbf{z}} = \hat{\mathbf{z}}_{H_\alpha}(\mathbf{s})$ is shorthand for the solution vector. For $\alpha = 1$ this yields the softmax Jacobian, while for $\alpha > 1$ it is a sparse matrix with zeros for all rows and columns outside of the support.

Correia *et al.* (2019) further derive derivatives with respect to α in order to tune this additional parameter via gradient methods.

Example 14. One-of-K and neural attention. Attention mechanisms in deep learning are differentiable components that are employed for extracting a single contextual representation of a set of K objects (*e.g.*, words) individually encoded as vectors. In particular, in the key-value attention formulation (Vaswani *et al.*, 2017), each object is encoded by a key vector \mathbf{k}_i and a value vector \mathbf{v}_i , row-stacked in the

matrices \mathbf{K} and \mathbf{V} . Given a query \mathbf{q} , attention is

$$\text{Attn}(\mathbf{q}, \mathbf{V}, \mathbf{K}) = \sum_{i=1}^K \alpha_i \mathbf{v}_i, \quad \text{where } \boldsymbol{\alpha} = \text{softmax}(\mathbf{K}\mathbf{q}).$$

Consider a deterministic key-value lookup, in which we would simply retrieve the value corresponding to the key with highest dot product with the query. Writing $s_i = \langle \mathbf{k}_i, \mathbf{q} \rangle$,

$$\text{Lookup}(\mathbf{q}, \mathbf{V}, \mathbf{K}) = \mathbf{v}_{i^*}, \quad \text{where } i^* = \arg \max_{i=1, \dots, K} s_i.$$

Using the set of one-of- K indicator vectors $\mathcal{Z} = \{\mathbf{e}_1, \dots, \mathbf{e}_K\}$, we have

$$\text{Lookup}(\mathbf{q}, \mathbf{V}, \mathbf{K}) = \mathbf{V}^\top \hat{\mathbf{z}}_0, \quad \text{where } \hat{\mathbf{z}}_0 = \arg \max_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{z}, \mathbf{s} \rangle.$$

Attention can then be seen to be the entropy-regularized relaxation of lookup, replacing $\hat{\mathbf{z}}_0$ with $\hat{\mathbf{z}}_{H_1} = \text{softmax}(\mathbf{s})$ as in Eq. 3.11. The entropy regularization forces all attention weights to be nonzero. However, sparse attention using H_α , as shown in figure 3.1, can bridge this gap and induce combinations between only a few objects.

3.4 Global Structured Relaxations and Structured Attention

In general structure prediction settings, $\text{conv}(\mathcal{Z})$ is a polytope that lacks a compact description. There are two ways to generalize softmax and sparsemax-style mappings in the structured case: *probabilistic regularization*, in which we rely on the interpretation of H as an entropy, measuring the flatness of probability distributions, and *mean regularization*, in which we think of H just as a regularizer on relaxed marginal assignments of parts.




Figure 3.1: Sparse attention in a sequence-to-sequence machine translation model. Each row is an independent sparse and differentiable selection \hat{z} over the German source sentence. Reproduced with permission from Niculae and Blondel (2017).

Probabilistic regularization. Consider the structured entropy

$$H(\mathbf{z}) = \sup_{\boldsymbol{\alpha}} \left\{ - \sum_i \alpha_i \log \alpha_i : \boldsymbol{\alpha} \in \Delta^{|\mathcal{Z}|}, \mathbb{E}_{\boldsymbol{\alpha}}[\mathbf{Z}] = \mathbf{z} \right\}. \quad (3.18)$$

H measures the highest entropy among all decompositions of \mathbf{z} as convex combinations of the elements of \mathcal{Z} . If \mathbf{z} is a vertex \mathbf{z}_0 , then the unique $\boldsymbol{\alpha}$ is an indicator vector and $H(\mathbf{z}) = 0$. If \mathbf{z} is at barycenter of the polytope, the max-entropy decomposition is $\boldsymbol{\alpha} = [1/|\mathcal{Z}|, \dots, 1/|\mathcal{Z}|]$, but other decompositions with lower entropy also may exist. The main motivation for using this regularizer is that the solution \mathbf{z}_H is none other than the average structure under the Gibbs distribution, and computing it is known as the *marginal inference* problem, which we have seen in §2.3 as a fundamental primitive for probabilistic structured output prediction, and now we see it also plays a central role in relaxed learning of latent structures:

$$\hat{\mathbf{z}}_H = \mathbb{E}[\mathbf{Z}] \quad \text{under} \quad p(\mathbf{z} | x) \propto \exp \langle \mathbf{z}, \mathbf{s} \rangle. \quad (3.19)$$

In the one-of-K model, this recovers softmax, because $\mathcal{Z} = \{\mathbf{e}_1, \dots, \mathbf{e}_K\}$, so $H(\mathbf{z}) = -\sum_i z_i \log z_i$. More generally, the solution takes the form $\hat{\mathbf{z}}_H = \mathbf{E}_{\mathbf{Z} \sim \boldsymbol{\alpha}_*}[\mathbf{Z}]$, where $\boldsymbol{\alpha}_*$ is the maximizer inside Eq. 3.18. The distribution $\boldsymbol{\alpha}_*$ cannot usually be computed, or even stored in memory, however, the mean $\hat{\mathbf{z}}_H$ can be computed efficiently for some select models of interest, *e.g.*, the ones in Table 2.2.

The marginal $\hat{\mathbf{z}}_H$ is continuously differentiable as a function of \mathbf{s} : this follows from the fact that $\hat{\mathbf{z}}_H$ is itself the gradient of the normalization constant $\log \sum_{\mathbf{z} \in \mathcal{Z}} \exp \langle \mathbf{z}, \mathbf{s} \rangle$ (Eq. 2.11), which is analytic. The gradient is thus available through automatic differentiation whenever marginal computation is available: as we have seen from Chapter 2 this is the case for global structure models representable as paths in (hyper)graphs, and for a select number of other global structure models. The next paragraph surveys research using this type of relaxation.

Structured attention networks. Kim *et al.* (2017) made use of the differentiability of marginal inference to propose structured attention mechanisms, including an attention mechanism with first-order Markov potentials to encourage contiguous selection, as well as a tree-structured attention mechanism demonstrated on arithmetic expression parsing. Both models are based on differentiating through dynamic programs (Li and Eisner, 2009; Mensch and Blondel, 2018). Liu and Lapata (2018) employ a similar strategy to learn tree-structured document representations, this time using unconstrained arborescences via the matrix-tree theorem. Liu *et al.* (2018) develop a structured attention network for alignment using a variant of an inside-outside parsing algorithm.

3.5 Mean Structure Regularization: Sinkhorn and SparseMAP

In some cases, marginal inference is not available for the structure we are interested in modeling. A more tractable regularizer in structured case can be obtained by applying a penalty directly on the marginal vector \mathbf{z} , rather than on the decomposition α . While this approach abandons probabilistic interpretations in terms of maximum entropy, it leads to useful algorithms.

Linear assignment and Sinkhorn. For certain structured representation, one may apply an entropy-inspired regularizer onto the coordinates of \mathbf{z} directly:

$$S(\mathbf{z}) = - \sum_r z_r \log z_r \neq H(\mathbf{z}). \quad (3.20)$$

If the representation \mathbf{z} consists of a concatenation of one-of-K variables, this can be seen as regularizing the (independent) entropies of each of those variables. This must not be confused with the Shannon entropy

regularizer $H(\mathbf{z})$ of Eq. 3.18. The main example for which this construction leads to efficient computation is linear assignment (matching), introduced in Example 3 and figures 2.1 and 2.2, which we now formalize as a special case of the transportation problem (Hitchcock, 1941; Kantorovich, 1942; Peyré and Cuturi, 2019).

Given two sets of size m with pairwise affinities between them $s_{i \leftrightarrow j}$ the task is to find a maximum-scoring assignment that maps each object in a set to exactly one counterpart in the other. Solutions can be represented as *permutation matrices*:

$$\mathcal{Z} = \{\mathbf{z} = \text{vec}(\mathbf{Z}) : \mathbf{Z} \in \{0, 1\}^{m \times m} : \mathbf{1}^\top \mathbf{Z} = \mathbf{Z} \mathbf{1} = \mathbf{1}\}. \quad (3.21)$$

and the marginal polytope is known as the Birkhoff polytopes of bistochastic matrices (Birkhoff, 1946):

$$\text{conv}(\mathcal{Z}) = \{\mathbf{z} = \text{vec}(\mathbf{Z}) : \mathbf{Z} \in \mathbb{R}_+^{m \times m}, \mathbf{1}^\top \mathbf{Z} = \mathbf{Z} \mathbf{1} = \mathbf{1}\}. \quad (3.22)$$

It is possible to generalize this formulation to allow *unbalanced* assignments between different-sized sets, by replacing the equality constraint along the larger dimension with an inequality constraint. Generalizing the row and column sums so distributions yields the transportation problem.

As per Table 2.2, the maximization problem

$$\arg \max_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{z}, \mathbf{s} \rangle \quad (3.23)$$

can be solved in polynomial time by algorithms such as Kuhn-Munkres, Jonkers-Volgenant, auction, however, the corresponding marginal inference problem is #P-complete (Garey and Johnson, 1979). However, the mean-regularized problem

$$\arg \max_{\mathbf{z} \in \text{conv}(\mathcal{Z})} \langle \mathbf{z}, \mathbf{s} \rangle + S(\mathbf{z}) \quad (3.24)$$

can be efficiently computed by an iterative renormalization algorithm due to Sinkhorn (1964), popular in optimal transport (Cuturi, 2013).

Adams and Zemel (2011) differentiate through Sinkhorn for learning to rank. For neural latent structures, we highlight Mena *et al.* (2018), who employ Sinkhorn within an approximate strategy for probabilistic permutations learning and Tay *et al.* (2020), who employ Sinkhorn attention for reordering inside Transformers.

Quadratic regularization: SparseMAP. Taking $Q(\mathbf{z}) = -\frac{1}{2}\|\mathbf{z}\|_2^2$ yields a projection onto the marginal polytope:

$$\hat{\mathbf{z}}_Q = \arg \max_{\mathbf{z} \in \text{conv}(\mathcal{Z})} \langle \mathbf{z}, \mathbf{s} \rangle - \frac{1}{2}\|\mathbf{z}\|_2^2 = \arg \min_{\mathbf{z} \in \text{conv}(\mathcal{Z})} \|\mathbf{z} - \mathbf{s}\|_2^2. \quad (3.25)$$

This quadratic penalty was used by Meshi *et al.* (2015) as a loss function.

The SparseMAP strategy for structure learning (Niculae *et al.*, 2018; Niculae and Martins, 2020) relies on this quadratic regularization, along with an efficient general-purpose algorithm for computing both the solution and its gradients, and observes that, analogously to sparsemax in the unstructured case, solutions tend to be sparse, *i.e.*, consist of combinations of only a few global structures. For computing the solution, Niculae *et al.* (2018) apply the general-purpose *active set* algorithm (Vinyes and Obozinski, 2017; Nocedal and Wright, 1999, Ch. 16.4 & 16.5), also known as Wolfe’s *min-norm point* algorithm, (Wolfe, 1976), part of a Frank-Wolfe family of algorithms (Frank and Wolfe, 1956; Lacoste-Julien and Jaggi, 2015) which do not require any specific form of $\text{conv}(\mathcal{Z})$ but instead only interact with the constraint via an *linear oracle* $\arg \max_{\mathbf{z} \in \text{conv}(\mathcal{Z})} \langle \mathbf{z}, \mathbf{s} \rangle$. In the context of structured prediction, this linear oracle corresponds to finding the highest-scoring structure, and thus algorithms are often well-studied (Table 2.2). The backward pass has an efficient low-rank structure that benefits from sparsity.

Example 15 (Dense vs. sparse linear assignment). Figure 3.2 shows




Figure 3.2: Relaxed bipartite matching (linear assignment) between tokens in two English sentences. Left: Sinkhorn, right: SparseMAP.

“soft matchings” between two English paraphrases, as induced by Sinkhorn and by SparseMAP. In both cases, the pairwise affinity scores are computed as

$$[s(x, \theta_f)]_{i \leftrightarrow j} = \langle [h(x_1; \theta_f)]_i, [h(x_2; \theta_f)]_j \rangle, \quad (3.26)$$

where $x = (x_1, x_2)$ is the tuple of input sentences, and h is a neural network encoder; in this case, an **albert-base** pretrained transformer (Lan *et al.*, 2020). Both induced alignments are differentiable. Analogous to the unstructured case, SparseMAP tends to make “hard assignments” for unambiguous variables; in this case, for identical words.

Implicit perturbation-based regularization. In some situations, the choice of regularizer itself is not necessarily important. Berthet *et al.* (2020) propose to implicitly induce a regularized argmax from the

smoothing generated by a random noise variable \mathbf{U} , defining

$$\widehat{\mathbf{z}}_{\Omega_{\mathbf{U}}}(\mathbf{s}) = \mathbb{E}_{\mathbf{U}} \left[\arg \max_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{z}, \mathbf{s} + \mathbf{U} \rangle \right]. \quad (3.27)$$

The expression of $\Omega_{\mathbf{U}}$ depends on the distribution of \mathbf{U} and is not generally available, but nevertheless $\widehat{\mathbf{z}}_{\Omega_{\mathbf{U}}}$ and its gradient can be estimated from samples. Despite the stochastic nature of this method, we catalogue it as a relaxation since it is based on computing $g(\mathbb{E}[\mathbf{Z}])$ and not $\mathbb{E}[g(\mathbf{Z})]$: the downstream objective is not optimized in expectation. Like SparseMAP, this strategy only requires calls to the maximization oracle $\arg \max_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{z}, \mathbf{s} \rangle$. Unlike SparseMAP, the maximizations can be performed in parallel; but in exchange one loses finite convergence properties and adaptive sparsity (more samples can always be drawn to improve the estimate). This strategy is also related to marginal inference, as for certain structures \mathbf{U} can be chosen such that $p(\widehat{\mathbf{z}}_{\Omega_{\mathbf{U}}}(\mathbf{s})) \propto \exp \langle \mathbf{z}, \mathbf{s} \rangle$: in the case of the one-of-K model \mathbf{U} would have to be the standard Gumbel variable, a connection discussed in more detail in Chapter 5.

Picking a regularizer. When an efficient and numerically stable algorithm for marginal inference is available, and sparsity of the induced structures is not desired,

marginal inference (*i. e.*, the structured entropy regularizer) is a good choice. When only a maximization algorithm is available, or sparsity is desired, or marginal inference is numerically unstable, SparseMAP provides a solution. If not even a maximization algorithm is available, Stoyanov *et al.* (2011) and Domke (2013) study differentiation of approximate marginal inference via belief propagation for pairwise Markov random fields, and Niculae and Martins (2020) propose an extension of SparseMAP to arbitrary factor graphs, supporting coarse decomposition when parts of the problem are tractable.

3.6 Summary

In this chapter we explored ways to relax a discrete variable selection into a continuous, differentiable mapping, while still retaining some of the desired properties from the discrete mapping. We have seen how results from convex analysis allow us to frame this relaxation as a smoothing, and lead to exact gradient computations. Examples of technique falling into this class include attention mechanisms and their structured counterparts, sparse transformations such as sparsemax and α -entmax, as well as their structured counterparts, such as Sinkhorn-based methods and SparseMAP, and perturbation-based regularization methods.

When a relaxation compromise is feasible, due to the model design or the application scenario, relaxation methods work generally well and are among the most stable to train, as they are essentially just standard deterministic deep models. While sparsity can bridge the gap between discrete and continuous and lead to integral outputs some of the time, in some applications discrete mappings are strictly necessary. The next two chapters focus on methods that do not compromise on discreteness.

4 Surrogate Gradients

4.1 Straight-Through Gradients

In §3.1, we have seen that a discrete latent variable assignment $\hat{z}(x, \theta_f)$ cannot simply be plugged into a downstream model

$$\tilde{g}(x, y, \theta_f, \theta_g) = g(x, y, \hat{z}(x, \theta_f); \theta_g),$$

as the gradient through a discrete encoder (Eq. 3.3) is either null or undefined, preventing gradient-based learning.

As before, invoking Assumption 2, we can focus on a model whose encoder generates factorized scores $\mathbf{s}(x, \theta_f)$, and the discrete variable assignment is $\hat{z}_0(\mathbf{s}) = \arg \max_{z \in \mathcal{Z}} \langle z, \mathbf{s} \rangle$. Instead of relaxing \hat{z}_0 , as in Chapter 3, in this chapter we keep the discrete encoder in the forward pass, but in the backward pass we replace its problematic derivative with a different function. A possible surrogate gradient is the identity




Figure 4.1: Selecting a hidden representation \hat{z} using a discrete argmax introduces a break (dotted line) in the backpropagation chain (thick purple arrows). The surrogate gradient approach keeps the forward pass unchanged, but replaces the dotted arrow with a different gradient, such as an identity function.

that leads to the following mapping \widehat{z} (Hinton, 2012; Bengio *et al.*, 2013b):

$$\begin{aligned} z_{\text{ST}}(\mathbf{s}) &= \arg \max_{z \in \mathcal{Z}} \langle z, \mathbf{s} \rangle, \\ \partial z_{\text{ST}}(\mathbf{s}) &:= \text{Id}. \end{aligned} \tag{4.1}$$

That is, the backward pass multiplies by an identity matrix regardless of the value of \mathbf{s} . This is the origin of the name “straight-through,” as during backpropagation the downstream gradient passes directly through the argmax mapping as if it weren’t there. The second line in Eq. 4.1 is not an assertion, but a construction: we replace the gradient of a function with another function. From an optimization point of view, this is incorrect, minimizing a different function than the intended objective. However, the update can be informative enough for effective learning in practical applications, and straight-through strategies are easy to implement.

4.2 Straight-Through Variants

Softmax-ST. In the unstructured model where $\mathcal{Z} = \{e_1, \dots, e_K\}$, recall that softmax gives a differentiable relaxation of the argmax operator, per §3.3. Therefore, as an alternative of using the identity matrix as the Jacobian in the backward pass, the derivative of softmax is a sensible alternative, corresponding to “pretending” we had used softmax in the forward pass instead of argmax (Jang *et al.*, 2017, Sec. 2.2):

$$\begin{aligned} \widehat{z}_{\text{Softmax-ST}}(\mathbf{s}) &:= \arg \max_{z \in \mathcal{Z}} \langle z, \mathbf{s} \rangle, \\ \partial \widehat{z}_{\text{Softmax-ST}}(\mathbf{s}) &:= \partial \text{softmax}(\mathbf{s}) \\ &= \text{diag}(\text{softmax}(\mathbf{s})) - \text{softmax}(\mathbf{s}) \text{softmax}(\mathbf{s})^\top, \end{aligned}$$

see Eq. 3.16. Since softmax can be seen as a regularized argmax, this may be interpreted to only turning on the regularization in the backward

pass, and so it may lead to a smaller mismatch between the forward and backward pass. In particular, if the maximum value is reached by a single element of \mathbf{s} , $\lim_{\gamma \rightarrow 0} \text{softmax}(\mathbf{s}/\gamma) = \widehat{\mathbf{z}}_{\text{Softmax-ST}}(\mathbf{s})$.

When tackling latent structured representations, surrogate straight-through methods fit well within the incremental formalism, as each step in the prediction is a one-of-K choice. For learning latent global structure, a natural extension of Softmax-ST to the structured case is given by the marginals $\widehat{\mathbf{z}}_H = \mathbb{E}[\mathbf{Z}]$, since, as discussed in §3.4, the softmax is the expected vector for the case of the one-of-K structure. As we know from Chapter 2, computing the marginals is an important computational primitive for structured models, but it is not always tractable. In cases when marginals are not available, alternative relaxations such as SparseMAP (§3.4) can be used, or one may consider the approaches described next.

Structured Projection of Intermediate Gradients (SPIGOT). Motivated by allowing hard decisions in the hidden layers of neural networks, while respecting the constraints of a structured arg max operation in the backward pass, SPIGOT (Peng *et al.*, 2018). computes an intermediate approximation of the gradient that it then projects onto the marginal polytope:

$$\begin{aligned} \widehat{\mathbf{z}}_{\text{SPIGOT}}(\mathbf{s}) &:= \widehat{\mathbf{z}}_0(\mathbf{s}) = \arg \max_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{z}, \mathbf{s} \rangle, \\ \partial \widehat{\mathbf{z}}_{\text{SPIGOT}}(\mathbf{s})(\mathbf{v}) &:= \widehat{\mathbf{z}}_0(\mathbf{s}) - \mathbf{proj}_{\text{conv}(\mathcal{Z})}(\widehat{\mathbf{z}}_0(\mathbf{s}) - \eta \mathbf{v}), \end{aligned} \tag{4.2}$$

where η is the step size. We provide an intuitive explanation of the above update, unifying it with STE, in §4.4. In the work originally proposing SPIGOT, Peng *et al.* (2018) compute the projection approximately using a further relaxation. Mihaylova *et al.* (2020) remark that this projection is the same computational primitive as the one in the SparseMAP relaxation (§3.4) and propose to compute it using the active

set algorithm.

Linear Interpolation (LI). Pogančić *et al.* (2020) study a surrogate gradient defined as a linear interpolation of the downstream value at two points:

$$\begin{aligned}\hat{z}_{\text{LI}}(\mathbf{s}) &:= \hat{z}_0(\mathbf{s}) \\ \partial \hat{z}_{\text{LI}}(\mathbf{s})(\mathbf{v}) &:= \frac{1}{\eta} (\hat{z}_0(\mathbf{s} + \eta \mathbf{v}) - \hat{z}_0(\mathbf{s})) .\end{aligned}\tag{4.3}$$

Notice that in the limit $\eta \rightarrow 0$ this recovers the definition of the directional derivative of \hat{z}_0 , which in this case would be ill-behaved. However, Pogančić *et al.* (2020) study this mapping in the regime of substantially larger $\eta \gg 0$ and argue for its desirable smoothing effect.

Unlike STE and Softmax-STE, SPIGOT and LI result in a derivative that is not a linear operator, it is therefore not the derivative of any function. For this reason we must write it as a nonlinear function of the direction vector \mathbf{v} and not as a Jacobian (in practice, this will be the downstream derivative of the loss *w.r.t.* \hat{z}).

Implementation Details. Surrogate gradient methods, by definition, override the actual gradient computation with something else. In PyTorch (PyTorch, 2017), at present we recommend an implementation using custom functions (inheriting from `torch.autograd.Function`). Such objects implement a `forward(self, s)` method which should return $\hat{z}_0(\mathbf{s})$, and a `backward(self, v)` method which returns the vector-Jacobian product at \mathbf{s} (saved from the forward pass) with \mathbf{v} , *i.e.*, it returns our construction of choice $\partial \hat{z}_{\text{ST}}(\mathbf{s})(\mathbf{v})$, $\partial \hat{z}_{\text{SPIGOT}}(\mathbf{s})(\mathbf{v})$, etc. Such an implementation cannot be tested using automatic gradient checks, since the gradient will be technically incorrect. We therefore recommend using custom unit tests against the theoretically expected answers, *e.g.*, using an identity encoder $\mathbf{s}(x, \boldsymbol{\theta}_f) := \boldsymbol{\theta}_f$.

4.3 Quantization: Straight-Through Friendly Models

In certain applications, it may be possible to rearrange the architecture of a neural network with discrete latent structure so that the discrete mapping is in some sense “close” to linear, thereby making the STE a more sensible choice. While these approaches require specific model choices, these choices can contain beneficial inductive biases and perform well.

4.3.1 Rounding

Consider a discrete latent variable whose domain is the integers, *i.e.* $\mathcal{Z} = \mathbb{Z} := \{\dots, -2, -1, 0, 1, \dots\}$. Our usual construction that assigns a score to each integer is infeasible since there are infinitely many options. However, the natural ordering of \mathbb{Z} suggests we can use a single real-valued score $s(x, \boldsymbol{\theta}_f) \in \mathbb{R}$ alongside a discretizing operation:

$$\widehat{z}_{\text{int}}(x, \boldsymbol{\theta}_f) = \lfloor s(x, \boldsymbol{\theta}_f) \rfloor.$$

Applying the straight-through idea here results in ignoring the floor operation in the backward pass:

$$\begin{aligned}\widehat{z}_{\text{int-ST}}(s) &:= \lfloor s \rfloor \\ \partial_s \widehat{z}_{\text{int-ST}}(s) &:= 1.\end{aligned}$$

The floor function is always within ± 1 of the identity function, and this strategy has proven useful in constructing integer normalizing flow models (Hoogeboom *et al.*, 2019) as well as in learning efficient neural networks with fixed precision weights (Gholami *et al.*, 2022). Note, however, that arbitrarily close functions may have arbitrarily far away derivatives.

4.3.2 Vector Quantization

Extending the scalar idea to a vector space leads to the VQ-VAE construction introduced by Oord *et al.* (2017). This model draws ideas from vector quantization, as used in clustering and compression. It similarly uses an “encoder-decoder” architecture, but unlike our framework, the VQ-VAE supports only discrete categorical variables represented by embeddings $\{\mathbf{z}_i \in \mathbb{R}^d : i = 1, \dots, K\}$. In its general form, the VQ-VAE uses an encoder $\tilde{f} : \mathcal{X} \rightarrow \mathbb{R}^D$, $\tilde{f}(\mathbf{x}) = \mathbf{z}$ and a decoder $\tilde{g} : \mathcal{X} \times \mathcal{Y} \times \mathbb{R}^D \rightarrow \mathbb{R}$ where for instance $g(\mathbf{x}, \mathbf{z}, \mathbf{y})$ models the probability $p(Y = \mathbf{y} \mid X = \mathbf{x}, Z = \mathbf{z})$. But instead of directly computing $\tilde{g}(\mathbf{x}, \mathbf{y}, \tilde{f}(\mathbf{x}))$, *i. e.*, letting \mathbf{z} be a continuous hidden vector, we instead employ a quantization step:

$$\tilde{g}(\mathbf{x}, \mathbf{y}, Q(\tilde{f}(\mathbf{x}))) \quad \text{where} \quad Q(\mathbf{s}) = \arg \min_{\mathbf{z} \in \mathcal{Z}} \|\mathbf{s} - \mathbf{z}\|_2. \quad (4.4)$$

The mapping $Q : \mathbb{R}^D \rightarrow \mathcal{Z}$ essentially *snaps-to-grid* the encoder output, by returning the nearest neighbor anchor point. Importantly, the decoder’s third argument (corresponding to \mathbf{z}) is parametrized by a neural network that can accept any continuous embedding in \mathbb{R}^D , and it is therefore possible to apply the STE to the quantization mapping Q . Additionally, the VQ-VAE loss includes a term $\|\tilde{f}(x) - Q(\tilde{f}(x))\|^2$ which simultaneously learns the embeddings and encourages the encoder to output values as close as possible to the set of allowed embeddings \mathcal{Z} : this term can be seen as a penalty on the STE approximation error, and if it were zero, the STE would be exact.

4.4 Interpretation via Pulled-Back Labels

STE and its variants are mainly motivated as a practical solution for learning problematic models with null gradients. In this section, we




Figure 4.2: Top: If ground truth supervision were available for the latent variable z , θ_f could be trained jointly with an auxiliary loss L_Z . Bottom: As such supervision is not available, we induce a best-guess *pulled-back label* z_* by pulling back the downstream loss. This strategy recovers the STE and SPIGOT estimators.

provide a possible interpretation of the end-to-end objective optimized when using STE or SPIGOT.

Consider a multitask learning scenario (Caruana, 1997), where we predict the latent variable as a separate task in addition to the downstream task. In this case, z would be fully observed, *i.e.*, we would have supervised data in the form of triplets (x, y, z) . We could therefore, in addition to the downstream loss on y , also incur a loss $L_Z(s, z)$ which would allow us to update the parameters θ_f . We will focus for now on the simple perceptron loss:

$$L_Z(s, z) := \langle s, z \rangle - \arg \max_{z' \in \mathcal{Z}} \langle s, z' \rangle. \quad (4.5)$$

In the latent variable case, we do not have access to this supervision and instead induce a *pulled-back label* as a best-effort guess at what z should be, given the available data, as illustrated in figure 4.2.

Let us define the pulled-back label as a sensible approximation to

the value for \mathbf{z} that minimizes the downstream loss:

$$\mathbf{z}_* \approx \arg \min_{\mathbf{z} \in \text{conv}(\mathcal{Z})} g(x, y, \mathbf{z}). \quad (4.6)$$

Notice that the optimization is with respect to \mathbf{z} and not the model parameters; the encoder f therefore does not enter this problem. Note also that this is a continuous optimization problem due to our choice of using the convex hull of \mathcal{Z} as domain. Searching only over the discrete set \mathcal{Z} would be intractable in the structured case; the continuous version can be approximated with iterative methods.

Still, this optimization problem is challenging and non-convex for many interesting models, such as deep neural networks. We may approach it using the *projected gradient algorithm* (Goldstein, 1964; Levitin and Polyak, 1966), which takes a step in the direction of the gradient (with step size η), and then projects the updated guess to the constraint set, yielding the iteration

$$\mathbf{z}^{(t+1)} = \mathbf{proj}_{\text{conv}(\mathcal{Z})}(\mathbf{z}^{(t)} - \eta \partial_3 g(x, y, \mathbf{z}^{(t)})).$$

If g was convex *w.r.t.* its third argument, we could repeat this iteration until convergence to the optimal \mathbf{z}_* , regardless of initialization. However, this is infeasible and may be an unwise use of resources, at least until the model is trained well.

Consider instead starting from the sensible initialization $\mathbf{z}^{(0)} := \hat{\mathbf{z}}_0(\mathbf{s}) = \arg \max_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{z}, \mathbf{s} \rangle$ and applying a single iteration of projected gradient, giving:

$$\mathbf{z}_* := \mathbf{proj}_{\text{conv}(\mathcal{Z})}(\hat{\mathbf{z}}_0 - \eta \partial_3 g(x, y, \hat{\mathbf{z}}_0)). \quad (4.7)$$

Using this pulled-back label to update the encoder by the perceptron

loss from Eq. 4.5, we get the following gradient *w.r.t.* \mathbf{s} :

$$\begin{aligned}
 \partial_s L_Z(\mathbf{s}, \mathbf{z}_\star) &= \widehat{\mathbf{z}}_0 - \mathbf{z}_\star \\
 &= \widehat{\mathbf{z}}_0 - \mathbf{proj}_{\text{conv}(\mathcal{Z})}(\widehat{\mathbf{z}}_0 - \eta \partial_3 g(x, y, \widehat{\mathbf{z}}_0)) \\
 &= \partial \widehat{\mathbf{z}}_{\text{SPIGOT}}(\mathbf{s}) \circ \partial_3 g(x, y, \widehat{\mathbf{z}}_{\text{SPIGOT}}(\mathbf{s})) \\
 &= \partial_s g(x, y, \widehat{\mathbf{z}}_{\text{SPIGOT}}(\mathbf{s})),
 \end{aligned} \tag{4.8}$$

recognizing the SPIGOT surrogate gradient from Eq. 4.2 and undoing the chain rule. This shows that SPIGOT minimizes the **perceptron loss** *w.r.t.* a pulled-back label computed by **one projected gradient step** on the downstream loss starting at $\widehat{\mathbf{z}}_0$.

Relaxing the constraints. We can relax the constraints for \mathbf{z}_\star and instead optimize

$$\mathbf{z}_\star := \underset{\mathbf{z} \in \mathbb{R}^D}{\text{arg min}} g(x, y, \mathbf{z}). \tag{4.9}$$

This optimization problem can again be solved iteratively, but without the projection step. One step of gradient descent starting at $\widehat{\mathbf{z}}_0$ yields:

$$\mathbf{z}_\star := \widehat{\mathbf{z}}_0 - \eta \partial_3 g(x, y, \widehat{\mathbf{z}}_0). \tag{4.10}$$

Applying the perceptron loss again, assuming $\eta = 1$,

$$\begin{aligned}
 \partial_s L_Z(\mathbf{s}, \mathbf{z}_\star) &= \widehat{\mathbf{z}}_0 - \mathbf{z}_\star \\
 &= \widehat{\mathbf{z}}_0 - (\widehat{\mathbf{z}}_0 - \partial_3 g(x, y, \widehat{\mathbf{z}}_0)) \\
 &= \partial_3 g(x, y, \widehat{\mathbf{z}}_0) \\
 &= \partial \widehat{\mathbf{z}}_{\text{STE}}(\mathbf{s}) \circ \partial_3 g(x, y, \widehat{\mathbf{z}}_{\text{STE}}(\mathbf{s})) \\
 &= \partial_s g(x, y, \widehat{\mathbf{z}}_{\text{STE}}(\mathbf{s})).
 \end{aligned} \tag{4.11}$$

This reveals that indeed STE can also be interpreted as minimizing a perceptron loss on a pulled-back label, and that SPIGOT is indeed a version of STE that takes into account the structure of the marginal polytope $\text{conv}(\mathcal{Z})$.

This view was proposed by Mihaylova *et al.* (2020) and employed to develop new surrogate gradient methods by changing other components, such as the loss function (from perceptron to cross-entropy loss) or the optimization algorithm (mirror descent instead of projected gradient).

4.5 Summary

Surrogate gradient methods provide a way to train deterministic models without compromising on the use of discrete mappings in the forward pass (*e.g.*, argmax or quantization layers). Whereas in Chapter 3 we relaxed the forward pass mapping and computed its gradient exactly, in this chapter we leave the forward pass unchanged, but only employ a relaxation in the backward pass. While strictly speaking incorrect, surrogate gradients can often be justified in principled ways and can work well in practice. Mihaylova *et al.* (2020) empirically compared surrogate gradient methods against relaxations on a task where \mathbf{z} is known, and found that both classes of methods are able to reach similar downstream performance. Relaxation methods seem to converge faster and are more robust to initialization, but are worse than straight-through methods at recovering a discrete latent \mathbf{z} that agrees with the experiment design.

All methods seen so far require assumption 2: continuity of the decoder g *w.r.t.* \mathbf{z} . Even if surrogate gradient methods use only discrete \mathbf{z} in the forward pass, the backward pass requires the gradient $\partial_{\mathbf{z}}g(x, y, \hat{\mathbf{z}})$. In the next chapter, we give the discrete latent a probabilistic treatment, allowing us to have correct (unbiased) estimation without compromising discreteness and without any requirements on g .

5

Probabilistic Latent Variables

5.1 Formulating the Probabilistic Model.

In this section, we directly tackle the probabilistic form of discrete latent variable learning problem. Instead of the downstream model seeing a single choice of \hat{z} , we will instead consider the expectation of the downstream model over all possible assignments of the latent variable:

$$\begin{aligned} \bar{g}(x, y) &:= \mathbb{E}_Z [g(x, y, Z)] \\ &= \sum_{z \in \mathcal{Z}} g(x, y, z) p(Z = z | x), \end{aligned} \tag{5.1}$$

The distribution $p(Z = z | x)$ is defined in terms of the encoder $f(x, z)$. A standard way is $p(z | x) \propto \exp f(x, z)$, but we shall see in this chapter that other ways to define this distribution are possible.

In Chapter 3, we have avoided the challenging sum and relaxed the problem by pushing the expectation inside: $\mathbb{E}_Z [g(x, y, Z)] \approx g(x, y, \mathbb{E}[Z])$. In Chapter 4, we eschewed any probabilistic interpretation and instead tried to approximate the discrete mapping $g(x, y, (\arg \max_{z \in \mathcal{Z}} f(x, z)))$. We now finally come face-to-face with this computation.

The probabilistic framework is challenging because, per Chapter 2, \mathcal{Z} may be an exponentially large combinatorial set. In general, this comes at a cost; the payoff however is that we can get well-principled models with much fewer assumptions on the downstream model g .

While many implementation of probabilistic models rely on stochastic computation (sampling), this is not necessarily the case. The first and the last strategy we present in this chapter are deterministic.

5.2 Explicit Marginalization by Enumeration

In the case where $|\mathcal{Z}|$ is of manageable size and when the decoder is not too expensive, it is of course possible to tackle Eq. 5.1 explicitly by enumerating all possibilities, which requires evaluating the decoder $|\mathcal{Z}|$ times. This is for instance possible for categorical (one-of- K) discrete latents with K not too large, or for bit vectors of manageable size. For example, $\mathcal{Z} = \{0, 1\}^D$ has 2^D elements, but for $D = 3$ it is not necessarily prohibitive to evaluate the decoder 8 times.

At this point we may see why explicit marginalization does not require assumption 2, and \mathbf{z} can be a completely discrete index: gradients of the form $\partial_3 g(x, y, \mathbf{z})$ are not required or used, because in Eq. 5.1 \mathbf{z} is not a transformed function of any parameters: it is an abstract summation index.

When one can afford to explicitly compute this objective, results tend to be unsurprisingly good in comparison to approximate approaches or relaxations. Of course, for large \mathcal{Z} enumeration is prohibitive.

Example 16. Neural attention as latent alignment. Recall from Example 14 that the neural attention mechanisms can be seen as relaxed Categorical (one-of- K) variable. Using a Categorical latent variable instead is known as hard attention and was proposed by Xu *et al.* (2015) for image captioning. Deng *et al.* (2018) observed that in the context of machine translation the model can be tackled with explicit enumeration. They proposed an explicitly-marginalized variational attention which performed the best on machine translation.

Example 17. Variational Auto-Encoder with a discrete latent variable. Kingma *et al.* (2014) propose a method for semi-supervised learning based on a generative neural network. Their model includes a categorical latent variable representing the target class. To train the model, they build a Variational Auto-Encoder, that is the posterior distribution over latent variables is approximated with a learned encoder network. To ensure that the encoder is correctly trained, they explicitly marginalize over the categorical latent variable, in contrast to the standard approach in this setting (reparameterization trick for continuous latent variables).

Even when not affordable on realistic problems, we recommend implementing explicit marginalization for testing purposes.

5.3 Monte Carlo Gradient Estimation

When $p(\mathbf{z} | x) \propto \exp f(x, \mathbf{z})$ and an algorithm is available to obtain samples from $p(\mathbf{z} | x) \propto \exp f(x, \mathbf{z})$, expectations such as Eq. 5.1 can be approximated via the Monte Carlo method (Metropolis and Ulam, 1949):

$$\mathbb{E}_{\mathbf{Z}}[g(x, y, \mathbf{Z})] \approx \frac{1}{S} \sum_{i=1}^S g(x, y, \mathbf{z}^{(i)}). \quad (5.2)$$

For learning the decoder parameters, we can directly differentiate Eq. 5.2, since the expectation operator is linear and so is the differentiation operator (Deisenroth *et al.*, 2020, § 6.4.1). For any function $F(\mathbf{z}, \boldsymbol{\theta})$ and

distribution $p(\mathbf{Z} = \mathbf{z})$ that does not depend on $\boldsymbol{\theta}$, we have

$$\begin{aligned} \partial_{\boldsymbol{\theta}} \mathbb{E}_{\mathbf{Z}}[F(\mathbf{Z}, \boldsymbol{\theta})] &= \partial_{\boldsymbol{\theta}} \sum_{\mathbf{z} \in \mathcal{Z}} F(\mathbf{z}, \boldsymbol{\theta}) p(\mathbf{z}) \\ &= \sum_{\mathbf{z} \in \mathcal{Z}} (\partial_{\boldsymbol{\theta}} F(\mathbf{z}, \boldsymbol{\theta})) p(\mathbf{z}) \\ &= \mathbb{E}_{\mathbf{Z}} [\partial_{\boldsymbol{\theta}} F(\mathbf{Z}, \boldsymbol{\theta})]. \end{aligned} \tag{5.3}$$

However, if $p(\mathbf{Z} \mid \boldsymbol{\theta})$ does depend on $\boldsymbol{\theta}$ (even if F does not) then differentiating is harder, as simply applying the chain rule does not directly lead to an expectation that can be estimated from samples $\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(S)}$:

$$\begin{aligned} \partial_{\boldsymbol{\theta}} \mathbb{E}_{\mathbf{Z} \sim p(\mathbf{z} \mid \boldsymbol{\theta})}[F(\mathbf{Z})] &= \partial_{\boldsymbol{\theta}} \sum_{\mathbf{z}} F(\mathbf{z}) p(\mathbf{z} \mid \boldsymbol{\theta}) \\ &= \sum_{\mathbf{z} \in \mathcal{Z}} F(\mathbf{z}) (\partial_{\boldsymbol{\theta}} p(\mathbf{z} \mid \boldsymbol{\theta})). \end{aligned} \tag{5.4}$$

For this reason, applying the Monte Carlo method to gradient estimation requires nontrivial techniques. The next two sections describe two main strategies used for Monte Carlo gradient estimation with discrete and structured latent variables. A more general and thorough review can be found in (Mohamed *et al.*, 2020).

5.4 Path Gradient Estimator (The Reparametrization Trick)

Estimating gradients of expectations is challenging not only in the discrete setting, but also in the continuous one. An alternative to the SFE is given by the *Path Gradient Estimator* (PGE), introduced by Ho and Cao (1983) and also known as the

push-out gradient method (Rubinstein, 1992), process derivative (Pflug, 2012), perturbation analysis (Glasserman, 1990), or the reparametrization trick (Rezende *et al.*, 2014; Kingma and Welling, 2014; Titsias and Lázaro-Gredilla, 2014).

The challenge, identified in Eq. 5.4, stems from the fact that the distribution of the latent variable, over which the expectation is taken, depends on the parameters $\boldsymbol{\theta}$. The PGE circumvents this problem by extracting the source of randomness into a parameter-free random variable \mathbf{U} , using a transformation ϕ that ensures $\mathbf{Z} = \phi(\mathbf{U}, \boldsymbol{\theta})$, where equality is in distribution. For any function $F(\mathbf{z})$ we may then write (Deisenroth *et al.*, 2020, Def. 6.3)

$$\mathbb{E}_{\mathbf{Z} \sim p(\mathbf{Z}|\boldsymbol{\theta})}[F(\mathbf{Z})] = \mathbb{E}_{\mathbf{U} \sim p(\mathbf{U})}[F(\phi(\mathbf{U}, \boldsymbol{\theta}))].$$

The right-hand side is an expectation over a distribution that does not depend on $\boldsymbol{\theta}$: we have moved the parameter into the argument of the expectation. We can then obtain the Monte Carlo gradient estimator:

$$\partial_{\boldsymbol{\theta}} \mathbb{E}_{\mathbf{Z} \sim p(\mathbf{Z}|\boldsymbol{\theta})}[F(\mathbf{Z})] = \mathbb{E}_{\mathbf{U} \sim p(\mathbf{U})}[\partial_{\boldsymbol{\theta}} F(\phi(\mathbf{U}, \boldsymbol{\theta}))]. \quad (5.5)$$

If \mathbf{U} is a continuous random variable, the right-hand-side expectation is an integral. In order for the exchange of integration and differentiation to be valid, the following conditions must hold (L'Ecuyer, 1995):

- $\partial_{\boldsymbol{\theta}} F(\phi(\mathbf{u}, \boldsymbol{\theta}))$ exists;
- $\mathbb{E}_{\mathbf{U}}[F(\phi(\mathbf{U}, \boldsymbol{\theta}))] < \infty$ for any $\boldsymbol{\theta}$.
- There exists an integrable function $h(\mathbf{u})$ such that

$$\sup_{\boldsymbol{\theta}} \|p(\mathbf{u}) \cdot \partial_{\boldsymbol{\theta}} F(\phi(\mathbf{u}, \boldsymbol{\theta}))\| \leq h(\mathbf{u}).$$

These assumptions are for instance satisfied if $F \circ \phi$ is Lipschitz.

The PGE finds many successful applications in continuous variables. A prominent example is the reparametrization of a Gaussian random variable: If $\mathbf{Z} \sim \text{Normal}(m, \sigma^2)$ then \mathbf{Z} is identical in distribution to a shifted and scaled standard normal:

$$\mathbf{Z} = \phi(\mathbf{U}, m, \sigma) = m + \sigma \mathbf{U}, \quad \text{where } \mathbf{U} \sim \text{Normal}(0, 1).$$

The Gumbel reparametrization. In the discrete setting, applying the PGE hits further roadblocks. However, this line of thinking opens the door to some valuable approximations. A fundamental result for the reparametrization of discrete random variables is the Gumbel-argmax construction for the categorical distribution. Taking $\mathcal{Z} = \{\mathbf{e}_1, \dots, \mathbf{e}_K\}$, the categorical distribution parametrized by a score vector $\mathbf{s} \in \mathbb{R}^K$ is

$$p(\mathbf{Z} = \mathbf{z}) = \exp \langle \mathbf{z}, \mathbf{s} \rangle / \sum_j \exp \langle \mathbf{e}_j, \mathbf{s} \rangle,$$

in other words, it is the distribution induced by the *softmax* of the scores. It turns out that \mathbf{Z} is identical in distribution to the perturbed argmax (Gumbel, 1954; Yellott, 1977):

$$\mathbf{Z} \stackrel{d}{=} \phi(\mathbf{U}, \mathbf{s}) = \arg \max_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{z}, \mathbf{s} + \mathbf{U} \rangle, \quad (5.6)$$

where \mathbf{U} is a vector of i.i.d. standard Gumbel random variables $U_i \sim \text{Gumbel}(0)$, and $\mathcal{Z} = \{\mathbf{e}_1, \dots, \mathbf{e}_K\}$. This reparametrization is, alas, not suitable for the PGE as in Eq. 5.5, because ϕ is discontinuous and therefore so is $F \circ \phi$. However, this reformulation paves the way for a relaxation we discuss next.

Gumbel-Softmax and the Concrete distribution. As in Chapter 3, we may relax the argmax in the Gumbel-Max reparametrization to a softmax mapping, yielding the Gumbel-Softmax (Jang *et al.*, 2017), also known as the Concrete distribution (Maddison *et al.*, 2017). Consider the parametrization

$$\phi_{\gamma H_1}(\mathbf{U}, \mathbf{s}) := \widehat{\mathbf{z}}_{\gamma H_1}(\mathbf{s} + \mathbf{U}) = \text{softmax}((\mathbf{s} + \mathbf{U})/\gamma),$$

which is a perturbed version of the softmax relaxation from §3.3 where $\gamma > 0$ is a regularization strength, or “temperature” hyperparameter. Let $\mathbf{Z}_{\gamma H_1} \stackrel{d}{=} \phi_{\gamma H_1}(\mathbf{U}, \mathbf{s})$. Since $\mathbf{Z}_{\gamma H_1}$ takes values in $\text{relint conv}(\mathcal{Z})$ and

\mathbf{Z} takes values in \mathcal{Z} , *i.e.*, their supports do not overlap, the two are not only distinct, but difficult to compare. (We revisit this issue at the end of this section.) Nevertheless, the approximation becomes exact in the $\gamma \rightarrow 0$ limit, suggesting a meaningful approximation.¹ Plugging in the Gumbel-Softmax approximation into our framework, we could fully embrace the relaxation and consider a continuous latent variable model:

$$\begin{aligned}\bar{g}(x, y, \boldsymbol{\theta}_f, \boldsymbol{\theta}_g) &= \mathbb{E}_{\mathbf{Z}_{\gamma H_1}} [g(x, y, \mathbf{Z}_{\gamma H_1}, \boldsymbol{\theta}_g)] \\ &= \mathbb{E}_{\mathbf{U}} [g(x, y, \hat{\mathbf{z}}_{\gamma H_1}(\mathbf{s}(x, \boldsymbol{\theta}_f) + \mathbf{U}), \boldsymbol{\theta}_g)]\end{aligned}$$

using the PGE.

For global-structured latent variables \mathbf{Z} whose structure is characterized or representable by a dynamic program, *i.e.*, equivalent to paths in a (hyper)graph (§2.3.3), sampling from $p(\mathbf{z} \mid x) \propto \exp f(x, \mathbf{z})$ is possible via the Forward-Filtering Backward-Sampling (FFBS) algorithm (Frühwirth-Schnatter, 1994). In this case, the Backward-Sampling pass can be modified with a Gumbel-Softmax distribution, leading to differentiable samples (Fu *et al.*, 2020). Similarly, in the temperature limit, the approximation becomes exact.

Structured perturbation methods. Extending the Gumbel reparametrization to structures is challenging: the direct approach would require perturbing the score of each structure independently, which is intractable. For directed models (*e.g.*, incremental structure prediction), Kool *et al.* (2020b) propose a top-down sampling approach based on sampling Gumbels conditional on their sum. For undirected models (MRF), Hazan *et al.* (2013) propose an algorithm based on factor-decomposed perturbations and Gibbs sampling.

¹This temperature-based argument requires care. Indeed, in the $\gamma \rightarrow 0$ limit, $\mathbf{Z}_{\gamma H_1} \rightarrow \mathbf{Z}$, but also in this limit ϕ approaches a threshold function, so its gradients approach zero (or infinity near discontinuities). In practice, γ must be set far away from zero to ensure numerical stability and meaningful gradients, but not too far, or else the approximation gap is too big, as reported, *e.g.*, by Baziotis *et al.* (2019).

The Perturb-and-MAP approach (Papandreou and Yuille, 2011) instead applies perturbations directly to the scores \mathbf{s} of an arbitrary structured model: this induces a random field \mathbf{Z} but the probability distribution is not the Gibbs one. Hazan and Jaakkola (2012) derive some bounds on the relationship between Perturb-and-MAP and the true log-partition function. However, in many cases it is not essential to prescribe \mathbf{Z} to have the Gibbs distribution, and the Perturb-and-MAP approach can thus be very useful. To ensure differentiability, the role of the softmax is taken by marginal inference for structured latent variables, and we can therefore use this strategy whenever an algorithm for marginal inference is available (Corro and Titov, 2019a; Corro and Titov, 2019b).

Straight-Through Gumbel (Gumbel-ST). Combining ideas of the PGE with the surrogate gradient strategies of Chapter 4, we may use the true discrete \mathbf{Z} in the forward pass, and only apply a reparametrized relaxation in the gradient estimation for the backward pass. This strategy, known as the Straight-Through Gumbel-Softmax estimator (Jang *et al.*, 2017), corresponds to using the construction $\hat{\mathbf{z}}_{\text{Softmax-ST}}$ from §4.2:

$$\mathbf{Z}_{\text{Softmax-ST}} := \hat{\mathbf{z}}_{\text{Softmax-ST}}(\mathbf{s} + \mathbf{U}),$$

leading to the following approximate gradient:

$$\begin{aligned} \partial_{\theta_f} \bar{g}(x, y, \theta_f, \theta_g) &\approx \partial_{\theta_f} \mathbb{E}_{\mathbf{Z}_{\text{Softmax-ST}}} [g(x, y, \mathbf{Z}_{\text{Softmax-ST}}, \theta_g)] \\ &= \partial_{\theta_f} \mathbb{E}_{\mathbf{U}} [g(x, y, \hat{\mathbf{z}}_{\text{Softmax-ST}}(\mathbf{s}(x, \theta_f) + \mathbf{U}), \theta_g)] \\ &= \mathbb{E}_{\mathbf{U}} \left[\partial_{\theta_f} g(x, y, \hat{\mathbf{z}}_{\text{Softmax-ST}}(\mathbf{s}(x, \theta_f) + \mathbf{U}), \theta_g) \right]. \end{aligned}$$

This time, in the forward pass only discrete values $\mathbf{z} \in \mathcal{Z}$ are passed to the downstream model. Still, due to the use of relaxations in both Gumbel-Softmax and Gumbel-ST, Assumption 2 is still necessary. Paulus *et al.* (2020) catalogue and study a number of structured Gumbel-based models, and Huijben *et al.* (2022) provide a broader survey of

gradient estimation with Gumbel-Max and Gumbel-Softmax is available. The Gumbel-ST approximation is well-suited and popular for learning models with incremental latent structure, for instance *sequence-to-sequence-to-sequence* models (Baziotis *et al.*, 2019) and related (Xu and Lapata, 2022).

Implicit Maximum Likelihood Estimation (I-MLE). One interesting approach is to combine the structured perturbation methods described above with surrogate gradient approaches (§4.4). An implicit maximum likelihood estimation approach with perturbation-based implicit differentiation has been proposed by Niepert *et al.* (2021) which combines these two ideas. I-MLE uses a surrogate gradient construction defined as

$$\begin{aligned}\widehat{\mathbf{z}}_{\text{I-MLE}}(\mathbf{s}) &:= \widehat{\mathbf{z}}_0(\mathbf{s}) = \arg \max_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{s}, \mathbf{z} \rangle, \\ \partial \widehat{\mathbf{z}}_{\text{I-MLE}}(\mathbf{s})(\mathbf{v}) &:= \widehat{\mathbf{z}}_0(\mathbf{s}) - \widehat{\mathbf{z}}_0(\mathbf{s} - \eta \mathbf{v}).\end{aligned}\tag{5.7}$$

This is similar to the linear interpolation surrogate of Pogančić *et al.* (2020), see Eq. 4.3, with η absorbed into the global learning rate. In addition, I-MLE uses a perturb-and-MAP construction

$$g(x, y, \widehat{\mathbf{z}}_{\text{I-MLE}}(\mathbf{s} + \mathbf{U})),$$

where \mathbf{U} is a noise variable. Niepert *et al.* (2021) point out that if \mathbf{U} is taken to always be zero, I-MLE recovers the deterministic linear interpolation approach; in their experiments, perturbations improve substantially. Remark that the noise sample from \mathbf{U} appears in both terms of the surrogate gradient.

Mixed Random Variables. As pointed out above, one disadvantage of the Gumbel-Softmax construction is that, for any temperature $\gamma > 0$, the resulting random variable $Z_{\gamma H_1}$ is continuous, taking values in $\text{relint conv}(\mathcal{Z})$, whereas the discrete random variable that one would




Figure 5.1: Projected distribution onto the unit interval. In this case, the base distribution is a Gaussian.

like to approximate, Z , takes values in \mathcal{Z} – their supports have empty intersection. Although in the zero temperature limit ($\gamma \rightarrow 0_+$) $Z_{\gamma H_1}$ approaches a categorical distribution, in practice it is hard to learn with low temperatures, as the gradients will be very small and numerical instabilities often arise.

This observation led several researchers to propose *truncated* distributions which can be regarded as hybrids of discrete and continuous random variables, called *mixed* random variables (Farinhas *et al.*, 2022). For the binary case ($|\mathcal{Z}| = 2$), the probability simplex is isomorphic to the unit interval, $\Delta_1 \simeq [0, 1]$, so a simple procedure to obtain a mixed random variable Z is first sampling a continuous variable U using an auxiliary distribution $p(U)$ whose supports contains $[0, 1]$ properly, and then defining Z as the projection onto the unit interval,

$$Z = \max\{0, \min\{1, U\}\}. \quad (5.8)$$

The resulting probability density function of Z will be the combination of two Dirac delta functions (point masses) placed at 0 and 1 with the probability density function of U in $]0, 1[$, as shown in figure 5.1 An




Figure 5.2: Multivariate distributions over Δ_{K-1} . Standard distributions, like the Logistic-Normal (left), assign zero probability to all faces except the maximal one $\text{relint}(\Delta_{K-1})$. In contrast, mixed distributions support assigning probability to the *full* simplex, including its boundary: the Gaussian-Sparsemax (right) induces a distribution over the 1-dimensional edges (shown as a histogram), and assigns $p(\{(1, 0, 0)\}) = .022$. Reproduced with permission from Farinhas *et al.* (2022).

example of this construction is the *Hard Concrete distribution* (Louizos *et al.*, 2018), where \mathbf{U} is a “stretched” Concrete distribution and \mathbf{Z} is defined as in Eq. 5.8. Of course, this rectification can be applied to other continuous distributions besides the (stretched) Concrete and has in fact been so; a simple choice is the Gaussian distribution, to which one-sided (Hinton and Ghahramani, 1997) and two-sided rectifications (Palmer *et al.*, 2017) have been proposed. Another example is the Hard-Kuma distribution (Bastings *et al.*, 2019), which uses the Kumaraswamy (Kumaraswamy, 1980) as the base distribution. These “stretch-and-rectify” techniques enable assigning probability mass to the boundary of Δ_1 and are similar in spirit to the spike-and-slab feature selection method (Mitchell and Beauchamp, 1988; Ishwaran, Rao, *et al.*, 2005) and for sparse codes in variational auto-encoders (Rolfe, 2017; Vahdat *et al.*, 2018). Since they apply to binary random variables, they can be used to approximate “bit-vector” latent variables, where $\mathcal{Z} = \{0, 1\}^K$ for some $K \in \mathbb{N}$.

We show next how to generalize the idea above to other choices

of \mathcal{Z} , for example to approximate a multivariate categorical variable Z . We note first that the hard sigmoid transformation of Eq. 5.8 is just a particular case of the sparsemax transformation when $|\mathcal{Z}| = 2$. This immediately suggests an extension of the Hard Concrete and rectified Gaussian distributions for the multivariate case, which has been proposed and studied in detail by Farinhas *et al.* (2022) under the names Gumbel-Sparsemax and Gaussian-Sparsemax, respectively. Figure 5.2 shows illustrates the latter – the Gaussian-Sparsemax can also be seen as a “mixed” counterpart of the multivariate logistic normal distribution, where the softmax is replaced by a sparsemax transformation. It has been shown by Farinhas *et al.* (2022) that these mixed distributions have solid mathematical properties which make them amenable for information-theoretic treatment, if one chooses as the underlying base measure the “direct-sum measure” – a combination of the counting measure (which underlies discrete variables) with Lebesgue-Borel measures of multiple dimensionalities, associated to each face of the probability simplex. By doing so, one avoids the cumbersome practice (which is commonplace when the Gumbel-Softmax relaxation is used with discrete latent variables) of treating a distribution as discrete or continuous as convenient in different computations. Indeed, differential entropy and discrete Shannon entropy have different properties, starting with the fact that the former can be negative. The full effect of this mismatch is not clear, but Farinhas *et al.* (2022) report improved performance when using mixed random variables under the appropriate direct-sum measure.

5.5 Score Function Estimator

The Score Function Estimator (SFE, Rubinstein, 1976; Paisley *et al.*, 2012), dubbed in some contexts REINFORCE (Williams, 1992) or the likelihood ratio method (Glynn, 1990), relies on the fact that $(\log t)' =$

$1/t$ implies that $f'(t) = f(t)(\log f(t))'$ to rewrite the expectation in Eq. 5.4 as

$$\begin{aligned} \partial_{\boldsymbol{\theta}} \mathbb{E}_{\mathbf{Z} \sim p(\mathbf{z} | \boldsymbol{\theta})} [F(\mathbf{Z})] &= \sum_{\mathbf{z} \in \mathcal{Z}} F(\mathbf{z}) (\partial_{\boldsymbol{\theta}} p(\mathbf{z} | \boldsymbol{\theta})) \\ &= \sum_{\mathbf{z} \in \mathcal{Z}} F(\mathbf{z}) (p(\mathbf{z} | \boldsymbol{\theta}) \partial_{\boldsymbol{\theta}} \log p(\mathbf{z} | \boldsymbol{\theta})) \quad (5.9) \\ &= \mathbb{E}_{\mathbf{Z}} [F(\mathbf{Z}) (\partial_{\boldsymbol{\theta}} \log p(\mathbf{Z} | \boldsymbol{\theta}))]. \end{aligned}$$

In this case, since \mathbf{Z} is discrete there are no requirements on F needed to swap differentiation and expectation.² This means we can use a Monte-Carlo approximation to train both the encoder and the decoder in a discrete latent variable model. With the notation $\bar{g}(x, y) := \mathbb{E}_{\mathbf{Z} \sim p(\mathbf{Z} | x; \boldsymbol{\theta}_f)} [g(x, y, \mathbf{Z}; \boldsymbol{\theta}_g)]$, we have

$$\begin{aligned} \partial_{\boldsymbol{\theta}_g} \bar{g}(x, y) &\approx \frac{1}{S} \sum_{i=1}^S \partial_{\boldsymbol{\theta}_g} g(x, y, \mathbf{z}^{(i)}, \boldsymbol{\theta}_g), \\ \partial_{\boldsymbol{\theta}_f} \bar{g}(x, y) &\approx \frac{1}{S} \sum_{i=1}^S g(x, y, \mathbf{z}^{(i)}, \boldsymbol{\theta}_g) \left(\partial_{\boldsymbol{\theta}_f} \log p(\mathbf{z}^{(i)} | x; \boldsymbol{\theta}_f) \right). \end{aligned} \quad (5.10)$$

Altogether, this leads to a strategy for gradient-based learning with discrete or structured latent variables, with no constraints at all on the dependency of g on \mathbf{z} . This explains the origins and success of the SFE in the reinforcement learning community, when often the rewards are computed via simulations and interactions with an environment that cannot always be relaxed.

Moreover, in contrast to Chapters 3 and 4, even if Eqs. 5.2 and 5.10 are approximate, they are **unbiased estimators** of the underlying stochastic quantities: they are equal to the true values in expectation (w.r.t. the sampling process).

²The SFE can also be applied in continuous cases, and in that case the requirements are more stringent (Mohamed *et al.*, 2020, §4.3).

The most prominent limitation is that SFE has rather high variance; the remainder of this section discusses variance reduction strategies. For structured latent variables, there is another, less-discussed limitation: the SFE has among the strictest computational requirements, needing a sampling oracle for $p(\mathbf{Z} \mid x)$, as well as access to $\partial_{\theta} \log p(\mathbf{Z} \mid x, \theta)$.

Implementation details. The SFE is by definition a custom gradient estimator, and thus not immediately compatible with the automatic differentiation language of current deep learning frameworks, in which we specify a computation (or estimator of) the loss, and the gradients are handled for us. The most common approach is to define a surrogate function $g_{\text{SFE}}(x, y)$ which, when automatically differentiated, yields the appropriate gradients from Eq. 5.10. This can be accomplished using the stopgrad (stop gradient) operator, which returns its argument in the forward pass, but a zero gradient in the backward pass. Then, the following surrogate function is suitable:

$$g_{\text{SFE}}(x, y) = \frac{1}{S} \sum_{i=1}^S \left[g(x, y, z^{(i)}, \theta_g) + \text{stopgrad} \left(g(x, y, z^{(i)}, \theta_g) \right) \log p(z^{(i)} \mid x; \theta_f) \right].$$

Since the value of g_{SFE} in the forward pass is arbitrary, it should not be used for reporting or monitoring performance. An alternative and perhaps more transparent implementation strategy can be to define a custom operation within the deep learning framework by explicitly specifying the forward and backward passes, as described in Chapter 4.

5.5.1 Variance Reduction

A drawback of the SFE is its high variance: unless many samples are taken, the estimator might be too noisy for effective learning. As a

result, there is a substantial body of work devoted to variance reduction techniques. We briefly discuss three main directions and their application for discrete latent variables.

Control variates. The control variate construction (Mohamed *et al.*, 2020, §7.4) is one of the most fruitful variance reduction strategies. Given some $h : \mathcal{Z} \rightarrow \mathbb{R}$ whose expectation under $p(\mathbf{z} \mid \boldsymbol{\theta})$ is known, and some $\beta \in \mathbb{R}$, we can construct

$$F_{\beta,h}(\mathbf{z}) = F(\mathbf{z}) - \beta(h(\mathbf{z}) - \mathbb{E}_{\mathbf{Z}}[h(\mathbf{Z})]).$$

Taking expectations,

$$\mathbb{E}_{\mathbf{Z}}[F_{\beta,h}(\mathbf{Z})] = \mathbb{E}_{\mathbf{Z}}[F(\mathbf{Z})],$$

so any unbiased gradient estimator will remain correct. It remains to choose h and β may lead to lower variance. We may take h to be the score function $\partial_{\boldsymbol{\theta}} \log p(\mathbf{z} \mid \boldsymbol{\theta})$, because

$$\begin{aligned} \mathbb{E}[\partial_{\boldsymbol{\theta}} \log p(\mathbf{Z} \mid \boldsymbol{\theta})] &= \sum_{\mathbf{z} \in \mathcal{Z}} p(\mathbf{z} \mid \boldsymbol{\theta}) \frac{\partial_{\boldsymbol{\theta}} p(\mathbf{z} \mid \boldsymbol{\theta})}{p(\mathbf{z} \mid \boldsymbol{\theta})} \\ &= \partial_{\boldsymbol{\theta}} \sum_{\mathbf{z} \in \mathcal{Z}} p(\mathbf{z} \mid \boldsymbol{\theta}) \\ &= \partial_{\boldsymbol{\theta}} 1 = 0. \end{aligned}$$

This leads to the estimator

$$\partial_{\boldsymbol{\theta}} \mathbb{E}_{\mathbf{Z} \sim p(\mathbf{z} \mid \boldsymbol{\theta})}[F(\mathbf{Z})] = \mathbb{E}_{\mathbf{Z}}[(F(\mathbf{Z}) - \beta)(\partial_{\boldsymbol{\theta}} \log p(\mathbf{Z} \mid \boldsymbol{\theta}))]. \quad (5.11)$$

A simple but effective way to choose β is as a value close to $F(\mathbf{Z})$, in order to keep the magnitude of the loss term small. This includes **running average** and **exponentially moving average** (EMA) of previously-seen values of $F(\mathbf{Z})$ during learning. Running averages give equal weight to all values, while EMA can prioritize recent observations,

which is useful in a machine learning context where the loss should decrease. From the same motivation, another good choice is the model prediction itself $\beta = F(\mathbf{z}')$ for some \mathbf{z}' that does not depend on \mathbf{Z} . Using the highest-scoring \mathbf{z}' is known as **self-critic baseline** (Rennie *et al.*, 2017), another good choice is to draw \mathbf{z}' as an independent sample from $p(\mathbf{Z} \mid \boldsymbol{\theta})$. These strategies are easy to implement and in practice tend to make a substantial difference compared to no variance reduction at all. We recommend never using SFE without at least one of these simple variance reduction strategies.

Learned data-driven baselines and NVIL. The EMA estimator for β corresponds to learning β as a scalar parameter using gradient descent with a squared loss $(\beta - F(\mathbf{z}))^2$. If the problem also has some contextualizing input x then one may consider a **learned data-driven baseline** $\beta(x; \boldsymbol{\theta}_\beta)$, with its own parameters to learn (Mnih and Gregor, 2014).

Relaxation-based control variates. MuProp (Gu *et al.*, 2016) approximates the objective using a deterministic, relaxed network (replacing \mathbf{Z} with $\mathbb{E}[\mathbf{Z}]$); but, in contrast to Chapter 3, this relaxation is only used for variance reduction when optimizing the true stochastic model. The derivation is based on a Taylor expansion around $\mathbb{E}[\mathbf{Z}]$. REBAR (Tucker *et al.*, 2017) uses Gumbel-Softmax relaxations (§5.4) to compute a stochastic but reparametrizable baseline. RELAX (Grathwohl *et al.*, 2018) uses a trained differentiable network as a surrogate, trading off additional parameters for a smaller approximation gap. These methods require continuity and differentiability of the objective (or, in the case of RELAX, of the surrogate network) w.r.t. the latent variable. MuProp seems applicable to globally structured latent variables, but REBAR and RELAX rely on a deterministic mapping between the relaxed continuous variable (e.g. Gumbel-Softmax) and the corresponding discrete

sample. For binary and one-of-K variables such a mapping exists, but in the structured case this is not available.

Rao-Blackwellization. This strategy relies on tractable conditioning as a way to reduce variance (Casella and Robert, 1996), and has been used for black box variational inference (Ranganath *et al.*, 2014; Titsias and Lázaro-Gredilla, 2015). One particularly successful example is the sum-and-sample estimator (Liu *et al.*, 2019), which provides a trade-off between explicit marginalization and the SFE by summing over some values $\bar{\mathcal{Z}} \subset \mathcal{Z}$ and compensating with a sample from the remainder:

$$\begin{aligned} \mathbb{E} [\partial_{\theta} F(\mathbf{Z})] &= \partial_{\theta} \left(\sum_{z \in \bar{\mathcal{Z}}} p(z | \theta) F(z) \right) \\ &\quad + (1 - p(\bar{\mathcal{Z}} | \theta)) \mathbb{E}_{p_{\bar{\mathcal{Z}}}(Z | \theta)} [F(\mathbf{Z}) \partial_{\theta} \log p_{\bar{\mathcal{Z}}}(\mathbf{Z} | \theta)] \end{aligned}$$

where we defined the renormalized tail distribution obtained after discarding the summed-out values $\bar{\mathcal{Z}}$:

$$p_{\bar{\mathcal{Z}}}(\mathbf{z}) \propto \begin{cases} 0, & \mathbf{z} \in \bar{\mathcal{Z}}, \\ p(\mathbf{z} | \theta), & \mathbf{z} \notin \bar{\mathcal{Z}}, \end{cases}$$

A good choice for $\bar{\mathcal{Z}}$ is a set of highest-probability values. This strategy does not efficiently translate to the general structured case, as the second term requires computing and sampling from the renormalized tail distribution. Another example of Rao-Blackwellized estimators is the sample-without-replacement estimator (Kool *et al.*, 2020c).

5.6 Sparsifying the Distribution

The first method discussed in this section, exact marginalization, is entirely deterministic; there is no sampling and all computations are exact, but it is costly or even intractable for large \mathcal{Z} . The next two

methods invoke Monte Carlo estimation to achieve tractability, but this introduces randomness, hence variance, which may be undesirable and lead to unstable learning. In this section we present middle ground strategies based on sparsity.

Until now, we have assumed that Z follows a Gibbs distribution,

$$p(\mathbf{Z} = \mathbf{z} \mid x; \boldsymbol{\theta}_f) \propto \exp f(x, \mathbf{z}, \boldsymbol{\theta}_f). \quad (5.12)$$

This choice can be motivated by a maximum-entropy principle, as discussed in §3.2

Unlike the Monte Carlo strategies, which are based on sampling, we now seek a way to deterministically compute the full expectation over a discrete set,

$$\mathbb{E}[g(x, y, \mathbf{Z})] = \sum_{\mathbf{z} \in \mathcal{Z}} g(x, y, \mathbf{z}) p(\mathbf{z} \mid x, \boldsymbol{\theta}_f). \quad (5.13)$$

by choosing a parametrization in which $p(\mathbf{z} \mid x, \boldsymbol{\theta}_f)$ is set to exactly zero for many $\mathbf{z} \in \mathcal{Z}$. When this happens, the sum and its gradient can be computed efficiently with much fewer calls to the downstream model g . Importantly, the Gibbs parametrization in Eq. 5.12 can never give us this property, since $\exp(\cdot) > 0$.

In §3.2, we have explored sparse counterparts to the softmax transform in the context of replacing \mathbf{Z} by a relaxed mean counterpart $\hat{\mathbf{z}} = \mathbb{E}[\mathbf{Z}]$, and using the continuity of $\hat{\mathbf{z}}$ to learn end-to-end models with a single call to a continuous decoder. In this section, we show how sparse mappings can also induce a sparse distribution over \mathbf{Z} , leading to efficient marginalization, as proposed by Correia *et al.* (2020).

One-of-K (categorical) case. Let us set

$$p(\mathbf{Z} = \mathbf{z} \mid x) = [f(x, \mathbf{z}) - \tau]_+, \quad (5.14)$$

with τ such as the distribution sums to 1. This corresponds to applying sparsemax to the vector of scores \mathbf{s} (where $s_{\mathbf{z}} = f(x, \mathbf{z})$) to set the

probability distribution of \mathbf{Z} , therefore it will assign zero probabilities to many choices. The vector $\boldsymbol{\alpha}_* = \text{sparsemax}(\mathbf{s})$ can be computed in average time $\Theta(K)$; typical implementations $\Theta(K \log K)$ are empirically faster. We then can treat such a sparse $\boldsymbol{\alpha}$ as a discrete probability distribution, setting $p(\mathbf{Z}) = \boldsymbol{\alpha}$, or more explicitly, $p(\mathbf{Z} = \mathbf{e}_i) = \alpha_i$. We can then compute $\mathbb{E}_{\boldsymbol{\alpha}}[g(x, y, \mathbf{Z})]$ and its gradients with only $\|\boldsymbol{\alpha}_*\|_0$ calls to g . For large enough K and/or expensive g , this can be substantially faster in practice than K calls, while performing comparably on downstream tasks, as the following example due to Correia *et al.* (2020) shows.

Example 18. Single-token communication game. Emergent communication between interacting agents can be formulated as a game, or equivalently a latent variable model (Lazaridou *et al.*, 2017; Havrylov and Titov, 2017): given a dataset of labeled images, player 1 and player 2 are both shown the same set of N images, but arranged in a different order. Player 1 picks an image and then a code word from a vocabulary of size K ; they send the code word to player 2. The communication is successful if player 2 identifies the correct image, and unsuccessful otherwise. Lazaridou *et al.* (2017) observe near-perfect communication success with SFE and straight-through estimators for the case $N = 2$. Havrylov and Titov (2017) find that with more distractor images and a larger vocabulary, performance degrades. In particular, Correia *et al.* (2020) find that with $N = 16, K = 256$ none of the benchmarked MC-based methods solve the game. (The best performing MC model is based on SFE, reaching 55.36% communication success rate. Explicit marginalization can reach 93.37% but requires 256 decoder calls. Marginalization over the sparsmax distribution of Eq. 5.14, however, reaches 93.35% with only 3.13 decoder calls on average, illustrating the potential of sparse marginalization for learning.

Despite in practice demonstrating very sparse distributions on average, notice that when $f(x, \mathbf{z}) = f(x, \mathbf{z}')$ for any $\mathbf{z}, \mathbf{z}' \in \mathcal{Z}$, the vector \mathbf{s} is constant and so $\boldsymbol{\alpha}_\star = \frac{1}{|\mathcal{Z}|}(1, \dots, 1)^\top$. Thus, in the worst case, it is possible for sparsemax to be as dense as softmax, rendering this method slow or even intractable for large or structured \mathcal{Z} . Below we discuss two generalizations that help in this scenario.

Top- k sparsemax. A way to bound the sparsity of sparsemax solutions is by introducing a non-convex ℓ_0 constraint:

$$\text{sparsemax}_{(k)}(\mathbf{s}) := \arg \min_{\boldsymbol{\alpha} \in \Delta^K, \|\boldsymbol{\alpha}\|_0 \leq k} \frac{1}{2} \|\mathbf{s} - \boldsymbol{\alpha}\|^2. \quad (5.15)$$

Despite the non-convexity, Kyriillidis *et al.* (2013) show that an optimal solution can be found by applying sparsemax to the k -dimensional vector of highest scores. Formally we may define an operator $\text{top}_{(k)} : \mathbb{R}^K \rightarrow \mathbb{R}^K$ such that $[\text{top}_{(k)}(\mathbf{s})]_i = s_i$ if i is among a set of k indices with score greater than or equal to all others, and $-\infty$ otherwise. But in practice all but the top k indices can be ignored during computation. In the unstructured case, finding the top k indices can be done with a partial quicksort in $O(n + k \log k)$. In the structured case, there is no general strategy, but specialized algorithms are available for many structures of interest.

Crucially, if the ℓ_0 constraint is loose, (*i.e.*, the solution has strictly fewer than k nonzeros), then for this particular \mathbf{s} , $\text{sparsemax}_{(k)}(\mathbf{s}) = \text{sparsemax}(\mathbf{s})$, and thus we have found an optimal solution much faster. Remarkably, in the structured case, this means we can sometimes compute exact expectations of arbitrary functions over exponentially many structures.

On one hand, this observation can lead to another algorithm for sparsemax that is more efficient in the average case: start with k set to a guess of what $\|\text{sparsemax}(\mathbf{s})\|_0$ might be, and compute $\boldsymbol{\alpha} =$

$\text{sparsemax}_{(k)}(\mathbf{s})$. If $\|\boldsymbol{\alpha}\|_0 < k$, we found the solution. Otherwise, increase k (e.g., doubling) and repeat. This strategy was proposed in the unstructured case by Peters *et al.* (2019).

On the other hand, Correia *et al.* (2020) propose to directly use $\text{sparsemax}_{(k)}(\mathbf{s})$ for marginalizing over latent variables, without increasing k in the case of suboptimality. This way we may control the computational complexity, at the cost of biasing the estimation toward the top. In practice, Correia *et al.* (2020) find that, provided sufficiently high k , after some training iterations, the encoder f becomes more confident and the estimation transitions from approximate to exact. This algorithm can be applied to structured latent variables as well, as long as a structured top- k oracle is available.

SparseMAP marginalization. Under Assumption 1, recall we have

$$f(x, \mathbf{z}) = \langle \mathbf{z}, \mathbf{s}(x; \boldsymbol{\theta}_f) \rangle$$

where the vector-valued function \mathbf{s} returns a vector of part scores. As described in §3.5, SparseMAP uses a quadratic regularization on the mean vector $\boldsymbol{\mu} = \mathbb{E}_{\mathbf{Z} \sim \boldsymbol{\alpha}}[\mathbf{Z}]$, interpreting $\boldsymbol{\alpha} \in \Delta_{|\mathcal{Z}|}$ as the distribution of \mathbf{Z} . We may equivalently rewrite the SparseMAP optimization problem in terms of $\boldsymbol{\alpha}$:

$$\begin{aligned} \text{SparseMAP}(\mathbf{s}) &= \mathbb{E}_{\mathbf{Z} \sim \boldsymbol{\alpha}_\star}[\mathbf{Z}], \\ \text{where } \boldsymbol{\alpha}_\star &\in \arg \max_{\boldsymbol{\alpha} \in \Delta_{|\mathcal{Z}|}} \langle \mathbb{E}_{\mathbf{Z} \sim \boldsymbol{\alpha}}[\mathbf{Z}], \mathbf{s} \rangle - \frac{1}{2} \|\mathbb{E}_{\mathbf{Z} \sim \boldsymbol{\alpha}}[\mathbf{Z}]\|^2. \end{aligned} \quad (5.16)$$

The active set algorithm (Nocedal and Wright, 1999) proposed by Niculae *et al.* (2018) for SparseMAP returns not only the unique solution $\text{SparseMAP}(\mathbf{s})$ but also a corresponding sparse $\boldsymbol{\alpha}_\star$ and a sparse generalized Jacobian of $\boldsymbol{\alpha}_\star$ w.r.t. \mathbf{s} , (Correia *et al.*, 2020), allowing the use SparseMAP instead of sparsemax for sparse marginalizing over $\boldsymbol{\alpha}_\star$. The non-uniqueness of $\boldsymbol{\alpha}_\star$ leads to a bias inherent to the deterministic

active set algorithm, but this does not seem to impact performance. Alternatively, Correia *et al.* (2020) propose mitigating the bias by turning active set into a stochastic algorithm by initializing it with a random structure instead of the MAP structure, leading to a (non-uniform) sample from the set $\{\alpha : \mathbb{E}_{Z \sim \alpha}[Z] = \text{SparseMAP}(\mathbf{s})\}$. If we denote by $p(\alpha)$ the implicit sampling distribution of the Active Set algorithm, we can then optimize

$$\mathbb{E}_{\alpha \sim p(\alpha)} \mathbb{E}_{Z \sim \alpha} [g(x, y, Z)],$$

with a sampled α from the outer expectation and an explicit sum over the nonzero \mathbf{z} for the inner one. The two strategies perform comparably. Like the top-k sparsemax, SparseMAP exhibits adaptive sparsity: when the models are very confident, SparseMAP and MAP coincide and then there is only one possible α_* which is a one-hot vector. The set of possible α is larger earlier in training when models are less certain. A more specific characterization of the set of possible α for a given \mathbf{s} in general seems elusive but may be possible for specific structured models.

Note that in the categorical (one-of-K) case both sparsemax and SparseMAP are equivalent, since the possible values of Z are one-hot vectors, so $\mathbb{E}_{Z \sim \alpha}[Z] = \alpha$

5.7 Summary

The probabilistic formulation allows us to handle discrete latent structure in a way that is both fundamentally principled as well as very flexible, leading to some of the methods with the lowest assumptions in our text. There is a price to pay in complexity, as we must always consider entire distributions and not just punctual values. This translates either to computational complexity (for exact marginalization) or sample complexity (in the case of Monte Carlo methods). Correia *et al.* (2020) compare variance-reduced SFE, Gumbel-ST, top-k sparsemax and SparseMAP on a variational autoencoder with a bit-vector

latent variable, finding relatively similar results in terms of heldout log-likelihood, but a different distribution in terms of distortion/rate decomposition in which SFE-based methods perform worse. When applicable, Gumbel-ST and related methods seem to perform very well in practice. However, bit vectors are one of the few types of structure for which these methods all apply. Zantedeschi *et al.* (2022) study permutation latent variables, for which neither SFE nor Gumbel-ST apply, but SparseMAP does.

6

Conclusions

6.1 Overview

We survey in this text several strategies for learning models with discrete latent structure. This sort of structure is relevant for many fields and applications, including natural language processing, computer vision, and bioinformatics, where data is often well represented by compositional structures such as trees, sequences, or matchings. Latent structure models are a powerful tool for learning to extract or induce such representations, offering a way to incorporate structural bias, discover insight about the data, and interpret decisions.

We present a range of different strategies for learning latent discrete structures organized into three main groups:

- In Chapter 3 we cover techniques based on continuous relaxations, where the discrete structure \mathbf{z} , originally belonging to a discrete and combinatorial set \mathcal{Z} , and *relaxed* to be part of a larger continuous and polyhedral space, $\text{conv}(\mathcal{Z})$. These techniques relax the forward pass mapping by maximizing a regularized objective, so that they can compute its gradient exactly. Examples of technique falling into this class include attention mechanisms and their structured counterparts, sparse transformations such as sparsemax and α -entmax, as well as their structured counterparts, such as Sinkhorn-based methods and SparseMAP, and perturbation-based

regularization methods.

- In Chapter 4 we present several techniques that sidestep the problem of differentiating through discrete structures by building surrogate gradients. Surrogate gradient methods provide a way to train deterministic models without compromising on the use of discrete mappings in the forward pass (*e.g.*, argmax or quantization layers). Unlike the previous class of methods, the forward pass is left unchanged, but an approximation is made in the backward pass. This approximation can often be justified in principled ways and can work well in practice, and we attempt to shed some light on these principles. Surrogate gradient methods typically also require continuity of the decoder g *w.r.t.* \mathbf{z} : Even though they use only discrete \mathbf{z} in the forward pass, the backward pass requires the gradient $\partial_3 g(x, y, \hat{\mathbf{z}})$. Examples of methods falling into this class include straight-through gradient estimation, its structured variants, such as SPIGOT, and quantization methods, such as rounding and vector quantization.
- Finally, in Chapter 5 we focus on techniques based on probabilistic estimation. Unlike the previous methods, this class of techniques give the discrete latent a probabilistic treatment, allowing us to have correct (unbiased) estimation without compromising discreteness and without any requirements on g . This is the class of methods with the lowest assumptions in our text. There is a price to pay in complexity, as we must always consider entire distributions and not just punctual values. This translates either to computational complexity (for exact marginalization) or sample complexity (in the case of Monte Carlo methods). Examples of methods in this class include explicit marginalization and their sparsifications, the score function estimator (often called REINFORCE), the path gradient estimation (based on the reparametrization

trick), structured perturbation methods such as Perturb-and-MAP, the Gumbel-Softmax relaxation, mixed random variables such as Gaussian-Sparsemax, and I-MLE—some of these methods involve also gradient surrogates, as the previous class of methods.

Throughout the text, we guide the reader through a unifying view, while highlighting connections, strengths, weaknesses, assumptions, and emerging recommended practices. A bird’s eye view of the discussed classes of methods and relationships between them is presented in Table 6.1.

6.2 Implementations and Libraries

The flexible nature of structure representation makes it generally challenging to fit latent structure learning inside a neat API framework, and research using methods we cover usually relies on custom implementations from scratch. Readers might find useful the code released alongside the work of Kim *et al.* (2017), Baziotis *et al.* (2019), and Mihaylova *et al.* (2020), for instance. Nevertheless, a few general purpose libraries have been developed. We list some prominent ones.

- [torch-struct](#) (Rush, 2020): a PyTorch library for differentiable structure prediction, supporting differentiable implementations of marginal inference for a number of supported structure models. It therefore provides building blocks used in Chapters 3 to 5.
- [synjax](#) (Stanojević and Sartran, 2023): a jax library with similar scope as torch-struct; applicable to Chapters 3 to 5.
- [entmax](#) (Peters *et al.*, 2019): a library implementing efficient sparse argmax operators (Chapter 5) for the categorical (one-of-K) case, and thus applicable for incremental structures.

- [lp-sparsemap](#) (Niculae and Martins, 2020): a pytorch frontend for differentiable optimization in factor graphs. Supports a select set of structures (one-of-K, arborescences, matchings) and arbitrarily complex models formed by interactions between them. Provides some of the computational primitives discussed in Chapters 3 and 4.
- [cvxpylayers](#) (Agrawal *et al.*, 2019). A multi-backend library for differentiable convex optimization written in a flexible API (disciplined parametrized programs); useful for exploring and prototyping relaxations and projections from Chapters 3 and 4.
- [jaxopt](#) (Blondel *et al.*, 2021): a jax library for differentiable optimization. Per Chapter 3, relaxed structured latents (as well as some types of surrogate gradients per Chapter 4) can be formulated as regularized constrained optimization problems and thus manually implemented using this library.
- [stochastic](#) (Krieken *et al.*, 2021): a PyTorch library for gradient estimation with discrete latent variables, supporting most methods we discuss in Chapter 5 for the discrete case.

Table 6.1: Overview of methods discussed for latent structure learning. We abbreviate $\hat{\mathbf{z}}_0 = \hat{\mathbf{z}}_0(\mathbf{s}) := \arg \max_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{z}, \mathbf{s} \rangle$. Convergent (cvg.) learning means the approximation error can be controlled.

Method	Description	Learning	Computation oracle	Notes
Relaxation methods $g(x, y, \mathbb{E}[\mathbf{Z}])$				
Marginals	$\mathbb{E}_{p(\mathbf{Z})} \times \exp f(\cdot) \cdot [\mathbf{Z}]$			<i>Cha</i> $e.g.$, belief propagation
SparseMAP	$\mathbf{proj}_{\text{conv}(\mathcal{Z})}$	exact approx. exact cvg.	dedicated algorithm local subproblems dedicated algorithm argmax	Active Set algorithm dual decomposition
Perturbation	$\mathbb{E}[\hat{\mathbf{z}}_0(\mathbf{s} + \mathbf{U})]$	approx.	local subproblems	
Sinkhorn	$\hat{\mathbf{z}}_S(\mathbf{s})$	cvg.	argmax	
		cvg.	local subproblems	only assignment/tran
Surrogate gradient methods $g(x, y, \hat{\mathbf{z}}_0(\mathbf{s}))$				
ST	$\partial \hat{\mathbf{z}}(\mathbf{s}) := \text{Id}$		fwd pass / bwd pass	<i>Cha</i>
ST-Softmax	$\partial \hat{\mathbf{z}}(\mathbf{s}) := \partial \text{softmax}$	approx.	argmax/—	same oracle as marginals
SPIGOT	$\partial \hat{\mathbf{z}}(\mathbf{s})(\mathbf{v}) := \hat{\mathbf{z}}_0 - \mathbf{proj}_{\text{conv}(\mathcal{Z})}(\hat{\mathbf{z}}_0 - \eta \mathbf{v})$	approx.	argmax/ ∂ marginals argmax/projection	same oracle as SparseMAP
LI	$\partial \hat{\mathbf{z}}(\mathbf{s})(\mathbf{v}) := 1/\eta (\hat{\mathbf{z}}_0(\mathbf{s} + \eta \mathbf{v}) - \hat{\mathbf{z}}_0(\mathbf{s}))$	approx.	argmax/argmax	
Probabilistic latent variable methods $\mathbb{E}[g(x, y, \mathbf{Z})]$				
Path gradients	$\hat{\mathbf{z}}_{\text{surr}}(\mathbf{s} + \mathbf{U})$	approx.	same as chosen surrogate	<i>Cha</i> incl. ST-Gumbel, I-M
SFE		cvg.	sampling, marginals	A1, A2 not required
Explicit sum	$p(\mathbf{z}) := \text{softmax}$ $p(\mathbf{z}) := \text{sparsemax}$ $p(\mathbf{z}) := \text{SparseMAP}$	exact approx. approx.	— top-k argmax	small \mathcal{Z} . A1,A2 not req. A1,A2 not req. A2 not req.

Acknowledgements

This work was partly supported by the Dutch Research Council (NWO) via VI.Veni.212.228 and Zwaartekracht Hybrid Intelligence, the European Union's Horizon Europe research and innovation programme via UTTER 101070631, and the European Research Council (ERC) via StG DeepSPIN 758969.



References

- Ackley, D. H., G. E. Hinton, and T. J. Sejnowski. (1985). “A learning algorithm for Boltzmann machines”. *Cognitive science*. 9(1): 147–169.
- Adams, R. P. and R. S. Zemel. (2011). “Ranking via sinkhorn propagation”. *preprint arXiv:1106.1925*.
- Agrawal, A., B. Amos, S. Barratt, S. Boyd, S. Diamond, and Z. Kolter. (2019). “Differentiable Convex Optimization Layers”. In: *Advances in Neural Information Processing Systems*.
- Anil, C., Y. Wu, A. Andreassen, A. Lewkowycz, V. Misra, V. Ramasesh, A. Slone, G. Gur-Ari, E. Dyer, and B. Neyshabur. (2022). “Exploring Length Generalization in Large Language Models”. In: *Advances in Neural Information Processing Systems*. Ed. by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh. Vol. 35. Curran Associates, Inc. 38546–38556. URL: https://proceedings.neurips.cc/paper_files/paper/2022/file/fb7451e43f9c1c35b774bcfad7a5714b-Paper-Conference.pdf.
- Baker, J. K. (1979). “Trainable grammars for speech recognition”. *The Journal of the Acoustical Society of America*. 65(S1): S132–S132. DOI: 10.1121/1.2017061. eprint: <https://doi.org/10.1121/1.2017061>. URL: <https://doi.org/10.1121/1.2017061>.
- Barber, D. (2012). *Bayesian reasoning and machine learning*. Cambridge University Press.

- Bastings, J., W. Aziz, and I. Titov. (2019). “Interpretable Neural Predictions with Differentiable Binary Variables”. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics. 2963–2977. DOI: [10.18653/v1/P19-1284](https://doi.org/10.18653/v1/P19-1284). URL: <https://aclanthology.org/P19-1284>.
- Baum, L. E. (1972). “An inequality and associated maximization technique in statistical estimation of probabilistic functions of a Markov process”. In: URL: <https://api.semanticscholar.org/CorpusID:60804212>.
- Baziotis, C., I. Androutsopoulos, I. Konstas, and A. Potamianos. (2019). “SEQ³: Differentiable Sequence-to-Sequence-to-Sequence Autoencoder for Unsupervised Abstractive Sentence Compression”. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Ed. by J. Burstein, C. Doran, and T. Solorio. Minneapolis, Minnesota: Association for Computational Linguistics. 673–681. DOI: [10.18653/v1/N19-1071](https://doi.org/10.18653/v1/N19-1071). URL: <https://aclanthology.org/N19-1071>.
- Beal, M., Z. Ghahramani, and C. Rasmussen. (2001). “The Infinite Hidden Markov Model”. In: *Advances in Neural Information Processing Systems*. Ed. by T. Dietterich, S. Becker, and Z. Ghahramani. Vol. 14. MIT Press. URL: https://proceedings.neurips.cc/paper_files/paper/2001/file/e3408432c1a48a52fb6c74d926b38886-Paper.pdf.
- Bellman, R. (1954). “The theory of dynamic programming”. *Bulletin of the American Mathematical Society*. 60(6): 503–515.
- Bengio, Y., A. Courville, and P. Vincent. (2013a). “Representation Learning: A Review and New Perspectives”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 35(8): 1798–1828. DOI: [10.1109/TPAMI.2013.50](https://doi.org/10.1109/TPAMI.2013.50).

- Bengio, Y., N. Léonard, and A. Courville. (2013b). “Estimating or propagating gradients through stochastic neurons for conditional computation”. In: *Proc. of NIPS*.
- Berthet, Q., M. Blondel, O. Teboul, M. Cuturi, J.-P. Vert, and F. Bach. (2020). “Learning with differentiable perturbed optimizers”. *Advances in neural information processing systems*. 33: 9508–9519.
- Bertsekas, D. P. (1988). “The auction algorithm: A distributed relaxation method for the assignment problem”. *Annals of Operations Research*. 14(1): 105–123.
- Bertsekas, D. P. (1999). *Nonlinear Programming*. Athena Scientific Belmont.
- Bertsekas, D. P. and J. N. Tsitsiklis. (2000). “Gradient convergence in gradient methods with errors”. *SIAM Journal on Optimization*. 10(3): 627–642.
- Birkhoff, G. (1946). “Tres observaciones sobre el algebra lineal”. *Univ. Nac. Tucumán Rev. Ser. A*. 5: 147–151.
- Bishop, C. M. (1998). “Latent Variable Models”. In: *Learning in Graphical Models*. Ed. by M. I. Jordan. Dordrecht: Springer Netherlands. 371–403. ISBN: 978-94-011-5014-9. DOI: [10.1007/978-94-011-5014-9_13](https://doi.org/10.1007/978-94-011-5014-9_13). URL: https://doi.org/10.1007/978-94-011-5014-9_13.
- Blondel, M., Q. Berthet, M. Cuturi, R. Frostig, S. Hoyer, F. Linares-López, F. Pedregosa, and J.-P. Vert. (2021). “Efficient and Modular Implicit Differentiation”. *arXiv preprint arXiv:2105.15183*.
- Blondel, M., A. F. Martins, and V. Niculae. (2020). “Learning with Fenchel-Young losses”. *Journal of Machine Learning Research*. 21(35): 1–69.

- Bogin, B., S. Subramanian, M. Gardner, and J. Berant. (2021). “Latent Compositional Representations Improve Systematic Generalization in Grounded Question Answering”. *Transactions of the Association for Computational Linguistics*. 9: 195–210. Ed. by B. Roark and A. Nenkova. DOI: [10.1162/tacl_a_00361](https://doi.org/10.1162/tacl_a_00361). URL: <https://aclanthology.org/2021.tacl-1.12>.
- Borwein, J. and A. S. Lewis. (2010). *Convex analysis and nonlinear optimization: theory and examples*. Springer Science & Business Media.
- Brent, M. R. (1999). “An efficient, probabilistically sound algorithm for segmentation and word discovery”. *Machine Learning*. 34: 71–105.
- Bridle, J. (1989). “Training Stochastic Model Recognition Algorithms as Networks can Lead to Maximum Mutual Information Estimation of Parameters”. In: *Advances in Neural Information Processing Systems*. Ed. by D. Touretzky. Vol. 2. Morgan-Kaufmann. URL: https://proceedings.neurips.cc/paper_files/paper/1989/file/0336dcbab05b9d5ad24f4333c7658a0e-Paper.pdf.
- Brucker, P. (1984). “An $O(n)$ algorithm for quadratic knapsack problems”. *Operations Research Letters*. 3(3): 163–166.
- Caruana, R. (1997). “Multitask learning”. *Machine learning*. 28(1): 41–75.
- Casella, G. and C. P. Robert. (1996). “Rao-Blackwellisation of Sampling Schemes”. *Biometrika*. 83(1): 81–94. ISSN: 00063444, 14643510. URL: <http://www.jstor.org/stable/2337434> (accessed on 11/06/2024).
- Chen, H., S. Branavan, R. Barzilay, and D. R. Karger. (2009). “Content modeling using latent permutations”. *Journal of Artificial Intelligence Research*. 36: 129–163.

- Cho, K., B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. (2014). “Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation”. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Doha, Qatar: Association for Computational Linguistics. 1724–1734. DOI: [10.3115/v1/D14-1179](https://doi.org/10.3115/v1/D14-1179). URL: <https://aclanthology.org/D14-1179>.
- Chu, Y.-J. and T.-H. Liu. (1965). “On the Shortest Arborescence of a Directed Graph”. *Science Sinica*. 14: 1396–1400.
- Clark, A. (2001). “Unsupervised induction of stochastic context-free grammars using distributional clustering”. In: *Proceedings of the ACL 2001 Workshop on Computational Natural Language Learning (ConLL)*. URL: <https://aclanthology.org/W01-0713>.
- Cocke, J. and J. T. Schwartz. (1970). “Programming languages and their compilers: preliminary notes”. *Tech. rep.*
- Cohen, S. (2019). *Bayesian Analysis in Natural Language Processing. Synth. Lect. Human Lang. Technol.* Morgan & Claypool.
- Cohen, S. B., C. Gómez-Rodríguez, and G. Satta. (2012). “Elimination of Spurious Ambiguity in Transition-Based Dependency Parsing”. *preprint arXiv:1206.6735*.
- Condat, L. (2016). “Fast projection onto the simplex and the ℓ_1 ball”. *Mathematical Programming*. 158(1-2): 575–585.
- Correia, G. M., V. Niculae, and A. F. Martins. (2019). “Adaptively Sparse Transformers”. In: *Proceedings of EMNLP-IJCNLP*.
- Correia, G. M., V. Niculae, W. Aziz, and A. F. T. Martins. (2020). “Efficient Marginalization of Discrete and Structured Latent Variables via Sparsity”. In: *Proceedings of NeurIPS*. URL: <http://arxiv.org/abs/2007.01919>.
- Corro, C. and I. Titov. (2019a). “Differentiable Perturb-and-Parse: Semi-Supervised Parsing with a Structured Variational Autoencoder”. In: *Proc. of ICLR*.

- Corro, C. and I. Titov. (2019b). “Learning latent trees with stochastic perturbations and differentiable dynamic programming”. In: *Proc. of ACL*.
- Cuturi, M. (2013). “Sinkhorn distances: Lightspeed computation of optimal transport”. In: *Proceedings of NeurIPS*.
- Cuturi, M. and M. Blondel. (2017). “Soft-DTW: a Differentiable Loss Function for Time-Series”. In: *Proceedings of ICML*. Ed. by D. Precup and Y. W. Teh. Vol. 70. *Proceedings of Machine Learning Research*. PMLR. URL: <https://proceedings.mlr.press/v70/cuturi17a.html>.
- d’Aspremont, A., D. Scieur, and A. Taylor. (2021). “Acceleration Methods”. *Foundations and Trends® in Optimization*. 5(1-2): 1–245. ISSN: 2167-3888. DOI: [10.1561/24000000036](https://doi.org/10.1561/24000000036). URL: <http://dx.doi.org/10.1561/24000000036>.
- Dantzig, G. B., A. Orden, and P. Wolfe. (1955). “The generalized simplex method for minimizing a linear form under linear inequality restraints”. *Pacific Journal of Mathematics*. 5(2): 183–195.
- Deisenroth, M. P., A. A. Faisal, and C. S. Ong. (2020). *Mathematics for machine learning*. Cambridge University Press.
- Dempster, A. P., N. M. Laird, and D. B. Rubin. (1977). “Maximum Likelihood from Incomplete Data Via the EM Algorithm”. *Journal of the Royal Statistical Society: Series B (Methodological)*. 39(1): 1–22. DOI: <https://doi.org/10.1111/j.2517-6161.1977.tb01600.x>. eprint: <https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/j.2517-6161.1977.tb01600.x>. URL: <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/j.2517-6161.1977.tb01600.x>.
- Deng, Y., Y. Kim, J. Chiu, D. Guo, and A. Rush. (2018). “Latent alignment and variational attention”. *Advances in Neural Information Processing Systems*. 31.

- Domke, J. (2013). “Learning graphical model parameters with approximate marginal inference”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 35(10): 2454–2467.
- Drozdo, A. and S. Bowman. (2017). “The coadaptation problem when learning how and what to compose”. In: *Proc of ReplNLP*.
- Du, L., W. Buntine, and M. Johnson. (2013). “Topic Segmentation with a Structured Topic Model”. In: *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Ed. by L. Vanderwende, H. Daumé III, and K. Kirchhoff. Atlanta, Georgia: Association for Computational Linguistics. 190–200. URL: <https://aclanthology.org/N13-1019>.
- Duchi, J. C., S. Shalev-Shwartz, Y. Singer, and T. Chandra. (2008). “Efficient projections onto the ℓ_1 -ball for learning in high dimensions”. In: *Proc. of ICML*.
- Edmonds, J. (1967). “Optimum branchings”. *J. Res. Nat. Bur. Stand.* 71B: 233–240.
- Eisenstein, J. and R. Barzilay. (2008). “Bayesian Unsupervised Topic Segmentation”. In: *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*. Ed. by M. Lapata and H. T. Ng. Honolulu, Hawaii: Association for Computational Linguistics. 334–343. URL: <https://aclanthology.org/D08-1035>.
- Eisner, J. (2016). “Inside-Outside and Forward-Backward Algorithms Are Just Backprop (tutorial paper)”. In: *Proceedings of the Workshop on Structured Prediction for NLP*. Ed. by K.-W. Chang, M.-W. Chang, A. Rush, and V. Srikumar. Austin, TX: Association for Computational Linguistics. 1–17. DOI: [10.18653/v1/W16-5901](https://doi.org/10.18653/v1/W16-5901). URL: <https://aclanthology.org/W16-5901>.

- Elsayed, G., A. Mahendran, S. van Steenkiste, K. Greff, M. C. Mozer, and T. Kipf. (2022). “SAVi++: Towards End-to-End Object-Centric Learning from Real-World Videos”. In: *Advances in Neural Information Processing Systems*. Ed. by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh. Vol. 35. Curran Associates, Inc. 28940–28954. URL: https://proceedings.neurips.cc/paper_files/paper/2022/file/ba1a6ba05319e410f0673f8477a871e3-Paper-Conference.pdf.
- Farinhas, A., W. Aziz, V. Niculae, and A. F. Martins. (2022). “Sparse Communication via Mixed Distributions”. In: *International Conference on Learning Representations*.
- Forney, G. (1973). “The Viterbi algorithm”. *Proceedings of the IEEE*. 61(3): 268–278. DOI: [10.1109/PROC.1973.9030](https://doi.org/10.1109/PROC.1973.9030).
- Frank, M. and P. Wolfe. (1956). “An algorithm for quadratic programming”. *Nav. Res. Log.* 3(1-2): 95–110.
- Frühwirth-Schnatter, S. (1994). “Data augmentation and dynamic linear models”. *Journal of Time Series Analysis*. 15(2): 183–202. DOI: <https://doi.org/10.1111/j.1467-9892.1994.tb00184.x>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-9892.1994.tb00184.x>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9892.1994.tb00184.x>.
- Fu, Y., C. Tan, B. Bi, M. Chen, Y. Feng, and A. Rush. (2020). “Latent Template Induction with Gumbel-CRFs”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin. Vol. 33. Curran Associates, Inc. 20259–20271. URL: <https://proceedings.neurips.cc/paper/2020/file/ea119a40c1592979f51819b0bd38d39d-Paper.pdf>.
- Garey, M. R. and D. S. Johnson. (1979). *Computers and intractability*. Vol. 174. freeman San Francisco.

- Gholami, A., S. Kim, Z. Dong, Z. Yao, M. W. Mahoney, and K. Keutzer. (2022). “A Survey of Quantization Methods for Efficient Neural Network Inference”. In: *Low-Power Computer Vision*. Ed. by G. K. Thiruvathukal, Y.-H. Lu, J. Kim, Y. Chen, and B. Chen. Chapman and Hall/CRC.
- Glasserman, P. (1990). *Gradient estimation via perturbation analysis*. Vol. 116. Springer Science & Business Media.
- Glynn, P. W. (1990). “Likelihood ratio gradient estimation for stochastic systems”. *Communications of the ACM*. 33(10): 75–84.
- Goldstein, A. A. (1964). “Convex programming in Hilbert space”. *Bulletin of the American Mathematical Society*. 70(5): 709–710.
- Goldwater, S., T. L. Griffiths, and M. Johnson. (2006). “Contextual Dependencies in Unsupervised Word Segmentation”. In: *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*. Ed. by N. Calzolari, C. Cardie, and P. Isabelle. Sydney, Australia: Association for Computational Linguistics. 673–680. DOI: [10.3115/1220175.1220260](https://doi.org/10.3115/1220175.1220260). URL: <https://aclanthology.org/P06-1085>.
- Goldwater, S. and T. Griffiths. (2007). “A fully Bayesian approach to unsupervised part-of-speech tagging”. In: *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*. Ed. by A. Zaenen and A. van den Bosch. Prague, Czech Republic: Association for Computational Linguistics. 744–751. URL: <https://aclanthology.org/P07-1094>.
- Grathwohl, W., D. Choi, Y. Wu, G. Roeder, and D. Duvenaud. (2018). “Backpropagation through the void: optimizing control variates for black-box gradient estimation”. In: *Proc. ICLR*.

- Greff, K., R. L. Kaufman, R. Kabra, N. Watters, C. Burgess, D. Zoran, L. Matthey, M. Botvinick, and A. Lerchner. (2019). “Multi-Object Representation Learning with Iterative Variational Inference”. In: *Proceedings of the 36th International Conference on Machine Learning*. Ed. by K. Chaudhuri and R. Salakhutdinov. Vol. 97. *Proceedings of Machine Learning Research*. PMLR. 2424–2433. URL: <https://proceedings.mlr.press/v97/greff19a.html>.
- Griewank, A. and A. Walther. (2008). *Evaluating Derivatives*. Second. Society for Industrial and Applied Mathematics. DOI: [10.1137/1.9780898717761](https://doi.org/10.1137/1.9780898717761). eprint: <https://epubs.siam.org/doi/pdf/10.1137/1.9780898717761>. URL: <https://epubs.siam.org/doi/abs/10.1137/1.9780898717761>.
- Gu, S., S. Levine, I. Sutskever, and A. Mnih. (2016). “[MuProp: unbiased backpropagation for stochastic neural networks](#)”. In: *Proc. ICLR*.
- Gumbel, E. J. (1954). *Statistical theory of extreme values and some practical applications: a series of lectures*. Vol. 33. US Government Printing Office.
- Havrylov, S. and I. Titov. (2017). “[Emergence of language with multi-agent games: Learning to communicate with sequences of symbols](#)”. In: *Proc. NeurIPS*.
- Hazan, T. and T. Jaakkola. (2012). “On the partition function and random maximum a-posteriori perturbations”. In: *Proceedings of ICML*.
- Hazan, T., S. Maji, and T. Jaakkola. (2013). “On Sampling from the Gibbs Distribution with Random Maximum A-Posteriori Perturbations”. In: *Advances in Neural Information Processing Systems*. Ed. by C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger. Vol. 26. Curran Associates, Inc. URL: <https://proceedings.neurips.cc/paper/2013/file/443cb001c138b2561a0d90720d6ce111-Paper.pdf>.

- Held, M., P. Wolfe, and H. P. Crowder. (1974). “[Validation of subgradient optimization](#)”. *Mathematical Programming*. 6(1): 62–88.
- Hinton, G. E. (2012). “Neural networks for machine learning”. Coursera, video lectures.
- Hinton, G. E. and Z. Ghahramani. (1997). “Generative models for discovering sparse distributed representations”. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*. 352(1358): 1177–1190.
- Hinton, G. E. (2002). “Training Products of Experts by Minimizing Contrastive Divergence”. *Neural Computation*. 14(8): 1771–1800. ISSN: 0899-7667. DOI: [10.1162/089976602760128018](https://doi.org/10.1162/089976602760128018). eprint: <https://direct.mit.edu/neco/article-pdf/14/8/1771/815447/089976602760128018.pdf>. URL: <https://doi.org/10.1162/089976602760128018>.
- Hinton, G. E., S. Osindero, and Y.-W. Teh. (2006). “A Fast Learning Algorithm for Deep Belief Nets”. *Neural Computation*. 18(7): 1527–1554. ISSN: 0899-7667. DOI: [10.1162/neco.2006.18.7.1527](https://doi.org/10.1162/neco.2006.18.7.1527). eprint: <https://direct.mit.edu/neco/article-pdf/18/7/1527/816558/neco.2006.18.7.1527.pdf>. URL: <https://doi.org/10.1162/neco.2006.18.7.1527>.
- Hitchcock, F. L. (1941). “The distribution of a product from several sources to numerous localities”. *Journal of mathematics and physics*. 20(1–4): 224–230.
- Ho, Y. and X. Cao. (1983). “Optimization and perturbation analysis of queueing networks”. *Journal of Optimization Theory and Applications*. 40(4): 559–582.
- Hoogeboom, E., J. Peters, R. van den Berg, and M. Welling. (2019). “[Integer discrete flows and lossless compression](#)”. In: *Advances in Neural Information Processing Systems*.

- Huang, L. (2008). “Advanced Dynamic Programming in Semiring and Hypergraph Frameworks”. In: *Coling 2008: Advanced Dynamic Programming in Computational Linguistics: Theory, Algorithms and Applications - Tutorial notes*. Ed. by L. Huang. Manchester, UK: Coling 2008 Organizing Committee. 1–18. URL: <https://aclanthology.org/C08-5001>.
- Huijben, I. A., W. Kool, M. B. Paulus, and R. J. Van Sloun. (2022). “A Review of the Gumbel-max Trick and its Extensions for Discrete Stochasticity in Machine Learning”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Ishwaran, H., J. S. Rao, *et al.* (2005). “Spike and slab variable selection: frequentist and Bayesian strategies”. *Annals of Statistics*. 33(2): 730–773.
- Jang, E., S. Gu, and B. Poole. (2017). “Categorical reparameterization with Gumbel-Softmax”. In: *Proc. of ICLR*.
- Johnson, M., T. Griffiths, and S. Goldwater. (2007). “Bayesian Inference for PCFGs via Markov Chain Monte Carlo”. In: *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*. Ed. by C. Sidner, T. Schultz, M. Stone, and C. Zhai. Rochester, New York: Association for Computational Linguistics. 139–146. URL: <https://aclanthology.org/N07-1018>.
- Johnson, M. J., D. K. Duvenaud, A. Wiltchko, R. P. Adams, and S. R. Datta. (2016). “Composing graphical models with neural networks for structured representations and fast inference”. In: *Advances in Neural Information Processing Systems*. Ed. by D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett. Vol. 29. Curran Associates, Inc. URL: https://proceedings.neurips.cc/paper_files/paper/2016/file/7d6044e95a16761171b130dcb476a43e-Paper.pdf.

- Jonker, R. and A. Volgenant. (1987). “A shortest augmenting path algorithm for dense and sparse linear assignment problems”. *Computing*. 38(4): 325–340.
- Jurafsky, D. and J. H. Martin. (2018). *Speech and Language Processing (3rd ed.)* draft.
- Kakade, S. M., S. Shalev-Shwartz, and A. Tewari. (2012). “Regularization techniques for learning with matrices”. *Journal of Machine Learning Research*. 13: 1865–1890.
- Kantorovich, L. (1942). “On the translocation of masses”. *Dokl Akad. Nauk SSSR*. 37(7–8): 227–229.
- Kasai, J., K. Sakaguchi, R. Le Bras, D. Radev, Y. Choi, and N. A. Smith. (2024). “A Call for Clarity in Beam Search: How It Works and When It Stops”. In: *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*. Ed. by N. Calzolari, M.-Y. Kan, V. Hoste, A. Lenci, S. Sakti, and N. Xue. Torino, Italia: ELRA and ICCL. 77–90. URL: <https://aclanthology.org/2024.lrec-main.7>.
- Kasami, T. (1966). “An efficient recognition and syntax-analysis algorithm for context-free languages”. *Tech. rep.*
- Keith, K., S. L. Blodgett, and B. O’Connor. (2018). “Monte Carlo Syntax Marginals for Exploring and Using Dependency Parses”. In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. New Orleans, Louisiana: Association for Computational Linguistics. DOI: 10.18653/v1/N18-1084. URL: <https://aclanthology.org/N18-1084>.
- Kim, Y., C. Denton, L. Hoang, and A. M. Rush. (2017). “Structured attention networks”. In: *Proc. of ICLR*.
- Kim, Y., S. Wiseman, and A. M. Rush. (2018). “A tutorial on deep latent variable models of natural language”. *arXiv preprint arXiv:1812.06834*.

- Kingma, D. and J. Ba. (2015). “Adam: A method for stochastic optimization”. In: *Proc. of ICLR*.
- Kingma, D. P. and M. Welling. (2014). “Auto-encoding Variational Bayes”. In: *Proceedings of ICLR*.
- Kingma, D. P., S. Mohamed, D. Jimenez Rezende, and M. Welling. (2014). “Semi-supervised Learning with Deep Generative Models”. In: *Advances in Neural Information Processing Systems*. Ed. by Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger. Vol. 27. Curran Associates, Inc. URL: <https://proceedings.neurips.cc/paper/2014/file/d523773c6b194f37b938d340d5d02232-Paper.pdf>.
- Kiperwasser, E. and Y. Goldberg. (2016). “Simple and accurate dependency parsing using bidirectional LSTM feature representations”. *TACL*. 4: 313–327.
- Kirchhoff, G. (1847). “Ueber die Auflösung der Gleichungen, auf welche man bei der Untersuchung der linearen Vertheilung galvanischer Ströme geführt wird”. *Annalen der Physik*. 148(12): 497–508.
- Klein, D. and C. Manning. (2004). “Corpus-Based Induction of Syntactic Structure: Models of Dependency and Constituency”. In: *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL-04)*. Barcelona, Spain. 478–485. DOI: [10.3115/1218955.1219016](https://doi.org/10.3115/1218955.1219016). URL: <https://aclanthology.org/P04-1061>.
- Komodakis, N., N. Paragios, and G. Tziritas. (2007). “MRF optimization via dual decomposition: Message-passing revisited.” In: *ICCV*. Vol. 1. No. 3. 5.
- Kool, W., H. van Hoof, and M. Welling. (2020a). “Ancestral Gumbel-Top-k Sampling for Sampling Without Replacement”. *Journal of Machine Learning Research*. 21(47): 1–36. URL: <http://jmlr.org/papers/v21/19-985.html>.

- Kool, W., H. van Hoof, and M. Welling. (2020b). “Ancestral Gumbel-Top-k Sampling for Sampling Without Replacement”. *Journal of Machine Learning Research*. 21(47): 1–36. URL: <http://jmlr.org/papers/v21/19-985.html>.
- Kool, W., H. van Hoof, and M. Welling. (2020c). “Estimating Gradients for Discrete Random Variables by Sampling without Replacement”. In: *Proc. ICLR*.
- Kool, W., H. Van Hoof, and M. Welling. (2019). “Stochastic Beams and Where To Find Them: The Gumbel-Top-k Trick for Sampling Sequences Without Replacement”. In: *Proceedings of the 36th International Conference on Machine Learning*. Ed. by K. Chaudhuri and R. Salakhutdinov. Vol. 97. *Proceedings of Machine Learning Research*. PMLR. 3499–3508. URL: <https://proceedings.mlr.press/v97/kool19a.html>.
- Krieken, E. van, J. M. Tomczak, and A. t. Teije. (2021). “Stochastic: A Framework for General Stochastic Automatic Differentiation”. 34: 7574–7587. Ed. by M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan. URL: https://proceedings.neurips.cc/paper_files/paper/2021/file/3dfe2f633108d604df160cd1b01710db-Paper.pdf.
- Kuhn, H. W. (1955). “The Hungarian method for the assignment problem”. *Nav. Res. Log.* 2(1-2): 83–97.
- Kumaraswamy, P. (1980). “A generalized probability density function for double-bounded random processes”. *Journal of hydrology*. 46(1-2): 79–88.
- Kyrillidis, A., S. Becker, V. Cevher, and C. Koch. (2013). “Sparse projections onto the simplex”. In: *Proc. ICML*.
- L’Ecuyer, P. (1995). “Note: On the interchange of derivative and expectation for likelihood ratio derivative estimators”. *Management Science*. 41(4): 738–747.

- Lacoste-Julien, S. and M. Jaggi. (2015). “On the global linear convergence of Frank-Wolfe optimization variants”. In: *Proc. of NIPS*.
- Lafferty, J., A. McCallum, F. Pereira, *et al.* (2001). “Conditional random fields: Probabilistic models for segmenting and labeling sequence data”. In: *Icml*. Vol. 1. No. 2. Williamstown, MA. 3.
- Lan, Z., M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut. (2020). “ALBERT: A Lite BERT for Self-supervised Learning of Language Representations”. In: *Proceedings of ICLR*. URL: <https://openreview.net/forum?id=H1eA7AEtvS>.
- Lari, K. and S. Young. (1990). “The estimation of stochastic context-free grammars using the Inside-Outside algorithm”. *Computer Speech & Language*. 4(1): 35–56. ISSN: 0885-2308. DOI: [https://doi.org/10.1016/0885-2308\(90\)90022-X](https://doi.org/10.1016/0885-2308(90)90022-X). URL: <https://www.sciencedirect.com/science/article/pii/088523089090022X>.
- Lazaridou, A., A. Peysakhovich, and M. Baroni. (2017). “Multi-agent cooperation and the emergence of (natural) language”. In: *Proc. ICLR*.
- Levitin, E. S. and B. T. Polyak. (1966). “Constrained minimization methods”. *USSR Computational mathematics and mathematical physics*. 6(5): 1–50.
- Li, Z. and J. Eisner. (2009). “First-and second-order expectation semirings with applications to minimum-risk training on translation forests”. In: *Proc. of EMNLP*.
- Liang, P., M. I. Jordan, and D. Klein. (2010). “Learning programs: A hierarchical Bayesian approach”. In: *Proceedings of ICML*.
- Lin, W., M. E. Khan, and N. Hubacher. (2018). “Variational Message Passing with Structured Inference Networks”. In: *International Conference on Learning Representations*. URL: <https://openreview.net/forum?id=HyH9lbZAW>.

- Linnainmaa, S. (1970). “The representation of the cumulative rounding error of an algorithm as a Taylor expansion of the local rounding errors”. *PhD thesis*. Master’s Thesis (in Finnish), Univ. Helsinki.
- Liu, C., S. An, Z. Lin, Q. Liu, B. Chen, J.-G. Lou, L. Wen, N. Zheng, and D. Zhang. (2021). “Learning Algebraic Recombination for Compositional Generalization”. In: *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*. Ed. by C. Zong, F. Xia, W. Li, and R. Navigli. Online: Association for Computational Linguistics. 1129–1144. DOI: [10.18653/v1/2021.findings-acl.97](https://doi.org/10.18653/v1/2021.findings-acl.97). URL: <https://aclanthology.org/2021.findings-acl.97>.
- Liu, D., M. Lapata, and F. Keller. (2023). “Visual Storytelling with Question-Answer Plans”. In: *Findings of the Association for Computational Linguistics: EMNLP 2023*. Ed. by H. Bouamor, J. Pino, and K. Bali. Singapore: Association for Computational Linguistics. 5800–5813. DOI: [10.18653/v1/2023.findings-emnlp.386](https://doi.org/10.18653/v1/2023.findings-emnlp.386). URL: <https://aclanthology.org/2023.findings-emnlp.386>.
- Liu, R., J. Regier, N. Tripuraneni, M. Jordan, and J. Mcauliffe. (2019). “Rao-Blackwellized stochastic gradients for discrete distributions”. In: *Proc. ICML*.
- Liu, Y., M. Gardner, and M. Lapata. (2018). “Structured Alignment Networks for Matching Sentences”. In: *Proceedings of EMNLP*. Association for Computational Linguistics. DOI: [10.18653/v1/D18-1184](https://doi.org/10.18653/v1/D18-1184). URL: <https://aclanthology.org/D18-1184>.
- Liu, Y. and M. Lapata. (2018). “Learning structured text representations”. *TACL*. 6: 63–75.

- Locatello, F., D. Weissenborn, T. Unterthiner, A. Mahendran, G. Heigold, J. Uszkoreit, A. Dosovitskiy, and T. Kipf. (2020). “Object-Centric Learning with Slot Attention”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin. Vol. 33. Curran Associates, Inc. 11525–11538. URL: https://proceedings.neurips.cc/paper_files/paper/2020/file/8511df98c02ab60aea1b2356c013bc0f-Paper.pdf.
- Loshchilov, I. and F. Hutter. (2019). “Decoupled Weight Decay Regularization”. In: *International Conference on Learning Representations*. URL: <https://openreview.net/forum?id=Bkg6RiCqY7>.
- Louizos, C., M. Welling, and D. P. Kingma. (2018). “Learning Sparse Neural Networks through L_0 Regularization”. In: *International Conference on Learning Representations*. URL: <https://openreview.net/forum?id=H1Y8hhg0b>.
- Lowerre, B. T. (1976). “The HARP Y speech recognition system”. *PhD*. Carnegie-Mellon University.
- Maddison, C. J., A. Mnih, and Y. W. Teh. (2017). “The concrete distribution: A continuous relaxation of discrete random variables”. In: *Proc. of ICLR*.
- Martins, A., N. Smith, and E. Xing. (2009a). “Concise Integer Linear Programming Formulations for Dependency Parsing”. In: *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*. Ed. by K.-Y. Su, J. Su, J. Wiebe, and H. Li. Suntec, Singapore: Association for Computational Linguistics. 342–350. URL: <https://aclanthology.org/P09-1039>.
- Martins, A. F. and R. F. Astudillo. (2016). “From softmax to sparsemax: A sparse model of attention and multi-label classification”. In: *Proc. of ICML*.

- Martins, A. F., M. A. Figueiredo, P. M. Aguiar, N. A. Smith, and E. P. Xing. (2015). “AD3: Alternating directions dual decomposition for MAP inference in graphical models”. *JMLR*. 16(1): 495–545.
- Martins, A. F., N. A. Smith, and E. P. Xing. (2009b). “Concise integer linear programming formulations for dependency parsing”. In: *Proc. of ACL-IJCNLP*.
- Meister, C., A. Amini, T. Vieira, and R. Cotterell. (2021). “Conditional Poisson Stochastic Beams”. In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Ed. by M.-F. Moens, X. Huang, L. Specia, and S. W.-t. Yih. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics. 664–681. DOI: [10.18653/v1/2021.emnlp-main.52](https://doi.org/10.18653/v1/2021.emnlp-main.52). URL: <https://aclanthology.org/2021.emnlp-main.52>.
- Meister, C., R. Cotterell, and T. Vieira. (2020a). “If beam search is the answer, what was the question?” In: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Ed. by B. Webber, T. Cohn, Y. He, and Y. Liu. Online: Association for Computational Linguistics. 2173–2185. DOI: [10.18653/v1/2020.emnlp-main.170](https://doi.org/10.18653/v1/2020.emnlp-main.170). URL: <https://aclanthology.org/2020.emnlp-main.170>.
- Meister, C., T. Vieira, and R. Cotterell. (2020b). “Best-First Beam Search”. *Transactions of the Association for Computational Linguistics*. 8: 795–809. Ed. by M. Johnson, B. Roark, and A. Nenkova. DOI: [10.1162/tacl_a_00346](https://doi.org/10.1162/tacl_a_00346). URL: <https://aclanthology.org/2020.tacl-1.51>.
- Mena, G., D. Belanger, S. Linderman, and J. Snoek. (2018). “Learning latent permutations with Gumbel-Sinkhorn networks”. In: *Proc. of ICLR*.
- Mensch, A. and M. Blondel. (2018). “Differentiable dynamic programming for structured prediction and attention”. In: *Proc. of ICML*.

- Meshi, O., M. Mahdavi, and A. G. Schwing. (2015). “Smooth and strong: MAP inference with linear convergence”. In: *Proc. of NIPS*.
- Metropolis, N. and S. Ulam. (1949). “The Monte Carlo Method”. *Journal of the American Statistical Association*. 44(247): 335–341. DOI: 10.1080/01621459.1949.10483310. eprint: <https://www.tandfonline.com/doi/pdf/10.1080/01621459.1949.10483310>. URL: <https://www.tandfonline.com/doi/abs/10.1080/01621459.1949.10483310>.
- Mihaylova, T., V. Niculae, and A. F. T. Martins. (2020). “Understanding the Mechanics of SPIGOT: Surrogate Gradients for Latent Structure Learning”. In: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Online: Association for Computational Linguistics. DOI: 10.18653/v1/2020.emnlp-main.171. URL: <https://aclanthology.org/2020.emnlp-main.171>.
- Mitchell, T. J. and J. J. Beauchamp. (1988). “Bayesian variable selection in linear regression”. *Journal of the American Statistical Association*. 83(404): 1023–1032.
- Mnih, A. and K. Gregor. (2014). “Neural Variational Inference and Learning in Belief Networks”. In: *Proceedings of ICML*.
- Mohamed, S., M. Rosca, M. Figurnov, and A. Mnih. (2020). “Monte Carlo Gradient Estimation in Machine Learning.” *J. Mach. Learn. Res.* 21(132): 1–62.
- Mohri, M. (2002). “Semiring frameworks and algorithms for shortest-distance problems”. *J. Autom. Lang. Comb.* 7(3): 321–350. ISSN: 1430-189X.
- Murphy, K. P. (2022). *Probabilistic Machine Learning: An introduction*. MIT Press. URL: probml.ai.

- Narayan, S., J. Maynez, R. K. Amplayo, K. Ganchev, A. Louis, F. Huot, A. Sandholm, D. Das, and M. Lapata. (2023). “Conditional Generation with a Question-Answering Blueprint”. *Transactions of the Association for Computational Linguistics*. 11: 974–996. DOI: [10.1162/tacl_a_00583](https://doi.org/10.1162/tacl_a_00583). URL: <https://aclanthology.org/2023.tacl-1.55>.
- Neal, R. M. (1992). “Connectionist learning of belief networks”. *Artificial Intelligence*. 56(1): 71–113. ISSN: 0004-3702. DOI: [https://doi.org/10.1016/0004-3702\(92\)90065-6](https://doi.org/10.1016/0004-3702(92)90065-6). URL: <https://www.sciencedirect.com/science/article/pii/0004370292900656>.
- Needleman, S. B. and C. D. Wunsch. (1970). “A general method applicable to the search for similarities in the amino acid sequence of two proteins”. *Journal of Molecular Biology*. 48(3): 443–453. ISSN: 0022-2836. DOI: [https://doi.org/10.1016/0022-2836\(70\)90057-4](https://doi.org/10.1016/0022-2836(70)90057-4). URL: <https://www.sciencedirect.com/science/article/pii/0022283670900574>.
- Nesterov, Y. (1983). “A method for solving the convex programming problem with convergence rate $O(1/k^2)$ ”. In: *Dokl akad nauk Sssr*. Vol. 269. 543.
- Niculae, V. and M. Blondel. (2017). “A regularized framework for sparse and structured neural attention”. In: *Proc. of NIPS*.
- Niculae, V. and A. Martins. (2020). “LP-SparseMAP: Differentiable Relaxed Optimization for Sparse Structured Prediction”. In: *Proceedings of ICML*. Ed. by H. D. III and A. Singh. Vol. 119. *Proceedings of Machine Learning Research*. PMLR. URL: <https://proceedings.mlr.press/v119/niculae20a.html>.
- Niculae, V., A. F. Martins, M. Blondel, and C. Cardie. (2018). “SparseMAP: Differentiable sparse structured inference”. In: *Proc. of ICML*.
- Niepert, M., P. Minervini, and L. Franceschi. (2021). “Implicit MLE: backpropagating through discrete exponential family distributions”. *Advances in Neural Information Processing Systems*. 34: 14567–14579.

- Nocedal, J. and S. Wright. (1999). *Numerical Optimization*. Springer New York.
- Oord, A. van den, O. Vinyals, and k. kavukcuoglu koray. (2017). “Neural Discrete Representation Learning”. In: *Advances in Neural Information Processing Systems*. Ed. by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett. Vol. 30. Curran Associates, Inc. URL: <https://proceedings.neurips.cc/paper/2017/file/7a98af17e63a0ac09ce2e96d03992fbc-Paper.pdf>.
- Paisley, J., D. M. Blei, and M. I. Jordan. (2012). “Variational bayesian inference with stochastic search”. In: *Proc. ICML*.
- Palmer, A. W., A. J. Hill, and S. J. Scheduling. (2017). “Methods for Stochastic Collection and Replenishment (SCAR) optimisation for persistent autonomy”. *Robotics and Autonomous Systems*. 87: 51–65.
- Papandreou, G. and A. L. Yuille. (2011). “Perturb-and-map random fields: Using discrete optimization to learn and sample from energy models”. In: *2011 International Conference on Computer Vision*. IEEE. 193–200.
- Parisi, G. and R. Shankar. (1988). “Statistical field theory”.
- Paulus, M., D. Choi, D. Tarlow, A. Krause, and C. J. Maddison. (2020). “Gradient estimation with stochastic softmax tricks”. *Advances in Neural Information Processing Systems*. 33: 5691–5704.
- Pearce, M. (2020). “The Gaussian Process Prior VAE for Interpretable Latent Dynamics from Pixels”. In: *Proceedings of The 2nd Symposium on Advances in Approximate Bayesian Inference*. Ed. by C. Zhang, F. Ruiz, T. Bui, A. B. Dieng, and D. Liang. Vol. 118. *Proceedings of Machine Learning Research*. PMLR. 1–12. URL: <https://proceedings.mlr.press/v118/pearce20a.html>.

- Peng, H., S. Thomson, and N. A. Smith. (2018). “Backpropagating through Structured Argmax using a SPIGOT”. In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Melbourne, Australia: Association for Computational Linguistics. 1863–1873. DOI: [10.18653/v1/P18-1173](https://doi.org/10.18653/v1/P18-1173). URL: <https://aclanthology.org/P18-1173>.
- Peter L Bartlett, M. I. J. and J. D. McAuliffe. (2006). “Convexity, Classification, and Risk Bounds”. *Journal of the American Statistical Association*. 101(473): 138–156. DOI: [10.1198/016214505000000907](https://doi.org/10.1198/016214505000000907). eprint: <https://doi.org/10.1198/016214505000000907>. URL: <https://doi.org/10.1198/016214505000000907>.
- Peters, B., V. Niculae, and A. F. Martins. (2019). “Sparse sequence-to-sequence models”. In: *Proc. ACL*.
- Peterson, C. (1987). “A mean field theory learning algorithm for neural network”. *Complex systems*. 1: 995–1019.
- Peyré, G. and M. Cuturi. (2019). “Computational optimal transport”. *Foundations and Trends® in Machine Learning*. 11(5-6): 355–607.
- Pflug, G. C. (2012). *Optimization of stochastic models: the interface between simulation and optimization*. Vol. 373. Springer Science & Business Media.
- Pogančić, M. V., A. Paulus, V. Musil, G. Martius, and M. Rolinek. (2020). “Differentiation of Blackbox Combinatorial Solvers”. In: *International Conference on Learning Representations*. URL: <https://openreview.net/forum?id=BkevoJSYPB>.
- Punyakanok, V., D. Roth, W.-t. Yih, and D. Zimak. (2004). “Semantic Role Labeling Via Integer Linear Programming Inference”. In: *Proceedings of the 20th International Conference on Computational Linguistics*. 1346–1352.
- PyTorch. (2017). “<http://pytorch.org>”.

- Rabiner, L. R. (1989). “A tutorial on Hidden Markov Models and selected applications in speech recognition”. *P. IEEE*. 77(2): 257–286.
- Ranganath, R., S. Gerrish, and D. Blei. (2014). “Black Box Variational Inference”. In: *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*. Ed. by S. Kaski and J. Corander. Vol. 33. *Proceedings of Machine Learning Research*. Reykjavik, Iceland: PMLR. 814–822. URL: <https://proceedings.mlr.press/v33/ranganath14.html>.
- Read, J., B. Pfahringer, G. Holmes, and E. Frank. (2011). “Classifier chains for multi-label classification”. *Machine learning*. 85(3): 333–359.
- Reid, M. D. and R. C. Williamson. (2010). “Composite Binary Losses”. *Journal of Machine Learning Research*. 11(83): 2387–2422. URL: <http://jmlr.org/papers/v11/reid10a.html>.
- Rennie, S. J., E. Marcheret, Y. Mroueh, J. Ross, and V. Goel. (2017). “Self-critical sequence training for image captioning”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 7008–7024.
- Rezende, D. J., S. Mohamed, and D. Wierstra. (2014). “Stochastic Backpropagation and Approximate Inference in Deep Generative Models”. In: *Proceedings of the 31st International Conference on Machine Learning*. Ed. by E. P. Xing and T. Jebara. Vol. 32. *Proceedings of Machine Learning Research*. No. 2. Beijing, China: PMLR. 1278–1286. URL: <https://proceedings.mlr.press/v32/rezende14.html>.
- Robbins, H. and S. Monro. (1951). “A Stochastic Approximation Method”. *The Annals of Mathematical Statistics*. 22(3): 400–407. DOI: [10.1214/aoms/1177729586](https://doi.org/10.1214/aoms/1177729586). URL: <https://doi.org/10.1214/aoms/1177729586>.

- Rolfe, J. T. (2017). “Discrete Variational Autoencoders”. In: *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net. URL: <https://openreview.net/forum?id=ryMxXPfex>.
- Roth, D. and W.-t. Yih. (2005). “Integer linear programming inference for conditional random fields”. In: *Proceedings of the 22nd international conference on Machine learning*. 736–743.
- Rubinstein, R. Y. (1992). “Sensitivity analysis of discrete event systems by the “push out” method”. *Annals of Operations Research*. 39(1): 229–250.
- Rubinstein, R. (1976). “A Monte Carlo method for estimating the gradient in a stochastic network”. *Unpublished manuscript, Technion, Haifa, Israel*.
- Rush, A. M., D. Sontag, M. Collins, and T. Jaakkola. (2010). “On dual decomposition and linear programming relaxations for natural language processing”. In: Association for Computational Linguistics.
- Rush, A. M. (2020). “Torch-Struct: Deep Structured Prediction Library”. arXiv: 2002.00876 [cs.CL].
- Sagae, K. and A. Lavie. (2005). “A Classifier-Based Parser with Linear Run-Time Complexity”. In: *Proceedings of the Ninth International Workshop on Parsing Technology*. Vancouver, British Columbia: Association for Computational Linguistics. 125–132.
- Sakoe, H. and S. Chiba. (1978). “Dynamic programming algorithm optimization for spoken word recognition”. *IEEE Transactions on Acoustics, Speech, and Signal Processing*. 26(1): 43–49. DOI: [10.1109/TASSP.1978.1163055](https://doi.org/10.1109/TASSP.1978.1163055).

- Salakhutdinov, R. and G. Hinton. (2009). “Deep Boltzmann Machines”. In: *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*. Ed. by D. van Dyk and M. Welling. Vol. 5. *Proceedings of Machine Learning Research*. Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA: PMLR. 448–455. URL: <https://proceedings.mlr.press/v5/salakhutdinov09a.html>.
- Saul, L. K., T. Jaakkola, and M. I. Jordan. (1996). “Mean field theory for sigmoid belief networks”. *Journal of artificial intelligence research*. 4: 61–76.
- Sinkhorn, R. (1964). “A relationship between arbitrary positive matrices and doubly stochastic matrices”. *The annals of mathematical statistics*. 35(2): 876–879.
- Stanojević, M. (2022). “Unbiased and Efficient Sampling of Dependency Trees”. In: *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Ed. by Y. Goldberg, Z. Kozareva, and Y. Zhang. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics. 1691–1706. DOI: [10.18653/v1/2022.emnlp-main.110](https://doi.org/10.18653/v1/2022.emnlp-main.110). URL: <https://aclanthology.org/2022.emnlp-main.110>.
- Stanojević, M. and L. Sartran. (2023). “SynJax: Structured Probability Distributions for JAX”. *arXiv preprint arXiv:2308.03291*. URL: <https://arxiv.org/abs/2308.03291>.
- Steenkiste, S. van, M. Chang, K. Greff, and J. Schmidhuber. (2018). “Relational Neural Expectation Maximization: Unsupervised Discovery of Objects and their Interactions”. In: *International Conference on Learning Representations*. URL: <https://openreview.net/forum?id=ryH20GbRW>.
- Stoyanov, V., A. Ropson, and J. Eisner. (2011). “Empirical risk minimization of graphical model parameters given approximate inference, decoding, and model structure”. In: *Proc. of AISTATS*.

- Sutskever, I., O. Vinyals, and Q. V. Le. (2014). “Sequence to Sequence Learning with Neural Networks”. In: *Advances in Neural Information Processing Systems*. Ed. by Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger. Vol. 27. Curran Associates, Inc. URL: <https://proceedings.neurips.cc/paper/2014/file/a14ac55a4f27472c5d894ec1c3c743d2-Paper.pdf>.
- Taskar, B. (2004). “Learning structured prediction models: A large margin approach”. *PhD thesis*. Stanford University.
- Tay, Y., D. Bahri, L. Yang, D. Metzler, and D.-C. Juan. (2020). “Sparse Sinkhorn attention”. In: *Proceedings of ICML*.
- Titsias, M. and M. Lázaro-Gredilla. (2014). “Doubly Stochastic Variational Bayes for non-Conjugate Inference”. In: *Proceedings of the 31st International Conference on Machine Learning*. Ed. by E. P. Xing and T. Jebara. Vol. 32. *Proceedings of Machine Learning Research*. No. 2. Beijing, China: PMLR. 1971–1979. URL: <https://proceedings.mlr.press/v32/titsias14.html>.
- Titsias, M. K. and M. Lázaro-Gredilla. (2015). “Local Expectation Gradients for Black Box Variational Inference”. In: *Advances in Neural Information Processing Systems*. Ed. by C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett. Vol. 28. Curran Associates, Inc. URL: https://proceedings.neurips.cc/paper_files/paper/2015/file/1373b284bc381890049e92d324f56de0-Paper.pdf.
- Tsoumakas, G. and I. Katakis. (2007). “Multi-Label Classification: An Overview”. *International Journal of Data Warehousing and Mining (IJDWM)*. 3(3): 1–13. URL: <https://EconPapers.repec.org/RePEc:igg:jdwm00:v:3:y:2007:i:3:p:1-13>.
- Tucker, G., A. Mnih, C. J. Maddison, D. Lawson, and J. Sohl-Dickstein. (2017). “REBAR: Low-variance, unbiased gradient estimates for discrete latent variable models”. In: *Proceedings of NeurIPS*.
- Tutte, W. T. (2001). *Graph theory*. Vol. 21. Cambridge university press.

- Vahdat, A., W. Macreedy, Z. Bian, A. Khoshaman, and E. Andriyash. (2018). “DVAE++: Discrete Variational Autoencoders with Overlapping Transformations”. In: *Proceedings of the 35th International Conference on Machine Learning*. Ed. by J. Dy and A. Krause. Vol. 80. *Proceedings of Machine Learning Research*. PMLR. 5035–5044. URL: <https://proceedings.mlr.press/v80/vahdat18a.html>.
- Valiant, L. (1979). “The complexity of computing the permanent”. *Theoretical Computer Science*. 8(2): 189–201. ISSN: 0304-3975. DOI: [https://doi.org/10.1016/0304-3975\(79\)90044-6](https://doi.org/10.1016/0304-3975(79)90044-6). URL: <https://www.sciencedirect.com/science/article/pii/0304397579900446>.
- Vapnik, V. (1991). “Principles of Risk Minimization for Learning Theory”. In: *Advances in Neural Information Processing Systems*. Ed. by J. Moody, S. Hanson, and R. Lippmann. Vol. 4. Morgan-Kaufmann. URL: https://proceedings.neurips.cc/paper_files/paper/1991/file/ff4d5fbbafdf976cfdc032e3bde78de5-Paper.pdf.
- Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. (2017). “Attention is All you Need”. In: *Advances in Neural Information Processing Systems*. Ed. by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett. Vol. 30. Curran Associates, Inc. URL: <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>.
- Venkataraman, A. (2001). “A Statistical Model for Word Discovery in Transcribed Speech”. *Computational Linguistics*. 27(3): 351–372. Ed. by J. Hirschberg. DOI: [10.1162/089120101317066113](https://doi.org/10.1162/089120101317066113). URL: <https://aclanthology.org/J01-3002>.
- Vintsyuk, T. K. (1968). “Speech discrimination by dynamic programming”. *Cybernetics*. 4(1): 52–57.
- Vinyes, M. and G. Obozinski. (2017). “Fast column generation for atomic norm regularization”. In: *Proc. of AISTATS*.

- Viterbi, A. (1967). “Error bounds for convolutional codes and an asymptotically optimum decoding algorithm”. *IEEE Transactions on Information Theory*. 13(2): 260–269. DOI: [10.1109/TIT.1967.1054010](https://doi.org/10.1109/TIT.1967.1054010).
- Wagner, R. A. and M. J. Fischer. (1974). “The String-to-String Correction Problem”. *J. ACM*. 21(1): 168–173. ISSN: 0004-5411. DOI: [10.1145/321796.321811](https://doi.org/10.1145/321796.321811). URL: <https://doi.org/10.1145/321796.321811>.
- Wainwright, M. J. and M. I. Jordan. (2008). “Graphical models, exponential families, and variational inference.” *Found. Trends Mach. Learn.* 1(1–2): 1–305.
- Williams, R. J. (1992). “Simple statistical gradient-following algorithms for connectionist reinforcement learning”. *Machine Learning*. 8(3-4): 229–256.
- Williamson, R. C., E. Vernet, and M. D. Reid. (2016). “Composite multiclass losses”. *Journal of Machine Learning Research*.
- Wolfe, P. (1976). “Finding the nearest point in a polytope”. *Mathematical Programming*. 11(1): 128–149.
- Wong, C., K. M. Ellis, J. Tenenbaum, and J. Andreas. (2021). “Leveraging language to learn program abstractions and search heuristics”. In: *Proceedings of ICML*. PMLR.
- Xu, K., J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio. (2015). “Show, Attend and Tell: Neural Image Caption Generation with Visual Attention”. In: *Proceedings of the 32nd International Conference on Machine Learning*. Ed. by F. Bach and D. Blei. Vol. 37. *Proceedings of Machine Learning Research*. Lille, France: PMLR. 2048–2057. URL: <https://proceedings.mlr.press/v37/xuc15.html>.
- Xu, Y. and M. Lapata. (2022). “Document Summarization with Latent Queries”. *Transactions of the Association for Computational Linguistics*. 10: 623–638. Ed. by B. Roark and A. Nenkova. DOI: [10.1162/tacl_a_00480](https://doi.org/10.1162/tacl_a_00480). URL: <https://aclanthology.org/2022.tacl-1.36>.

- Yao, Y. and A. Koller. (2022). “Structural generalization is hard for sequence-to-sequence models”. In: *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Ed. by Y. Goldberg, Z. Kozareva, and Y. Zhang. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics. 5048–5062. DOI: [10.18653/v1/2022.emnlp-main.337](https://doi.org/10.18653/v1/2022.emnlp-main.337). URL: <https://aclanthology.org/2022.emnlp-main.337>.
- Yellott Jr., J. I. (1977). “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.
- Younger, D. H. (1967). “Recognition and parsing of context-free languages in time n^3 ”. *Information and Control*. 10(2): 189–208. ISSN: 0019-9958. DOI: [https://doi.org/10.1016/S0019-9958\(67\)80007-X](https://doi.org/10.1016/S0019-9958(67)80007-X). URL: <https://www.sciencedirect.com/science/article/pii/S001999586780007X>.
- Zălinescu, C. (2002). *Convex Analysis in General Vector Spaces*. World Scientific.
- Zanon Boito, M., B. Yusuf, L. Ondel, A. Villavicencio, and L. Besacier. (2022). “Unsupervised Word Segmentation from Discrete Speech Units in Low-Resource Settings”. In: *Proceedings of the 1st Annual Meeting of the ELRA/ISCA Special Interest Group on Under-Resourced Languages*. Ed. by M. Meler, S. Sakti, and C. Soria. Marseille, France: European Language Resources Association. 1–9. URL: <https://aclanthology.org/2022.sigul-1.1>.
- Zantedeschi, V., J. Kaddour, L. Franceschi, M. Kusner, and V. Niculae. (2022). “DAG learning on the Permutahedron”. In: *ICLR2022 Workshop on the Elements of Reasoning: Objects, Structure and Causality*.

- Zhao, Y. and S. Linderman. (2023). “Revisiting Structured Variational Autoencoders”. In: *Proceedings of the 40th International Conference on Machine Learning*. Ed. by A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett. Vol. 202. *Proceedings of Machine Learning Research*. PMLR. 42046–42057. URL: <https://proceedings.mlr.press/v202/zhao23c.html>.
- Zhou, H., A. Bradley, E. Littwin, N. Razin, O. Saremi, J. M. Susskind, S. Bengio, and P. Nakkiran. (2024). “What Algorithms can Transformers Learn? A Study in Length Generalization”. In: *The Twelfth International Conference on Learning Representations*. URL: <https://openreview.net/forum?id=AssIuHnmHX>.
- Zmigrod, R., T. Vieira, and R. Cotterell. (2021). “Efficient Sampling of Dependency Structure”. In: *Proceedings of EMNLP*. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics. DOI: [10.18653/v1/2021.emnlp-main.824](https://doi.org/10.18653/v1/2021.emnlp-main.824). URL: <https://aclanthology.org/2021.emnlp-main.824>.