Deep Discrete Latent Variable Models

Deep Learning 2 – 2022

Wilker Aziz
w.aziz@uva.nl


University of Amsterdam
Institute for Logic, Language and Computation

Previously

Probabilistic models prescribe the probability measure of a random experiment

- one of the many ways to achieve that is to specify the pdf (or pmf) of a collection of random variables

Previously

Probabilistic models prescribe the probability measure of a random experiment

- one of the many ways to achieve that is to specify the pdf (or pmf) of a collection of random variables
- a deep probabilistic model uses NNs to parameterise this pdf

Previously

Probabilistic models prescribe the probability measure of a random experiment

- one of the many ways to achieve that is to specify the pdf (or pmf) of a collection of random variables
- a deep probabilistic model uses NNs to parameterise this pdf

For parameter estimation, we decided to employ MLE

Previously

Probabilistic models prescribe the probability measure of a random experiment

- one of the many ways to achieve that is to specify the pdf (or pmf) of a collection of random variables
- a deep probabilistic model uses NNs to parameterise this pdf

For parameter estimation, we decided to employ MLE

- with a tractable and differentiable likelihood function, gradient-based search for NN parameters gives us a general purpose mechanism to approach supervised learning

Previously

Probabilistic models prescribe the probability measure of a random experiment

- one of the many ways to achieve that is to specify the pdf (or pmf) of a collection of random variables
- a deep probabilistic model uses NNs to parameterise this pdf

For parameter estimation, we decided to employ MLE

- with a tractable and differentiable likelihood function, gradient-based search for NN parameters gives us a general purpose mechanism to approach supervised learning

Finally, once the probabilistic model is fully specified (which includes parameter estimation), together with a decision rule, it can power applications (tasks).

Goals for this session

1. Prescribe joint distribution involving discrete and unobserved random variables
2. Estimate parameter via gradient-based MLE
3. Recognise the role of a model's posterior distribution in parameter estimation
4. Apply amortised variational inference
5. Derive gradient estimators for deep discrete latent variable models

Outline

Modelling observed random variables

Our goal is to learn a distribution over a set of **observed** random variables.

Observed random variables are the result of random experiments that have already happened: e.g., sentences in a collection of news articles, number of stars in a product review.

Modelling unobserved random variables

Unobserved random variables are variables that are

- observable in principle, but not available for observation (e.g., the topic of a piece of text, a semantic graph)
- unobservable (e.g., a 100-dimensional sentence embedding)

they help us prescribe and even estimate our models.

Our goal is to learn a distribution over *observed and unobserved rvs*

- make explicit assumptions about statistical dependence
- discover hidden structure
- mimic intuitions/knowledge about the data generating process
- deal with missing data
- estimate uncertainty about predictions

Deterministic predictors may also be available.

Unobserved random variables are also called latent variables.

For those interested in Bayesian statistics, note that the presence of unobserved random variables does not imply Bayesian modelling. Bayesian principles are a collection of ideas organised in what is called the Bayesian Theory (or Bayesian Decision Theory) for rational decision making under uncertainty (Bernardo and Smith, 2009). These ideas may cross paths with many aspects of our ML solutions.

Outline

Latent Structure and Over-Dispersion

We found some old manuscripts in an excavation site, historians began labelling them for publication date.

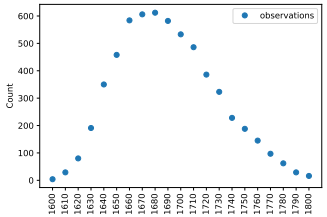

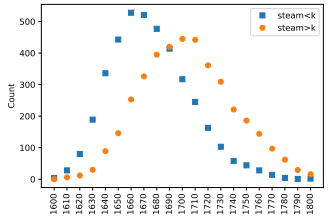

Marginally (left), it looked like our observations could have been drawn from a Poisson distribution.

Left: publication date of documents in the labelled collection.

Our historians claim that above some threshold k a document was likely written after 1699 (when Thomas Savery demonstrated his first steam engine to the British Royal society).

Right: data as a function of the frequency of the word *steam*.

Plotting two streams of data under such criterion reveals what could have been 2 different Poisson distributions.

Latent Structure and Over-Dispersion

We found some old manuscripts in an excavation site, historians began labelling them for publication date.

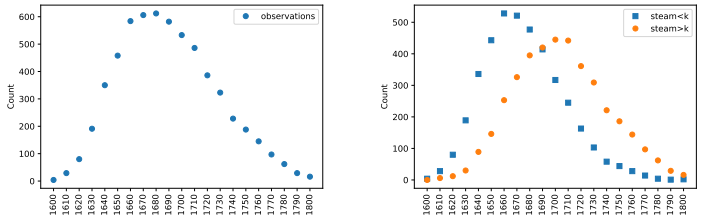

But they might also have been the result of mixing (right) into one population draws from two different Poisson distributions.

Left: publication date of documents in the labelled collection.

Our historians claim that above some threshold k a document was likely written after 1699 (when Thomas Savery demonstrated his first steam engine to the British Royal society).

Right: data as a function of the frequency of the word *steam*.

Plotting two streams of data under such criterion reveals what could have been 2 different Poisson distributions.

Latent Structure and Over-Dispersion

We found some old manuscripts in an excavation site, historians began labelling them for publication date.

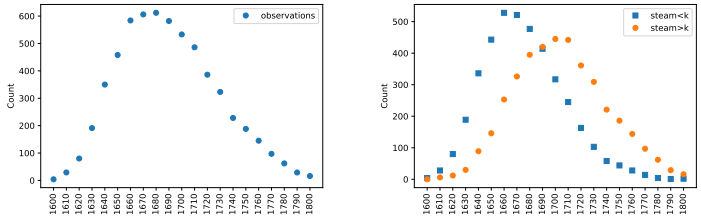

The way a model *views* the data tells us something about latent factors that account for (cause or correlate with) observed variance.

Left: publication date of documents in the labelled collection.

Our historians claim that above some threshold k a document was likely written after 1699 (when Thomas Savery demonstrated his first steam engine to the British Royal society).

Right: data as a function of the frequency of the word *steam*.

Plotting two streams of data under such criterion reveals what could have been 2 different Poisson distributions.

"Oh, latent variables are like hidden units, right?"

Latent structure here has to do with a partitioning of the probability space in terms of intermediate outcomes that depend on one another.

Hidden layers in an NN output deterministic transformations of their observed inputs. They are not statements about statistical dependence.

Example:

$$Y|h \sim \text{Poisson}\left(\underbrace{\exp\left(\sum_{i=1}^{D} w_i h_i\right)}_{\text{shallow NN}}\right)$$

any one draw comes from the exact same Poisson.

To reveal latent structure that is likely supported by observations, we need to postulate a joint distribution where observations and latent variables interact.

'Interacting' is a matter of **statistical dependence**.

Latent Structure and Multimodality

For a few manuscripts, we obtained labels from multiple experts.

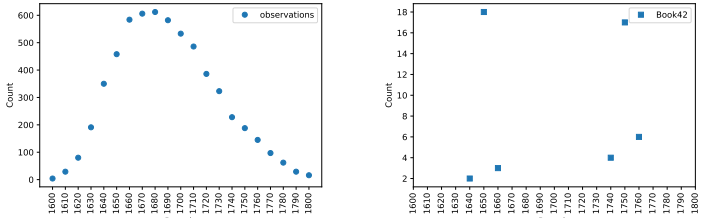

Left: observations for Y across the collection. Right: observations for Y given doc is book-42.

a unimodal conditional $Y|\mathrm{doc}$ seems inappropriate.

When we plot observations for Y (e.g., left) we see the data *marginally*.

If we intend to model the data conditionally, that plot won't help us pick a family. That is because our choice should be informed by plots of the kind $Y|\mathrm{doc}$. If we group our data into bins, where bin membership depends on matching a specific value of doc, more often than not our bins will each contain a single data point. Should we conclude that *conditionally* our data can be seen as deterministic? By no means!

Be aware of sneaky modelling assumptions. The combination of *'1 bin per unique document'* and *'one plot per bin'* is a modelling choice (for visualisation purposes, but still). One that suffers from data sparsity so tremendously that it makes a random variable look deterministic. Concluding that we can model the data deterministically is in fact an instance of overfitting (by humans).

Note that sometimes we can *construct* more meaningful $Y|\mathrm{doc}$ plots that reveal the stochastic nature of the data. For example, if we have direct access to the mechanism by which observations are generated, we can fix doc and draw Y multiple times (rightmost plot on the slide).

"But NNs can, in principle, learn anything, right?"

Not quite. We identify a *probability measure* by parameterising a joint pdf (or pmf). Thus our models are limited by the expressiveness of the families we choose.

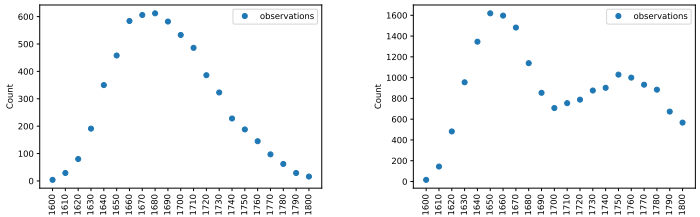

Left: data look unimodal and could have been drawn from a Poisson.
Right: data look bimodal and a single Poisson seems less likely.

"But NNs can, in principle, learn anything, right?"

Not quite. We identify a *probability measure* by parameterising a joint pdf (or pmf). Thus our models are limited by the expressiveness of the families we choose.

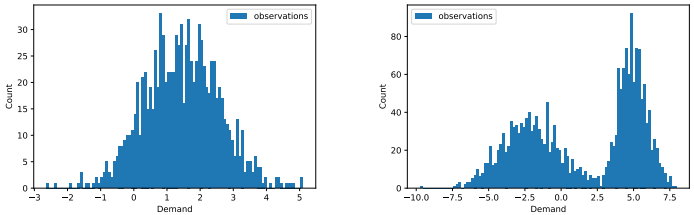

Left: data look unimodal and could have been drawn from a Gaussian.
Right: data look bimodal and a single Gaussian seems less likely.

Can we combine simple distributions?

We can however *mix* K members of each family to get a good fit:

For example, with $K = 2$

$$Z|b \sim \text{Bernoulli}(b)$$
$$Y|\lambda, b \sim \text{Poisson}(\lambda_z)$$

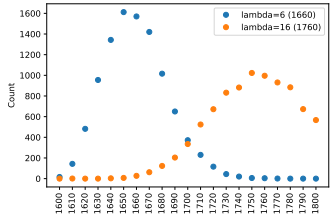

$$Z|b \sim \text{Bernoulli}(b)$$
$$Y|\mu, \sigma, b \sim \mathcal{N}(\mu_z, \sigma_z^2)$$

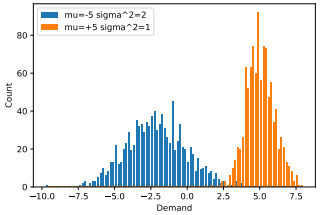

This is known as a **mixture model**. The specific ones on the slide combines two conditional distributions, namely, $Y|Z = 0$ and $Y|Z = 1$. The model mixes its conditional *components* stochastically, a process controlled by a distribution over components, whose probabilities $p(z|\theta)$ are known as *mixing weights*. That is, with probability $p(z|\theta)$ the component $Y|Z = z$ generates a draw in $\mathcal{Y}$. In this example, $p(Z = 1|\theta) = b$ and $p(Z = 0|\theta) = 1 - b$.

For $K > 2$ components, $Z \sim \text{Cat}(\pi_1, \ldots, \pi_K)$, thus $p(z|\pi) = \pi_z$.

Mixture model

A **mixture model** assigns probability density

$$p_{ZY}(z, y|\theta) = p_Z(z|\theta)_{Y|Z} p(y|z, \theta)$$

to joint outcomes in $\mathcal{Z} \times \mathcal{Y}$. That is, it prescribes a *joint distribution* over observed and unobserved random variables.

A mixture model encodes the assumption that data points are each drawn from one of a finite number of independent distributions.

The latent variable Z captures this *unobserved component assignment*. It is governed by a distribution we call the *prior*. Oftentimes this is as simple as a uniform distribution over the sample space $\mathcal{Z}$.

Given an observation y drawn from the mixture, we can *infer* a distribution over component assignments by basic probability calculus, this very famous result is known as Bayes rule:

$$p(z|y, \theta) = \frac{p(z, y|\theta)}{p(y|\theta)} = \frac{p(y|z, \theta)p(z|\theta)}{\sum_{z' \in \mathcal{Z}} p(y|z', \theta)p(z'|\theta)}$$

Note that this *posterior pdf* $p(z|y, \theta)$ involves the *marginal pdf* $p(y|\theta)$, which we discuss next.

Prescribing Flexible Distributions

The **marginal** distribution of the mixture model is potentially multimodal and exhibits richer covariance structure. It assigns probability density

$$p(y|\theta) = \sum_{z=1}^{K} p(z|\theta)p(y|z,\theta)$$

to an outcome $y \in \mathcal{Y}$ by *marginalisation* of assignments $z \in \mathcal{Z}$ of the latent random variable.

Marginal inference for the MM scales linearly in the number of *components*. This is a key type of computation, for example, indispensable for parameter estimation / learning via maximum likelihood.

Say we have a dataset $\mathcal{D}$ of N i.i.d. observations for Y. MLE depends on the log-likelihood function, which in turn depends on assessments of the *marginal probability* of each observation:

$$\begin{aligned}
\mathcal{L}_{\mathcal{D}}(\theta) &= \sum_{n=1}^{N} \log p(y^{(n)}|\theta) \\
&= \sum_{n=1}^{N} \log \sum_{z^{(n)}=1}^{K} p(y^{(n)}, z^{(n)}|\theta) \\
&= \sum_{n=1}^{N} \log \sum_{z^{(n)}=1}^{K} p(y^{(n)}|z^{(n)}, \theta)p(z^{(n)}|\theta)
\end{aligned}$$

Posterior component assignment

Given an observation y we can infer a distribution over component assignments via Bayes rule

$$p(z|y,\theta) = \frac{p(z,y|\theta)}{p(y|\theta)} = \frac{p(y|z,\theta)p(z|\theta)}{\sum_{z' \in \mathcal{Z}} p(y|z',\theta)p(z'|\theta)}$$

MMs are one of the first options when it comes to organising massive collections of unlabelled data into smaller groups (clustering).

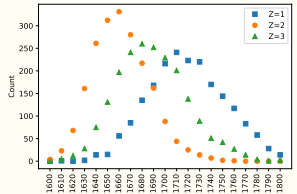

Here is a mixture model (3 Poisson components) of our historical data.

Components in a mixture model are not *labelled* with self-evident information such as *'pre-steam-engine'* and *'post-steam-engine'*, but sometimes by inspecting likely component assignments we can recognise some salient features data bring data points together under a certain component. We can also use it to target annotation efforts, for example, to avoid under-representing certain decades (in our running example).

Labelling components with self-evident information can be done by experts with assistance of posterior queries or even semi-automatically by extending mixture models in interesting ways. See LDA (Blei et al., 2003), for example.

Predictors are welcome

It is also possible to use mixture models in conditional models:

$$p(z, y|x, \theta) = p(z|x, \theta)p(y|x, z, \theta)$$

and we may use x in different ways, e.g.

$$p(z, y|x, \theta) = p(z|\theta)p(y|x, z, \theta)$$
$$p(z, y|x, \theta) = p(z|x, \theta)p(y|z, \theta)$$

Whether to use predictors to parameterise mixing weights, the conditional model, or both will depend on the application.

'Parameterising mixing weights' means specifying a distribution over K mixture components, e.g.

$$Z|x \sim \text{Cat}(g(x; \theta))$$

An alternative to giving control of mixing weights to a neural network, or fixing the weights to something superficially intuitive (like a uniform distribution), is to prescribe a *prior* distribution over the mixing coefficients. This would get you very close to Bayesian realms. Do you know any distribution which has the space of K-dimensional probability vectors as support?

Semi-supervised learning

Suppose some documents are annotated and others are not (as in the example), and say we model generatively.

For labelled documents, we observe (x, y) whose joint probability is

$$p_{XY}(x, y|\theta) = p_Y(y|\theta)_{X|Y} p(x|y, \theta)$$

and the marginal probability of an unlabelled document x is

$$p_X(x|\theta) = \sum_{y \in \mathcal{Y}} p_Y(y|\theta) p_{X|Y}(x|y, \theta)$$

For a countably finite set $\mathcal{Y}$, this is a **mixture model**!

This is a very special mixture model for its components are *labelled* with self-evident information (e.g., decades).

A generative model of this kind can be thought of as a classifier (the *task* point of view), we need only apply Bayes rule to obtain a conditional $p(y|x_*, \theta)$ that can power a decision rule for a novel document x_*.

But, above all, a generative model of this kind is a model of all of our observations (the *random experiment* point of view). Our observations are indeed a collection of documents, where some documents (very few) are labelled for decade. When we model conditionally we call the labelled instances *training data* and ignore all unlabelled instances (the vast majority of our observations).

Besides powering a classification rule, the generative formulation could be used to shed light onto vocabulary shifts over the decades. One way to specify the component $p(x|y, \theta)$ is to assume it generates a document by drawing words independently given a decade-specific parameter θ_y. That is, $X_i|Y = y \sim \text{Cat}(\theta_y)$ for $i = 1, \ldots, |x|$.

Competition or cooperation?

In a mixture model the components *compete* to generate a data point. This means they cannot *cooperate* to account for some observed variance.

Sometimes, however, we want to stipulate the presence of a number of latent factors that together contribute to our observations distributing the way they do. Think of it in terms of clustering: sometimes we need overlapping clusters, or rather, *attributes*.

For example, our documents are scientific documents, and the period in consideration covers the European Scientific Revolution, as it came to be named. A number of inventions and new ideas marked this period. Documents were likely influenced by subsets of those ideas, rather than any singe idea in particular.

Like in mixture models we can recognise two roles for the class of models we are about to develop.

They can serve task-driven goals and power models that can predict attributes of an input (e.g., attributes of product, aspects of review, morphological features of a word).

They can serve knowledge-seeking goals and power inferences about latent structure that account (cause or correlate with) observed variance (e.g., in what latent aspects/dimensions are data points related).

Latent factor document model

Let us consider a latent factor model for document modelling:

- a document $x = (x_1, \ldots, x_I)$ consists of I i.i.d. categorical draws from that model
- the categorical distribution in turn depends on binary latent factors $z = (z_1, \ldots, z_D)$ which are also i.i.d.

$$\begin{aligned} Z_j &\sim \text{Bernoulli}(\alpha) && (1 \leq d \leq D) \\ X_i|z &\sim \text{Categorical}\left(f(z;\theta)\right) && (1 \leq i \leq I) \end{aligned}$$

Here $f(\cdot;\theta)$ is an NN and θ its parameters.

To keep the model simple we will assume $X_i \perp X_j|Z$ for $i \neq j$. We could, however, relax this conditional independence if we wanted. For example, we could model $X_i|z, x_{<i} \sim \text{Cat}(f(z, x_{<i};\theta))$.

Graphical model

Joint distribution: independent latent variables

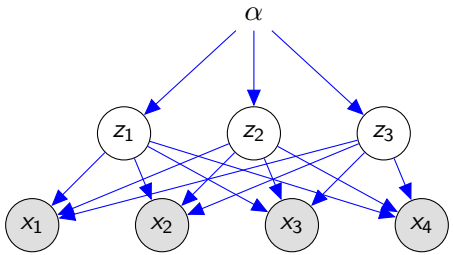

I omit θ from the graphical model, but every $X_i|z$ depends on it.

Suppose, for example, $D = 3$ and $I = 4$.

Posterior

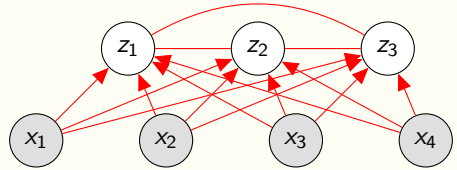

Marginal

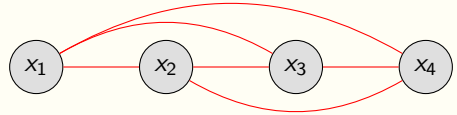

I'm omitting θ and α from the graphical models.

Intractable Marginals

In the latent factor model, marginalisation takes time $\mathcal{O}(2^D)$

$$
\begin{aligned}
p(x|\alpha,\theta) &= \sum_{z\in\{0,1\}^D} p(z|\alpha)p(x|z,\theta) \\
&= \sum_{z\in\{0,1\}^D} \prod_{d=1}^{D} \text{Bern}(z_d|\alpha) \prod_{i=1}^{I} \text{Cat}(x_i|f(z;\theta))
\end{aligned}
$$

As a consequence, we cannot assess $\log p(x|\alpha,\theta)$ nor its gradient.

Thinking ahead: if we cannot assess $\log p(x|\alpha,\theta)$ for an observation x, nor its gradient, how are we going to estimate parameters for this model?

Combinatorial Latent Structure

The posterior of the latent factor model reveals attributes that are relevant to an observation.

Sampling from it can help discover discrete factors of variation (e.g. morphological attributes of a word).

Unfortunately, the posterior in this case is a Gibbs distribution whose parameter is intractable to compute

- its natural parameter has length 2^D
- its log-normaliser requires a summation over $z \in \{0, 1\}^D$

Applications

Alignment Learn to match two data structures (e.g., word alignment, phrase alignment, visual question answering).

Data: $\langle x_1, \ldots, x_I \rangle$ and $\langle y_1, \ldots, y_J \rangle$

Generate each part of y using a subset of the parts of x.

- this can be a mixture model
- or a latent factor model
- and there can be constraints on the parts (e.g., disjoint)

(Rios et al., 2018; Deng et al., 2018; Kawakami et al., 2019)

Applications

Latent attribution What parts of the input (or of a computation graph) affect predictions.

Data: $\langle x_1, \ldots, x_I \rangle$ and y

$$Z_i \sim \text{Bern}(\alpha)$$
$$Y|x, z \sim \text{Cat}(f(x \odot z; \theta))$$

x could also be every hidden state internal to a given NN, and y could be that NN's output $y = g(x; \phi)$

(Lei et al., 2016; Bastings et al., 2019; Cao et al., 2020)

Applications

Compositionality Learn a computation graph.

Sample a structure z from a prior or conditional distribution $Z|x$ and let this structure determine a composition function to parameterise a distribution $Y|x, z$. This can be used for semi-supervised learning of syntactic/semantic representations, for learning to solve arithmetic expressions, interpretable text classifiers, etc.

(Yogatama et al., 2017; Corro and Titov, 2018; Niculae et al., 2018; Havrylov et al., 2019)

Applications

Controllable generation Learn to affect a conditional generator by controlling a latent prompt.

For example, translation models parameterise a conditional distribution $Y|x, \theta$ over translations of a given input x.

The source may contain a certain word (e.g., doctor), and the target language gender-marks nouns. The source sentence *does not* contain enough information to resolve the ambiguity, wouldn't it be nice to have a mechanism, other than requiring the user to produce a less ambiguous x, to control inflections?

(Hu et al., 2017; Zhou and Neubig, 2017; Ataman et al., 2020)

Applications

(Hu et al., 2017; Zhou and Neubig, 2017; Ataman et al., 2020)

Summary

Mixture model ('learning clusters')

Latent factor model ('learning attributes or overlapping clusters')

Applications:

- unsupervised learning (e.g., word alignments, LDA, IBP)
- semi-supervised learning (e.g., generative classifiers, disentanglement learning)
- transparency (e.g., latent rationales)

Examples:

- word alignments (Brown et al., 1993; Vogel et al., 1996; Rios et al., 2018)
- LDA (Blei et al., 2003)
- IBP (Ghahramani and Griffiths, 2006)
- semi-supervised deep generative models (Kingma et al., 2014; **?**)
- latent rationales (Lei et al., 2016; Bastings et al., 2019)

Outline

Latent Variable Models

When talking about some generic model I will follow this convention

- X is an rv taking on values in $\mathcal{X}$
- $x \in \mathcal{X}$ is an observation
- Z is a *discrete* rv taking on values in $\mathcal{Z}$
- $z \in \mathcal{Z}$ is a latent assignment
- the joint pdf factorises as $p(x, z|\theta) = p(z|\theta)p(x|z, \theta)$
- $p(z|\theta)$ is called the *prior*
- $p(x|z, \theta)$ is called the *observational model*
- $p(x|\theta)$ is the *marginal* (or *evidence*)
- and $p(z|x, \theta) = \frac{p(x,z|\theta)}{p(x|\theta)}$ is the posterior
- anything in the model can be parameterised by NNs

Throughout, we shall assume we have N i.i.d. observations. With deterministic parameters θ, we can make all our arguments in terms of a single observation x. Recall, the likelihood-function $\mathcal{L}_{\mathcal{D}}(\theta)$ is just $\sum_{x\sim\mathcal{D}} \log p(x|\theta)$.

By the way, can you draw a plate diagram for our generic latent variable model?

Many models admit exact marginals

Examples (and the algorithms for marginalisation)

- Mixture models (enumeration)
- HMMs (forward algorithm)
- CFGs (inside algorithm)
- Spanning-tree random fields (matrix-tree theorem)

Tractable marginalisation depends on the conditional independence assumptions of a model (e.g., in an HMM a hidden state is independent of all but its preceding state), not on how that model's probability distributions are parameterised (e.g., a transition distribution in the HMM may be stored in a table, predicted by a log-linear model or by an NN).

Marginalisation algorithms are generally harder to parallelise on GPUs.

Recall that to use NNs in probabilistic models we converged to two constraints on the log-likelihood function:

- differentiability with respect to parameters
- and tractability

If $p(z|\theta)$ and $p(x|z,\theta)$ are differentiable functions of their parameters, there is no impediment to gradient-based parameter estimation. Can you show that to yourself? Hint: expand $\nabla_\theta \log p(x|\theta)$.

Tractability depends on whether $p(x|\theta) = \sum_{z \in \mathcal{Z}} p(x, z|\theta)$, or its logarithm, can be evaluated in feasible time. Though it may seem so, this is not always a matter of cardinality of $\mathcal{Z}$.

For example, there is a Catalan number of trees in a CFG, yet because of the strong independence assumptions in the model, the marginal $p(x|\theta)$ is computable in cubic-time (w.r.t. sequence length) via the inside algorithm. Similarly, there is an exponential number of state sequences in an HMM, but its marginal is computable in linear-time (w.r.t. sequence length) via the forward algorithm.

Neural {MM, HMM, CFG, CRF, ... }

An NN-parameterisation of a classic discrete LVM, for which exact marginals are tractable, still needs to preserve all of that model's statistical assumptions about unobserved random variables.

We won't necessarily achieve a more complex distribution.

Though we may condition on complex data more effectively.

A neural HMM could look like:

$$p(x|\theta) = \sum_{z \in \{1,\dots,K\}^{|x|}} \prod_{i=1}^{|x|} \underbrace{p(z_i|z_{i-1}, x_{<i}, \theta)}_{\text{Cat}(z_i|g(x_{<i}, z_{i-1}; \theta))} \underbrace{p(x_i|z_i, x_{<i})}_{\text{Cat}(x_i|f(x_{<i}, z_i; \theta))}$$

the entire history of already generated words is available for conditioning

NNs allow us to condition on complex observations, like a long history of words $x_{<i}$ in unsupervised part-of-speech tagging.

We cannot, as easily, exploit that power to relax statistical conditional independence assumptions, for those assumptions are crucial in order to maintain *exact and tractable* access to marginal probabilities.

Think of it this way, what makes the HMM the HMM is the first-order (or n-order) Markov assumption $Z_i \perp Z_j | Z_{i-1}$ for j other than i and $i-1$. Relaxing that turns the HMM into something else, for which exact inference is likely impossible. See Wang et al. (2018) for a neural HMM.

Gradient-based MLE

What happens when we autodiff the quantity $\log p(x|\theta)$, which we computed exactly and tractably?

Let's inspect this gradient ourselves $\nabla_\theta \log p(x|\theta)$

Gradient of log-marginal

- It all starts with the derivative of log, followed by chain rule again.

Gradient-based MLE

What happens when we autodiff the quantity $\log p(x|\theta)$, which we computed exactly and tractably?

Let's inspect this gradient ourselves $\nabla_\theta \log p(x|\theta)$

$$= \frac{1}{p(x|\theta)} \nabla_\theta p(x|\theta) =$$

Gradient of log-marginal

- It all starts with the derivative of log, followed by chain rule again.
- The next step requires marginalisation.

Gradient-based MLE

What happens when we autodiff the quantity $\log p(x|\theta)$, which we computed exactly and tractably?

Let's inspect this gradient ourselves $\nabla_\theta \log p(x|\theta)$

$$= \frac{1}{p(x|\theta)} \nabla_\theta p(x|\theta) = \frac{1}{p(x|\theta)} \nabla_\theta \sum_{z \in \mathcal{Z}} p(z, x|\theta)$$

Gradient of log-marginal

- It all starts with the derivative of log, followed by chain rule again.
- The next step requires marginalisation.
- Now we need the gradient of a big sum.

Gradient-based MLE

What happens when we autodiff the quantity $\log p(x|\theta)$, which we computed exactly and tractably?

Let's inspect this gradient ourselves $\nabla_\theta \log p(x|\theta)$

$$= \frac{1}{p(x|\theta)} \nabla_\theta p(x|\theta) = \frac{1}{p(x|\theta)} \nabla_\theta \sum_{z \in \mathcal{Z}} p(z, x|\theta)$$

$$= \frac{1}{p(x|\theta)} \sum_{z \in \mathcal{Z}} \nabla_\theta p(z, x|\theta) =$$

Gradient of log-marginal

- It all starts with the derivative of log, followed by chain rule again.
- The next step requires marginalisation.
- Now we need the gradient of a big sum.
- Derivatives are linear, so we can sum gradients instead.

Gradient-based MLE

What happens when we autodiff the quantity $\log p(x|\theta)$, which we computed exactly and tractably?

Let's inspect this gradient ourselves $\nabla_\theta \log p(x|\theta)$

$$= \frac{1}{p(x|\theta)} \nabla_\theta p(x|\theta) = \frac{1}{p(x|\theta)} \nabla_\theta \sum_{z \in \mathcal{Z}} p(z, x|\theta)$$

$$= \frac{1}{p(x|\theta)} \sum_{z \in \mathcal{Z}} \nabla_\theta p(z, x|\theta) = \frac{1}{p(x|\theta)} \sum_{z \in \mathcal{Z}} p(z, x|\theta) \nabla_\theta \log p(z, x|\theta)$$

Gradient of log-marginal

- It all starts with the derivative of log, followed by chain rule again.
- The next step requires marginalisation.
- Now we need the gradient of a big sum.
- Derivatives are linear, so we can sum gradients instead.
- Sums are fine, but let's use the log identity $f' = f(\log f)'$

Gradient-based MLE

What happens when we autodiff the quantity $\log p(x|\theta)$, which we computed exactly and tractably?

Let's inspect this gradient ourselves $\nabla_\theta \log p(x|\theta)$

$$
\begin{aligned}
&= \frac{1}{p(x|\theta)} \nabla_\theta p(x|\theta) = \frac{1}{p(x|\theta)} \nabla_\theta \sum_{z \in \mathcal{Z}} p(z, x|\theta) \\
&= \frac{1}{p(x|\theta)} \sum_{z \in \mathcal{Z}} \nabla_\theta p(z, x|\theta) = \frac{1}{p(x|\theta)} \sum_{z \in \mathcal{Z}} p(z, x|\theta) \nabla_\theta \log p(z, x|\theta) \\
&= \sum_{z \in \mathcal{Z}} \frac{p(z, x|\theta)}{p(x|\theta)} \nabla_\theta \log p(z, x|\theta) =
\end{aligned}
$$

Gradient of log-marginal

- It all starts with the derivative of log, followed by chain rule again.
- The next step requires marginalisation.
- Now we need the gradient of a big sum.
- Derivatives are linear, so we can sum gradients instead.
- Sums are fine, but let's use the log identity $f' = f(\log f)'$
- The marginal is constant for $z \in \mathcal{Z}$, so distribute it over the sum.

Gradient-based MLE

What happens when we autodiff the quantity $\log p(x|\theta)$, which we computed exactly and tractably?

Let's inspect this gradient ourselves $\nabla_\theta \log p(x|\theta)$

$$
\begin{aligned}
&= \frac{1}{p(x|\theta)} \nabla_\theta p(x|\theta) = \frac{1}{p(x|\theta)} \nabla_\theta \sum_{z \in \mathcal{Z}} p(z, x|\theta) \\
&= \frac{1}{p(x|\theta)} \sum_{z \in \mathcal{Z}} \nabla_\theta p(z, x|\theta) = \frac{1}{p(x|\theta)} \sum_{z \in \mathcal{Z}} p(z, x|\theta) \nabla_\theta \log p(z, x|\theta) \\
&= \sum_{z \in \mathcal{Z}} \frac{p(z, x|\theta)}{p(x|\theta)} \nabla_\theta \log p(z, x|\theta) = \sum_{z \in \mathcal{Z}} p(z|x, \theta) \nabla_\theta \log p(z, x|\theta)
\end{aligned}
$$

Gradient of log-marginal

- It all starts with the derivative of log, followed by chain rule again.
- The next step requires marginalisation.
- Now we need the gradient of a big sum.
- Derivatives are linear, so we can sum gradients instead.
- Sums are fine, but let's use the log identity $f' = f(\log f)'$
- The marginal is constant for $z \in \mathcal{Z}$, so distribute it over the sum.
- This gives us a recognisable object! Joint probability, divided by evidence, that's the posterior! And we have a weighted average, coefficients given by a pmf, and we sum over the entire support $\mathcal{Z}$.

Gradient-based MLE

What happens when we autodiff the quantity $\log p(x|\theta)$, which we computed exactly and tractably?

Let's inspect this gradient ourselves $\nabla_\theta \log p(x|\theta)$

$$
\begin{aligned}
&= \frac{1}{p(x|\theta)} \nabla_\theta p(x|\theta) = \frac{1}{p(x|\theta)} \nabla_\theta \sum_{z \in \mathcal{Z}} p(z, x|\theta) \\
&= \frac{1}{p(x|\theta)} \sum_{z \in \mathcal{Z}} \nabla_\theta p(z, x|\theta) = \frac{1}{p(x|\theta)} \sum_{z \in \mathcal{Z}} p(z, x|\theta) \nabla_\theta \log p(z, x|\theta) \\
&= \sum_{z \in \mathcal{Z}} \frac{p(z, x|\theta)}{p(x|\theta)} \nabla_\theta \log p(z, x|\theta) = \sum_{z \in \mathcal{Z}} p(z|x, \theta) \nabla_\theta \log p(z, x|\theta) \\
&= \mathbb{E}_{p(z|x,\theta)} \left[\nabla_\theta \log p(z, x|\theta) \right]
\end{aligned}
$$

Gradient of log-marginal

- It all starts with the derivative of log, followed by chain rule again.
- The next step requires marginalisation.
- Now we need the gradient of a big sum.
- Derivatives are linear, so we can sum gradients instead.
- Sums are fine, but let's use the log identity $f' = f(\log f)'$
- The marginal is constant for $z \in \mathcal{Z}$, so distribute it over the sum.
- This gives us a recognisable object! Joint probability, divided by evidence, that's the posterior! And we have a weighted average, coefficients given by a pmf, and we sum over the entire support $\mathcal{Z}$.
- We have an expectation! The gradient of the log-marginal of x is the expected gradient of the log joint probability of x and z, where x is observed and z is a draw from the posterior distribution $Z|x, \theta$. Dependency on Z makes the gradient of log-joint $G(Z) = \nabla_\theta \log P(Z, X = x)$ a random variable.

Gradient-based MLE

What happens when we autodiff the quantity $\log p(x|\theta)$, which we computed exactly and tractably?

Let's inspect this gradient ourselves $\nabla_\theta \log p(x|\theta)$

$$
\begin{aligned}
&= \frac{1}{p(x|\theta)} \nabla_\theta p(x|\theta) = \frac{1}{p(x|\theta)} \nabla_\theta \sum_{z \in \mathcal{Z}} p(z, x|\theta) \\
&= \frac{1}{p(x|\theta)} \sum_{z \in \mathcal{Z}} \nabla_\theta p(z, x|\theta) = \frac{1}{p(x|\theta)} \sum_{z \in \mathcal{Z}} p(z, x|\theta) \nabla_\theta \log p(z, x|\theta) \\
&= \sum_{z \in \mathcal{Z}} \frac{p(z, x|\theta)}{p(x|\theta)} \nabla_\theta \log p(z, x|\theta) = \sum_{z \in \mathcal{Z}} p(z|x, \theta) \nabla_\theta \log p(z, x|\theta) \\
&= \mathbb{E}_{p(z|x,\theta)} \left[\nabla_\theta \log p(z, x|\theta) \right]
\end{aligned}
$$

Autodiff performs **exact posterior inference** for us!

Gradient of log-marginal

- It all starts with the derivative of log, followed by chain rule again.
- The next step requires marginalisation.
- Now we need the gradient of a big sum.
- Derivatives are linear, so we can sum gradients instead.
- Sums are fine, but let's use the log identity $f' = f(\log f)'$
- The marginal is constant for $z \in \mathcal{Z}$, so distribute it over the sum.
- This gives us a recognisable object! Joint probability, divided by evidence, that's the posterior! And we have a weighted average, coefficients given by a pmf, and we sum over the entire support $\mathcal{Z}$.
- We have an expectation! The gradient of the log-marginal of x is the expected gradient of the log joint probability of x and z, where x is observed and z is a draw from the posterior distribution $Z|x, \theta$. Dependency on Z makes the gradient of log-joint $G(Z) = \nabla_\theta \log P(Z, X = x)$ a random variable.
- The gradient of the log-marginal $\nabla_\theta \log p(x|\theta)$ is deterministic, it is the expected value $\mathbb{E}_{Z|X=x,\theta}[G(Z)]$. You evaluate the marginal, autodiff evaluates the expectation.

Gradient-based MLE

What happens when we autodiff the quantity $\log p(x|\theta)$, which we computed exactly and tractably?

Let's inspect this gradient ourselves $\nabla_\theta \log p(x|\theta)$

$$
\begin{aligned}
&= \frac{1}{p(x|\theta)} \nabla_\theta p(x|\theta) = \frac{1}{p(x|\theta)} \nabla_\theta \sum_{z \in \mathcal{Z}} p(z, x|\theta) \\
&= \frac{1}{p(x|\theta)} \sum_{z \in \mathcal{Z}} \nabla_\theta p(z, x|\theta) = \frac{1}{p(x|\theta)} \sum_{z \in \mathcal{Z}} p(z, x|\theta) \nabla_\theta \log p(z, x|\theta) \\
&= \sum_{z \in \mathcal{Z}} \frac{p(z, x|\theta)}{p(x|\theta)} \nabla_\theta \log p(z, x|\theta) = \sum_{z \in \mathcal{Z}} p(z|x, \theta) \nabla_\theta \log p(z, x|\theta) \\
&= \mathbb{E}_{p(z|x,\theta)} \left[\nabla_\theta \log p(z, x|\theta) \right]
\end{aligned}
$$

Autodiff performs **exact posterior inference** for us!

Gradient of log-marginal

- It all starts with the derivative of log, followed by chain rule again.
- The next step requires marginalisation.
- Now we need the gradient of a big sum.
- Derivatives are linear, so we can sum gradients instead.
- Sums are fine, but let's use the log identity $f' = f(\log f)'$
- The marginal is constant for $z \in \mathcal{Z}$, so distribute it over the sum.
- This gives us a recognisable object! Joint probability, divided by evidence, that's the posterior! And we have a weighted average, coefficients given by a pmf, and we sum over the entire support $\mathcal{Z}$.
- We have an expectation! The gradient of the log-marginal of x is the expected gradient of the log joint probability of x and z, where x is observed and z is a draw from the posterior distribution $Z|x, \theta$. Dependency on Z makes the gradient of log-joint $G(Z) = \nabla_\theta \log P(Z, X = x)$ a random variable.
- The gradient of the log-marginal $\nabla_\theta \log p(x|\theta)$ is deterministic, it is the expected value $\mathbb{E}_{Z|X=x,\theta}[G(Z)]$. You evaluate the marginal, autodiff evaluates the expectation.

Summary

Many discrete LVMs admit tractable marginalisation

Assessing the gradient of the log-marginal probability of an observation corresponds to assessing an expectation under the posterior distribution over latent variables. Think of it this way:

- we need posterior inference to compute the gradient
- and we need the gradient for parameter estimation
- with exact marginals, autodiff assesses the gradient thus abstracting posterior inference away

What happens when we cannot solve $\sum_{z \in \mathcal{Z}} p(x, z|\theta)$?

Many interesting models are such that the exact marginal is intractable. We've seen, for example, the case where $p(x, z|\theta)$ is a latent factor model.

Autodiff cannot differentiate a quantity that cannot be assessed. So if we cannot compute the exact log-marginal probability of an observation, we won't get automatic posterior inference for free. We will have to resort to rather explicit approaches to approximate inference.

Outline

Latent factor document model

Let us consider a latent factor model for document modelling:

- a document $x = (x_1, \ldots, x_n)$ consists of n i.i.d. categorical draws from that model
- the categorical distribution in turn depends on binary latent factors $z = (z_1, \ldots, z_D)$ which are also i.i.d.

$$\begin{aligned} Z_j &\sim \text{Bernoulli}(\alpha) && (1 \leq d \leq D) \\ X_i|z &\sim \text{Categorical}\left(f(z;\theta)\right) && (1 \leq i \leq n) \end{aligned}$$

Here $0 < \alpha < 1$ specifies a Bernoulli prior
and $f(\cdot;\theta)$ is a function computed by an NN, e.g.:

$$\begin{aligned} f(z;\theta) &= \text{softmax}(Wz + b) \\ \theta &= \{W, b\} \end{aligned}$$

To keep the model simple we will assume $X_i \perp X_j|Z$ for $i \neq j$. We could, however, relax this conditional independence if we wanted. For example, we could model $X_i|z, x_{<i} \sim \text{Cat}(f(z, x_{<i};\theta))$.

We could model *any* multivariate and/or structured data: e.g., an image x as a collection of pixel intensities along different colours, a molecule as a graph, a gene as a sequence of symbols, etc.

Graphical model

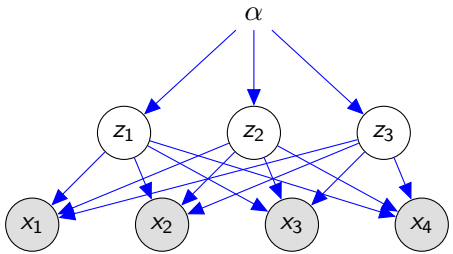

Joint distribution: independent latent variables

Suppose, for example, $D = 3$ and $n = 4$.

I omit θ from the graphical model, but every $X_i|z$ depends on it.

Intractable Marginals

In the latent factor model, marginalisation takes time $\mathcal{O}(2^D)$

$$p(x|\alpha,\theta) = \sum_{z\in\{0,1\}^D} p(z|\alpha)p(x|z,\theta)$$
$$= \sum_{z\in\{0,1\}^D} \prod_{d=1}^{D} \text{Bern}(z_d|\alpha) \prod_{i=1}^{n} \text{Cat}(x_i|f(z;\theta))$$

As a consequence, we cannot assess $\log p(x|\alpha,\theta)$ nor its gradient.

But we know that

$$\nabla_\theta \log p(x|\alpha,\theta) = \mathbb{E}_{p(z|x,\alpha,\theta)}[\nabla_\theta \log p(x,z|\alpha,\theta)]$$

Unfortunately, we cannot count on autodiff to solve the expectation for us in this case.

Perhaps we can estimate the gradient?

Gradient estimates?

Monte Carlo to the rescue?

$$\nabla_\theta \log p(x|\theta) = \mathbb{E}_{p(z|x,\theta)} \left[\nabla_\theta \log p(x, z|\theta)\right]$$

$$\overset{\text{MC}}{\approx} \frac{1}{K} \sum_{k=1}^{K} \nabla_\theta \log p(x, z^{(k)}|\theta) \quad \text{where } z^{(k)} \sim p(z|x, \theta)$$

I am omitting α, the prior parameter. For simplicity, assume it is fixed.

Gradient estimates?

Monte Carlo to the rescue?

$$\nabla_\theta \log p(x|\theta) = \mathbb{E}_{p(z|x,\theta)}\left[\nabla_\theta \log p(x,z|\theta)\right]$$
$$\overset{\text{MC}}{\approx} \frac{1}{K}\sum_{k=1}^{K} \nabla_\theta \log p(x,z^{(k)}|\theta) \quad \text{where } z^{(k)} \sim p(z|x,\theta)$$

Hold on, the posterior is not available either!

$$p(z|x,\theta) = \frac{p(x,z|\theta)}{p(x|\theta)}$$

I am omitting α, the prior parameter. For simplicity, assume it is fixed.

The Basic Problem

We want to compute the posterior over latent variables $p(z|x, \theta)$. This involves computing the marginal likelihood

$$p(x|\theta) = \sum_{z \in \mathcal{Z}} p(x, z|\theta)$$

which is generally **intractable**. This problem motivates the use of **approximate inference** techniques.

You may think all we really need is the marginal and we shouldn't bother about the posterior. I'd argue there's barely a difference. Probabilistic inference (i.e., computations involving probability calculus) is the core of the problem. Sometimes we see it as the need for a marginal, sometimes as the need for a posterior.

Strategy

Variational Inference

- Accept that $p(z|x,\theta)$ is not computable.
- Approximate it by an auxiliary distribution $q(z)$ that is computable!
- Choose $q(z)$ as close as possible to $p(z|x,\theta)$ to obtain a faithful approximation.

We are going to derive VI's objective from two points of view

- first, we will concentrate on the intractable log-marginal, and attempt to bound it by a tractable quantity;
- then we will concentrate on the intractable posterior, and attempt to learn a tractable approximation to it;

The two views will turn out intimately related.

This is the outline for **variational inference** (VI; Jordan et al., 1999; Blei et al., 2017).

There are alternatives to VI, but they are not covered in this course. Here are some pointers:

- Markov chain Monte Carlo (MCMC). Here is an excellent material by Michael Betancourt: `https://betanalpha.github.io/assets/case_studies/markov_chain_monte_carlo.html`.
- Expectation propagation (EP; Minka, 2001; Vehtari et al., 2020)

Outline

VI derivation I

Deriving VI from the log-evidence:

$$\log p(x|\theta) = \log \sum_{z \in \mathcal{Z}} p(x, z|\theta)$$

The evidence lower-bound (ELBO):

- Let's start from the log marginal.

VI derivation I

Deriving VI from the log-evidence:

$$\log p(x|\theta) = \log \sum_{z \in \mathcal{Z}} p(x, z|\theta)$$

$$= \log \sum_{z \in \mathcal{Z}} q(z) \frac{p(x, z|\theta)}{q(z)}$$

The evidence lower-bound (ELBO):

- Let's start from the log marginal.
- And introduce $q(z)$ such that $q(z) > 0$ if $p(z|x, \theta) > 0$.

VI derivation I

Deriving VI from the log-evidence:

$$\begin{aligned}\log p(x|\theta) &= \log \sum_{z \in \mathcal{Z}} p(x, z|\theta) \\ &= \log \sum_{z \in \mathcal{Z}} q(z) \frac{p(x, z|\theta)}{q(z)} \\ &= \log \mathbb{E}_{q(z)} \left[\frac{p(x, z|\theta)}{q(z)} \right]\end{aligned}$$

The evidence lower-bound (ELBO):

- Let's start from the log marginal.
- And introduce $q(z)$ such that $q(z) > 0$ if $p(z|x, \theta) > 0$.
- And note we got an expectation w.r.t. $q(z)$, and recall that, unlike $p(z|x, \theta)$, we know $q(z)$, as we chose it.

VI derivation I

Deriving VI from the log-evidence:

$$\begin{aligned}
\log p(x|\theta) &= \log \sum_{z \in \mathcal{Z}} p(x, z|\theta) \\
&= \log \sum_{z \in \mathcal{Z}} q(z) \frac{p(x, z|\theta)}{q(z)} \\
&= \log \mathbb{E}_{q(z)} \left[\frac{p(x, z|\theta)}{q(z)} \right] \\
&\overset{\text{JI}}{\geq} \mathbb{E}_{q(z)} \left[\log \frac{p(x, z|\theta)}{q(z)} \right]
\end{aligned}$$

The evidence lower-bound (ELBO):

- Let's start from the log marginal.
- And introduce $q(z)$ such that $q(z) > 0$ if $p(z|x, \theta) > 0$.
- And note we got an expectation w.r.t. $q(z)$, and recall that, unlike $p(z|x, \theta)$, we know $q(z)$, as we chose it.
- And Jensen's inequality allows us to push the log through the expectation.

VI derivation I

Deriving VI from the log-evidence:

$$\begin{aligned}\log p(x|\theta) &= \log \sum_{z\in\mathcal{Z}} p(x,z|\theta) \\ &= \log \sum_{z\in\mathcal{Z}} q(z)\frac{p(x,z|\theta)}{q(z)} \\ &= \log \mathbb{E}_{q(z)}\left[\frac{p(x,z|\theta)}{q(z)}\right] \\ &\overset{\text{JI}}{\geq} \mathbb{E}_{q(z)}\left[\log \frac{p(x,z|\theta)}{q(z)}\right]\end{aligned}$$

This is the *lowerbound on the log-evidence*, also known as ELBO. Crucially, it **does not require the true posterior**!

The evidence lower-bound (ELBO):

- Let's start from the log marginal.
- And introduce $q(z)$ such that $q(z) > 0$ if $p(z|x,\theta) > 0$.
- And note we got an expectation w.r.t. $q(z)$, and recall that, unlike $p(z|x,\theta)$, we know $q(z)$, as we chose it.
- And Jensen's inequality allows us to push the log through the expectation.

VI derivation I

Let's gain insight about this bound

$$\log p(x|\theta) \geq \mathbb{E}_{q(z)}\left[\log \frac{p(x,z|\theta)}{q(z)}\right]$$

What can be said about $q(z)$?

- Let's start from the ELBO.

VI derivation I

Let's gain insight about this bound

$$\log p(x|\theta) \geq \mathbb{E}_{q(z)}\left[\log \frac{p(x,z|\theta)}{q(z)}\right]$$
$$= \mathbb{E}_{q(z)}\left[\log \frac{p(z|x,\theta)p(x|\theta)}{q(z)}\right]$$

What can be said about $q(z)$?

- Let's start from the ELBO.
- Now let's factorise the joint probability using the marginal and the posterior (these are clearly not available to us, but they will help us understand what is going on).

VI derivation I

Let's gain insight about this bound

$$\begin{aligned}\log p(x|\theta) &\geq \mathbb{E}_{q(z)}\left[\log \frac{p(x,z|\theta)}{q(z)}\right] \\ &= \mathbb{E}_{q(z)}\left[\log \frac{p(z|x,\theta)p(x|\theta)}{q(z)}\right] \\ &= \log p(x|\theta) + \sum_{z\in\mathcal{Z}} q(z) \log \frac{p(z|x,\theta)}{q(z)}\end{aligned}$$

What can be said about $q(z)$?

- Let's start from the ELBO.
- Now let's factorise the joint probability using the marginal and the posterior (these are clearly not available to us, but they will help us understand what is going on).
- The log-marginal is constant w.r.t. z, thus its expected valued under $q(z)$ is itself, i.e., $\log p(x|\theta)$.

VI derivation I

Let's gain insight about this bound

$$\begin{aligned}
\log p(x|\theta) &\geq \mathbb{E}_{q(z)}\left[\log \frac{p(x,z|\theta)}{q(z)}\right] \\
&= \mathbb{E}_{q(z)}\left[\log \frac{p(z|x,\theta)p(x|\theta)}{q(z)}\right] \\
&= \log p(x|\theta) + \sum_{z\in\mathcal{Z}} q(z) \log \frac{p(z|x,\theta)}{q(z)} \\
&= \log p(x|\theta) - \sum_{z\in\mathcal{Z}} q(z) \log \frac{q(z)}{p(z|x,\theta)}
\end{aligned}$$

What can be said about $q(z)$?

- Let's start from the ELBO.
- Now let's factorise the joint probability using the marginal and the posterior (these are clearly not available to us, but they will help us understand what is going on).
- The log-marginal is constant w.r.t. z, thus its expected valued under $q(z)$ is itself, i.e., $\log p(x|\theta)$.
- We can apply a property of logs to rearrange the fraction.

VI derivation I

Let's gain insight about this bound

$$
\begin{aligned}
\log p(x|\theta) &\geq \mathbb{E}_{q(z)}\left[\log \frac{p(x,z|\theta)}{q(z)}\right] \\
&= \mathbb{E}_{q(z)}\left[\log \frac{p(z|x,\theta)p(x|\theta)}{q(z)}\right] \\
&= \log p(x|\theta) + \sum_{z\in\mathcal{Z}} q(z) \log \frac{p(z|x,\theta)}{q(z)} \\
&= \log p(x|\theta) - \sum_{z\in\mathcal{Z}} q(z) \log \frac{q(z)}{p(z|x,\theta)} \\
&= \log p(x|\theta) - \underbrace{\mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right)}_{\geq 0}
\end{aligned}
$$

What can be said about $q(z)$?

- Let's start from the ELBO.
- Now let's factorise the joint probability using the marginal and the posterior (these are clearly not available to us, but they will help us understand what is going on).
- The log-marginal is constant w.r.t. z, thus its expected valued under $q(z)$ is itself, i.e., $\log p(x|\theta)$.
- We can apply a property of logs to rearrange the fraction.
- Which gives us the KL divergence from $p(z|x,\theta)$ to $q(z)$. Recall, $\mathrm{KL}\left(q \,||\, p\right) \geq 0$ and equality holds only if $q = p$.

VI derivation I

Let's gain insight about this bound

$$\begin{aligned}
\log p(x|\theta) &\geq \mathbb{E}_{q(z)}\left[\log \frac{p(x,z|\theta)}{q(z)}\right] \\
&= \mathbb{E}_{q(z)}\left[\log \frac{p(z|x,\theta)p(x|\theta)}{q(z)}\right] \\
&= \log p(x|\theta) + \sum_{z\in\mathcal{Z}} q(z) \log \frac{p(z|x,\theta)}{q(z)} \\
&= \log p(x|\theta) - \sum_{z\in\mathcal{Z}} q(z) \log \frac{q(z)}{p(z|x,\theta)} \\
&= \log p(x|\theta) - \underbrace{\mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right)}_{\geq 0}
\end{aligned}$$

We have derived a lower bound on the log-evidence whose gap is exactly $\mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right)$.

What can be said about $q(z)$?

- Let's start from the ELBO.
- Now let's factorise the joint probability using the marginal and the posterior (these are clearly not available to us, but they will help us understand what is going on).
- The log-marginal is constant w.r.t. z, thus its expected valued under $q(z)$ is itself, i.e., $\log p(x|\theta)$.
- We can apply a property of logs to rearrange the fraction.
- Which gives us the KL divergence from $p(z|x,\theta)$ to $q(z)$. Recall, $\mathrm{KL}\left(q \,||\, p\right) \geq 0$ and equality holds only if $q = p$.

It looks like $q(z)$ should be as close as possible to $p(z|x,\theta)$!

VI derivation I

Let's gain insight about this bound

$$\begin{aligned}
\log p(x|\theta) &\geq \mathbb{E}_{q(z)}\left[\log \frac{p(x,z|\theta)}{q(z)}\right] \\
&= \mathbb{E}_{q(z)}\left[\log \frac{p(z|x,\theta)p(x|\theta)}{q(z)}\right] \\
&= \log p(x|\theta) + \sum_{z\in\mathcal{Z}} q(z) \log \frac{p(z|x,\theta)}{q(z)} \\
&= \log p(x|\theta) - \sum_{z\in\mathcal{Z}} q(z) \log \frac{q(z)}{p(z|x,\theta)} \\
&= \log p(x|\theta) - \underbrace{\mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right)}_{\geq 0}
\end{aligned}$$

We have derived a lower bound on the log-evidence whose gap is exactly $\mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right)$.

What can be said about $q(z)$?

- Let's start from the ELBO.
- Now let's factorise the joint probability using the marginal and the posterior (these are clearly not available to us, but they will help us understand what is going on).
- The log-marginal is constant w.r.t. z, thus its expected valued under $q(z)$ is itself, i.e., $\log p(x|\theta)$.
- We can apply a property of logs to rearrange the fraction.
- Which gives us the KL divergence from $p(z|x,\theta)$ to $q(z)$. Recall, $\mathrm{KL}\left(q \,||\, p\right) \geq 0$ and equality holds only if $q = p$.

It looks like $q(z)$ should be as close as possible to $p(z|x,\theta)$!

Outline

VI derivation II

Derive VI by optimisation:

$$\arg\max_{q(z)} - \mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right)$$

The previous derivation suggests that we should attempt to choose $q(z)$ such that the gap relative to the true posterior $\mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right)$ is as small as possible.

- Let's state that objective explicitly and seek some optimum $q(z)$.

VI derivation II

Derive VI by optimisation:

$$\arg\max_{q(z)} - \mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right)$$

$$= \arg\max_{q(z)} \sum_{z \in \mathcal{Z}} q(z) \log \frac{p(z|x,\theta)}{q(z)}$$

The previous derivation suggests that we should attempt to choose $q(z)$ such that the gap relative to the true posterior $\mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right)$ is as small as possible.

- Let's state that objective explicitly and seek some optimum $q(z)$.
- We do not have access to the true posterior probability of any z, thus let's decompose it via Bayes rule.

VI derivation II

Derive VI by optimisation:

$$\begin{aligned}
&\arg\max_{q(z)} - \mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right) \\
&= \arg\max_{q(z)} \sum_{z\in\mathcal{Z}} q(z) \log \frac{p(z|x,\theta)}{q(z)} \\
&= \arg\max_{q(z)} \sum_{z\in\mathcal{Z}} q(z) \log \frac{p(x,z|\theta)}{q(z)p(x|\theta)}
\end{aligned}$$

The previous derivation suggests that we should attempt to choose $q(z)$ such that the gap relative to the true posterior $\mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right)$ is as small as possible.

- Let's state that objective explicitly and seek some optimum $q(z)$.
- We do not have access to the true posterior probability of any z, thus let's decompose it via Bayes rule.
- Bayes rule reveals the marginal. Note that $p(x|\theta)$ does not depend on z.

VI derivation II

Derive VI by optimisation:

$$\begin{aligned}
&\arg\max_{q(z)} - \mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right) \\
&= \arg\max_{q(z)} \sum_{z\in\mathcal{Z}} q(z) \log \frac{p(z|x,\theta)}{q(z)} \\
&= \arg\max_{q(z)} \sum_{z\in\mathcal{Z}} q(z) \log \frac{p(x,z|\theta)}{q(z)p(x|\theta)} \\
&= \arg\max_{q(z)} \sum_{z\in\mathcal{Z}} q(z) \log \frac{p(x,z|\theta)}{q(z)} - \overbrace{\log p(x|\theta)}^{\textit{constant}}
\end{aligned}$$

The previous derivation suggests that we should attempt to choose $q(z)$ such that the gap relative to the true posterior $\mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right)$ is as small as possible.

- Let's state that objective explicitly and seek some optimum $q(z)$.
- We do not have access to the true posterior probability of any z, thus let's decompose it via Bayes rule.
- Bayes rule reveals the marginal. Note that $p(x|\theta)$ does not depend on z.
- Nor it depends on our choice of $q(z)$.

VI derivation II

Derive VI by optimisation:

$$
\begin{aligned}
&\arg\max_{q(z)} - \mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right) \\
&= \arg\max_{q(z)} \sum_{z\in\mathcal{Z}} q(z) \log \frac{p(z|x,\theta)}{q(z)} \\
&= \arg\max_{q(z)} \sum_{z\in\mathcal{Z}} q(z) \log \frac{p(x,z|\theta)}{q(z)p(x|\theta)} \\
&= \arg\max_{q(z)} \sum_{z\in\mathcal{Z}} q(z) \log \frac{p(x,z|\theta)}{q(z)} - \overbrace{\log p(x|\theta)}^{\text{constant}} \\
&= \arg\max_{q(z)} \underbrace{\sum_{z\in\mathcal{Z}} q(z) \log p(x,z|\theta)}_{\mathbb{E}_{q(z)}[\log p(x,z|\theta)]} - \underbrace{\sum_{z\in\mathcal{Z}} q(z) \log q(z)}_{-\mathbb{H}(q(z))}
\end{aligned}
$$

The previous derivation suggests that we should attempt to choose $q(z)$ such that the gap relative to the true posterior $\mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right)$ is as small as possible.

- Let's state that objective explicitly and seek some optimum $q(z)$.
- We do not have access to the true posterior probability of any z, thus let's decompose it via Bayes rule.
- Bayes rule reveals the marginal. Note that $p(x|\theta)$ does not depend on z.
- Nor it depends on our choice of $q(z)$.
- The final objective involves only joint densities (always tractable, by assumption), and $q(z)$.

VI derivation II

Derive VI by optimisation:

$$
\begin{aligned}
&\arg\max_{q(z)} - \mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right) \\
&= \arg\max_{q(z)} \sum_{z\in\mathcal{Z}} q(z) \log \frac{p(z|x,\theta)}{q(z)} \\
&= \arg\max_{q(z)} \sum_{z\in\mathcal{Z}} q(z) \log \frac{p(x,z|\theta)}{q(z)p(x|\theta)} \\
&= \arg\max_{q(z)} \sum_{z\in\mathcal{Z}} q(z) \log \frac{p(x,z|\theta)}{q(z)} - \overbrace{\log p(x|\theta)}^{\text{constant}} \\
&= \arg\max_{q(z)} \underbrace{\sum_{z\in\mathcal{Z}} q(z) \log p(x,z|\theta)}_{\mathbb{E}_{q(z)}[\log p(x,z|\theta)]} - \underbrace{\sum_{z\in\mathcal{Z}} q(z) \log q(z)}_{-\mathbb{H}(q(z))}
\end{aligned}
$$

The previous derivation suggests that we should attempt to choose $q(z)$ such that the gap relative to the true posterior $\mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right)$ is as small as possible.

- Let's state that objective explicitly and seek some optimum $q(z)$.
- We do not have access to the true posterior probability of any z, thus let's decompose it via Bayes rule.
- Bayes rule reveals the marginal. Note that $p(x|\theta)$ does not depend on z.
- Nor it depends on our choice of $q(z)$.
- The final objective involves only joint densities (always tractable, by assumption), and $q(z)$.

ELBO

The evidence lowerbound (ELBO) is the optimisation objective in **variational inference**.

$$\begin{aligned}&\arg\max_{q(z)} \mathbb{E}_{x\sim\mathcal{D}}\left[\mathbb{E}_{q(z)}\left[\log p(x,z|\theta)\right]+\mathbb{H}\left(q(z)\right)\right]\\ =&\arg\max_{q(z)} \mathbb{E}_{x\sim\mathcal{D}}\left[\mathbb{E}_{q(z)}\left[\log p(x|z,\theta)\right]-\mathrm{KL}\left(q(z)\,||\,p(z)\right)\right]\end{aligned}$$

The ELBO circumvents intractable posterior inference by optimisation: we search the approximate posterior that is closest to the true posterior in terms of $\mathrm{KL}\left(q(z)\,||\,p(z|x,\theta)\right)$. For example, if $q(z|\lambda)$ is in a certain parametric family, we search for its parameter.

ELBO highlights

- we get to design $q(z)$, so for example, while the true posterior of a latent factor model depends on an intractable marginalisation, the approximate posterior $q(z)$ might simply combine D independent Bernoulli distributions (one per latent factor);
- as we get to pick $q(z)$, we get to choose a family that's convenient, for example, one for which we can obtain independent samples;
- tractable samples from $q(z)$ means that we can obtain MC estimates of the ELBO;
- that's because for some given $z \in \mathcal{Z}$, the ELBO only requires assessing the joint density $\log p(x,z|\theta)$ and $\log q(z)$;
- ideally, we would choose a family for which the entropy is also tractable.

Can you show to yourself that the second expression is true?

- if $p(z)$ and $q(z)$ are in the same (exponential) family, chances are the term $\mathrm{KL}\left(q(z)\,||\,p(z)\right)$ is known in closed form;

Designing a tractable approximation

Mean field approximation:

- make **all** latent variables independent under $q(z)$.
- pick a parametric family with tractable pmf.

For example, in the latent factor model this takes the form:

$$
\begin{aligned}
q(z|\lambda) &= \prod_{d=1}^{D} q(z_d|\lambda_d) \\
&= \prod_{d=1}^{D} \text{Bern}(z_d|\lambda_d)
\end{aligned}
$$

where λ is a vector that specifies D Bernoulli parameters.

Mean Field Latent Factor Model Inference

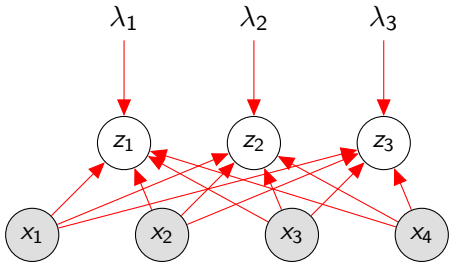

$Z_d|\lambda \sim \text{Bernoulli}(\lambda_d)$

Instead of inferring the true posterior $p_{Z|X=x}$, a computation that takes assessing the marginal probability $p_X(x|\theta)$, and thus requires all 2^D assessments of the joint density $p_{XZ}(x, z|\theta)$, **we optimise** exactly D parameters. One per Bernoulli factor in the posterior approximation $q(z|\lambda)$.

Clearly, such independence assumption is a strong simplification. In some cases we need to design *structured* approximate posteriors, that is, approximations that can correlate latent variables (we will hear more about those later).

Amortised variational inference

Amortise the cost of inference using NNs

$$q(z_1, \ldots, z_D|\lambda, x) = \prod_{d=1}^{D} q_\lambda(z_d|\lambda, x)$$

still mean field

$$Z_d|\lambda, x \sim \text{Bernoulli}(b_d)$$

but with a shared set of parameters

- where $b_1^D = \text{NN}(x; \lambda)$

The true posterior $p_{Z|X=x}$ follows from conditioning on observed x.

With NNs, we can condition on complex data efficiently, thus it seems like an interesting idea to jointly parameterise the independent factors of the posterior approximation $q(z|x, \lambda)$.

This leads to fewer parameters (more latent variables will not demand more parameters) and has a potentially useful by-product: an *inference model*, that is, a *model* of the distribution of the latent variable.

Recall our notion of model: a mechanism to predict the outcomes of a random experiment. So far, we've been attempting to model outcomes of some joint distribution $p(x, z|\theta)$. In variational inference, we introduce a rather unusual model, i.e., $q(z|x, \lambda)$, it predicts another model's posterior inferences.

Overview

Joint distribution

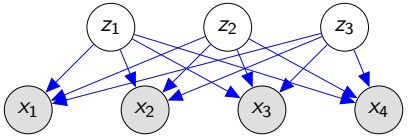

Posterior

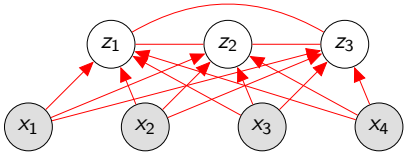

Mean field

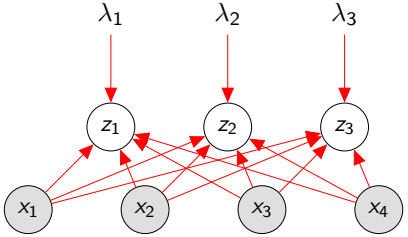

Amortised VI

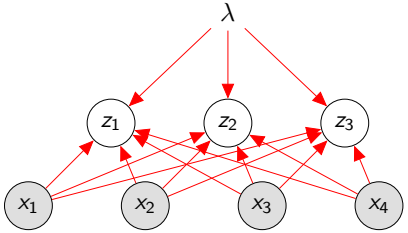

Joint distribution: latent variables are independent a priori. This is a model assumption.

Posterior: latent variables are correlated. That is because for any $z \in \mathcal{Z}$ the value $p(z|x,\theta)$ depends on $p(x,z'|\theta)$ for all $z' \in \mathcal{Z}$ via $p(x|\theta)$.

Mean field approximation: we postulate a simple distribution over latent variables, e.g., where every variable is controlled by an independent distribution. The parameters of these distributions are chosen to maximise the ELBO.

Amortised VI: we design a probabilistic model of the latent variables. That is, we design a tractable model that maps from observations to an approximation of the true posterior distribution. This inference model is typically parameterised by an inference (neural) network, its parameters are too estimated to maximise the ELBO.

Summary

- Posterior inference is often **intractable** because the marginal (or **evidence**) $p(x|\theta)$ cannot be computed efficiently.
- Variational inference approximates the posterior $p(z|x,\theta)$ with a simpler distribution $q(z)$.
- The variational objective is the **evidence lower bound (ELBO)**:

$$\mathbb{E}_{q(z)}\left[\log p(x,z|\theta)\right] + \mathbb{H}\left(q(z)\right)$$

- The solution to the ELBO minimises $\mathrm{KL}\left(q(z) \,||\, p(z|x,\theta)\right)$

There's an interesting special case of VI which is likely familiar to you. When $q(z) = p(z|x,\theta)$ we recover EM. Check the (optional) Appendix.

Summary

- We design $q(z)$ to be simple
- A common approximation is the **mean field** approximation which assumes that all latent variables are independent:

$$q(z|\lambda) = \prod_{d=1}^{D} q(z_d|\lambda_d)$$

- In amortised VI, we condition on a data point x to parameterise a collection of variational factors $\prod_{d=1}^{D} q(z_d|x, \lambda)$ and we typically use NNs for that.

Outline

Variational Inference Learning (NVIL)

Train a deep latent variable model with using amortised variational inference.

$$\lambda^*, \theta^* = \arg\max_{\lambda,\theta} \mathbb{E}_{x\sim\mathcal{D}} \left[\underbrace{\mathbb{E}_{q(z|x,\lambda)} \left[\log \frac{p(x,z|\theta)}{q(z|x,\lambda)} \right]}_{\text{ELBO}_x(\lambda,\theta)} \right]$$

Approach parameter estimation via stochastic gradient-based optimisation.

Now we discuss the concrete case of training deep discrete latent variable models with amortised variational inference.

The main difference with respect to VI as we saw is that we will be learning the inference model $q(z|x,\lambda)$ along with the joint distribution $p(x,z|\theta)$.

Concretely, we will use gradient-based optimisation to update λ and θ towards a (local) maximum of the ELBO.

The $\text{ELBO}_{\mathcal{D}}(\lambda,\theta)$, just like the log-likelihood function $\mathcal{L}_{\mathcal{D}}(\theta)$, factorises as a sum over i.i.d. observations.

Generative model

Again, let's take the latent factor document model as an example:

- a document $x = (x_1, \ldots, x_n)$ consists of n i.i.d. categorical draws from that model
- the categorical distribution in turn depends on binary latent factors $z = (z_1, \ldots, z_D)$ which are also i.i.d.

$$Z_d \sim \text{Bernoulli}(\alpha) \qquad (1 \leq d \leq D)$$
$$X_i | z \sim \text{Categorical}(f(z; \theta)) \qquad (1 \leq i \leq n)$$

Here $0 < \alpha < 1$ specifies a Bernoulli prior (assume fixed) and $f(\cdot; \theta)$ is a function computed by an NN

$$f(z; \theta) = \text{softmax}(Wz + b)$$
$$\theta = \{W, b\}$$

We've chosen a very shallow NN for the observational model. It's just an affine transformation and softmax (a log-linear model).

Nothing prevents us from using a more complex model, both in terms of parameterisation (e.g., a deeper FFNN) and statistical assumptions (e.g., a factorisation of the sequence x without Markov assumptions).

Example Model

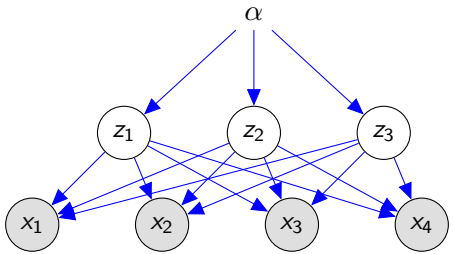

Joint distribution: independent latent variables

I omit θ from the graphical model, but recall that every $X_i|\theta, z$ in the joint distribution depends on it. Moreover, every $Z_d|\theta, x$ in the true posterior distribution also depends on it.

Example Model

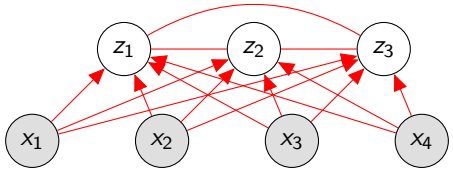

Posterior: latent variables are marginally dependent.

For our variational distribution we are going to assume that they are not (recall: mean field assumption).

I omit θ from the graphical model, but recall that every $X_i|\theta, z$ in the joint distribution depends on it. Moreover, every $Z_d|\theta, x$ in the true posterior distribution also depends on it.

Mean Field Inference

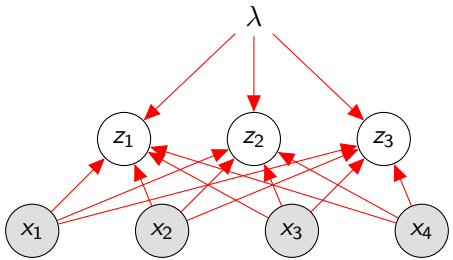

The inference network needs to predict D Bernoulli parameters b_1^D. Any neural network with sigmoid output will do that job.

The inference model is *independent* of θ.

That is the whole point, rather than actually inferring the true posterior, we want to independently estimate a model to perform approximate inference.

Inference Model

Model

$$q(z|x,\lambda) = \prod_{d=1}^{D} \text{Bern}(z_d|b_d)$$
$$\text{where } b_1^D = g(x;\lambda)$$

Example architecture (inference network)

$$h = \frac{1}{n}\sum_{i=1}^{n} E_{x_i} \qquad b_1^D = \text{sigmoid}(Mh + c)$$

$\lambda = \{E, M, c\}$

Some will call $q(z|x,\lambda)$ the inference network or recognition network. To be consistent with the vocabulary we've developed so far, I prefer calling the distribution $Z|X = x, \lambda$ an *inference model*. The *inference network* then is the NN architecture that parameterises the inference model. The term *recognition network* comes from the literature around Wake-Sleep (WS; Hinton et al., 1995), a heuristic form of VI that we discuss in the (optional) Appendix.

In this example, the inference network is very shallow: we embed the words using an embedding matrix E, combine them into an average h, project that to D real values via an affine transformation $Mh + c$, and use elementwise sigmoid to map each of those to he interval $(0, 1)$, necessary for the Bernoulli distributions. Nothing prevents us from using a more complex architecture, for example: we could encode the entire document using an LSTM and use the LSTM's last hidden state instead of the average of embeddings.

Making the inference model more complex, for example to correlate the latent assignments, is harder. But, if we knew a more complex model whose pmf is tractable (to assess and to sample from) we could use it instead. Can you think of any?

Objective

Let's concentrate on a single observation $x \in \mathcal{D}$:

$$\begin{aligned}\text{ELBO}_x(\lambda, \theta) &= \mathbb{E}_{q(z|x,\lambda)}\left[\log \frac{p(x, z|\theta)}{q(z|x, \lambda)}\right] \\ &= \mathbb{E}_{q(z|x,\lambda)}\left[\log p(x, z|\theta)\right] + \mathbb{H}\left(q(z|x, \lambda)\right) \\ &= \mathbb{E}_{q(z|x,\lambda)}\left[\log p(x|z, \theta)\right] - \text{KL}\left(q(z|x, \lambda) \,||\, p(z)\right)\end{aligned}$$

Parameter estimation

$$\arg\max_{\theta,\lambda} \; \mathbb{E}_{q(z|x,\lambda)}\left[\log p(x|z, \theta)\right] - \text{KL}\left(q(z|x, \lambda) \,||\, p(z)\right)$$

Here I list all 3 forms of the ELBO for a single data point. Generally, we need to pick a family for which we can sample from $Z|X = x, \lambda$ and assess the probability of a sample.

1. The first form is convenient when that is precisely all we can do.
2. The second form is convenient when in addition to that we can assess the entropy $\mathbb{H}(Z|X = x, \lambda)$.
3. The last form is convenient when $p(z)$ and $q(z|x, \lambda)$ are in the same exponential family.

Let's focus on the 3rd form for now, as it suffices to illustrate all challenges and the solutions we will develop. We have two terms, let's call them *the expected likelihood term* and the *KL term*, for brevity.

Our **goal** for the rest of this section is to compute ∇_θ and ∇_λ, or at least unbiased estimates thereof, as we need those for optimisation.

KL term

KL between D independent Bernoulli distributions is tractable

$$\mathrm{KL}\left(q(z|x,\lambda) \,||\, p(z|\alpha)\right) = \sum_{d=1}^{D} \mathrm{KL}\left(q(z_d|x,\lambda) \,||\, p(z_d|\alpha)\right)$$

$$= \sum_{d=1}^{D} \mathrm{KL}\left(\mathrm{Bernoulli}\left(b_d\right)\right) \,||\, \mathrm{Bernoulli}\left(\alpha\right))$$

$$= \sum_{d=1}^{D} b_d \log \frac{b_d}{\alpha} + (1-b_d) \log \frac{1-b_d}{1-\alpha}$$

In our example, the prior is a product of D independent Bernoulli distributions. Similarly, the inference model is a product of D independent distributions. This means that the KL term is a sum of D independent KL terms. Moreover, each $\mathrm{KL}\left(Z_d|X=x,\lambda \,||\, Z_d|\alpha\right)$ is known analytically, since both distributions are in the same exponential family (i.e., the Bernoulli family).

Being able to solve this expression in closed-form and with a computation that scales linearly in D means that there's no challenge in representing the KL term in a computation graph, and autodiff will be able to differentiate it with respect to λ (and even with respect to α should our prior not be fixed).

Updating the generative model

$$\frac{\partial}{\partial \theta}\left(\mathbb{E}_{q(z|x,\lambda)}\left[\log p(x|z,\theta)\right] - \overbrace{\mathrm{KL}\left(q(z|x,\lambda) \,||\, p(z)\right)}^{\text{constant wrt } \theta}\right)$$

Updating the generative model is actually rather simple

- The second term is constant in this case, and poses no challenge. Even if it depend on θ, that is, if the prior depended on θ, as long as we can evaluate the KL term, autodiff would differentiate it for us. The first term seems less obvious, after all, we cannot solve the expected value in closed-form (it would take a sum over $z \in \mathcal{Z}$. Avoiding this sum is the whole point.

Updating the generative model

$$\frac{\partial}{\partial \theta}\left(\mathbb{E}_{q(z|x,\lambda)}\left[\log p(x|z,\theta)\right] - \overbrace{\mathrm{KL}\left(q(z|x,\lambda) \,||\, p(z)\right)}^{\text{constant wrt } \theta}\right)$$

$$= \underbrace{\mathbb{E}_{q(z|x,\lambda)}\left[\frac{\partial}{\partial \theta}\log p(x|z,\theta)\right]}_{\text{expected gradient :)}}$$

Updating the generative model is actually rather simple

- The second term is constant in this case, and poses no challenge. Even if it depend on θ, that is, if the prior depended on θ, as long as we can evaluate the KL term, autodiff would differentiate it for us. The first term seems less obvious, after all, we cannot solve the expected value in closed-form (it would take a sum over $z \in \mathcal{Z}$. Avoiding this sum is the whole point.
- But note that the distribution we take expectations with respect to is the inference model $q(z|x,\lambda)$, which does not depend on θ. As derivatives are linear, we compute an expected derivative instead of differentiating an expected value.

Updating the generative model

$$\frac{\partial}{\partial \theta}\left(\mathbb{E}_{q(z|x,\lambda)}\left[\log p(x|z,\theta)\right] - \overbrace{\mathrm{KL}\left(q(z|x,\lambda)\,||\,p(z)\right)}^{\text{constant wrt } \theta}\right)$$

$$= \underbrace{\mathbb{E}_{q(z|x,\lambda)}\left[\frac{\partial}{\partial \theta}\log p(x|z,\theta)\right]}_{\text{expected gradient :)}}$$

$$\overset{\text{MC}}{\approx} \frac{1}{S}\sum_{s=1}^{S}\frac{\partial}{\partial \theta}\log p(x|z^{(s)},\theta) \quad \text{where } z^{(s)} \sim q(z|x,\lambda)$$

Updating the generative model is actually rather simple

- The second term is constant in this case, and poses no challenge. Even if it depend on θ, that is, if the prior depended on θ, as long as we can evaluate the KL term, autodiff would differentiate it for us. The first term seems less obvious, after all, we cannot solve the expected value in closed-form (it would take a sum over $z \in \mathcal{Z}$. Avoiding this sum is the whole point.
- But note that the distribution we take expectations with respect to is the inference model $q(z|x,\lambda)$, which does not depend on θ. As derivatives are linear, we compute an expected derivative instead of differentiating an expected value.
- Expected values are great for we know how to estimate them without bias. More often than not we use a single sample per observation.

Updating the generative model

$$\frac{\partial}{\partial \theta}\left(\mathbb{E}_{q(z|x,\lambda)}\left[\log p(x|z,\theta)\right] - \overbrace{\mathrm{KL}\left(q(z|x,\lambda)\,||\,p(z)\right)}^{\text{constant wrt } \theta}\right)$$

$$= \underbrace{\mathbb{E}_{q(z|x,\lambda)}\left[\frac{\partial}{\partial \theta}\log p(x|z,\theta)\right]}_{\text{expected gradient :)}}$$

$$\overset{\text{MC}}{\approx} \frac{1}{S}\sum_{s=1}^{S}\frac{\partial}{\partial \theta}\log p(x|z^{(s)},\theta) \quad \text{where } z^{(s)} \sim q(z|x,\lambda)$$

Monte Carlo (MC) estimation gives us a gradient estimate with a computation that does not depend on the size of $\mathcal{Z}$.

Updating the generative model is actually rather simple

- The second term is constant in this case, and poses no challenge. Even if it depend on θ, that is, if the prior depended on θ, as long as we can evaluate the KL term, autodiff would differentiate it for us. The first term seems less obvious, after all, we cannot solve the expected value in closed-form (it would take a sum over $z \in \mathcal{Z}$. Avoiding this sum is the whole point.
- But note that the distribution we take expectations with respect to is the inference model $q(z|x,\lambda)$, which does not depend on θ. As derivatives are linear, we compute an expected derivative instead of differentiating an expected value.
- Expected values are great for we know how to estimate them without bias. More often than not we use a single sample per observation.

Updating the inference model

$$\frac{\partial}{\partial \lambda}\left(\mathbb{E}_{q(z|x,\lambda)}\left[\log p(x|z,\theta)\right] - \overbrace{\mathrm{KL}\left(q(z|x,\lambda) \,||\, p(z)\right)}^{\text{analytical}}\right)$$

Updating the inference model is not as simple

- The KL term is tractable to assess, thus autodiff will handle it, and we don't need to worry about the exact form of the gradient.

Updating the inference model

$$\frac{\partial}{\partial \lambda}\left(\mathbb{E}_{q(z|x,\lambda)}\left[\log p(x|z,\theta)\right] - \overbrace{\mathrm{KL}\left(q(z|x,\lambda) \,||\, p(z)\right)}^{\text{analytical}}\right)$$

$$= \frac{\partial}{\partial \lambda}\mathbb{E}_{q(z|x,\lambda)}\left[\log p(x|z,\theta)\right] - \underbrace{\frac{\partial}{\partial \lambda}\mathrm{KL}\left(q(z|x,\lambda) \,||\, p(z)\right)}_{\text{analytical computation}}$$

Updating the inference model is not as simple

- The KL term is tractable to assess, thus autodiff will handle it, and we don't need to worry about the exact form of the gradient.
- The first term requires an intractable sum over $z \in \mathcal{Z}$ which we mean to avoid. Unfortunately this time we cannot simply 'push' the derivative inside as the expectation is taken w.r.t. $q(z|x,\lambda)$, which clearly depends on λ.

Updating the inference model

$$\frac{\partial}{\partial \lambda}\left(\mathbb{E}_{q(z|x,\lambda)}\left[\log p(x|z,\theta)\right] - \overbrace{\mathrm{KL}\left(q(z|x,\lambda) \,||\, p(z)\right)}^{\text{analytical}}\right)$$

$$= \frac{\partial}{\partial \lambda}\mathbb{E}_{q(z|x,\lambda)}\left[\log p(x|z,\theta)\right] - \underbrace{\frac{\partial}{\partial \lambda}\mathrm{KL}\left(q(z|x,\lambda) \,||\, p(z)\right)}_{\text{analytical computation}}$$

The first term again requires approximation by sampling, but there is a problem

Updating the inference model is not as simple

- The KL term is tractable to assess, thus autodiff will handle it, and we don't need to worry about the exact form of the gradient.
- The first term requires an intractable sum over $z \in \mathcal{Z}$ which we mean to avoid. Unfortunately this time we cannot simply 'push' the derivative inside as the expectation is taken w.r.t. $q(z|x,\lambda)$, which clearly depends on λ.

MC is not differentiable

$$\frac{\partial}{\partial \lambda} \mathbb{E}_{q(z|x,\lambda)} \left[\log p(x|z,\theta) \right]$$
$$= \frac{\partial}{\partial \lambda} \sum_z q(z|x,\lambda) \log p(x|z,\theta)$$

Unfortunately, we cannot turn to MC either, as we can only MC estimate expected values, and the derivative of the expected likelihood term does not seem to be an expected value.

MC is not differentiable

$$\frac{\partial}{\partial \lambda} \mathbb{E}_{q(z|x,\lambda)} \left[\log p(x|z,\theta) \right]$$
$$= \frac{\partial}{\partial \lambda} \sum_z q(z|x,\lambda) \log p(x|z,\theta)$$
$$= \underbrace{\sum_z \frac{\partial}{\partial \lambda} (q(z|x,\lambda)) \log p(x|z,\theta)}_{\text{not an expectation}}$$

Unfortunately, we cannot turn to MC either, as we can only MC estimate expected values, and the derivative of the expected likelihood term does not seem to be an expected value.

- Writing the expected likelihood explicitly we can see that though we can sum derivatives, as differentiation is linear, we cannot hope to evaluate all $|\mathcal{Z}|$ terms in our lifetime.

MC is not differentiable

$$
\begin{aligned}
&\frac{\partial}{\partial \lambda} \mathbb{E}_{q(z|x,\lambda)} \left[\log p(x|z,\theta) \right] \\
&= \frac{\partial}{\partial \lambda} \sum_z q(z|x,\lambda) \log p(x|z,\theta) \\
&= \underbrace{\sum_z \frac{\partial}{\partial \lambda} (q(z|x,\lambda)) \log p(x|z,\theta)}_{\text{not an expectation}}
\end{aligned}
$$

- MC estimator is non-differentiable

Unfortunately, we cannot turn to MC either, as we can only MC estimate expected values, and the derivative of the expected likelihood term does not seem to be an expected value.

- Writing the expected likelihood explicitly we can see that though we can sum derivatives, as differentiation is linear, we cannot hope to evaluate all $|\mathcal{Z}|$ terms in our lifetime.
- This shows that in general we cannot differentiate an MC estimate.

MC is not differentiable

$$\begin{aligned}
&\frac{\partial}{\partial \lambda} \mathbb{E}_{q(z|x,\lambda)} \left[\log p(x|z,\theta)\right] \\
&= \frac{\partial}{\partial \lambda} \sum_z q(z|x,\lambda) \log p(x|z,\theta) \\
&= \underbrace{\sum_z \frac{\partial}{\partial \lambda} (q(z|x,\lambda)) \log p(x|z,\theta)}_{\text{not an expectation}}
\end{aligned}$$

- MC estimator is non-differentiable
- Differentiating the expression does not yield an expectation: cannot approximate via MC

Unfortunately, we cannot turn to MC either, as we can only MC estimate expected values, and the derivative of the expected likelihood term does not seem to be an expected value.

- Writing the expected likelihood explicitly we can see that though we can sum derivatives, as differentiation is linear, we cannot hope to evaluate all $|\mathcal{Z}|$ terms in our lifetime.
- This shows that in general we cannot differentiate an MC estimate.

Score Function Estimator

We can again use the log identity for derivatives

$$\frac{\partial}{\partial \lambda} \mathbb{E}_{q(z|x,\lambda)} \left[\log p(x|z,\theta) \right]$$
$$= \sum_z \frac{\partial}{\partial \lambda} (q(z|x,\lambda)) \log p(x|z,\theta)$$

It turns out we've already seen this form of gradient when we derived the general form of $\nabla_\theta \log p(x|\theta)$ for models with tractable marginals.

Score Function Estimator

We can again use the log identity for derivatives

$$\begin{aligned}
&\frac{\partial}{\partial \lambda} \mathbb{E}_{q(z|x,\lambda)} \left[\log p(x|z,\theta)\right] \\
&= \sum_z \frac{\partial}{\partial \lambda} (q(z|x,\lambda)) \log p(x|z,\theta) \\
&= \sum_z q(z|x,\lambda) \frac{\partial}{\partial \lambda} (\log q(z|x,\lambda)) \log p(x|z,\theta)
\end{aligned}$$

It turns out we've already seen this form of gradient when we derived the general form of $\nabla_\theta \log p(x|\theta)$ for models with tractable marginals.

- We can use the log identity for derivatives (i.e., $f' = f(\log f)'$) to re-express the sum as an expectation with respect to $q(z|x,\lambda)$.

Score Function Estimator

We can again use the log identity for derivatives

$$\begin{aligned}
&\frac{\partial}{\partial\lambda}\mathbb{E}_{q(z|x,\lambda)}\left[\log p(x|z,\theta)\right] \\
&= \sum_z \frac{\partial}{\partial\lambda}(q(z|x,\lambda))\log p(x|z,\theta) \\
&= \sum_z q(z|x,\lambda)\frac{\partial}{\partial\lambda}(\log q(z|x,\lambda))\log p(x|z,\theta) \\
&= \underbrace{\mathbb{E}_{q(z|x,\lambda)}\left[\log p(x|z,\theta)\frac{\partial}{\partial\lambda}\log q(z|x,\lambda)\right]}_{\text{expected gradient :)}}
\end{aligned}$$

It turns out we've already seen this form of gradient when we derived the general form of $\nabla_\theta \log p(x|\theta)$ for models with tractable marginals.

- We can use the log identity for derivatives (i.e., $f' = f(\log f)'$) to re-express the sum as an expectation with respect to $q(z|x,\lambda)$.
- This estimator is known as the *score function estimator*.

Score Function Estimator

We can again use the log identity for derivatives

$$\begin{aligned}
&\frac{\partial}{\partial\lambda}\mathbb{E}_{q(z|x,\lambda)}\left[\log p(x|z,\theta)\right] \\
&= \sum_z \frac{\partial}{\partial\lambda}(q(z|x,\lambda))\log p(x|z,\theta) \\
&= \sum_z q(z|x,\lambda)\frac{\partial}{\partial\lambda}(\log q(z|x,\lambda))\log p(x|z,\theta) \\
&= \underbrace{\mathbb{E}_{q(z|x,\lambda)}\left[\log p(x|z,\theta)\frac{\partial}{\partial\lambda}\log q(z|x,\lambda)\right]}_{\text{expected gradient :)}}
\end{aligned}$$

We turned the derivative of an expectation into the expected value of a derivative!

It turns out we've already seen this form of gradient when we derived the general form of $\nabla_\theta \log p(x|\theta)$ for models with tractable marginals.

- We can use the log identity for derivatives (i.e., $f' = f(\log f)'$) to re-express the sum as an expectation with respect to $q(z|x,\lambda)$.
- This estimator is known as the *score function estimator*.

Score Function Estimator

We can now build an MC estimator

$$\frac{\partial}{\partial \lambda} \mathbb{E}_{q(z|x,\lambda)} \left[\log p(x|z,\theta) \right]$$
$$= \mathbb{E}_{q(z|x,\lambda)} \left[\log p(x|z,\theta) \frac{\partial}{\partial \lambda} \log q(z|x,\lambda) \right]$$

And, as always, expected gradients can be estimated free of bias via MC.

Score Function Estimator

We can now build an MC estimator

$$\frac{\partial}{\partial \lambda} \mathbb{E}_{q(z|x,\lambda)} \left[\log p(x|z,\theta) \right]$$
$$= \mathbb{E}_{q(z|x,\lambda)} \left[\log p(x|z,\theta) \frac{\partial}{\partial \lambda} \log q(z|x,\lambda) \right]$$
$$\overset{\text{MC}}{\approx} \frac{1}{S} \sum_{s=1}^{S} \log p(x|z^{(s)},\theta) \frac{\partial}{\partial \lambda} \log q(z^{(s)}|x,\lambda)$$

where $z^{(s)} \sim q(z|x,\lambda)$

And, as always, expected gradients can be estimated free of bias via MC.

Computation Graph

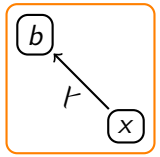

inference network

Let's put everything together in a computation graph

- we map an observation x to the parameters b of our inference model, this uses an NN with parameters λ;

Computation Graph

$z \sim \text{Bernoulli}(b)$

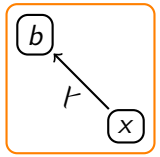

inference network

Let's put everything together in a computation graph

- we map an observation x to the parameters b of our inference model, this uses an NN with parameters λ;
- with b we can parameterise Bernoulli distributions (in our example), from which we know how to obtain independent samples;

Computation Graph

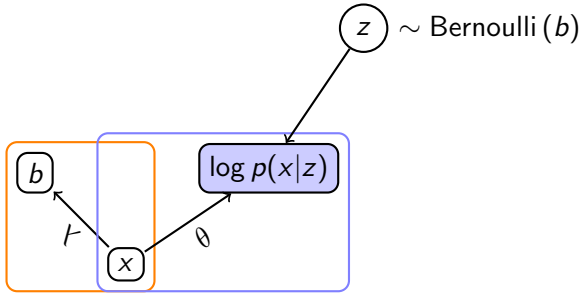

Let's put everything together in a computation graph

- we map an observation x to the parameters b of our inference model, this uses an NN with parameters λ;
- with b we can parameterise Bernoulli distributions (in our example), from which we know how to obtain independent samples;
- besides, we have our main neural network, which maps from z to the log-probability $\log p(x|z, \theta)$, this is a quantity that depends on θ and whose gradient we need in order to update the generative model;

Computation Graph

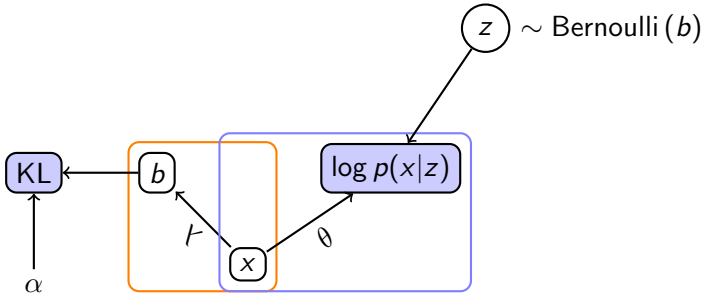

Let's put everything together in a computation graph

- we map an observation x to the parameters b of our inference model, this uses an NN with parameters λ;
- with b we can parameterise Bernoulli distributions (in our example), from which we know how to obtain independent samples;
- besides, we have our main neural network, which maps from z to the log-probability $\log p(x|z, \theta)$, this is a quantity that depends on θ and whose gradient we need in order to update the generative model;
- with b and the prior parameter α, we can assess $\mathrm{KL}\left(q(z|x, \lambda) \,||\, p(z|\alpha)\right)$, whose gradient we need in order to update the inference model;

Computation Graph

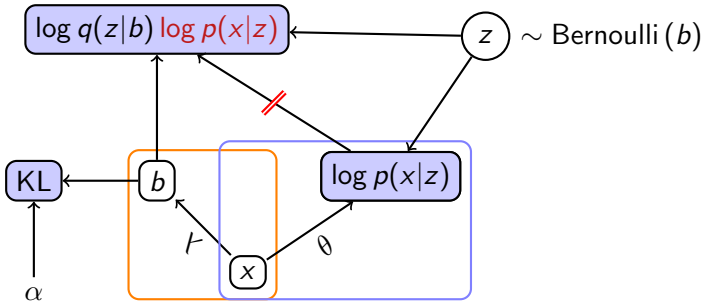

inference network

generative network

Let's put everything together in a computation graph

- we map an observation x to the parameters b of our inference model, this uses an NN with parameters λ;
- with b we can parameterise Bernoulli distributions (in our example), from which we know how to obtain independent samples;
- besides, we have our main neural network, which maps from z to the log-probability $\log p(x|z, \theta)$, this is a quantity that depends on θ and whose gradient we need in order to update the generative model;
- with b and the prior parameter α, we can assess $\mathrm{KL}\left(q(z|x, \lambda) \,||\, p(z|\alpha)\right)$, whose gradient we need in order to update the inference model;
- finally, to update the inference model we also need the score function estimator, which is $\log p(x|z, \theta) \nabla_\lambda \log q(z|x, \lambda)$; to obtain that gradient using autodiff we need to get a gradient for $\log q(z|x, \lambda)$ and scale it by the log-likelihood $\log p(x|z, \theta)$; we can do that if we can achieve that by multiplying $\log q(z|x, \lambda)$ and a 'detached' (constant) version of $\log p(x|z, \theta)$;

Computation Graph

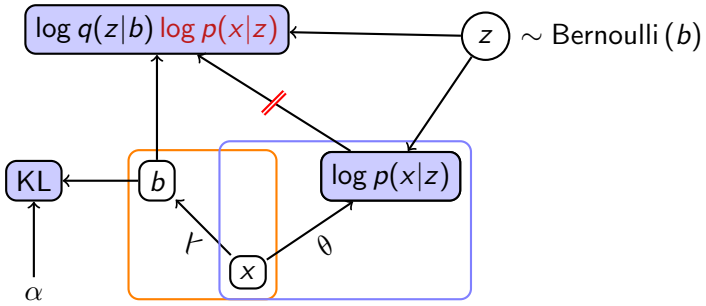

inference network

generative network

Let's put everything together in a computation graph

- we map an observation x to the parameters b of our inference model, this uses an NN with parameters λ;
- with b we can parameterise Bernoulli distributions (in our example), from which we know how to obtain independent samples;
- besides, we have our main neural network, which maps from z to the log-probability $\log p(x|z, \theta)$, this is a quantity that depends on θ and whose gradient we need in order to update the generative model;
- with b and the prior parameter α, we can assess $\mathrm{KL}\left(q(z|x, \lambda) \,||\, p(z|\alpha)\right)$, whose gradient we need in order to update the inference model;
- finally, to update the inference model we also need the score function estimator, which is $\log p(x|z, \theta) \nabla_\lambda \log q(z|x, \lambda)$; to obtain that gradient using autodiff we need to get a gradient for $\log q(z|x, \lambda)$ and scale it by the log-likelihood $\log p(x|z, \theta)$; we can do that if we can achieve that by multiplying $\log q(z|x, \lambda)$ and a 'detached' (constant) version of $\log p(x|z, \theta)$;

Stochastic surrogate objectives

A computation node whose gradient estimates the gradient we want:

$$\log p(x|z,\theta) - \mathrm{KL}\left(q(z|x,\lambda \mid\mid p(z|\alpha)\right) + \underbrace{\log p(x|z,\not{\theta})}_{\text{'detached'}} \log q(z|x,\lambda)$$

Can you verify $\nabla_{\theta,\lambda}$ of the surrogate objective yields the correct partials?

Implementation goal: we want a forward pass whose backward estimates $\nabla_{\lambda,\theta}\,\mathrm{ELBO}_x(\lambda,\theta)$.

That is, a quantity whose gradient w.r.t. λ,θ as computed by an automatic differentiation algorithm yields the correct partial derivatives for the generative and the inference model.

To implement this efficiently, we resort to the notion of a 'detached' computation node. That is, a node whose value is interpreted as a constant (its outputs are disconnected from NN parameters during back-propagation). For brevity, we will denote this by crossing the parameter out (e.g., $\not{\theta}$).

Score Function Estimator: Variance

$$\frac{\partial}{\partial \lambda} \mathbb{E}_{q(z|x,\lambda)} \left[\log p(x|z,\theta) \right] = \mathbb{E}_{q(z|x,\lambda)} \left[\log p(x|z,\theta) \frac{\partial}{\partial \lambda} \log q(z|x,\lambda) \right]$$

Empirically this estimator often exhibits high variance.

- the magnitude of $\log p(x|z,\theta)$ varies widely
- the model likelihood does not contribute to direction of gradient (it only scales the gradient)

We can get gradient estimates and they are unbiased, but they are too noisy to be useful out of the box.

How can we reduce the variance of an estimator?

Score Function Estimator: Variance

We could:

- sample more (better MC estimates)
- use variance reduction techniques (e.g. baselines and control variates)

Sampling more is not a very efficient way to reduce variance, as the variance drops with the square root of the number of samples.

Perhaps we can do better with less computation?

Score Function Estimator: Variance

Idea: standardise the "reward" $r(z) := \log p(x|z,\theta)$ to have a mean at 0 and a variance of 1

- Keep a moving average of the mean and variance $\log p(x|z,\theta)$: $\hat{\mu}$ and $\hat{\sigma^2}$.
- $\hat{r}(z) = \frac{\log p(x|z,\theta) - \hat{\mu}}{\hat{\sigma^2}}$

It can be shown that

$$\begin{aligned}\frac{\partial}{\partial\lambda}\mathbb{E}_{q(z|x,\lambda)}\left[\log p(x|z,\theta)\right] &= \mathbb{E}_{q(z|x,\lambda)}\left[\log p(x|z,\theta)\frac{\partial}{\partial\lambda}\log q(z|x,\lambda)\right] \\ &= \mathbb{E}_{q(z|x,\lambda)}\left[\hat{r}(z)\frac{\partial}{\partial\lambda}\log q(z|x,\lambda)\right]\end{aligned}$$

To understand why this is true, we need to learn more about control variates (Greensmith et al., 2004). You can see the (optional) Appendix.

In reinforcement learning, $\hat{\mu}$ is also known as a baseline. Score function estimation along with baselines is what is known as REINFORCE (Williams, 1992).

Score Function Estimator: Variance

- We can show that using these *baselines* does not bias the estimator.
- We can add more advanced *control variates* and other *baselines* to further reduce variance.
- More about this in the (optional) Appendix.

Back to the KL term

We can easily relax our constraints about the tractability of the KL term. In general, we could have the approximate posterior and the prior in different families, and the prior could even depend on θ.

Recall that

$$\mathrm{KL}\left(q(z|x,\lambda) \,||\, p(z|\theta)\right) = \mathbb{E}_{q(z|x,\lambda)}\left[\log \frac{q(z|x,\lambda)}{p(z|\theta)}\right]$$

If this quantity is not tractable we can work with gradient estimates of it:

$$\nabla_\theta \,\mathrm{KL}\left(q(z|x,\lambda) \,||\, p(z|\theta)\right) = \mathbb{E}_{q(z|x,\lambda)}\left[-\nabla_\theta \log p(z|\theta)\right]$$

$$\nabla_\lambda \,\mathrm{KL}\left(q(z|x,\lambda) \,||\, p(z|\theta)\right) = \mathbb{E}_{q(z|x,\lambda)}\left[\log \frac{q(z|x,\lambda)}{p(z|\theta)} \nabla_\lambda \log q(z|x,\lambda)\right]$$

By rewriting the KL term as an expectation we can see that its gradient w.r.t. θ is indeed the expected value of a gradient, which we can MC-estimate directly.

For the gradient w.r.t. λ, we again need to use the score function estimator, which re-expressed the gradient as an expected value, for which then MC estimation is possible.

It is an interesting exercise to show to yourself that the expression for $\nabla_\lambda \,\mathrm{KL}\left(q(z|x,\lambda) \,||\, p(z|\theta)\right)$ indeed holds.

Pros and Cons

Pros:

- Applicable to all distributions
- Many libraries come with samplers for common distributions

Cons:

- High Variance!

Unfortunately, for discrete latent variables there is not alternative. Combating the cons takes studying and deploying variance reduction techniques such as control variates (Gu et al., 2016; Tucker et al., 2017; Grathwohl et al., 2018), Rao-Blackwellization (Liu et al., 2019), as well as other techniques developed in reinforcement learning literature (Rennie et al., 2017; Schulman et al., 2017).

Mohamed et al. (2019) present an extensive survey.

NVIL's original paper (Mnih and Gregor, 2014). The same ideas power black-box inference outside the context of deep learning (Ranganath et al., 2014). Mnih and Rezende (2016) present an extension based on multiple-sample MC estimates.

Summary

- One objective, two purposes.
- Use ELBO to estimate the parameters of the inference model (i.e., for approximate posterior inference)
- Use ELBO—as a proxy to the log-likelihood function—to estimate the parameters of the generative model.
- Gradient estimation for the generative model is simple.
- Gradient estimation for the inference model requires score function estimation, which is cursed with high variance.
- Stochastic surrogate objectives give us a forward node whose backward estimates the gradient we need for optimisation.
- Reward standardisation helps obtain a useful gradient estimator.

Outline

I trained my first discrete LVM, but does it work?

Well, you gotta know where to look :-)

- Track validation ELBO, distortion (D), and rate (R).
- $I(X; Z) \leq R$, if rate is low the latent is not informative.
- For a generative model, $E_{x\sim\mathcal{D}}[p(z|x)]$ should match the prior. Thus for an approximate posterior it is worth tracking K, which should be small. A visual check often suffices (plot a histogram of prior samples against posterior samples for the entire dev set).
- Track importance sampling estimates of the model's log-likelihood (after all, this is what MLE would have optimised).
- Good NLL does not say much: certain likelihood functions do not need any additional flexibility to model the data well.

We are looking for models that have large ELBO, small NLL, large R, and very small K.

$$
\begin{aligned}
\mathrm{D} &= -\mathbb{E}_{x\sim\mathcal{D}}[\log p(x|z,\theta)] \\
\mathrm{R} &= -\mathbb{E}_{x\sim\mathcal{D}}[\mathrm{KL}\left((||\ q)\left(z|x,\lambda\right)||p(z)\right)] \\
\mathrm{ELBO} &= -D - R \\
q(z|\lambda) &= \mathbb{E}_{x\sim\mathcal{D}}[q(z|x,\lambda)] \\
\mathrm{K} &= \mathrm{KL}\left(q(z|\lambda)\ ||\ p(z)\right)
\end{aligned}
$$

Hoffman and Johnson (2016); Alemi et al. (2018); Poole et al. (2019)

Look for Failure Modes

Generating from prior samples should cover enough of the data space.

Generating from (approximate) posterior samples should preserve some appreciable aspect of the seed data.

Conversely, $\mathrm{KL}\left(q(z|x,\lambda) \,||\, q(z|x',\lambda)\right)$ should be smaller the more related the two data samples are.

You can use nearest neighbour retrieval to help automate some of this analysis.

Final Remarks

- Probabilistic models are extremely flexible tools.
- They are interesting precisely because we can make choices about unobserved aspects of the data.
- Discrete latent variables are oftentimes key to revealing interpretable structure, or to imposing some interpretable structure on a joint distribution.
- Learning discrete LVMs is challenging, but recent years have seen amazing progress.
- Join the party! Apply these models, extend them, discover problems with their estimation/evaluation, investigate solutions.
- Avoid approaching LVMs wondering whether they will beat some non-LVM NN. If such NN exists, then you are probably looking at an aspect of the problem that does not require latent variables.

What next?

Next week we talk about deep latent variable models with continuous random variables (a.k.a. VAEs).

References I

Alexander Alemi, Ben Poole, Ian Fischer, Joshua Dillon, Rif A. Saurous, and Kevin Murphy. Fixing a Broken ELBO. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 159–168, Stockholmsmässan, Stockholm Sweden, July 2018. PMLR. URL http://proceedings.mlr.press/v80/alemi18a.html.

Duygu Ataman, Wilker Aziz, and Alexandra Birch. A Latent Morphology Model for Open-Vocabulary Neural Machine Translation. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=BJxSI1SKDH.

References II

Jasmijn Bastings, Wilker Aziz, and Ivan Titov. Interpretable neural predictions with differentiable binary variables. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 2963–2977, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1284. URL https://www.aclweb.org/anthology/P19-1284.

José M Bernardo and Adrian FM Smith. *Bayesian theory*, volume 405. John Wiley & Sons, 2009.

David M. Blei, Andrew Y. Ng, and Michael I. Jordan. Latent Dirichlet Allocation. *J. Mach. Learn. Res.*, 3:993–1022, 2003. ISSN 1532-4435. URL http://dl.acm.org/citation.cfm?id=944919.944937. Publisher: JMLR.org.

References III

David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518):859–877, 2017. Publisher: Taylor & Francis.

Peter F. Brown, Stephen A. Della Pietra, Vincent J. Della Pietra, and Robert L. Mercer. The Mathematics of Statistical Machine Translation: Parameter Estimation. *Computational Linguistics*, 19(2):263–311, 1993. URL https://www.aclweb.org/anthology/J93-2003.

Nicola De Cao, Michael Schlichtkrull, Wilker Aziz, and Ivan Titov. How do Decisions Emerge across Layers in Neural Models? Interpretation with Differentiable Masking. In *EMNLP*, 2020.

Caio Corro and Ivan Titov. Differentiable Perturb-and-Parse: Semi-Supervised Parsing with a Structured Variational Autoencoder. In *ICLR*, September 2018. URL https://openreview.net/forum?id=BJlgNh0qKQ.

References IV

Yuntian Deng, Yoon Kim, Justin Chiu, Demi Guo, and Alexander Rush. Latent Alignment and Variational Attention. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 9712–9724. Curran Associates, Inc., 2018. URL http://papers.nips.cc/paper/8179-latent-alignment-and-variational-attention.pdf.

Zoubin Ghahramani and Thomas L. Griffiths. Infinite latent feature models and the Indian buffet process. In Y. Weiss, B. Schölkopf, and J. C. Platt, editors, *Advances in Neural Information Processing Systems 18*, pages 475–482. MIT Press, 2006.

References V

Will Grathwohl, Dami Choi, Yuhuai Wu, Geoff Roeder, and David Duvenaud. Backpropagation through the Void: Optimizing control variates for black-box gradient estimation. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=SyzKd1bCW.

Evan Greensmith, Peter L. Bartlett, and Jonathan Baxter. Variance Reduction Techniques for Gradient Estimates in Reinforcement Learning. *Journal of Machine Learning Research*, 5(Nov):1471–1530, 2004. ISSN ISSN 1533-7928. URL https://www.jmlr.org/papers/v5/greensmith04a.html.

Shixiang Gu, Sergey Levine, Ilya Sutskever, and Andriy Mnih. MuProp: Unbiased backpropagation for stochastic neural networks. In *ICLR (poster)*, 2016. URL http://arxiv.org/abs/1511.05176. tex.cdate: 1451606400000 tex.crossref: conf/iclr/2016.

References VI

Serhii Havrylov, Germán Kruszewski, and Armand Joulin. Cooperative learning of disjoint syntax and semantics. 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)*, pages 1118–1128, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1115. URL https://www.aclweb.org/anthology/N19-1115.

G. E. Hinton, P. Dayan, B. J. Frey, and R. M. Neal. The Wake-Sleep Algorithm for Unsupervised Neural Networks. *Science*, 268:1158–1161, 1995.

References VII

Matthew D Hoffman and Matthew J Johnson. Elbo surgery: yet another way to carve up the variational evidence lower bound. In *Workshop in Advances in Approximate Bayesian Inference, NIPS*, volume 1, page 2, 2016.

Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P. Xing. Toward controlled generation of text. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1587–1596. PMLR, 06–11 Aug 2017. URL http://proceedings.mlr.press/v70/hu17e.html.

Michael I. Jordan, Zoubin Ghahramani, Tommi S. Jaakkola, and Lawrence K. Saul. An Introduction to Variational Methods for Graphical Models. *Machine Learning*, 37(2):183–233, November 1999. ISSN 1573-0565. doi: 10.1023/A:1007665907178. URL https://doi.org/10.1023/A:1007665907178.

References VIII

Kazuya Kawakami, Chris Dyer, and Phil Blunsom. Learning to discover, ground and use words with segmental neural language models. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6429–6441, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1645. URL https://www.aclweb.org/anthology/P19-1645.

Durk P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised Learning with Deep Generative Models. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3581–3589. Curran Associates, Inc., 2014.

References IX

Tao Lei, Regina Barzilay, and Tommi Jaakkola. Rationalizing Neural Predictions. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 107–117, Austin, Texas, November 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1011. URL https://www.aclweb.org/anthology/D16-1011.

Runjing Liu, Jeffrey Regier, Nilesh Tripuraneni, Michael Jordan, and Jon Mcauliffe. Rao-blackwellized stochastic gradients for discrete distributions. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *ICML*, volume 97 of *Proceedings of machine learning research*, pages 4023–4031, Long Beach, California, USA, June 2019. PMLR. URL http://proceedings.mlr.press/v97/liu19c.html. tex.pdf: http://proceedings.mlr.press/v97/liu19c/liu19c.pdf.

References X

Thomas P. Minka. Expectation propagation for approximate bayesian inference. In *Proceedings of the seventeenth conference on uncertainty in artificial intelligence*, UAI'01, pages 362–369, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. ISBN 1-55860-800-1. Number of pages: 8 Place: Seattle, Washington.

Andriy Mnih and Karol Gregor. Neural Variational Inference and Learning in Belief Networks. In *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32*, ICML'14, pages II–1791–II–1799. JMLR.org, 2014. event-place: Beijing, China.

Andriy Mnih and Danilo Rezende. Variational inference for monte carlo objectives. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *ICML*, volume 48 of *Proceedings of machine learning research*, pages 2188–2196, New York, New York, USA, June 2016. PMLR. URL http://proceedings.mlr.press/v48/mnihb16.html. tex.pdf: http://proceedings.mlr.press/v48/mnihb16.pdf.

References XI

Shakir Mohamed, Mihaela Rosca, Michael Figurnov, and Andriy Mnih. Monte Carlo Gradient Estimation in Machine Learning. *CoRR*, abs/1906.10652, 2019. URL http://arxiv.org/abs/1906.10652.

Vlad Niculae, André F. T. Martins, and Claire Cardie. Towards Dynamic Computation Graphs via Sparse Latent Structure. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 905–911, Brussels, Belgium, October 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1108. URL https://www.aclweb.org/anthology/D18-1108.

Ben Poole, Sherjil Ozair, Aaron Van Den Oord, Alex Alemi, and George Tucker. On variational bounds of mutual information. In *International Conference on Machine Learning*, pages 5171–5180. PMLR, 2019.

References XII

Rajesh Ranganath, Sean Gerrish, and David Blei. Black Box Variational Inference. In Samuel Kaski and Jukka Corander, editors, *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, volume 33 of *Proceedings of Machine Learning Research*, pages 814–822, Reykjavik, Iceland, April 2014. PMLR. URL http://proceedings.mlr.press/v33/ranganath14.html.

Steven J. Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. Self-Critical Sequence Training for Image Captioning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 1179–1195. IEEE Computer Society, 2017. doi: 10.1109/CVPR.2017.131. URL https://doi.org/10.1109/CVPR.2017.131.

References XIII

Miguel Rios, Wilker Aziz, and Khalil Sima'an. Deep Generative Model for Joint Alignment and Word Representation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1011–1023, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1092. URL https://www.aclweb.org/anthology/N18-1092.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal Policy Optimization Algorithms. *CoRR*, abs/1707.06347, 2017. URL http://arxiv.org/abs/1707.06347.

References XIV

George Tucker, Andriy Mnih, Chris J Maddison, John Lawson, and Jascha Sohl-Dickstein. REBAR: Low-variance, unbiased gradient estimates for discrete latent variable models. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 2627–2636. Curran Associates, Inc., 2017.

Aki Vehtari, Andrew Gelman, Tuomas Sivula, Pasi Jylänki, Dustin Tran, Swupnil Sahai, Paul Blomstedt, John P Cunningham, David Schiminovich, and Christian P Robert. Expectation propagation as a way of life: A framework for bayesian inference on partitioned data. *Journal of Machine Learning Research*, 21(17):1–53, 2020.

References XV

Stephan Vogel, Hermann Ney, and Christoph Tillmann. HMM-Based word alignment in statistical translation. In *COLING 1996 volume 2: The 16th international conference on computational linguistics*, 1996. URL https://www.aclweb.org/anthology/C96-2141.

Weiyue Wang, Derui Zhu, Tamer Alkhouli, Zixuan Gan, and Hermann Ney. Neural Hidden Markov Model for Machine Translation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 377–382, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-2060. URL https://www.aclweb.org/anthology/P18-2060.

References XVI

Ronald J. Williams. Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning. *Machine Learning*, 8(3-4): 229–256, May 1992. ISSN 0885-6125. doi: 10.1007/BF00992696. URL https://doi.org/10.1007/BF00992696.

Dani Yogatama, Phil Blunsom, Chris Dyer, Edward Grefenstette, and Wang Ling. Learning to compose words into sentences with reinforcement learning. In *ICLR*, 2017.

Chunting Zhou and Graham Neubig. Multi-space variational encoder-decoders for semi-supervised labeled sequence transduction. In *Proceedings of the 55th annual meeting of the association for computational linguistics (volume 1: Long papers)*, pages 310–320, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1029. URL https://www.aclweb.org/anthology/P17-1029.