Variational Inference using Implicit Distributions

Ferenc Huszár [1]

Abstract

Generative adversarial networks (GANs) have given us a great tool to fit implicit generative models to data. Implicit distributions are ones we can sample from easily, and take derivatives of samples with respect to model parameters. These models are highly expressive and we argue they can prove just as useful for variational inference (VI) as they are for generative modelling. Several papers have proposed GAN-like algorithms for inference, however, connections to the theory of VI are not always well understood. This paper provides a unifying review of existing algorithms establishing connections between variational autoencoders, adversarially learned inference, operator VI, GAN-based image reconstruction, and more. Secondly, the paper provides a framework for building new algorithms: depending on the way the variational bound is expressed we introduce prior-contrastive and joint-contrastive methods, and show practical inference algorithms based on either density ratio estimation or denoising.

1. Introduction

Implicit distributions are probability models whose probability density function may be intractable, but there is a way to

1. sample from them exactly and/or calculate and approximate expectations under them, and
2. calculate or estimate gradients of such expectations with respect to model parameters.

A popular example of implicit models are stochastic generative networks: samples from a simple distribution - such as uniform or Gaussian - are transformed nonlinearly and non-invertably by a deep neural network. Such networks can flexibly parametrise a wide range of probability distributions, including even degenerate ones which may not

even have a continuous density.

Implicit models have been successfully applied to generative modelling in generative adversarial networks (GANs Goodfellow et al., 2014) and subsequent work (Salimans et al., 2016; Radford et al., 2016; Donahue et al., 2017; Dumoulin et al., 2017). They work particularly well for visual data, partly because they can exploit the inductive biases of convolutional neural networks, and partly because they can flexibly model potentially degenerate, manifold-like distributions which natural images are assumed to follow.

This note is about using implicit distributions in another important probabilistic machine learning problem: approximate inference in latent variable models. Unlike in the first applications of GANs to generative modelling where an implicit model directly models the distribution of observed data, in approximate inference we are interested in modelling the posterior distribution of latent variables given observed data. Direct generative modelling and approximate inference are very different problems indeed: in the former we are provided with samples x_i from the distribution to be modelled, in the latter we are given a joint distribution of latents z and observations x, and a set of observed samples x_i, but no samples from the posterior itself.

In this note we focus on variational inference (VI) which works by minimising a divergence between the approximate and real posteriors. More precisely we follow the usual KL-divergence formulation, but other, more general variational methods exist (e.g. Li & Turner, 2016; Ranganath et al., 2016). VI also provides a lower bound to the marginal likelihood or model evidence - the *evidence lower bound* or ELBO - which can be maximised with respect to parameters of the latent variable model to approximate maximum likelihood learning. It's important to keep in mind that despite several algorithms in this note look and feel like adversarial training procedures, the way the model is fit to observed data is more akin to variational auto-encoders (VAE, Kingma & Welling, 2014; Rezende & Mohamed, 2015) than to GANs.

There are several reasons to explore implicit distributions in the context of variational inference. Firstly, explicit VI is often limited to exponential family distributions or other distributions with tractable densities (Rezende & Mohamed, 2015; Kingma et al., 2016) which may not be ex-

[1] Twitter, London, U.K. Correspondence to: Ferenc Huszár <fhuszar@twitter.com>.

pressive enough to capture non-trivial dependencies that the real posterior exhibits. A flexible, implicit distribution may provide a better approximation to the posterior and a sharper lower bound. Secondly, It may be desirable to use an implicit likelihood as the resulting latent variable model may fit the data better. For example, the likelihood or forward model might be described as a probabilistic program (Vajda, 2014) whose density is intractable or unknown. Finally, sometimes we may want to use implicit priors over latent variables. For example in a deep hierarchical latent variable model the prior for a layer may be a complicated probabilistic model with an intractable density. Or, when solving inference problems in computational photography, the prior may be the empirical distribution of natural images as in (Sønderby et al., 2017). In summary, any or all of the prior, the likelihood and the approximate posterior may have to be modelled implicitly, and we need VI procedures that are ready to tackle these situations.

In this note we present two sets of tools to handle implicit distributions in variational inference: GAN-like *adversarial algorithms* which rely on density ratio estimation, and *denoising-based algorithms* which build a representation of the gradients of each implicit distribution's log-density and use these gradient estimates directly in a stochastic gradient descent (SGD) algorithm. We further classify algorithms as *prior-contrastive* and *joint-contrastive* depending on which form of the variational bound they use. Prior-contrastive methods only deal with implicit distributions over latent variables (i. e. the prior or approximate posterior), while joint-contrastive methods can handle fully implicit models where none of the distributions involved has a tractable density. This classification gives rise to a range of algorithms listed in Table 1, alongside related algorithms from prior work. All of the algorithms presented here can perform variational approximate inference, which is the main focus of this note, but not all of them can perform learning unless the likelihood is explicitly defined.

1.1. Overview of prior work

Several of the algorithms proposed here have been discovered in some form before. However, their connections to variational inference is rarely made explicit. In this section we review algorithms for inference and feature learning which use implicit distributions or adversarial training. As we will see, several of these admit a variational interpretation or can be rather straightforwardly modified to fit the variational framework.

GANs have been used rather successfully to solve inverse problems in computer vision. These inverse problems can be cast as a special case of approximate inference. Dosovitskiy & Brox (2016) used GANs to reconstruct and generate images from non-linear feature representations. As

ALGORITHM	IMPLICIT			VI
	$p_\theta(z)$	$p_\theta(x\|z)$	$q_\psi(z\|x)$	
VAE (KINGMA & WELLING, 2014) NF (REZENDE & MOHAMED, 2015)				✓
PC-ADV, ALGORITHM 1 AffGAN† (SØNDERBY ET AL., 2017) AVB (MESCHEDER ET AL., 2017)	I		✓	✓
OPVI (RANGANATH ET AL., 2016)	I		✓	✓
PC-DEN, ALGORITHM 3	I		✓	✓
JC-ADV, ALGORITHM 2	I	I	✓	✓
JC-DEN	I	I	✓	✓
JC-ADV-RMD‡	✓	✓	✓	✓
AAE (MAKHZANI ET AL., 2016)	I		✓	
DEEPSiM (DOSOVITSKIY & BROX, 2016)	I		✓	
ALI (DUMOULIN ET AL., 2017) BiGAN (DONAHUE ET AL., 2017)	✓	✓	✓	

Table 1. Summary of algorithms for variational inference and learning in latent variable models using implicit distributions. Columns 2-4 indicate whether the component distributions: the prior, likelihood or approximate posterior are handled implicitly by the algorithm. "I" denotes inference only - the parameters θ of these distributions cannot be learned unless they are explicitly defined, but inference can still be performed. The last column indicates whether the algorithm has a variational interpretation, i. e. minimises an approximation to ELBO. In naming the algorithms `PC` and `JC` stand for prior-contrastive or joint-contrastive, `Adv` and `Den` stand for adversarial (Section 3) or denoiser-based (Section 4). Algorithms that share a row are equivalent or special cases of each other. †`AffGAN` is specialised to the task of image super-resolution where the likelihood is degenerate and linear. ‡ The reverse-mode differentiation-based `JC-Adv-RMD` algorithm has not been validated experimentally.

pointed out later by Sønderby et al. (2017), this method, DeePSiM, can be interpreted as a special case of amortised maximum a posteriori (MAP) or variational inference with a Gaussian observation model. GANs have also been used for inference in image super-resolution Ledig et al. (2016); Sønderby et al. (2017). Connections between GANs and VI in this context were first pointed out in (Sønderby et al., 2017, Appendix F). Sønderby et al. (2017) also introduced a modified objective function for the GAN generator which ensures that the algorithm minimises the Kullback-Leibler divergence as opposed to the Jensen-Shannon divergence, an essential step to using GANs for VI. The AffGAN algorithm presented there is highly application specific, thus, it does not solve VI in general.

In more recent, parallel work, (Mescheder et al., 2017) proposes *adversarial variational Bayes*(AVB), perhaps the best description of the use of GANs for variational inference. AVB is a general algorithm that allows for implicit variational distributions and is in fact equivalent to the prior-contrastive adversarial algorithm (`PC-Adv`, Algorithm 1) described in Section 3. Operator variational in-

ference (OPVI, Ranganath et al., 2016) formulates a general class of variational lower bounds based on operator divergences, resulting in a practical algorithm for training implicit inference networks without a tractable density. As is shown in the paper, the KL-divergence-based variational bound used here and in (Mescheder et al., 2017) is a special case of OPVI. Adversarial autoencoders (AAE, Makhzani et al., 2016) are similar to variational autoencoders where the KL-divergence term is replaced by an adversarial objective. However, AAEs do not use the KL-divergence formulation of the adversarial loss and their discriminator is independent of the encoder's input, thus they are not a true variational method. Finally, Karaletsos (2016) proposed variational message passing, in which adversaries are employed to minimise local Jensen-Shannon divergences in an algorithm more akin to expectation propagation (Minka, 2001) than to variational inference.

Another line of research extends GANs to latent variable models by training the discriminator on the joint distribution of latent and observed variables. This technique has been independently discovered as bi-directional GAN (BiGAN, Dumoulin et al., 2017) and adversarially learned inference (ALI, Donahue et al., 2017). These algorithms are closely related to the joint-contrastive adversarial algorithm (`JC-Adv`, Algorithm 2). ALI and BiGAN use the Jensen-Shannon formulation of GANs rather than the KL-divergence ones used here. On the one hand, this means that the Jensen-Shannon variants aren't technically VI algorithms. On the other hand, the symmetry of the Jensen-Shannon divergence makes ALI and BiGAN completely symmetric, enabling not only approximate inference but also learning in the same algorithm. Unfortunately, this is no longer true when KL divergences are used: `JC-Adv` is an algorithm for inference only.

The algorithms mentioned so far are examples of adversarial techniques which rely on density ratio estimation as the primary tool for dealing with implicit distributions (Mohamed & Lakshminarayanan, 2016). Sønderby et al. (2017) and Warde-Farley & Bengio (2017) demonstrated an alternative or complementary technique based on denoising autoencoders. As shown by (Alain & Bengio, 2014) the optimal denoising function learns to represent gradients of the log data density - which in turn can be used in an inference method. Sønderby et al. (2017) used this insight to build a denoiser-based inference algorithm for image super-resolution and connected it to amortised maximum a posteriori (MAP) inference. The extension from MAP to variational inference is straightforward and this method is closely related to the prior-contrastive denoising VI (`PCDen`, Algorithm 3) algorithm presented here.

2. Variational Inference: Two forms

In this section, we give a lightweight overview of amortised variational inference (VI) in a latent variable model, in a model similar to e.g. variational autoencoders (Kingma & Welling, 2014). We observe an i.i.d. sequence of N observed data $\mathcal{D} = \{x_n, n = 1 \ldots N\}$. For each data point there exists an associated latent variable $z_n, n = 1 \ldots N$. We specify a prior $p_\theta(z)$ over latent variables and a forward model $p_\theta(x|y)$ which describes how the observations are related to latents. In such model we are interested in maximum likelihood learning, which maximises the marginal likelihood or model evidence $\sum_{n=1}^N \log p_\theta(x_n)$ with respect to parameters θ, and inference which involves calculating the posterior $p_\theta(z|x)$. We assume that neither the marginal likelihood or the posterior are tractable.

In amortized VI we introduce an auxiliary probability distribution $q(z|x_n; \psi)$, known as the recognition model, inference network or approximate posterior. Using q_ψ we define the evidence lower bound (ELBO) as follows:

$$\mathcal{L}(\theta, \psi) = \sum_{n=1}^N \{\log p_\theta(x_n) - \mathrm{KL}\left[q_\psi(z|x_n) \| p_\theta(z|x_n)\right]\} \tag{1}$$

As the name suggests, ELBO is a lower bound to the model evidence $p_\theta(\mathcal{D})$ and it is exact when q_ψ matches the true posterior $p_\theta(z|x)$ exactly. Maximising ELBO with respect to ψ is known as *variational inference*. This minimises the KL divergence $\mathrm{KL}\left[\left[\|q\right]_\psi |p_\theta(z|x)\right]$ thus moving the q_ψ closer to the posterior. Conversely, maximising ELBO with respect to θ is known as *variational learning* which approximates maximum likelihood learning.

The ELBO can be calculated exactly for many combinations of p_θ and q_ψ, whose densities are tractable. VAEs use a re-parametrisation trick to construct a low variance estimator to ELBO, but still require tractable densities for both the model p_θ and recognition model q_ψ. If p_θ and/or q_ψ are implicit the ELBO needs to be approximated differently. As we will see in the next sections, it is useful to formulate ELBO in terms of density ratios.

There are two main forms considered here. Firstly, the *prior-contrastive* form used also by VAEs (Kingma & Welling, 2014):

$$\mathcal{L} = \sum_{n=1}^N \mathbb{E}_{z \sim q_\psi(z|x_n)} \left[\log p_\theta(x_n|z) - r_{\theta,\psi}(z, x_n)\right] \tag{2}$$

$$= \sum_{n=1}^N \left[\mathbb{E}_{z \sim q_\psi(z|x_n)} \log p_\theta(x_n|z) - \mathrm{KL}\left[q_\psi(z|x_n) \| p_\theta(z)\right]\right],$$

where we introduced notation for the logarithmic density ratio $r_{\theta,\psi} = \log \frac{q_\psi(z|x_n)}{p_\theta(z)}$.

We call Eqn. (2.1) the *prior-contrastive* expression as the KL term contrasts the approximate posterior q_ψ with the prior $p_\theta(z)$. Alternatively, we can write ELBO in a *joint-contrastive* form as follows:

$$\mathcal{L}(\theta, \psi) = -N \cdot (\text{KL}\left[q_\psi(z|x)p_\mathcal{D}(x)\|p_\theta(x, z)\right] - \mathbb{H}[p_\mathcal{D}])$$

$$= -\sum_{n=1}^{N} \mathbb{E}_{z \sim q_\psi(z|x_n)} s_{\theta, \psi}(x_n, z) - N \cdot \mathbb{H}[p_\mathcal{D}], \qquad (3)$$

where we introduced notation $p_\mathcal{D}$ to denote the real data distribution and $\mathbb{H}[p_\mathcal{D}]$ denotes its entropy[1]. $\mathbb{H}[p_\mathcal{D}]$ can be ignored as it is constant with respect to both θ and ψ. We also introduced notation $s_{\theta, \psi}$ to denote the logarithmic density ratio $s_{\theta, \psi}(x_n, z) = \log \frac{q_\psi(z|x_n)p_\mathcal{D}(x)}{p_\theta(x_n, z)}$. Note that while $r_{\theta, \psi}$ was a log-ratio between densities over z, $s_{\theta, \psi}$ is the ratio of joint densities over the tuple (x, z). As this form contrasts joint distributions, we call Eqn. (2.1) the *joint-contrastive* expression.

When using implicit models the density ratios $r_{\theta, \psi}$ and $s_{\theta, \psi}$ cannot be computed analytically. Indeed, even if all distributions involved are explicitly defined, $p_\mathcal{D}$ is only available as an empirical distribution, thus $s_{\theta, \psi}$ cannot be calculated even if the densities of other distributions are tractable. In this note we rely on techniques for estimating $r_{\theta, \psi}$ or $s_{\theta, \psi}$, or their gradients, directly from samples. For this to work we need to deal with a final difficulty: that $r_{\theta, \psi}$ or $s_{\theta, \psi}$ themselves implicitly depend on the parameter ψ which we would like to optimise.

2.1. Dependence of $r_{\theta, \psi}$ and $s_{\theta, \psi}$ on ψ

The KL-divergences in equations and depend on ψ in two ways: first, an expectation is taken with respect to q_ψ - this is fine as we assumed expectations under implicit distributions and their gradients can be approximated easily. Secondly, the ratios $r_{\theta, \psi}$ and $s_{\theta, \psi}$ themselves depend on ψ, which may cause difficulties. If one optimised ELBO na ively via gradient descent, one should back-propagate through both of these dependencies. Fortunately, the second dependence can be ignored:

$$\left.\frac{\partial}{\partial \psi} \mathbb{E}_{z \sim q_\psi} r_{\theta, \psi}(z)\right|_{\psi = \psi_0} = \left.\frac{\partial}{\partial \psi} \mathbb{E}_{z \sim q_\psi} r_{\theta, \psi_0}(z)\right|_{\psi = \psi_0} \qquad (4)$$

The only difference between the LHS and RHS of the equation is in the subscripts $r_{\theta, \psi}$ v. s. r_{θ, ψ_0}. As r_{θ, ψ_0} is a con-

stant with respect to ψ, Eqn. (4) reduces to the gradient of an expectation with respect to q_ψ, which we assumed we can approximate if q_ψ is an implicit distribution. The detailed proof of Eqn. (4) is in Appendix A, the key idea of which is the observation that for any ψ_0

$$\mathbb{E}_{z \sim q_\psi} r_{\theta, \psi}(z) = \mathbb{E}_{x, z \sim q_{\psi_0}} r_{\theta, \psi}(x, z) + \text{KL}\left[q_\psi \| q_{\psi_0}\right]$$

A similar equation analogously holds for $s_{\theta, \psi}$ in Eqn. (2.1), or indeed, any other KL divergence as well.

2.2. Approximate SGD Algorithms for VI

In the following sections we outline algorithms for VI which allow for implicit distributions. These algorithms can generally described as two nested loops of the following nature:

- the outer loop performs stochastic gradient descent (SGD) on an approximation to ELBO with respect to ϕ, using gradient estimates obtained by the inner loop

- in each iteration of outer loop, with $\psi = \psi_t$ fixed, the inner loop constructs an estimate to r_{θ, ψ_t}, $s\theta, \psi_t$, or more generally to the gradient in Eqn. (4)

As long as the gradient estimates provided by the inner loop has no memory between subsequent iterations of the outer loop, and the gradient estimates provided by the inner loop on average constitute a conservative vector field, the algorithms can be seen as instances of SGD, and as such, should have the convergence properties of SGD.

3. Direct Density Ratio Estimation

Direct density ratio estimation, also known as direct importance estimation, is the task of estimating the ratio between the densities of two probability distribution given only i. i. d. samples from each of the distributions (Kanamori et al., 2009; Sugiyama et al., 2008; Mohamed & Lakshminarayanan, 2016, see e. g.). This task is relevant in many machine learning applications, such as dealing with covariate shift or domain adaptation. A range of methods have been introduced to learn density ratios from samples, here we focus on *adversarial* techniques which employ a discriminator trained via logistic regression. We note that other methods such as KLIEP (Sugiyama et al., 2008; Mohamed & Lakshminarayanan, 2016) or LSIF (Kanamori et al., 2009; Uehara et al., 2016) could be used just as well.

3.1. Adversarial approach using discriminators

(Bickel et al., 2007) proposed estimating density ratios by training a logistic regression classifier between samples from the two distributions. Assuming the classifier is close

---

[1] In practice, $p_\mathcal{D}$ is an empirical distribution of samples, so technically it does not have a continuous density or differential entropy $\mathbb{H}[p_\mathcal{D}]$. We still use this notation liberally to avoid unnecessarily complicating the derivations.

to a unique Bayes-optimum, it can then be used directly to provide an estimate of the logarithmic density ratio. This approach has found great application in generative adversarial networks (Sønderby et al., 2017; Mohamed & Lakshminarayanan, 2016; Uehara et al., 2016), which work particularly well for generative modelling of images (see e. g. Salimans et al., 2016).

Let us use this to construct an approximation r_ϕ to the logarithmic density ratio $r_{\psi,\theta}$ from Eqn. (2.1). We can do this by minimising the following objective function, typically via SGD:

$$\ell_{\psi,\theta}^{\text{PC-Adv}}(\phi) = \sum_{n=1}^{N} \mathbb{E}_{z \sim p_\theta} \xi^+ \left(r_\phi(x_n, z) \right)$$
$$- \sum_{n=1}^{N} \mathbb{E}_{z \sim q_\psi(z|x_n)} \xi^- \left(r_\phi(x_n, z) \right), \quad (5)$$

where $\xi^+(t) = \log(1 + \exp^t)$ and $\xi^-(t) = t - \xi^+(t)$ are the softplus and softminus functions, respectively. Once the approximate log ratio r_ϕ is found, we can use it to take a gradient descent step along the approximate negative ELBO:

$$-\mathcal{L}_\phi(\psi, \theta) = \sum_{n=1}^{N} \mathbb{E}_{\epsilon \sim \mathcal{N}(0,I)} r_\phi \left(x_n, g_\psi(x_n, \epsilon) \right)$$
$$- \sum_{n=1}^{N} \mathbb{E}_{\epsilon \sim \mathcal{N}(0,I)} \log p_\theta \left(x_n | g_\psi(x_n, \epsilon) \right), \quad (6)$$

where we re-parametrised sampling from q_ψ in terms of a generator function g_ψ and noise ϵ. When q_ψ is explicitly defined, this re-parametrisation is the same as the re-parametrisation in VAEs. When q_ψ is an implicit distribution, it often already defined in terms of a non-linear function g_ψ and a noise variable ϵ which it transforms.

Equations and are analogous to the discriminator and generator losses in generative adversarial networks, with $1/(1 + \exp(-r_\phi))$ taking the role of the discriminator. Optimising the two losses in tandem gives rise to the prior-contrastive adversarial algorithm (PC-Adv, Algorithm 1) for variational inference. This algorithm is equivalent to the independently developed *adversarial variational Bayes* (Mescheder et al., 2017, AVB).

As the likelihood $p_\theta(x|z)$ appears in Eqn. (3.1), in Algorithm 1 the forward model has to be explicitly defined, but the prior $p_\theta(z)$ and approximate posterior q_ψ can be implicit. Algorithm 1 only describes variational inference - finding ψ given θ - but the approximate ELBO in Eqn. can

---

Algorithm 1 PC-Adv: prior-contrastive adversarial VI

Input: data $\mathcal{D}$, model p_θ, batchsize B, iter. count K
repeat
  for $k = 1$ **to** K **do**
    $\{z_{k,b}^p, b = 1 \dots B\} \leftarrow$ sample B items from p_θ
    $\{x_{k,b}, b = 1 \dots B\} \leftarrow$ sample B items from $\mathcal{D}$
    for all $x_{k,b}$ **do**
      $z_{k,b}^q \leftarrow$ sample from $q_\psi(z|x_{k,b})$
    end for
    update ϕ by gradient descent step on
    $$\sum_{b=1}^{B} \left[\xi^+ \left(r_\phi(x_{k,b}, z_{k,b}^p) \right) - \xi^- \left(r_\phi(x_{k,b}, z_{k,b}^q) \right) \right]$$
  end for
  $\{x_b, b = 1 \dots B\} \leftarrow$ sample B items from $\mathcal{D}$
  $\{\epsilon_b, b = 1 \dots B\} \leftarrow$ sample B items from $\mathcal{N}(0, I)$
  update ψ by gradient descent step on
  $$\sum_{b=1}^{B} r_\phi(x_b, g_\psi(x_b, \epsilon_b)) - \log p_\theta(x_{n,b} | g_\psi(x_b, \epsilon_b))$$
until change in ψ is negligible

---

be used for variational learning of θ as well, with the exception for parameters of the prior $p_\theta(z)$.

Learning prior parameters involves minimising the KL-divergence $\text{KL} \left[\sum_{n=1}^{N} q_\psi(z|x_n) \middle\| p_\theta(z) \right]$ which is akin to fitting p_θ to samples from the aggregate posterior $\sum_{n=1}^{N} q_\psi(z|x_n)$ via maximum likelihood. If the prior has a tractable density, this may be an easy task to do. A more interesting case is though when the prior $p_\theta(z)$ itself is a latent variable model, in which case we can lower bound the said KL divergence with another ELBO, thereby stacking multiple models on top of each other in a hierarchical fashion (Kingma & Welling, 2014; Rezende & Mohamed, 2015; Sønderby et al., 2016).

A similar adversarial algorithm (JC-Adv, Algorithm 2) can be constructed to target $s_{\theta,\psi}$ in the joint-contrastive formulation of ELBO (Eqn. 2.1). JC-Adv is very similar to ALI (Dumoulin et al., 2017) and BiGAN (Donahue et al., 2017) in that it learns to discriminate between the joint distributions $p_\mathcal{D}(x)q_\psi(z|x)$ and $p_\theta(x, z)$. Unlike these methods, however, JC-Adv uses the correct loss functions so it maximises an approximation to the ELBO. Unlike in PC-Adv, which required a tractable likelihood $p_\theta(x|z)$, JC-Adv also works with completely implicitly defined models. As a downside, JC-Adv provides no direct way for variational learning of θ. ALI and BiGAN exploit the symmetry of the Jensen-Shannon divergence to optimise for θ, but as JC-Adv uses the asymmetric KL-divergence, this is not an option. Section 7 explores an idea for fixing this drawback of JC-Adv.

Algorithm 2 `JC-Adv`: joint-contrastive adversarial VI

> **Input:** data $\mathcal{D}$, model p_θ, batchsize B, iter. count K
> **repeat**
>> **for** $k = 1$ **to** K **do**
>>> $\{x_{k,b}^q, b = 1 \ldots B\} \leftarrow$ sample B items from $\mathcal{D}$
>>> **for all** $x_{k,b}^q$ **do**
>>>> $z_{k,b}^q \leftarrow$ sample from $q_\psi(z|x_{k,b}^q)$
>>> **end for**
>>> $\{z_{k,b}^p, b = 1 \ldots B\} \leftarrow$ sample B items from p_θ
>>> **for all** $z_{k,b}^p$ **do**
>>>> $x_{k,b}^p \leftarrow$ sample from $p_\theta(x|z_{k,b}^p)$
>>> **end for**
>>> **update** ϕ by gradient descent step on
>>> $$\sum_{b=1}^{B} \left[\xi^+ \left(s_\phi(x_{k,b}^p, z_{k,b}^p) \right) - \xi^- \left(s_\phi(x_{k,b}^q, z_{k,b}^q) \right) \right]$$
>> **end for**
>> $\{x_b, b = 1 \ldots B\} \leftarrow$ sample B items from $\mathcal{D}$
>> $\{\epsilon_b, b = 1 \ldots B\} \leftarrow$ sample B items from $\mathcal{N}(0, I)$
>> **update** ψ by gradient descent step on
>> $$\sum_{b=1}^{B} s_\phi(x_b, g_\psi(x_b, \epsilon_b))$$
> **until** change in ψ is negligible

Algorithm 3 `PC-Den`: prior-contrastive denoising VI

> **Input:** data $\mathcal{D}$, p_θ, batchsize B, iter. count K, σ_n
> **repeat**
>> **for** $k = 1$ **to** K **do**
>>> $\{x_{k,b}, b = 1 \ldots B\} \leftarrow B$ samples from $\mathcal{D}$
>>> **for all** $x_{k,b}$ **do**
>>>> $z_{k,b} \leftarrow$ sample from $q_\psi(z|x_{k,b})$
>>> **end for**
>>> $\eta_{k,b}, b = 1 \ldots B\} \leftarrow B$ samples from $\mathcal{N}(0, \sigma_n I)$
>>> **update** ϕ by gradient descent step on
>>> $$\sum_{b=1}^{B} \|u_\phi(z_{k,b} + \eta_{k,b}) - z_{k,b}\|^2$$
>> **end for**
>> $\{x_b, b = 1 \ldots B\} \leftarrow$ sample B items from $\mathcal{D}$
>> $\{\epsilon_b, b = 1 \ldots B\} \leftarrow$ sample B items from $\mathcal{N}(0, I)$
>> **for all** x_b **do**
>>> $z_b \leftarrow g_\psi(x_b, \epsilon_b)$
>> **end for**
>> **update** ψ by gradient descent using gradient
>> $$\sum_{b=1}^{B} \frac{\partial g_\psi(x_b, \epsilon_b)}{\partial \psi} \left[\left. \frac{\partial p_\theta(x_b, z)}{\partial z} \right|_{z = z_b} + \frac{z_b - u_\phi(z_b)}{\sigma_n^2} \right]$$
> **until** change in ψ is negligible

---

4. Denoiser-guided learning

Although most versions of GAN use an adversarial discriminator based on logistic regression, there are other ways one can tackle learning and inference with implicit distributions. One interesting tool that has emerged in recent papers (Sønderby et al., 2017; Warde-Farley & Bengio, 2017) is the use of denoising autoencoders (DAEs, Vincent et al., 2008) or reconstruction-contractive autoencoders (RCAEs, Alain & Bengio, 2014).

The key observation for using denoising is that the Bayes-optimal denoiser function captures gradients of the log density of the data generating distribution:

$$v^*(x, z) = \text{argmin}_v \, \mathbb{E}_{x, z \sim q(x,z), \eta \sim \mathcal{N}_{\sigma_n}} \|v(x, z + \eta) - z\|^2$$

$$\approx z + \sigma_n^2 \frac{\partial \log q(z|x)}{\partial z} \qquad \text{as } \sigma_n \to 0 \qquad (7)$$

This allows us to construct an estimator to the *score* of a distribution by fitting a DAE to samples. We note that it is possible to obtain a more precise analytical expression for for the optimal denoising function (Valpola et al., 2016).

Let's see how one can use this in the prior-contrastive scenario to deal with an implicit q_θ. First, we fit a denoising

function u_ϕ by minimising the following loss:

$$\ell_{\psi, \sigma_n}^{\text{PC-Den}}(\phi) = \sum_{n=1}^{N} \mathbb{E}_{z \sim q_\theta(z|x_n), \eta \sim \mathcal{N}_{\sigma_n}} \|v_\phi(z + \eta, x_n) - z\|^2,$$
$$(8)$$

which can then be used to approximate the gradient of ELBO (Eqn. (4)) with respect to ψ as follows:

$$\frac{\partial \mathcal{L}(\theta, \psi)}{\partial \psi} \approx \sum_{n=1}^{N} \mathbb{E}_{\epsilon \sim \mathcal{N}} \frac{\partial g_\psi(x_n, \epsilon)}{\partial \psi} \left. \frac{\partial \log p_\theta(x_n, z)}{\partial z} \right|_{z = g_\psi(x_n, \epsilon)}$$
$$+ \sum_{n=1}^{N} \mathbb{E}_{\epsilon \sim \mathcal{N}} \frac{\partial g_\psi(x_n, \epsilon)}{\partial \psi} \frac{g_\psi(x_n, \epsilon) - u_\phi(g_\psi(x_n, \epsilon))}{\sigma_n^2}. \qquad (9)$$

Several SGD methods only require gradients of the objective function as input, this gradient estimate can be readily used to optimise an approximate ELBO. The resulting iterative algorithm, prior-contrastive denoising VI (`PC-Den`, Algorithm 3) updates the denoiser and the variational distribution in tandem. Following similar derivation one can construct several other variants of the algorithm. The denoiser approach is more flexible than the adversarial approach as one can pick and choose which individual distributions are modelled explitly, and which ones are implicit. For example, when the prior p_θ is implicit, we can train a denoiser to represent its score function. Or, one can start from the joint-contrastive formulation of the ELBO and train a DAE over joint distribution of x and z, giving rise to the joint-contrastive denoising VI (`JC-Den`). In the

interest of space the detailed description of these variants is omitted here.

As the denoising criterion estimates the gradients of ELBO but not the value itself, the denoising approach does not provide a direct way to learn model parameters θ. The denoising method may be best utilised in conjunction with an adversarial algorithm such as a combination of `PC-Den` and `PC-Adv`. The denoising method works better early on in the training when q_ψ and p_θ are very different, and therefore the discrimination task is too easy. Conversely, as q_ψ approaches p_θ, the discriminator can focus its efforts on modelling the residual differences between them rather than trying to model everything about q_θ in isolation as the denoiser in Algorithm 3 does. Warde-Farley & Bengio (2017) already used such combination of adversarial training with a denoising criterion for generative modelling. However, the additional nonlinear transformation before denoising introduced in that work breaks the mathematical connections to KL divergence minimisation.

5. Summary

To summarise, we have presented two main ways to formulate ELBO in terms of logarithmic density ratios r and s. We called these prior-contrastive (PC) and joint-contrastive (JC). We have then described two techniques by which these density ratios, or their gradients, can be estimated if the distributions involved are implicit: adversarial methods (`PC-Adv` and `JC-Adv`, Algorithms 1&2) directly estimate density ratios via logistic regression, denoising methods (`DC-Den` and `JC-Den`, Algorithm 3) estimate gradients of the log densities via denoising autoencoders. We have mentioned that these methods can be combined, and that such combination may improve convergence.

While all of these algorithms can perform variational inference - fitting the variational parameters ψ - not all of them can perform full variational learning of model parameters θ if the model itself is implicitly defined. In Section 7 we outline an idea based on reverse mode differentiation (RMD) idea by Maclaurin et al. (2015) which, giving rise to an algorithm we refer to as `JC-Adv-RMD`, which can in theory perform fully variational inference and learning in a model where all distributions involved are implicit.

The capabilities of algorithms presented here and in related work are summarised in Table 1. The adversarial variational Bayes (Mescheder et al., 2017) is equivalent to PC-Adv, while ALI (Dumoulin et al., 2017) and BiGAN (Donahue et al., 2017) are closely related to JC-Adv. (Sønderby et al., 2017) and (Dosovitskiy & Brox, 2016) are closely related to PC-Adv, although the former solves a limited special case and the latter uses the Jensen-Shannon formulation and hence is not fully variational.

6. Experiments

Several related papers have already demonstrated the success of methods surveyed here on real world datasets, see for example (Mescheder et al., 2017; Dosovitskiy & Brox, 2016) for `PC-Adv`, (Dumoulin et al., 2017; Donahue et al., 2017) for `JC-Adv` and (Sønderby et al., 2017; Warde-Farley & Bengio, 2017) for denoiser-based techniques. Experiments in these papers typically focus on the models' ability to learn θ, and the quality of samples from the learnt generative model $p_\theta(x)$.

As the focus here is on inference rather than learning, the goal of this section is to validate the algorithms' ability to perform inference. To this end, we have devised a simple toy problem loosely based on the "sprinkle" example which exhibits *explaining away* (Pearl, 1988). In our "continuous sprinkler" model, two independent scalar hidden variables z_1 and z_2 are combined nonlinearly to produce a univariate observation x:

$$(z_1, z_2) \sim \mathcal{N}(0, \sigma^2 I_{2 \times 2})$$
$$x \sim \mathrm{EXP}(3 + \max(0, z_1)^3 + \max(0, z_2)^3)$$

Although z_1 and z_2 are a priori independent, the likelihood introduces dependence between the two variables once conditioned on data: either latent variable taking a large value can explain a large observed x. This is an example of explaining away which is an important phenomenon in latent variable models that is known to be hard to model with simple, unimodal distributions.

Column A in Figure 5 illustrates the joint posterior density of z_1 and z_2 for various values of x. The subsequent columns show the posterior approximations by `PC-Adv`, `JC-Adv` and `PC-Den`, respectively. q_θ is implemented as a stochastic generative network where the observation x and Gaussian noise variables ϵ are fed as input to a multilayer perceptron g_ψ. The discriminators and denoisers were implemented as multilayer perceptrons as well. Columns C and D illustrate the limiting behaviour of the discriminator in the `PC-Adv` algorithm: as q_ψ converges to the true posterior, r_ψ is expected resemble the likelihood $p_\theta(x|z)$ up to an additive constant. In `JC-Adv` the discriminator eventually converges to the flat $s_\phi = 0$ solution.

7. Discussion and future work

Are adversaries really needed? When using adversarial techniques for VI, we model the distribution of latent variables rather than observations. The distributions we encounter in VI are usually thought of as simpler than the distribution of observed data, so the question arises whether the flexibility of the adversarial framework is really needed.

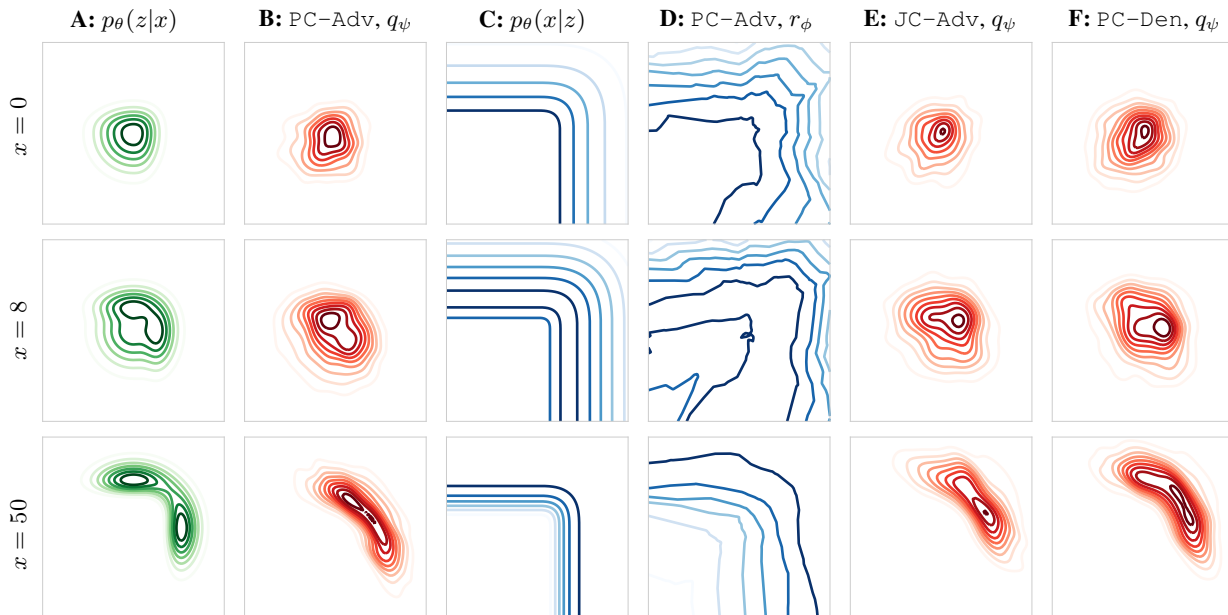**Is the prior-contrastive too much like noise-contrastive?**

Figure 1. Approximate inference in the "continuous sprinkler" toy model from Section 6. Rows correspond to different values of x. Column A shows contours of the real posterior $p_\theta(z|x)$. With increasing values of x the conditional dependence between z_1 and z_2 increases. Columns B, E and F show the approximate posterior q_ψ obtained by PC-Adv (Alg. 1), JC-Adv (Alg. 2) and PC-Den (Alg. 3), respectively. We could not observe any systematic difference in the algorithms' final estimates. The quality of approximation seems to be predominantly influenced by the choice of architecture used to implement q_ψ. Column D shows the final estimate r_ϕ in the PC-Adv algorithm. If the approximate posterior was perfect, r_ϕ should match the likelihood $p_\theta(x|z)$, shown in column C, up to an additive constant.

In the PC-Adv algorithm, the discriminator compares samples from the approximate posterior to the prior, and the prior is often high-dimensional Gaussian noise. Even at convergence, the two distributions the discriminator sees never overlap, and this may slow down training. This can be remedied by observing that as q_ψ converges to the true posterior, the discriminator will converge to the log-likelihood plus constant $r_\phi(x, z) \approx \log p(x|y) + c$. Hence, the task of the discriminator can be made easier by forming an ensemble between a neural network and the actual log-likelihood.

Aren't denoising methods imprecise? The main criticism of denoiser-based methods is that the gradient estimates are imprecise. As (Valpola et al., 2016) pointed out, the optimal denoising function represents the gradients of the noise-corrupted distribution rather than the original, and in practical cases the noise level σ_n may not be small enough for this effect to be negligible. Sønderby et al. (2017) observed that denoiser-based methods can not produce results as sharp as adversarial counterparts. Finally, for the outer loop SGD to work consistently, the gradient estimates provided by the inner loop have to form a conservative vector field. While the Bayes-optimal denoiser function satisfies this, it is unclear to what degree this property is preserved

when using suboptimal denoisers (Im et al., 2016). We believe that an alternative approach based on score matching (Hyvärinen, 2005) - a task intimately related to denoising (Vincent, 2011) - might overcome both of these issues.

How to learn θ? The focus of this note is on variational inference, which is finding ψ. However, it is equally important to think about learning θ. Unfortunately, none of the algorithms presented here allow for fully variational learning of model parameters θ when $p_\theta(x|z)$ is implicit. ALI and BiGAN do provide an algorithm, but as we mentioned, they are not fully variational. We close by highlighting one possible avenue for future work to enable this: differentiating the inner loop of the JC-Adv algorithm via reverse mode differentiation (RMD, Maclaurin et al., 2015). To learn θ via SGD, one only needs an estimate of the gradient $\frac{\partial s_{\psi,\theta}}{\partial \theta}$. We can't compute $s_{\psi,\theta}$, only an approximation s_ϕ which is reached via SGD. Each step of SGD depends implicitly on θ. Following Maclaurin et al. (2015) we can algorithmically differentiate the SGD algorithm in a memory-efficient way to obtain an estimate of the gradient we need for learning θ. We have not validated this approach experimentally, but included it as JC-Adv-RMD in Table 1.

References

Alain, Guillaume and Bengio, Yoshua. What regularized auto-encoders learn from the data-generating distribution? *Journal of Machine Learning Research*, 15(1): 3563–3593, 2014.

Bickel, Steffen, Brückner, Michael, and Scheffer, Tobias. Discriminative learning for differing training and test distributions. In *Proceedings of the 24th International Conference on Machine learning*, pp. 81–88. ACM, 2007.

Donahue, Jeff, Krhenb uhl, Philipp, and Darrell, Trevor. Adversarially feature learning. In *International Conference on Learning Representations*, 2017. URL https://arxiv.org/abs/1605.09782.

Dosovitskiy, Alexey and Brox, Thomas. Generating images with perceptual similarity metrics based on deep networks. In *Advances in Neural Information Processing Systems*, pp. 658–666, 2016.

Dumoulin, Vincent, Belghazi, Ishmael, Poole, Ben, Lamb, Alex, Arjovsky, Martin, Mastropietro, Olivier, and Courville, Aaron. Adversarially learned inference. In *International Conference on Learning Representations*, 2017. URL https://arxiv.org/abs/1606.00704.

Goodfellow, Ian, Pouget-Abadie, Jean, Mirza, Mehdi, Xu, Bing, Warde-Farley, David, Ozair, Sherjil, Courville, Aaron, and Bengio, Yoshua. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 2672–2680, 2014.

Hyvärinen, Aapo. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(Apr):695–709, 2005.

Im, Daniel Jiwoong, Belghazi, Mohamed Ishmael, and Memisevic, Roland. Conservativeness of untied auto-encoders. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pp. 1694–1700. AAAI Press, 2016.

Kanamori, Takafumi, Hido, Shohei, and Sugiyama, Masashi. A least-squares approach to direct importance estimation. *Journal of Machine Learning Research*, 10 (Jul):1391–1445, 2009.

Karaletsos, Theofanis. Adversarial message passing for graphical models. *arXiv preprint arXiv:1612.05048*, 2016. URL https://arxiv.org/abs/1612.05048.

Kingma, Diderik and Welling, Max. Auto-encoding variational Bayes. In *International Conference on Learning Representations*, 2014. URL https://arxiv.org/abs/1312.6114.

Kingma, Diederik P, Salimans, Tim, and Welling, Max. Improving variational inference with inverse autoregressive flow. *arXiv preprint arXiv:1606.04934*, 2016.

Ledig, Christian, Theis, Lucas, Huszár, Ferenc, Caballero, Jose, Cunningham, Andrew, Acosta, Alejandro, Aitken, Andrew, Tejani, Alykhan, Totz, Johannes, Wang, Zehan, et al. Photo-realistic single image super-resolution using a generative adversarial network. *arXiv preprint arXiv:1609.04802*, 2016.

Li, Yingzhen and Turner, Richard E. Rényi divergence variational inference. In *Advances in Neural Information Processing Systems*, pp. 1073–1081, 2016.

Maclaurin, Dougal, Duvenaud, David K, and Adams, Ryan P. Gradient-based hyperparameter optimization through reversible learning. In *International Conference on Machine Learning*, pp. 2113–2122, 2015.

Makhzani, Alireza, Shlens, Jonathon, Jaitly, Navdeep, and Goodfellow, Ian. Adversarial autoencoders. In *International Conference on Learning Representations*, 2016. URL http://arxiv.org/abs/1511.05644.

Mescheder, Lars M., Nowozin, Sebastian, and Geiger, Andreas. Adversarial variational Bayes: Unifying variational autoencoders and generative adversarial networks. *CoRR*, abs/1701.04722, 2017. URL http://arxiv.org/abs/1701.04722.

Minka, Thomas P. Expectation propagation for approximate bayesian inference. In *Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence*, pp. 362–369. Morgan Kaufmann Publishers Inc., 2001.

Mohamed, Shakir and Lakshminarayanan, Balaji. Learning in implicit generative models. *arXiv preprint arXiv:1610.03483*, 2016.

Pearl, Judea. Embracing causality in default reasoning. *Artificial Intelligence*, 35(2):259–271, 1988.

Radford, Alec, Metz, Luke, and Chintala, Soumith. Unsupervised representation learning with deep convolutional generative adversarial networks. In *International Conference on Learning Representations*, 2016. URL https://arxiv.org/abs/1511.06434.

Ranganath, Rajesh, Tran, Dustin, Altosaar, Jaan, and Blei, David. Operator variational inference. In *Advances in Neural Information Processing Systems*, pp. 496–504, 2016.

Rezende, Danilo and Mohamed, Shakir. Variational inference with normalizing flows. In Blei, David and Bach, Francis (eds.), *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pp. 1530–1538. JMLR Workshop and Conference Proceedings, 2015. URL http://jmlr.org/proceedings/papers/v37/rezende15.pdf.

Salimans, Tim, Goodfellow, Ian, Zaremba, Wojciech, Cheung, Vicki, Radford, Alec, and Chen, Xi. Improved techniques for training gans. In *Advances in Neural Information Processing Systems*, pp. 2226–2234, 2016.

Sønderby, Casper Kaae, Raiko, Tapani, Maaløe, Lars, Sønderby, Søren Kaae, and Winther, Ole. Ladder variational autoencoders. In *Advances in Neural Information Processing Systems*, pp. 3738–3746, 2016.

Sønderby, Casper Kaae, Caballero, Jose, Theis, Lucas, Shi, Wenzhe, and Huszár, Ferenc. Amortised MAP inference for image super-resolution. In *International Conference on Learning Representations*, 2017.

Sugiyama, Masashi, Nakajima, Shinichi, Kashima, Hisashi, Buenau, Paul V, and Kawanabe, Motoaki. Direct importance estimation with model selection and its application to covariate shift adaptation. In *Advances in neural information processing systems*, pp. 1433–1440, 2008.

Uehara, Masatoshi, Sato, Issei, Suzuki, Masahiro, Nakayama, Kotaro, and Matsuo, Yutaka. Generative adversarial nets from a density ratio estimation perspective. *arXiv preprint arXiv:1610.02920*, 2016.

Vajda, Steven. *Probabilistic programming*. Academic Press, 2014.

Valpola et al. Learning by denoising part 2. connection between data distribution and denoising function. https://thecuriousaicompany.com/connection-to-g/, 2016. retrieved on 24 February 2017.

Vincent, Pascal. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.

Vincent, Pascal, Larochelle, Hugo, Bengio, Yoshua, and Manzagol, Pierre-Antoine. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th International Conference on Machine learning*, pp. 1096–1103. ACM, 2008.

Warde-Farley, David and Bengio, Yoshua. Improving generative adversarial networks with denoising feature matching. In *International Conference on Learning Representations*, 2017. URL

https://openreview.net/forum?id=S1X7nhsxl&noteId=S1X7nhsxl.

A. Ignoring implicit dependence on ψ

Proof of Eqn. (4):

$$
\frac{\partial}{\partial \psi} \operatorname{KL}\left[q_\psi \| p_\theta\right]\Big|_{\psi=\psi_0} = \frac{\partial}{\partial \psi} \mathbb{E}_{z \sim q_\psi} r_{\theta,\psi}(z)\big|_{\psi=\psi_0}
$$

$$
= \frac{\partial}{\partial \psi} \mathbb{E}_{z \sim q_\psi} \log \frac{q_{\psi_0}(z)}{p(z)}\Bigg|_{\psi=\psi_0} + \frac{\partial}{\partial \psi} \operatorname{KL}\left[q_\psi \| q_{\psi_0}\right]\Big|_{\psi=\psi_0}
$$

$$
= \frac{\partial}{\partial \psi} \mathbb{E}_{z \sim q_\psi} \log \frac{q_{\psi_0}(z)}{p(z)}\Bigg|_{\psi=\psi_0}
$$

$$
= \frac{\partial}{\partial \psi} \mathbb{E}_{z \sim q_\psi} r_{\theta,\psi_0}(z)\Big|_{\psi=\psi_0}, \tag{10}
$$

where the third line is obtained by noting that ψ_0 is a local minimum of $\operatorname{KL}\left[q_\psi \| q_{\psi_0}\right]$, hence the second term in the second line is 0.