
Importance Weighting and Variational Inference

Justin Domke¹ and Daniel Sheldon^{1,2}

¹ College of Information and Computer Sciences, University of Massachusetts Amherst

² Department of Computer Science, Mount Holyoke College

Abstract

Recent work used importance sampling ideas for better variational bounds on likelihoods. We clarify the applicability of these ideas to pure probabilistic inference, by showing the resulting Importance Weighted Variational Inference (IWVI) technique is an instance of augmented variational inference, thus identifying the looseness in previous work. Experiments confirm IWVI’s practicality for probabilistic inference. As a second contribution, we investigate inference with elliptical distributions, which improves accuracy in low dimensions, and convergence in high dimensions.

1 Introduction

Probabilistic modeling is used to reason about the world by formulating a joint model $p(\mathbf{z}, \mathbf{x})$ for unobserved variables \mathbf{z} and observed variables \mathbf{x} , and then querying the posterior distribution $p(\mathbf{z} \mid \mathbf{x})$ to learn about hidden quantities given evidence \mathbf{x} . Common tasks are to draw samples from $p(\mathbf{z} \mid \mathbf{x})$ or compute posterior expectations. However, it is often intractable to perform these tasks directly, so considerable research has been devoted to methods for approximate probabilistic inference.

Variational inference (VI) is a leading approach for approximate inference. In VI, $p(\mathbf{z} \mid \mathbf{x})$ is approximated by a distribution $q(\mathbf{z})$ in a simpler family for which inference is tractable. The process to select q is based on the following decomposition [21, Eqs. 11-12]:

$$\log p(\mathbf{x}) = \underbrace{\mathbb{E}_{q(\mathbf{z})} \log \frac{p(\mathbf{z}, \mathbf{x})}{q(\mathbf{z})}}_{\text{ELBO}[q(\mathbf{z}) \parallel p(\mathbf{z}, \mathbf{x})]} + \underbrace{\text{KL}[q(\mathbf{z}) \parallel p(\mathbf{z} \mid \mathbf{x})]}_{\text{divergence}}. \quad (1)$$

The first term is a lower bound of $\log p(\mathbf{x})$ known as the “evidence lower bound” (ELBO). Selecting q to make the ELBO as big as possible simultaneously obtains a lower bound of $\log p(\mathbf{x})$ that is as tight as possible and drives q close to p in KL-divergence.

The ELBO is closely related to importance sampling. For fixed q , let $R = p(\mathbf{z}, \mathbf{x})/q(\mathbf{z})$ where $\mathbf{z} \sim q$. This random variable satisfies $p(\mathbf{x}) = \mathbb{E} R$, which is the foundation of importance sampling. Similarly, we can write by Jensen’s inequality that $\log p(\mathbf{x}) \geq \mathbb{E} \log R = \text{ELBO}[q \parallel p]$, which is the foundation of modern “black-box” versions of VI (BBVI) [19] in which Monte Carlo samples are used to estimate $\mathbb{E} \log R$, in the same way that IS estimates $\mathbb{E} R$.

Critically, the *only* property VI uses to obtain a lower bound is $p(\mathbf{x}) = \mathbb{E} R$. Further, it is straightforward to see that Jensen’s inequality yields a tighter bound when R is more concentrated about its mean $p(\mathbf{x})$. So, it is natural to consider different random variables with the same mean that are more concentrated, for example the sample average $R_M = \frac{1}{M} \sum_{m=1}^M R_m$. Then, by identical reasoning, $\log p(\mathbf{x}) \geq \mathbb{E} \log R_M$. The last quantity is the objective of importance-weighted auto-encoders [5]; we call it the *importance weighted ELBO (IW-ELBO)*, and the process of selecting q to maximize it *importance-weighted VI (IWVI)*.

However, at this point we should pause. The decomposition in Eq. 1 makes it clear exactly in what sense standard VI, when optimizing the ELBO, makes q close to p . By switching to the

one-dimensional random variable R_M , we derived the IW-ELBO, which gives a tighter bound on $\log p(\mathbf{x})$. For learning applications, this may be all we want. But for probabilistic inference, we are left uncertain exactly in what sense q "is close to" p , and how we should use q to approximate p , say, for computing posterior expectations.

Our first contribution is to provide a new perspective on IWVI by highlighting a precise connection between IWVI and *self-normalized importance sampling* (NIS) [17], which instructs us how to use IWVI for "pure inference" applications. Specifically, IWVI is an instance of augmented VI. Maximizing the IW-ELBO corresponds exactly to minimizing the KL divergence between joint distributions q_M and p_M , where q_M is derived from NIS over a batch of M samples from q , and p_M is the joint distribution obtained by drawing one sample from p and $M - 1$ "dummy" samples from q . This has strong implications for probabilistic inference (as opposed to learning) which is our primary focus. After optimizing q , one should compute posterior expectations using NIS. We show that not only does IWVI significantly tighten bounds on $\log p(\mathbf{x})$, but, by using q this way at test time, it significantly reduces estimation error for posterior expectations.

Previous work has connected IWVI and NIS by showing that the importance weighted ELBO is a lower bound of the ELBO applied to the NIS distribution [6, 16, 2]. Our work makes this relationship precise as an instance of augmented VI, and exactly quantifies the gap between the IW-ELBO and conventional ELBO applied to the NIS distribution, which is a conditional KL divergence.

Our second contribution is to further explore the connection between variational inference and importance sampling by adapting ideas of "defensive sampling" [17] to VI. Defensive importance sampling uses a widely dispersed q distribution to reduce variance by avoiding situations where q places essentially no mass in an area with p has density. This idea is incompatible with regular VI due to its "mode seeking" behavior, but it is quite compatible with IWVI. We show how to use elliptical distributions and reparameterization to achieve a form of defensive sampling with almost no additional overhead to black-box VI (BBVI). "Elliptical VI" provides small improvements over Gaussian BBVI in terms of ELBO and posterior expectations. In higher dimensions, these improvements diminish, but elliptical VI provides significant improvement in the convergence reliability and speed. This is consistent with the notion that using a "defensive" q distribution is advisable when it is not well matched to p (e.g., before optimization has completed).

2 Variational Inference

Consider again the "ELBO decomposition" in Eq. 1. Variational inference maximizes the "evidence lower bound" (ELBO) over q . Since the divergence is non-negative, this tightens a lower-bound on $\log p(\mathbf{x})$. But, of course, since the divergence and ELBO vary by a constant, maximizing the ELBO is equivalent to minimizing the divergence. Thus, variational inference can be thought of as simultaneously solving two problems:

- **"probabilistic inference"** or finding a distribution $q(\mathbf{z})$ that is close to $p(\mathbf{z}|\mathbf{x})$ in KL-divergence.
- **"bounding the marginal likelihood"** or finding a lower-bound on $\log p(\mathbf{x})$.

The first problem is typically used with Bayesian inference: A user specifies a model $p(\mathbf{z}, \mathbf{x})$, observes some data \mathbf{x} , and is interested in the posterior $p(\mathbf{z}|\mathbf{x})$ over the latent variables. While Markov chain Monte Carlo is most commonly for these problems [9, 22], the high computational expense motivates VI [11, 3]. While a user might be interested in any aspect of the posterior, for concreteness, we focus on "posterior expectations", where the user specifies some arbitrary $t(\mathbf{z})$ and wants to approximate $\mathbb{E}_{p(\mathbf{z}|\mathbf{x})} t(\mathbf{z})$.

The second problem is typically used to support maximum likelihood learning. Suppose that $p_\theta(\mathbf{z}, \mathbf{x})$ is some distribution over observed data \mathbf{x} and hidden variables \mathbf{z} . In principle, one would like to set θ to maximize the marginal likelihood over the observed data. When the integral $p_\theta(\mathbf{x}) = \int p_\theta(\mathbf{z}, \mathbf{x}) d\mathbf{z}$ is intractable, one can optimize the lower-bound $\mathbb{E}_{q(\mathbf{z})} \log(p_\theta(\mathbf{z}, \mathbf{x})/q(\mathbf{z}))$ instead [21], over both θ and the parameters of q . This idea has been used to great success recently with variational auto-encoders (VAEs) [10].

3 Importance Weighting

Recently, ideas from importance sampling have been applied to obtain tighter ELBOs for learning in VAEs [5]. We review the idea and then draw novel connections to augmented VI that make it clear how adapt apply these ideas to probabilistic inference.

Take any random variable R such that $\mathbb{E} R = p(\mathbf{x})$, which we will think of as an ‘‘estimator’’ of $p(\mathbf{x})$. Then it’s easy to see via Jensen’s inequality that

$$\log p(\mathbf{x}) = \underbrace{\mathbb{E} \log R}_{\text{bound}} + \underbrace{\mathbb{E} \log \frac{p(\mathbf{x})}{R}}_{\text{looseness}}, \quad (2)$$

where the first term is a lower bound on $\log p(\mathbf{x})$, and the second (non-negative) term is the looseness. The bound will be tight if R is highly concentrated.

While Eq. 2 looks quite trivial, it is a generalization of the ‘‘ELBO’’ decomposition in Eq. 1. To see that, use the random variable

$$R = \omega(\mathbf{z}) = \frac{p(\mathbf{z}, \mathbf{x})}{q(\mathbf{z})}, \quad \mathbf{z} \sim q, \quad (3)$$

which clearly obeys $\mathbb{E} R = p(\mathbf{x})$, and for which Eq. 2 becomes Eq. 1.

The advantage of Eq. 2 over Eq. 1 is increased flexibility: alternative estimators R can give a tighter bound on $\log p(\mathbf{x})$. One natural idea is to draw multiple i.i.d. samples from q and average the estimates as in importance sampling (IS) . This gives the estimator

$$R_M = \frac{1}{M} \sum_{m=1}^M \frac{p(\mathbf{z}_m, \mathbf{x})}{q(\mathbf{z}_m)}, \quad \mathbf{z}_m \sim q. \quad (4)$$

It’s always true that $\mathbb{E} R_M = p(\mathbf{x})$, but the distribution of R_M places less mass near zero for larger M , which leads to a tighter bound (Fig. 1).

This leads to a tighter ‘‘importance weighted ELBO’’ (IW-ELBO) lower bound on $\log p(\mathbf{x})$, namely

$$\text{IW-ELBO}_M [q(\mathbf{z}) \| p(\mathbf{z}, \mathbf{x})] := \mathbb{E}_{q(\mathbf{z}_{1:M})} \log \frac{1}{M} \sum_{m=1}^M \frac{p(\mathbf{z}_m, \mathbf{x})}{q(\mathbf{z}_m)}, \quad (5)$$

where $\mathbf{z}_{1:M}$ is a shorthand for $(\mathbf{z}_1, \dots, \mathbf{z}_M)$ and $q(\mathbf{z}_{1:M}) = q(\mathbf{z}_1) \cdots q(\mathbf{z}_M)$. This bound was first introduced by Burda et al. [5] in the context of supporting maximum likelihood learning of a variational auto-encoder.

3.1 A generative process for the importance weighted ELBO

While Eq. 2 makes clear that optimizing the IW-ELBO tightens a bound on $\log p(\mathbf{x})$, it isn’t obvious what connection this has to probabilistic inference. Is there some divergence that is being minimized? Theorem 1 shows this can be understood by constructing ‘‘augmented’’ distributions $p_M(\mathbf{z}_{1:M}, \mathbf{x})$ and $q_M(\mathbf{z}_{1:M})$ and then applying the ELBO decomposition in Eq. 1 to the joint distributions.

Theorem 1 (IWVI). *Let $q_M(\mathbf{z}_{1:M})$ be the density of the generative process described by Alg. 1, which is based on self-normalized importance sampling over a batch of M samples from q . Let $p_M(\mathbf{z}_{1:M}, \mathbf{x}) = p(\mathbf{z}_1, \mathbf{x})q(\mathbf{z}_{2:M})$ be the density obtained by drawing \mathbf{z}_1 and \mathbf{x} from p and drawing the ‘‘dummy’’ samples $\mathbf{z}_{2:M}$ from q . Then*

$$q_M(\mathbf{z}_{1:M}) = \frac{p_M(\mathbf{z}_{1:M}, \mathbf{x})}{\frac{1}{M} \sum_{m=1}^M \omega(\mathbf{z}_m)}. \quad (6)$$

Further, the ELBO decomposition in Eq. 1 applied to q_M and p_M is

$$\log p(\mathbf{x}) = \text{IW-ELBO}_M [q(\mathbf{z}) \| p(\mathbf{z}, \mathbf{x})] + \text{KL}[q_M(\mathbf{z}_{1:M}) \| p_M(\mathbf{z}_{1:M} | \mathbf{x})]. \quad (7)$$

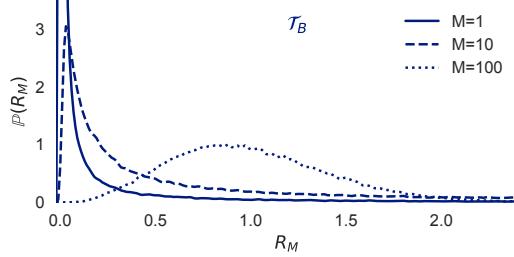


Figure 1: How the density of R_M changes with M . (Distribution and setting as in Fig. 2.)

Algorithm 1 A generative process for $q_M(\mathbf{z}_{1:M})$

1. Draw $\hat{\mathbf{z}}_1, \hat{\mathbf{z}}_2, \dots, \hat{\mathbf{z}}_M$ independently from $q(\mathbf{z})$.
2. Choose $m \in \{1, \dots, M\}$ with probability $\frac{\omega(\hat{\mathbf{z}}_m)}{\sum_{m'=1}^M \omega(\hat{\mathbf{z}}_{m'})}$.
3. Set $\mathbf{z}_1 = \hat{\mathbf{z}}_m$ and $\mathbf{z}_{2:M} = \hat{\mathbf{z}}_{-m}$ and return $\mathbf{z}_{1:M}$.

We will call the process of maximizing the IW-ELBO “Importance Weighted Variational Inference” (IWVI). (Burda et al. used “Importance Weighted Auto-encoder” for optimizing Eq. 5 as a bound on the likelihood of a variational auto-encoder, but this terminology ties the idea to a particular model, and is not suggestive of the probabilistic inference setting.)

The generative process for q_M in Alg. 1 is very similar to self-normalized importance sampling. The usual NIS distribution draws a batch of size M , and then “selects” a single variable with probability in proportion to its importance weight. NIS is exactly equivalent to the marginal distribution $q_M(\mathbf{z}_1)$. The generative process for $q_M(\mathbf{z}_{1:M})$ additionally keeps the *unselected* variables and relabels them as $\mathbf{z}_{2:M}$.

Previous work [6, 2, 16, 12] investigated a similar connection between NIS and the importance-weighted ELBO. In our notation, they showed that

$$\log p(\mathbf{x}) \geq \text{ELBO}[q_M(\mathbf{z}_1) \| p(\mathbf{z}_1, \mathbf{x})] \geq \text{IW-ELBO}_M[q(\mathbf{z}) \| p(\mathbf{z}, \mathbf{x})]. \quad (8)$$

That is, they showed that the IW-ELBO *lower bounds* the ELBO between the NIS distribution and p , without quantifying the gap in the second inequality. Our result makes it clear exactly what KL-divergence is being minimized by maximizing the IW-ELBO and in what sense doing this makes q “close to” p . As a corollary, we also quantify the gap in the inequality above (see Thm. 2 below).

A recent decomposition [12, Claim 1] is related to Thm. 1, but based on different augmented distributions q_M^{IS} and p_M^{IS} . This result is fundamentally different in that it holds q_M^{IS} “fixed” to be an independent sample of size M from q , and modifies p_M^{IS} so its marginals approach q . This does not inform inference. Contrast this with our result, where $q_M(\mathbf{z}_1)$ gets closer and closer to $p(\mathbf{z}_1 | \mathbf{x})$, and can be used for probabilistic inference. See appendix (Section A.3.2) for details.

Identifying the precise generative process is useful if IWVI will be used for general probabilistic queries, which is a focus of our work, and, to our knowledge, has not been investigated before. For example, the expected value of $t(\mathbf{z})$ can be approximated as

$$\mathbb{E}_{p(\mathbf{z}|\mathbf{x})} t(\mathbf{z}) = \mathbb{E}_{p_M(\mathbf{z}_1|\mathbf{x})} t(\mathbf{z}_1) \approx \mathbb{E}_{q_M(\mathbf{z}_1)} t(\mathbf{z}_1) = \mathbb{E}_{q(\mathbf{z}_{1:M})} \frac{\sum_{m=1}^M \omega(\mathbf{z}_m) t(\mathbf{z}_m)}{\sum_{m=1}^M \omega(\mathbf{z}_m)}. \quad (9)$$

The final equality is established by Lemma 4 in the Appendix. Here, the inner approximation is justified since IWVI minimizes the joint divergence between $q_M(\mathbf{z}_{1:M})$ and $p_M(\mathbf{z}_{1:M}|\mathbf{x})$. However, this is *not* equivalent to minimizing the divergence between $q_M(\mathbf{z}_1)$ and $p_M(\mathbf{z}_1|\mathbf{x})$, as the following result shows.

Theorem 2. *The marginal and joint divergences relevant to IWVI are related by*

$$\text{KL}[q_M(\mathbf{z}_{1:M}) \| p_M(\mathbf{z}_{1:M}|\mathbf{x})] = \text{KL}[q_M(\mathbf{z}_1) \| p(\mathbf{z}_1|\mathbf{x})] + \text{KL}[q_M(\mathbf{z}_{2:M}|\mathbf{z}_1) \| q(\mathbf{z}_{2:M})].$$

As a consequence, the gap in the first inequality of Eq. 8 is exactly $\text{KL}[q_M(\mathbf{z}_1) \| p(\mathbf{z}_1|\mathbf{x})]$ and the gap in the second inequality is exactly $\text{KL}[q_M(\mathbf{z}_{2:M}|\mathbf{z}_1) \| q(\mathbf{z}_{2:M})]$.

The first term is the divergence between the marginal of q_M , i.e., the “standard” NIS distribution, and the posterior. In principle, this is exactly the divergence we would like to minimize to justify Eq. 9. However, the second term is not zero since the selection phase in Alg. 1 leaves $\mathbf{z}_{2:M}$ distributed differently under q_M than under q . Since this term is irrelevant to the quality of the approximation in Eq. 9, IWVI truly minimizes an upper-bound. Thus, IWVI can be seen as an instance of auxiliary variational inference [1] where a joint divergence upper-bounds the divergence of interest.

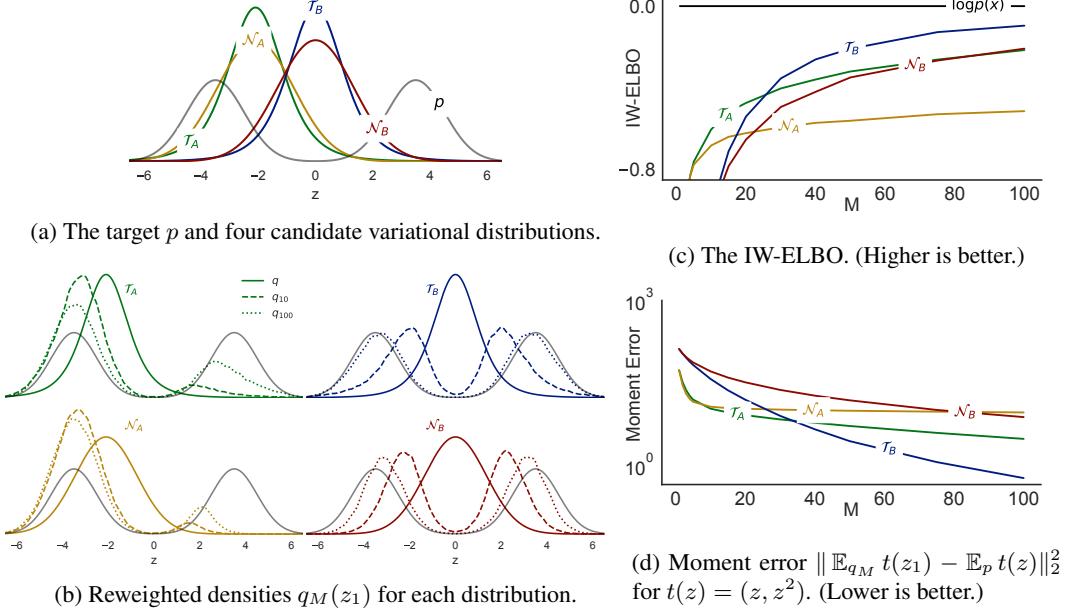


Figure 2: Two Gaussian (\mathcal{N}) and two Student-T (\mathcal{T}) variational distributions, all with constant variance and one of two means (A or B). For $M = 1$ it is better to use a mean closer to one mode of p . For large M , a mean in the center is superior, and the heavy tails of the Student T lead to better approximation of p and better performance both in terms of IW-ELBO and moment error.

4 Importance Sampling Variance

This section considers the family for the variational distribution. For small M , the mode-seeking behavior of VI will favor weak tails, while for large M , variance reduction provided by importance weighting will favor wider tails.

The most common variational distribution is the Gaussian. One explanation for this is the Bayesian central limit theorem, which, in many cases, guarantees that the posterior is asymptotically Gaussian. Another is that it's "safest" to have weak tails: since the objective is $\mathbb{E} \log R$, small values of R are most harmful. So, VI wants to avoid cases where $q(\mathbf{z}) \gg p(\mathbf{z}, \mathbf{x})$, which is difficult if q is heavy-tailed. (This is the "mode-seeking" behavior of the KL-divergence [23].)

With IWVI, the situation changes. Asymptotically in M , R_M in Eq. 4 concentrates around $p(\mathbf{x})$, and so it is the variance of R_M that matters, as formalized in the following result.

Theorem 3. *For large M , the looseness of the IW-ELBO is given by the variance of R . Formally, if there exists some $\alpha > 0$ such that $\mathbb{E} |R - p(\mathbf{x})|^{2+\alpha} < \infty$ and $\limsup_{M \rightarrow \infty} \mathbb{E}[1/R_M] < \infty$, then*

$$\lim_{M \rightarrow \infty} M \underbrace{\left(\log p(\mathbf{x}) - \text{IW-ELBO}_M [q(\mathbf{z}) \| p(\mathbf{z}, \mathbf{x})] \right)}_{\text{KL}[q_M \| p_M]} = \frac{\mathbb{V}[R]}{2p(\mathbf{x})^2}.$$

Maddison et al. [13] give a related result. Their Proposition 1 applied to R_M gives the same conclusion (after an argument based on the Marcinkiewicz-Zygmund inequality; see appendix) but requires the sixth central moment to exist, whereas we require only existence of $\mathbb{E} |R - p(\mathbf{x})|^{2+\alpha}$ for any $\alpha > 0$. The \limsup assumption on $\mathbb{E} 1/R_M$ is implied by assuming that $\mathbb{E} 1/R_M < \infty$ for any finite M (or for R itself). Rainforth et al. [18, Theorem 1 in Appendix] provide a related asymptotic for errors in gradient variance, assuming at least the third moment exists.

Directly minimizing the variance of R is equivalent to minimizing the χ^2 divergence between $q(\mathbf{z})$ and $p(\mathbf{z}|\mathbf{x})$, as explored by Dieng et al. [7]. Overdispersed VI [20] reduces the variance of score-function estimators using heavy-tailed distributions.

The quantity inside the parentheses on the left-hand side is exactly the KL-divergence between q_M and p_M in Eq. 7, and accordingly, even for constant q , this divergence asymptotically decreases at a $1/M$ rate.

The variance of R is a well-explored topic in traditional importance sampling. Here the situation is reversed from traditional VI—since R is non-negative, it is very *large* values of R that can cause high variance, which occurs when $q(\mathbf{z}) \ll p(\mathbf{z}, \mathbf{x})$. The typical recommendation is “defensive sampling” or using a widely-dispersed proposal [17]. For these reasons, we believe that the best form for q will vary depending on the value of M . Figure 1 explores a simple example of this in 1-D.

5 Elliptical Distributions

Elliptical distributions are a generalization of Gaussians that includes the Student-T, Cauchy, scale-mixtures of Gaussians, and many others. The following short review assumes a density function exists, enabling a simpler presentation than the typical one based on characteristic functions [8].

We first describe the special case of spherical distributions. Take some density $\rho(r)$ for a non-negative r with $\int_0^\infty \rho(r) = 1$. Define the **spherical random variable** ϵ corresponding to ρ as

$$\epsilon = r\mathbf{u}, \quad r \sim \rho, \quad \mathbf{u} \sim S, \quad (10)$$

where S represents the uniform distribution over the unit sphere in d dimensions. The density of ϵ can be found using two observations. First, it is constant for all ϵ with a fixed radius $\|\epsilon\|$. Second, if if $q_\epsilon(\epsilon)$ is integrated over $\{\epsilon : \|\epsilon\| = r\}$ the result must be $\rho(r)$. Using these, it is not hard to show that the density must be

$$q_\epsilon(\epsilon) = g(\|\epsilon\|_2^2), \quad g(a) = \frac{1}{S_{d-1}a^{(d-1)/2}} \rho(\sqrt{a}), \quad (11)$$

where S_{d-1} is the surface area of the unit sphere in d dimensions (and so $S_{d-1}a^{(d-1)/2}$ is the surface area of the sphere with radius a) and g is the **density generator**.

Generalizing, this, take some positive definite matrix Σ and some vector μ . Define the **elliptical random variable** \mathbf{z} corresponding to ρ, Σ , and μ by

$$\mathbf{z} = rA^\top \mathbf{u} + \mu, \quad r \sim \rho, \quad \mathbf{u} \sim S, \quad (12)$$

where A is some matrix such that $A^\top A = \Sigma$. Since \mathbf{z} is an affine transformation of ϵ , it is not hard to show by the “Jacobian determinant” formula for changes of variables that the density of \mathbf{z} is

$$q(\mathbf{z}|\mu, \Sigma) = \frac{1}{|\Sigma|^{1/2}} g\left((\mathbf{z} - \mu)^\top \Sigma^{-1} (\mathbf{z} - \mu)\right), \quad (13)$$

where g is again as in Eq. 11. The mean and covariance are $\mathbb{E}[\mathbf{z}] = \mu$, and $\mathbb{C}[\mathbf{z}] = (\mathbb{E}[r^2]/d) \Sigma$.

For some distributions, $\rho(r)$ can be found from observing that r has the same distribution as $\|\epsilon\|$. For example, with a Gaussian, $r^2 = \|\epsilon\|^2$ is a sum of d i.i.d. squared Gaussian variables, and so, by definition, $r \sim \chi_d$.

6 Reparameterization and Elliptical Distributions

Suppose the variational family $q(\mathbf{z}|w)$ has parameters w to optimize during inference. The reparameterization trick is based on finding some density $q_\epsilon(\epsilon)$ independent of w and a “reparameterization function” $\mathcal{T}(\epsilon; w)$ such that $\mathcal{T}(\epsilon; w)$ is distributed as $q(\mathbf{z}|w)$. Then, the ELBO can be re-written as

$$\text{ELBO}[q(\mathbf{z}|w) \| p(\mathbf{z}, \mathbf{x})] = \mathbb{E}_{q_\epsilon(\epsilon)} \log \frac{p(\mathcal{T}(\epsilon; w), \mathbf{x})}{q(\mathcal{T}(\epsilon; w)|w)}.$$

The advantage of this formulation is that the expectation is independent of w . Thus, computing the gradient of the term inside the expectation for a random ϵ gives an unbiased estimate of the gradient. By far the most common case is the multivariate Gaussian distribution, in which case the base density $q_\epsilon(\epsilon)$ is just a standard Gaussian and for some A_w such that $A_w^\top A_w = \Sigma_w$,

$$\mathcal{T}(\epsilon; w) = A_w^\top \epsilon + \mu_w. \quad (14)$$

6.1 Elliptical Reparameterization

To understand Gaussian reparameterization from the perspective of elliptical distributions, note the similarity of Eq. 14 to Eq. 12. Essentially, the reparameterization in Eq. 14 combines r and \mathbf{u} into $\epsilon = r\mathbf{u}$. This same idea can be applied more broadly: for any elliptical distribution, *provided the density generator g is independent of w* , the reparameterization in Eq. 14 will be valid, provided that ϵ comes from the corresponding spherical distribution.

While this independence is true for Gaussians, this is not the case for other elliptical distributions. If ρ_w itself is a function of w , Eq. 14 must be generalized. In that case, think of the generative process (for v sampled uniformly from $[0, 1]$)

$$\mathcal{T}(\mathbf{u}, v; w) = F_w^{-1}(v) A_w^T \mathbf{u} + \boldsymbol{\mu}_w, \quad (15)$$

where $F_w^{-1}(v)$ is the inverse CDF corresponding to the distribution $\rho_w(r)$. Here, we should think of the vector (\mathbf{u}, v) playing the role of ϵ above, and the base density as $q_{\mathbf{u}, v}(\mathbf{u}, v)$ being a spherical density for \mathbf{u} and a uniform density for v .

To calculate derivatives with respect to w , backpropagation through A_w and $\boldsymbol{\mu}_w$ is simple using any modern autodiff system. So, if the inverse CDF F_w^{-1} has a closed-form, autodiff can be directly applied to Eq. 15. If the inverse CDF does not have a simple closed-form, the following section shows that only the CDF is actually needed, provided that one can at least sample from $\rho(r)$.

6.2 Dealing CDFs without closed-form inverses

For many distributions ρ , the inverse CDF may not have a simple closed form, yet highly efficient samplers still exist (most commonly custom rejection samplers with very high acceptance rates). In such cases, one can still achieve the *effect* of Eq. 15 on a random v using only the CDF (not the inverse). The idea is to first directly generate $r \sim \rho_w$ using the specialized sampler, and only then find the corresponding $v = F_w(r)$ using the closed-form CDF. To understand this, observe that if $r \sim \rho$ and $v \sim \text{Uniform}[0, 1]$, then the pairs $(r, F_w(r))$ and $(F_w^{-1}(v), v)$ are identically distributed. Then, via the implicit function theorem, $\nabla_w F_w^{-1}(v) = -\nabla_w F_w(r)/\nabla_r F_w(r)$. All gradients can then be computed by “pretending” that one had started with v and computed r using the inverse CDF.

6.3 Student T distributions

The following experiments will consider student T distributions. The spherical T distribution can be defined as $\epsilon = \sqrt{\nu}\delta/s$ where $\delta \sim \mathcal{N}(0, I)$ and $s \sim \chi_\nu$ [8]. Equivalently, write $r = \|\epsilon\| = \sqrt{\nu t}/s$ with $t \sim \chi_d$. This shows that r is the ratio of two independent χ variables, and thus determined by an F-distribution, the CDF of which could be used directly in Eq. 15. We found a slightly “bespoke” simplification helpful. As there is no need for gradients with respect to d (which is fixed), we represent ϵ as $\epsilon = (\sqrt{\nu t}/s)\mathbf{u}$, leading to reparameterizing the elliptical T distribution as

$$\mathcal{T}(\mathbf{u}, t, v; w) = \frac{\sqrt{\nu t}}{F_\nu^{-1}(v)} A_w^T \mathbf{u} + \boldsymbol{\mu}_w,$$

where F_ν is the CDF for the χ_ν distribution. This is convenient since the CDF of the χ distribution is more widely available than that of the F distribution.

7 Experiments

All the following experiments compare “E-IWVI” using student T distributions to “IWVI” using Gaussians. Regular “VI” is equivalent to IWVI with $M = 1$.

We consider experiments on three distributions. In the first two, a computable $\log p(\mathbf{x})$ enables estimation of the KL-divergence and computable true mean and variance of the posterior enable a precise evaluation of test integral estimation. On these, we used a fixed set of $10,000 \times M$ random inputs to \mathcal{T} and optimized using batch L-BFGS, avoiding heuristic tuning of a learning rate sequence.

A first experiment considered random Dirichlet distributions $p(\boldsymbol{\theta}|\boldsymbol{\alpha})$ over the probability simplex in K dimensions, $\boldsymbol{\theta} \in \Delta^K$. Each parameter α_k is drawn i.i.d. from a Gamma distribution with a shape parameter of 10. Since this density is defined only over the probability simplex, we borrow from Stan

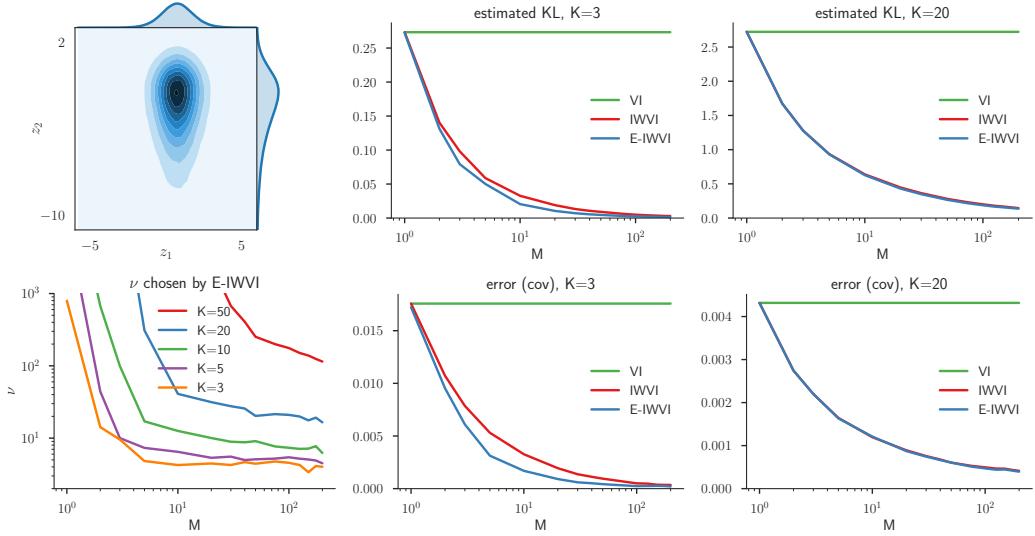


Figure 3: Random Dirichlets, averaged over 20 repetitions. Top left shows an example posterior for $K = 3$. The test-integral error is $\|\mathbb{C}[\theta] - \hat{\mathbb{C}}[\theta]\|_F$ where $\hat{\mathbb{C}}$ is the empirical covariance of samples drawn from $q_M(z_1)$ and then transformed to Δ^K . In all cases, IWVI is able to reduce the error of VI to negligible levels. E-IWVI provides an accuracy benefit in low dimensions but little when $K = 20$.

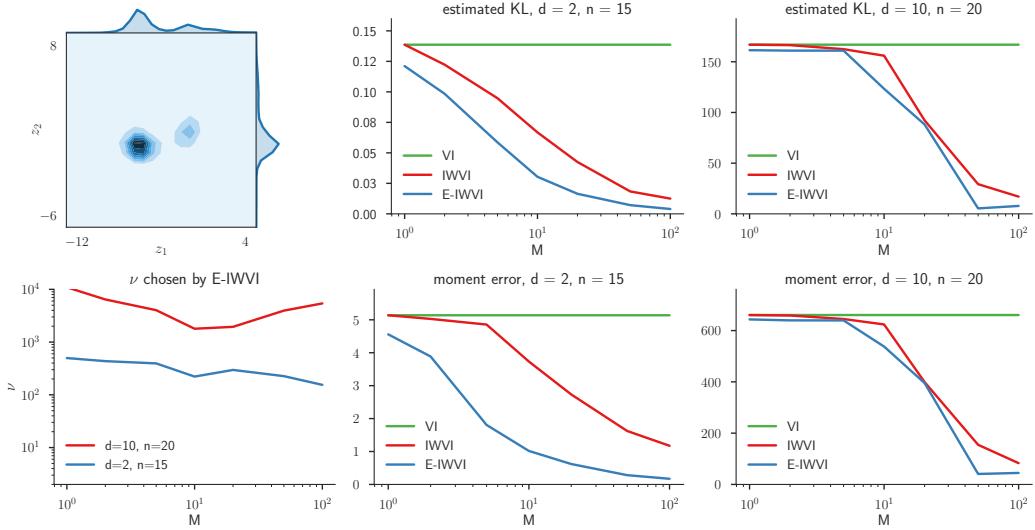


Figure 4: Clutter Distributions, averaged over 50 repetitions. The error shows the error in the estimated second moment $\mathbb{E}[\mathbf{z}\mathbf{z}^T]$. IWVI reduces the errors of VI by orders of magnitude. E-IWVI provides a diminishing benefit in higher dimensions.

the strategy of transforming to an unconstrained $\mathbf{z} \in \mathbb{R}^{K-1}$ space via a stick-breaking process [22]. To compute test integrals over variational distributions, the reverse transformation is used. Results are shown in Fig. 3.

A second experiment uses Minka’s “clutter” model [15]: $\mathbf{z} \in \mathbb{R}^d$ is a hidden object location, and $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ is a set of n noisy observations, with $p(\mathbf{z}) = \mathcal{N}(\mathbf{z}; \mathbf{0}, 100I)$ and $p(\mathbf{x}_i | \mathbf{z}) = 0.25 \mathcal{N}(\mathbf{x}_i; \mathbf{z}, I) + 0.75 \mathcal{N}(\mathbf{x}_i; \mathbf{0}, 10I)$. The posterior $p(\mathbf{z} | \mathbf{x})$ is a mixture of 2^n Gaussians, for which we can do exact inference for moderate n . Results are shown in Fig. 4.

Finally, we considered a (non-conjugate) logistic regression model with a Cauchy prior with a scale of 10, using stochastic gradient descent with various step sizes. On these higher dimensional problems, we found that when the step-size was perfectly tuned and optimization had many iterations, both methods performed similarly in terms of the IW-ELBO. E-IWVI never performed worse, and

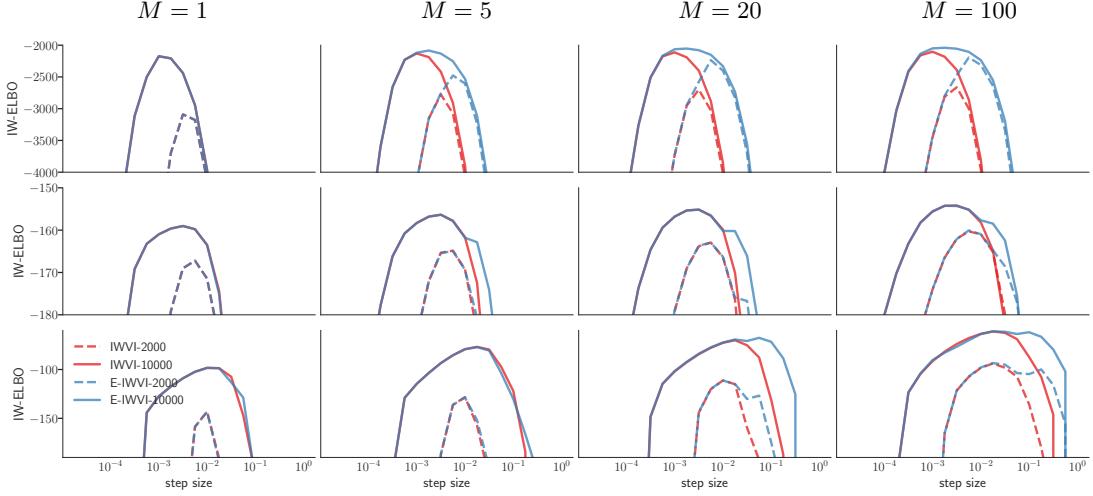


Figure 5: Logistic regression comparing IWVI (red) and E-IWVI (blue) with various M and step sizes. The IW-ELBO is shown after 2,000 (dashed lines) and 10,000 (solid) iterations. A larger M consistently improves both methods. E-IWVI converges more reliably, particularly on higher-dimensional data. From top: Madelon ($d = 500$) Sonar ($d = 60$), Mushrooms ($d = 112$).

sometimes performed very slightly better. E-IWVI exhibited superior convergence behavior and was easier to tune, as illustrated in Fig. 5, where E-IWVI converges at least as well as IWVI for *all* step sizes. We suspect this is because when w is far from optimal, both the IW-ELBO and gradient variance is better with E-IWVI.

Acknowledgements

We thank Tom Rainforth for insightful comments regarding asymptotics and Theorem 3. This material is based upon work supported by the National Science Foundation under Grant No. 1617533.

References

- [1] Felix V. Agakov and David Barber. An auxiliary variational method. In *Neural Information Processing*, Lecture Notes in Computer Science, pages 561–566. Springer, Berlin, Heidelberg, 2004.
- [2] Philip Bachman and Doina Precup. Training deep generative models: Variations on a theme. In *NIPS Workshop: Advances in Approximate Bayesian Inference*, 2015.
- [3] Robert Bamler, Cheng Zhang, Manfred Opper, and Stephan Mandt. Perturbative black box variational inference. In *NIPS*, 2017.
- [4] Peter J Bickel and Kjell A Doksum. *Mathematical statistics: basic ideas and selected topics, volume I*, volume 117. CRC Press, 2015.
- [5] Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov. Importance weighted autoencoders. 2015.
- [6] Chris Cremer, Quaid Morris, and David Duvenaud. Reinterpreting importance-weighted autoencoders. 2017.
- [7] Adji Boussou Dieng, Dustin Tran, Rajesh Ranganath, John Paisley, and David Blei. Variational inference via χ upper bound minimization. In *NIPS*, pages 2729–2738. 2017.
- [8] Kaitai Fang, Samuel Kotz, and Kai Wang Ng. *Symmetric multivariate and related distributions*. Number 36 in Monographs on statistics and applied probability. Chapman and Hall, 1990.

- [9] W. R. Gilks, A. Thomas, and D. J. Spiegelhalter. A language and program for complex bayesian modelling. 43(1):169–177, 1994.
- [10] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*.
- [11] Alp Kucukelbir, Dustin Tran, Rajesh Ranganath, Andrew Gelman, and David M. Blei. Automatic differentiation variational inference. 18(14):1–45, 2017.
- [12] Tuan Anh Le, Maximilian Igl, Tom Rainforth, Tom Jin, and Frank Wood. Auto-Encoding Sequential Monte Carlo. In *ICLR*, 2018.
- [13] Chris J Maddison, John Lawson, George Tucker, Nicolas Heess, Mohammad Norouzi, Andriy Mnih, Arnaud Doucet, and Yee Teh. Filtering variational objectives. In *NIPS*, pages 6576–6586. 2017.
- [14] Józef Marcinkiewicz and Antoni Zygmund. Quelques théoremes sur les fonctions indépendantes. *Fund. Math.*, 29:60–90, 1937.
- [15] Minka, Thomas. Expectation propagation for approximate bayesian inference. In *UAI*, 2001.
- [16] Christian A. Naesseth, Scott W. Linderman, Rajesh Ranganath, and David M. Blei. Variational sequential monte carlo. In *AISTATS*, 2018.
- [17] Art Owen. *Monte Carlo theory, methods and examples*. 2013.
- [18] Tom Rainforth, Adam R. Kosiorek, Tuan Anh Le, Chris J. Maddison, Maximilian Igl, Frank Wood, and Yee Whye Teh. Tighter variational bounds are not necessarily better.
- [19] Rajesh Ranganath, Sean Gerrish, and David M. Blei. Black box variational inference. In *AISTATS*, 2014.
- [20] Francisco J. R. Ruiz, Michalis K. Titsias, and David M. Blei. Overdispersed black-box variational inference. In *UAI*, 2016.
- [21] L. K. Saul, T. Jaakkola, and M. I. Jordan. Mean field theory for sigmoid belief networks. *Journal of Artificial Intelligence Research*, 4:61–76, 1996.
- [22] Stan Development Team. Modeling language user’s guide and reference manual, version 2.17.0, 2017.
- [23] Tom Minka. Divergence measures and message passing. 2005.

A Appendix

A.1 Additional Experimental Results

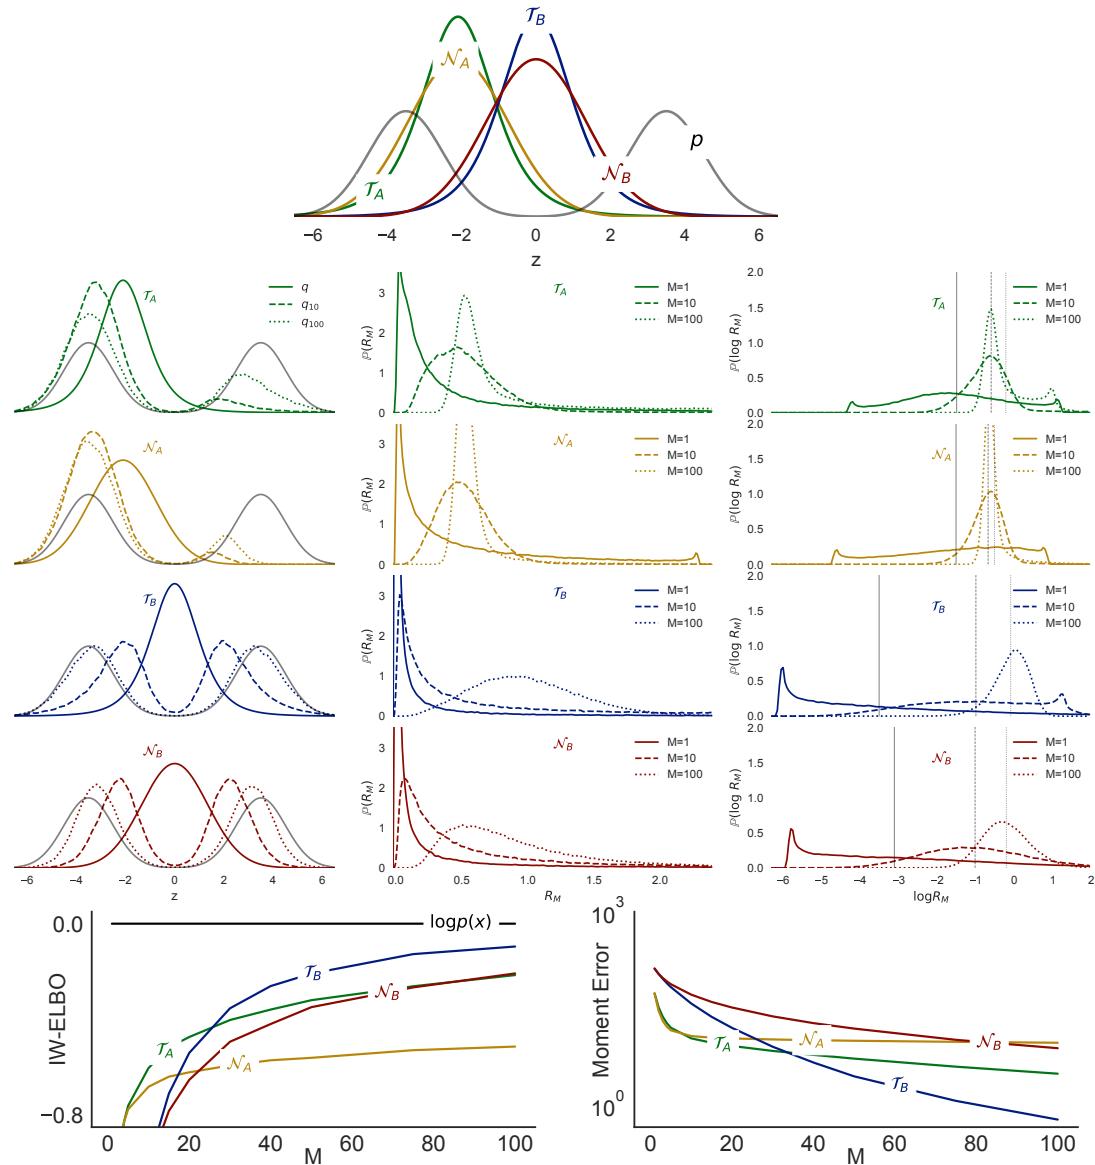


Figure 6: More figures corresponding to the 1-D example.

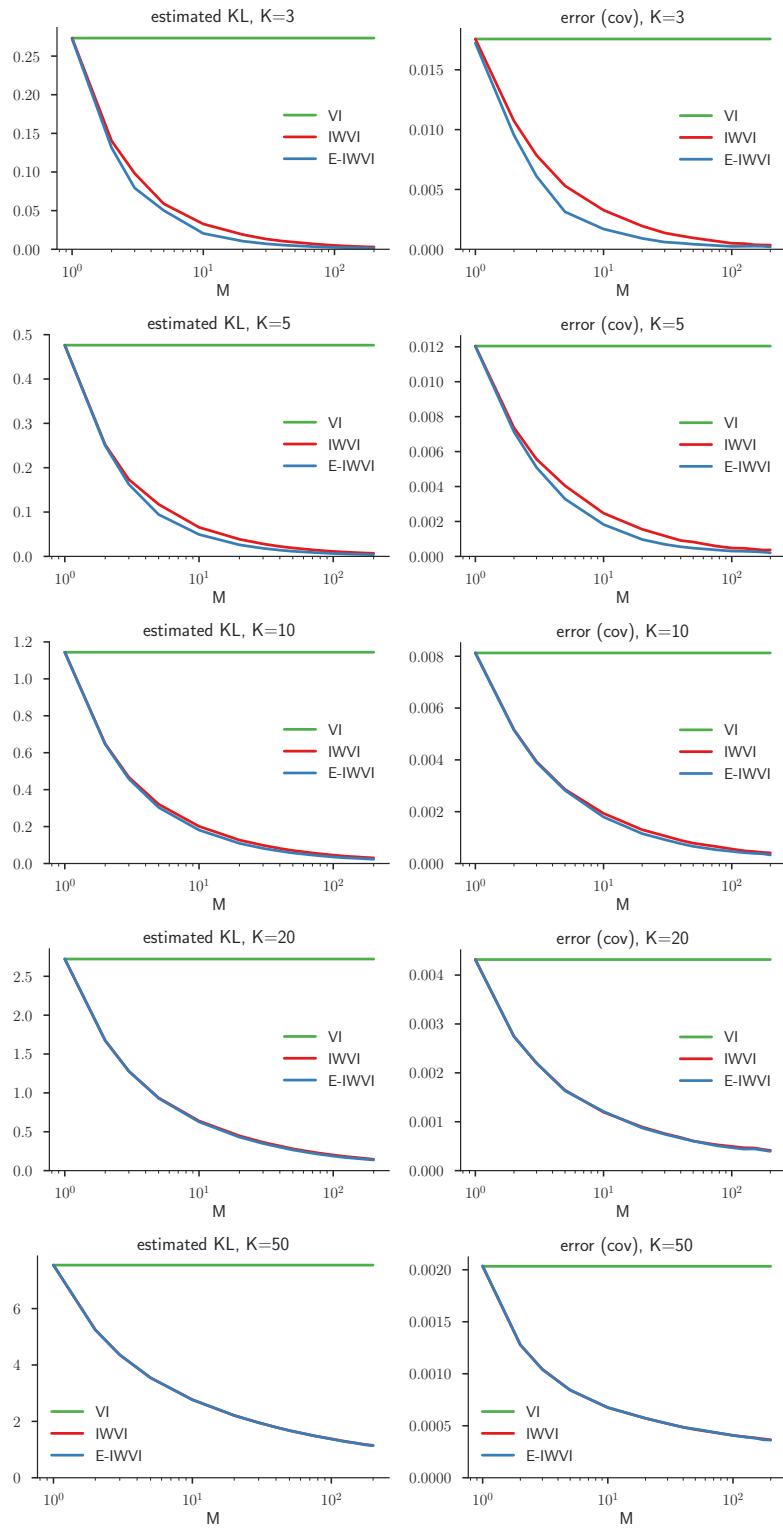


Figure 7: More Results on Inference with Dirichlet distributions

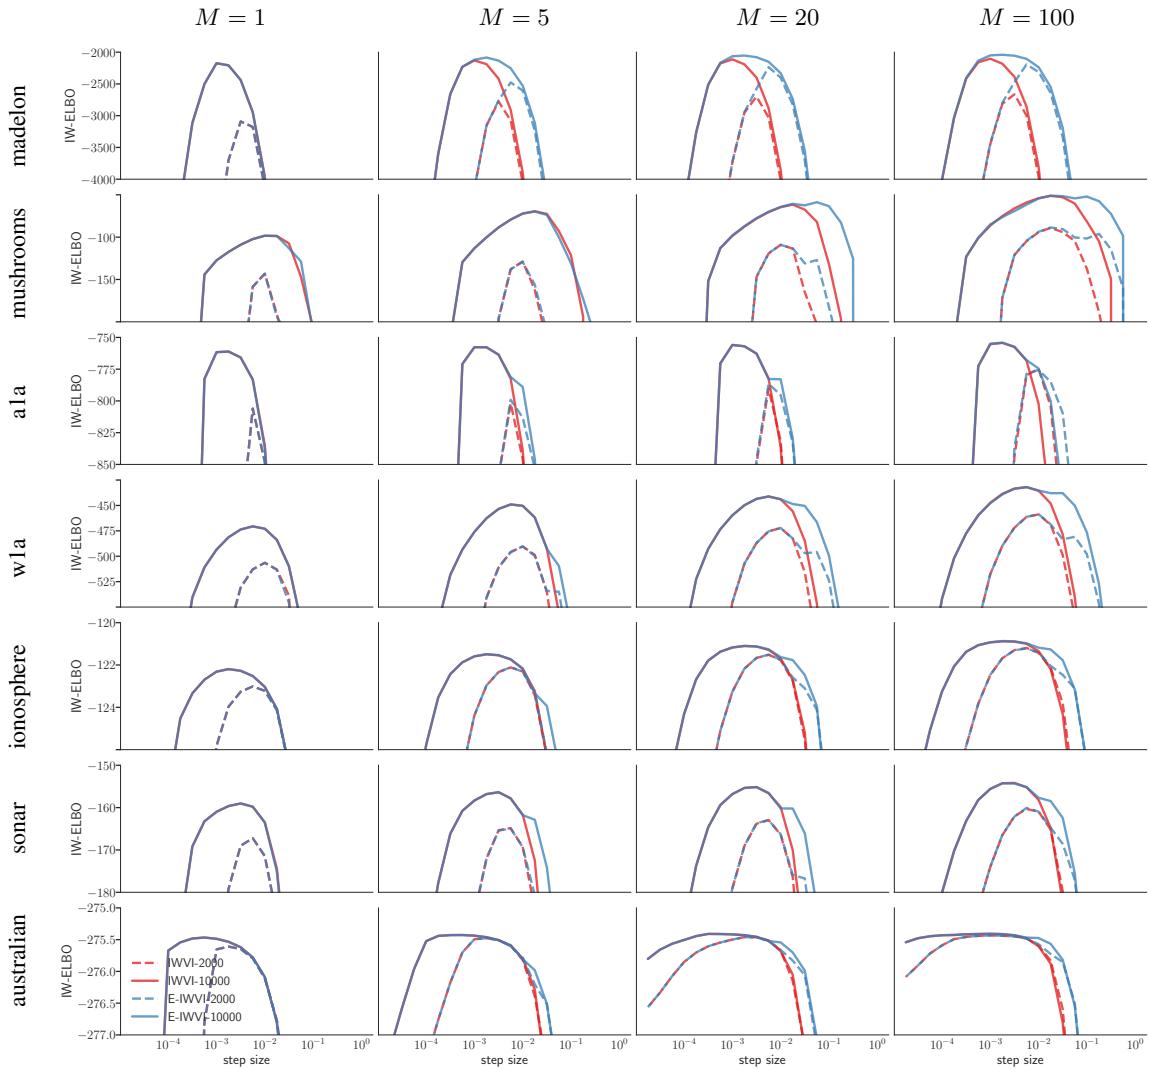


Figure 8: Logistic regression experiments (as in Fig. 5) on more datasets.

A.2 Proofs for Section 3

Theorem 1 (IWVI). Let $q_M(\mathbf{z}_{1:M})$ be the density of the generative process described by Alg. 1, which is based on self-normalized importance sampling over a batch of M samples from q . Let $p_M(\mathbf{z}_{1:M}, \mathbf{x}) = p(\mathbf{z}_1, \mathbf{x})q(\mathbf{z}_{2:M})$ be the density obtained by drawing \mathbf{z}_1 and \mathbf{x} from p and drawing the “dummy” samples $\mathbf{z}_{2:M}$ from q . Then

$$q_M(\mathbf{z}_{1:M}) = \frac{p_M(\mathbf{z}_{1:M}, \mathbf{x})}{\frac{1}{M} \sum_{m=1}^M \omega(\mathbf{z}_m)}. \quad (6)$$

Further, the ELBO decomposition in Eq. 1 applied to q_M and p_M is

$$\log p(\mathbf{x}) = \text{IW-ELBO}_M [q(\mathbf{z}) \| p(\mathbf{z}, \mathbf{x})] + \text{KL}[q_M(\mathbf{z}_{1:M}) \| p_M(\mathbf{z}_{1:M} | \mathbf{x})]. \quad (7)$$

Proof. For the density q_M , define the distribution

$$\begin{aligned} q_M(\hat{\mathbf{z}}_{1:M}, \mathbf{z}_{1:M}, h) &= q_M(\hat{\mathbf{z}}_{1:M}) q_M(h | \hat{\mathbf{z}}_{1:M}) q_M(\mathbf{z}_{1:M} | \hat{\mathbf{z}}_{1:M}, h) \\ q_M(\hat{\mathbf{z}}_{1:M}) &= q(\hat{\mathbf{z}}_{1:M}) \\ q_M(h | \hat{\mathbf{z}}_{1:M}) &= \frac{p(\hat{\mathbf{z}}_h, x) / q(\hat{\mathbf{z}}_h)}{\sum_{m=1}^M p(\hat{\mathbf{z}}_m, x) / q(\hat{\mathbf{z}}_m)} \\ q_M(\mathbf{z}_{1:M} | \hat{\mathbf{z}}_{1:M}, h) &= \delta(\mathbf{z}_1 - \hat{\mathbf{z}}_h) \delta(\mathbf{z}_{2:M} - \hat{\mathbf{z}}_{-h}). \end{aligned}$$

What is the marginal distribution over $\mathbf{z}_{1:M}$?

$$\begin{aligned} q_M(\mathbf{z}_{1:M}) &= \int \sum_{h=1}^M q_M(\hat{\mathbf{z}}_{1:M}) q_M(h | \hat{\mathbf{z}}_{1:M}) q_M(\mathbf{z}_{1:M} | \hat{\mathbf{z}}_{1:M}, h) d\hat{\mathbf{z}}_{1:M} \\ &= \int \sum_{h=1}^M q(\hat{\mathbf{z}}_{1:M}) \frac{p(\hat{\mathbf{z}}_h, \mathbf{x}) / q(\hat{\mathbf{z}}_h)}{\sum_{m=1}^M p(\hat{\mathbf{z}}_m, \mathbf{x}) / q(\hat{\mathbf{z}}_m)} \delta(\mathbf{z}_1 - \hat{\mathbf{z}}_h) \delta(\mathbf{z}_{2:M} - \hat{\mathbf{z}}_{-h}) d\hat{\mathbf{z}}_{1:M} \\ &= \sum_{h=1}^M \int q(\hat{\mathbf{z}}_{1:M}) \frac{p(\hat{\mathbf{z}}_h, \mathbf{x}) / q(\hat{\mathbf{z}}_h)}{\sum_{m=1}^M p(\hat{\mathbf{z}}_m, \mathbf{x}) / q(\hat{\mathbf{z}}_m)} \delta(\mathbf{z}_1 - \hat{\mathbf{z}}_h) \delta(\mathbf{z}_{2:M} - \hat{\mathbf{z}}_{-h}) d\hat{\mathbf{z}}_{1:M} \\ &= M \int q(\hat{\mathbf{z}}_{1:M}) \frac{p(\hat{\mathbf{z}}_1, \mathbf{x}) / q(\hat{\mathbf{z}}_1)}{\sum_{m=1}^M p(\hat{\mathbf{z}}_m, \mathbf{x}) / q(\hat{\mathbf{z}}_m)} \delta(\mathbf{z}_1 - \hat{\mathbf{z}}_1) \delta(\mathbf{z}_{2:M} - \hat{\mathbf{z}}_{2:M}) d\hat{\mathbf{z}}_{1:M} \\ &= M \int \frac{p(\hat{\mathbf{z}}_1, \mathbf{x}) q(\hat{\mathbf{z}}_{2:M})}{\sum_{m=1}^M p(\hat{\mathbf{z}}_m, \mathbf{x}) / q(\hat{\mathbf{z}}_m)} \delta(\mathbf{z}_1 - \hat{\mathbf{z}}_1) \delta(\mathbf{z}_{2:M} - \hat{\mathbf{z}}_{2:M}) d\hat{\mathbf{z}}_{1:M} \\ &= M \frac{p(\mathbf{z}_1, \mathbf{x}) q(\mathbf{z}_{2:M})}{\sum_{m=1}^M p(\mathbf{z}_m, \mathbf{x}) / q(\mathbf{z}_m)} \\ &= \frac{p(\mathbf{z}_1, \mathbf{x}) q(\mathbf{z}_{2:M})}{\frac{1}{M} \sum_{m=1}^M p(\mathbf{z}_m, \mathbf{x}) / q(\mathbf{z}_m)} \end{aligned}$$

For the decomposition, we have, by Eq. 1 that

$$\log p_M(\mathbf{x}) = \mathbb{E}_{q_M(\mathbf{z}_{1:M})} \log \frac{p_M(\mathbf{z}_{1:M}, \mathbf{x})}{q_M(\mathbf{z}_{1:M})} + \text{KL}[q_M(\mathbf{z}_{1:M}) \| p_M(\mathbf{z}_{1:M} | \mathbf{x})].$$

Now, by the definition of p_M , it's easy to see that $p(\mathbf{x}) = p_M(\mathbf{x})$.

Next, re-write the importance-weighted ELBO as

$$\begin{aligned} \mathbb{E}_{q_M(\mathbf{z}_{1:M})} \log \frac{p_M(\mathbf{z}_{1:M}, \mathbf{x})}{q_M(\mathbf{z}_{1:M})} &= \mathbb{E}_{q_M(\mathbf{z}_{1:M})} \log \frac{p(\mathbf{z}_1, \mathbf{x}) q(\mathbf{z}_{2:M})}{\frac{1}{M} \sum_{m=1}^M p(\mathbf{z}_m, \mathbf{x}) / q(\mathbf{z}_m)} \\ &= \mathbb{E}_{q_M(\mathbf{z}_{1:M})} \log \left(\frac{1}{M} \sum_{m=1}^M \frac{p(\mathbf{z}_m, \mathbf{x})}{q(\mathbf{z}_m)} \right). \end{aligned}$$

This gives that

$$\log p(\mathbf{x}) = \underbrace{\mathbb{E}_{q_M(\mathbf{z}_{1:M})} \log \left(\frac{1}{M} \sum_{m=1}^M \frac{p(\mathbf{z}_m, \mathbf{x})}{q(\mathbf{z}_m)} \right)}_{\text{importance weighted ELBO}} + \text{KL}[q_M(\mathbf{z}_{1:M}) \| p_M(\mathbf{z}_{1:M} | \mathbf{x})].$$

□

$$\text{Lemma 4. } \mathbb{E}_{q_M(\mathbf{z}_1)} t(\mathbf{z}_1) = \mathbb{E}_{q(\mathbf{z}_{1:M})} \frac{\sum_{m=1}^M \omega(\mathbf{z}_m) t(\mathbf{z}_m)}{\sum_{m=1}^M \omega(\mathbf{z}_m)}.$$

Proof.

$$\begin{aligned} \mathbb{E}_{q_M(\mathbf{z}_{1:M})} t(\mathbf{z}_1) &= \int \frac{t(\mathbf{z}_1) p(\mathbf{z}_1, \mathbf{x}) q(\mathbf{z}_{2:M})}{\frac{1}{M} \sum_{m=1}^M p(\mathbf{z}_m, \mathbf{x}) / q(\mathbf{z}_m)} d\mathbf{z}_{1:M} \\ &= \int q(\mathbf{z}_{1:M}) \frac{t(\mathbf{z}_1) p(\mathbf{z}_1, \mathbf{x}) / q(\mathbf{z}_1)}{\frac{1}{M} \sum_{m=1}^M p(\mathbf{z}_m, \mathbf{x}) / q(\mathbf{z}_m)} d\mathbf{z}_{1:M} \\ &= \mathbb{E}_{q(\mathbf{z}_{1:M})} \frac{t(\mathbf{z}_1) p(\mathbf{z}_1, \mathbf{x}) / q(\mathbf{z}_1)}{\frac{1}{M} \sum_{m=1}^M p(\mathbf{z}_m, \mathbf{x}) / q(\mathbf{z}_m)} \\ &= \mathbb{E}_{q(\mathbf{z}_{1:M})} \frac{\frac{1}{M} \sum_{m=1}^M t(\mathbf{z}_m) p(\mathbf{z}_m, \mathbf{x}) / q(\mathbf{z}_m)}{\frac{1}{M} \sum_{m=1}^M p(\mathbf{z}_m, \mathbf{x}) / q(\mathbf{z}_m)} \\ &= \mathbb{E}_{q(\mathbf{z}_{1:M})} \frac{\sum_{m=1}^M \omega(\mathbf{z}_m) t(\mathbf{z}_m)}{\sum_{m=1}^M \omega(\mathbf{z}_m)} \end{aligned}$$

□

Theorem 2. *The marginal and joint divergences relevant to IWVI are related by*

$$\text{KL}[q_M(\mathbf{z}_{1:M}) \| p_M(\mathbf{z}_{1:M} | \mathbf{x})] = \text{KL}[q_M(\mathbf{z}_1) \| p(\mathbf{z}_1 | \mathbf{x})] + \text{KL}[q_M(\mathbf{z}_{2:M} | \mathbf{z}_1) \| q(\mathbf{z}_{2:M})].$$

As a consequence, the gap in the first inequality of Eq. 8 is exactly $\text{KL}[q_M(\mathbf{z}_1) \| p(\mathbf{z}_1 | \mathbf{x})]$ and the gap in the second inequality is exactly $\text{KL}[q_M(\mathbf{z}_{2:M} | \mathbf{z}_1) \| q(\mathbf{z}_{2:M})]$.

Proof.

$$\begin{aligned} \text{KL}[q_M(\mathbf{z}_{1:M}) \| p_M(\mathbf{z}_{1:M} | \mathbf{x})] &= \text{KL}[q_M(\mathbf{z}_1) \| p_M(\mathbf{z}_1 | \mathbf{x})] + \text{KL}[q_M(\mathbf{z}_{2:M} | \mathbf{z}_1) \| p_M(\mathbf{z}_{2:M} | \mathbf{z}_1, \mathbf{x})] \\ &\quad \text{by the chain rule of KL-divergence} \\ &= \text{KL}[q_M(\mathbf{z}_1) \| p(\mathbf{z}_1 | \mathbf{x})] + \text{KL}[q_M(\mathbf{z}_{2:M} | \mathbf{z}_1) \| q(\mathbf{z}_{2:M})] \\ &\quad \text{since } p_M(\mathbf{z}_1 | \mathbf{x}) = p(\mathbf{z}_1 | \mathbf{x}) \text{ and } p_M(\mathbf{z}_{2:M} | \mathbf{z}_1, \mathbf{x}) = q(\mathbf{z}_{2:M}). \end{aligned}$$

The KL-divergences can be identified with the gaps in the inequalities in Eq. 8 through the application of Eq. 1 to give that

$$\log p(\mathbf{x}) - \text{ELBO}[q_M(\mathbf{z}_1) \| p(\mathbf{z}_1, \mathbf{x})] = \text{KL}[q_M(\mathbf{z}_1) \| p_M(\mathbf{z}_1 | \mathbf{x})]$$

which establishes the looseness of the first inequality. Then, Thm. 1 gives that

$$\log p(\mathbf{x}) - \text{IW-ELBO}_M[q(\mathbf{z}) \| p(\mathbf{z}, \mathbf{x})] = \text{KL}[q_M(\mathbf{z}_{1:M}) \| p_M(\mathbf{z}_{1:M} | \mathbf{x})].$$

The difference of the previous two equations gives that the looseness of the second inequality is

$$\begin{aligned} \text{ELBO}[q_M(\mathbf{z}_1) \| p(\mathbf{z}_1, \mathbf{x})] - \text{IW-ELBO}_M[q(\mathbf{z}) \| p(\mathbf{z}, \mathbf{x})] &= \text{KL}[q_M(\mathbf{z}_{1:M}) \| p_M(\mathbf{z}_{1:M} | \mathbf{x})] \\ &\quad - \text{KL}[q_M(\mathbf{z}_1) \| p_M(\mathbf{z}_1 | \mathbf{x})] \\ &= \text{KL}[q_M(\mathbf{z}_{2:M} | \mathbf{z}_1) \| q(\mathbf{z}_{2:M})]. \end{aligned}$$

□

A.3 Asymptotics

Theorem 3. For large M , the looseness of the IW-ELBO is given by the variance of R . Formally, if there exists some $\alpha > 0$ such that $\mathbb{E} |R - p(\mathbf{x})|^{2+\alpha} < \infty$ and $\limsup_{M \rightarrow \infty} \mathbb{E}[1/R_M] < \infty$, then

$$\lim_{M \rightarrow \infty} M \underbrace{\left(\log p(\mathbf{x}) - \text{IW-ELBO}_M [q(\mathbf{z}) \| p(\mathbf{z}, \mathbf{x})] \right)}_{\text{KL}[q_M \| p_M]} = \frac{\mathbb{V}[R]}{2p(\mathbf{x})^2}.$$

We first give more context for this theorem, and then its proof. Since $\text{IW-ELBO}_M [p \| q] = \mathbb{E} \log(R_M)$ where $\sqrt{M}(R_M - p(\mathbf{x}))$ converges in distribution to a Gaussian distribution, the result is *nearly* a straightforward application of the “delta method for moments” (e.g. [4, Chapter 5.3.1]). The key difficulty is that the derivatives of $\log(r)$ are unbounded at $r = 0$; bounded derivatives are typically required to establish convergence rates.

The assumption that $\limsup_{M \rightarrow \infty} \mathbb{E}[1/R_M] < \infty$ warrants further discussion. One (rather strong) assumption that implies this¹ would be that $\mathbb{E} 1/R < \infty$. However, this is not necessary. For example, if R were uniform on the $[0, 1]$ interval, then $\mathbb{E} 1/R$ does not exist, yet $\mathbb{E} 1/R_M$ does for any $M \geq 2$. It can be shown² that if $M \geq M_0$ and $\mathbb{E} 1/R_{M_0} < \infty$ then $\mathbb{E} 1/R_M \leq \mathbb{E} 1/R_{M_0}$. Thus, assuming only that there is some finite M such that $\mathbb{E} 1/R_M < \infty$ is sufficient for the \limsup condition.

Both Maddison et al. [13, Prop. 1] and Rainforth et al. [18, Eq. 7] give related results that control the rate of convergence. It can be shown that Proposition 1 of Maddison et al. implies the conclusion of Theorem 3 if $\mathbb{E} [(R - p(\mathbf{x}))^6] < \infty$. Their Proposition 1, specialized to our notation and setting, is:

Proposition 1 ([13]). *If $g(M) = \mathbb{E} [(R_M - p(\mathbf{x}))^6] < \infty$ and $\limsup_{M \rightarrow \infty} \mathbb{E}[1/R_M] < \infty$, then*

$$\log p(\mathbf{x}) - \mathbb{E} \log R_M = \frac{\mathbb{V}[R_M]}{2p(\mathbf{x})^2} + O(\sqrt{g(M)}).$$

In order to imply the conclusion of Theorem 3, it is necessary to bound the final term. To do this, we can use the following lemma, which is a consequence of the Marcinkiewicz–Zygmund inequality [14] and provides an asymptotic bound on the higher moments of a sample mean. We will also use this lemma in our proof of Theorem 3 below.

Lemma 5 (Bounds on sample moments). *Let U_1, \dots, U_M be i.i.d random variables with $\mathbb{E}[U_i] = 0$ and let $\bar{U}_M = \frac{1}{M} \sum_{i=1}^M U_i$. Then, for each $s \geq 2$ there is a constant $B_s > 0$ such that*

$$\mathbb{E} |\bar{U}_M|^s \leq B_s M^{-s/2} \mathbb{E} |U_1|^s.$$

We now show that if the assumptions of Prop. 1 are true, this lemma can be used to bound $g(M)$ and therefore imply the conclusion of Theorem 3. If $\mathbb{E} |R - p(\mathbf{x})|^6 < \infty$ then $g(M) = \mathbb{E} |R_M - p(\mathbf{x})|^6 \leq B_6 M^{-3} \mathbb{E} |R - p(\mathbf{x})|^6 \in O(M^{-3})$ and $\sqrt{g(M)} \in O(M^{-3/2})$. Then, since $\mathbb{V}[R_M] = \mathbb{V}[R]/M$, we can multiply by M in both sides of Prop. 1 to get

$$M(\log p(\mathbf{x}) - \mathbb{E} \log R_M) = \frac{\mathbb{V}[R]}{2p(\mathbf{x})^2} + O(M^{-1/2}),$$

which goes to $\frac{\mathbb{V}[R]}{2p(\mathbf{x})^2}$ as $M \rightarrow \infty$, as desired.

Proof of Theorem 3. Our proof will follow the same high-level structure as the proof of Prop. 1 from Maddison et al. [13], but we will more tightly bound the Taylor remainder term that appears below.

¹To see this, observe that since $1/r$ is convex over $r > 0$, Jensen’s inequality gives that $(\frac{1}{M} \sum_{m=1}^M r_m)^{-1} \leq \frac{1}{M} \sum_{m=1}^M r_m^{-1}$ and so $\mathbb{E} 1/R_M \leq \mathbb{E} 1/R$.

²Define σ to be a uniformly random over all permutations of $1, \dots, M$. Then, Jensen’s inequality gives that

$$\left(\frac{1}{M} \sum_{m=1}^M r_m \right)^{-1} = \left(\mathbb{E}_{\sigma} \frac{1}{M_0} \sum_{m=1}^{M_0} r_{\sigma(m)} \right)^{-1} \leq \mathbb{E}_{\sigma} \left(\frac{1}{M_0} \sum_{m=1}^{M_0} r_{\sigma(m)} \right)^{-1}.$$

Since R_M is a mean of i.i.d. variables, the permutation vanishes under expectations and so $\mathbb{E} 1/R_M \leq \mathbb{E} 1/R_{M_0}$.

Let $\theta = p(\mathbf{x}) = \mathbb{E} R$ and $\sigma^2 = \mathbb{V}[R]$. For any $r > 0$, let $\Delta = \Delta(r) = \frac{r-\theta}{\theta} = \frac{r}{\theta} - 1$. Then $\log \theta - \log r = -\log(1 + \Delta)$. Since $r > 0$, we only need to consider $-1 < \Delta < \infty$.

Consider the second-order Taylor expansion of $\log(1 + \Delta)$:

$$\log(1 + \Delta) = \Delta - \frac{1}{2}\Delta^2 + \int_0^\Delta \frac{x^2}{1+x} dx$$

Now, let $\Delta_M = \Delta(R_M)$. Then, since $\mathbb{E}[\Delta_M] = 0$ and $\mathbb{E}[\Delta_M^2] = \frac{1}{\theta^2} \frac{\sigma^2}{M}$,

$$\begin{aligned} \mathbb{E}(\log \theta - \log R_M) &= -\mathbb{E} \log(1 + \Delta_M) = \frac{1}{2} \mathbb{E} \Delta_M^2 - \mathbb{E} \int_0^{\Delta_M} \frac{x^2}{1+x} dx \\ &= \frac{\sigma^2/M}{2\theta^2} - \mathbb{E} \int_0^{\Delta_M} \frac{x^2}{1+x} dx \end{aligned}$$

Moving M and taking the limit, this is

$$\lim_{M \rightarrow \infty} M(\log \theta - \mathbb{E} \log R_M) = \frac{\sigma^2}{2\theta^2} - \lim_{M \rightarrow \infty} M \mathbb{E} \int_0^{\Delta_M} \frac{x^2}{1+x} dx.$$

Our desired result holds if and only if $\lim_{M \rightarrow \infty} \left| M \mathbb{E} \int_0^{\Delta_M} \frac{x^2}{1+x} dx \right| = 0$. Lemma 7 (proven in Section A.3.1 below) bounds the absolute value of this integral for fixed Δ . Choosing $\Delta = \Delta_M$, multiplying by M and taking the expectation of both sides of Lemma 7 is equivalent to the statement that, for any $\epsilon, \alpha > 0$:

$$M \mathbb{E} \left| \int_0^{\Delta_M} \frac{x^2}{1+x} dx \right| \leq M \mathbb{E} \left[C_\epsilon \left| \frac{1}{1+\Delta_M} \right|^{\frac{\epsilon}{1+\epsilon}} |\Delta_M|^{\frac{2+3\epsilon}{1+\epsilon}} \right] + M D_\alpha \mathbb{E} |\Delta_M|^{2+\alpha}. \quad (16)$$

Let α be as given in the conditions of the theorem, so that $\mathbb{E} |R - \theta|^{2+\alpha} < \infty$. We will show that both terms on the right-hand side of Eq. (16) have a limit of zero as $M \rightarrow \infty$ for suitable ϵ . For the second term, let $s = 2 + \alpha$. Then by Lemma 5,

$$\mathbb{E} |\Delta_M|^{2+\alpha} = \mathbb{E} |\Delta_M|^s = \theta^{-s} \mathbb{E} |R_M - \theta|^s \leq \theta^{-s} B_s M^{-s/2} \mathbb{E} |R - \theta|^s. \quad (17)$$

Since $s/2 > 1$ and $\mathbb{E} |R - \theta|^s < \infty$, this implies that the $D_\alpha \mathbb{E} |\Delta_M|^{2+\alpha}$ is $o(M^{-1})$ and so the limit of the second term on the right of Eq. 16 is zero.

For the first term on the right-hand side of Eq. (16), apply Holder's inequality with $p = \frac{1+\epsilon}{\epsilon}$ and $q = 1 + \epsilon$, to get that

$$M \mathbb{E} \left[C_\epsilon \left| \frac{1}{1+\Delta_M} \right|^{\frac{\epsilon}{1+\epsilon}} |\Delta_M|^{\frac{2+3\epsilon}{1+\epsilon}} \right] \leq M C_\epsilon \left(\mathbb{E} \left| \frac{1}{1+\Delta_M} \right| \right)^{\frac{\epsilon}{1+\epsilon}} \left(\mathbb{E} |\Delta_M|^{2+3\epsilon} \right)^{\frac{1}{1+\epsilon}}.$$

Now, use the fact that $\limsup(a_M b_M) \leq \limsup a_M \limsup b_M$ to get that

$$\begin{aligned} \limsup_{M \rightarrow \infty} M \mathbb{E} \left[C_\epsilon \left| \frac{1}{1+\Delta_M} \right|^{\frac{\epsilon}{1+\epsilon}} |\Delta_M|^{\frac{2+3\epsilon}{1+\epsilon}} \right] \\ \leq C_\epsilon \limsup_{M \rightarrow \infty} \left(\mathbb{E} \left| \frac{1}{1+\Delta_M} \right| \right)^{\frac{\epsilon}{1+\epsilon}} \limsup_{M \rightarrow \infty} M \left(\mathbb{E} |\Delta_M|^{2+3\epsilon} \right)^{\frac{1}{1+\epsilon}} \end{aligned} \quad (18)$$

We will now show that the first limit on the right of Eq. 18 is finite, while the second is zero. For the first limit, our assumption that $\limsup_{M \rightarrow \infty} \mathbb{E} \frac{1}{R_M} < \infty$, means that for sufficiently large M , $\mathbb{E} \frac{1}{R_M}$ is bounded by a constant. Thus, we have that regardless of ϵ , the first limit of

$$\limsup_{M \rightarrow \infty} \left(\mathbb{E} \left| \frac{1}{1+\Delta_M} \right| \right)^{\frac{\epsilon}{1+\epsilon}} = \limsup_{M \rightarrow \infty} \left(\mathbb{E} \left| \frac{\theta}{R_M} \right| \right)^{\frac{\epsilon}{1+\epsilon}}$$

is bounded by a constant.

Now, consider the second limit on the right of Eq. 18. Let $\epsilon = \alpha/3$ and $s' = \frac{2+\alpha}{1+\epsilon} > 2$. Then, using the bound we already established above in Eq. 17 we have that

$$\left(\mathbb{E} |\Delta_M|^{2+3\epsilon} \right)^{\frac{1}{1+\epsilon}} = \left(\mathbb{E} |\Delta_M|^{2+\alpha} \right)^{\frac{1}{1+\epsilon}} \leq \theta^{-s'} B_s^{\frac{1}{1+\epsilon}} M^{-s'/2} \left(\mathbb{E} |R - \theta|^s \right)^{\frac{1}{1+\epsilon}}.$$

Since $s' > 2$ and $\mathbb{E} |R - \theta|^s < \infty$, this proves that the second limit in Eq. (18) is zero. Since we already showed that the first limit on the right of Eq. 18 is finite we have that the limit of the first term on the right of Eq. (16) is zero, completing the proof. \square

A.3.1 Proofs of Lemmas

Lemma 6 (Bounds on sample moments). *Let U_1, \dots, U_M be i.i.d random variables with $\mathbb{E}[U_i] = 0$ and let $\bar{U}_M = \frac{1}{M} \sum_{i=1}^M U_i$. Then, for each $s \geq 2$ there is a constant $B_s > 0$ such that*

$$\mathbb{E} |\bar{U}_M|^s \leq B_s M^{-s/2} \mathbb{E} |U_1|^s.$$

Proof. The Marcinkiewicz–Zygmund inequality [14] states that, under the same conditions, for any $s \geq 1$ there exists $B_s > 0$ such that

$$\mathbb{E} \left(\left| \sum_{i=1}^M U_i \right|^s \right) \leq B_s \mathbb{E} \left(\left(\sum_{i=1}^M |U_i|^2 \right)^{s/2} \right)$$

Therefore,

$$\begin{aligned} \mathbb{E} \left(\left| \frac{1}{M} \sum_{i=1}^M U_i \right|^s \right) &= M^{-s} \mathbb{E} \left(\left| \sum_{i=1}^M U_i \right|^s \right) \\ &\leq B_s M^{-s} \mathbb{E} \left(\left(\sum_{i=1}^M |U_i|^2 \right)^{s/2} \right) \\ &= B_s M^{-s/2} \mathbb{E} \left(\left(\frac{1}{M} \sum_{i=1}^M |U_i|^2 \right)^{s/2} \right) \end{aligned}$$

Now, since $v \mapsto v^{s/2}$ is convex for $s \geq 2$

$$\left(\frac{1}{M} \sum_{i=1}^M |U_i|^2 \right)^{s/2} \leq \frac{1}{M} \sum_{i=1}^M |U_i|^s,$$

and $\mathbb{E} \left(\frac{1}{M} \sum_{i=1}^M |U_i|^s \right) = \mathbb{E} |U_1|^s$, so we have

$$\mathbb{E} \left(\left| \frac{1}{M} \sum_{i=1}^M U_i \right|^s \right) \leq B_s M^{-s/2} \mathbb{E} |U_1|^s.$$

\square

Lemma 7. *For every $\epsilon, \alpha > 0$ there exists constants C_ϵ, D_α such that, for all $\Delta > -1$,*

$$\left| \int_0^\Delta \frac{x^2}{1+x} dx \right| \leq C_\epsilon \left| \frac{1}{1+\Delta} \right|^{\frac{\epsilon}{1+\epsilon}} |\Delta|^{\frac{2+3\epsilon}{1+\epsilon}} + D_\alpha |\Delta|^{2+\alpha}.$$

Proof. We will treat positive and negative Δ separately, and show that:

1. If $-1 < \Delta < 0$, then for every $\epsilon > 0$, there exists $C_\epsilon > 0$ such that

$$\left| \int_0^\Delta \frac{x^2}{1+x} dx \right| \leq C_\epsilon \left| \frac{1}{1+\Delta} \right|^{\frac{\epsilon}{1+\epsilon}} |\Delta|^{\frac{2+3\epsilon}{1+\epsilon}}. \quad (19)$$

2. If $\Delta \geq 0$, then for every $\alpha > 0$,

$$\left| \int_0^\Delta \frac{x^2}{1+x} dx \right| \leq \underbrace{\frac{1}{2+\alpha}}_{D_\alpha} \Delta^{2+\alpha}. \quad (20)$$

Put together, these imply that for all $\Delta > -1$, the quantity $\left| \int_0^\Delta \frac{x^2}{1+x} dx \right|$ is no more than the maximum of the upper bounds in Eqs. (19) and (20). Since these are both non-negative, it is also no more than their sum, which will prove the lemma.

We now prove the bound in Eq. (19). For $-1 < \Delta < 0$, substitute $u = -x$ to obtain an integral with non-negative integrand and integration limits:

$$\int_0^\Delta \frac{x^2}{1+x} dx = - \int_0^{-\Delta} \frac{u^2}{1-u} du < 0$$

Therefore:

$$\left| \int_0^\Delta \frac{x^2}{1+x} dx \right| = \int_0^{-\Delta} \frac{u^2}{1-u} du$$

Now apply Holder's inequality with $p, q > 1$ such that $\frac{1}{p} + \frac{1}{q} = 1$:

$$\begin{aligned} \int_0^{-\Delta} \frac{u^2}{1-u} du &\leq \left(\int_0^{-\Delta} \frac{1}{(1-u)^p} du \right)^{1/p} \cdot \left(\int_0^{-\Delta} u^{2q} du \right)^{1/q} \\ &= \left(\frac{1}{p-1} \frac{1 - (1+\Delta)^{p-1}}{(1+\Delta)^{p-1}} \right)^{1/p} \cdot \left(\frac{1}{2q+1} (-\Delta)^{2q+1} \right)^{1/q} \\ &= C_{p,q} \cdot \left(\frac{1 - (1+\Delta)^{p-1}}{(1+\Delta)^{p-1}} \right)^{1/p} \cdot \left(|\Delta|^{2q+1} \right)^{1/q} \\ &\leq C_{p,q} \cdot \left(\frac{1}{(1+\Delta)^{p-1}} \right)^{1/p} \cdot \left(|\Delta|^{2q+1} \right)^{1/q} \\ &= C_{p,q} \cdot \left| \frac{1}{1+\Delta} \right|^{\frac{p-1}{p}} \cdot |\Delta|^{\frac{2q+1}{q}} \end{aligned}$$

In the fourth line, we used the fact that $0 < (1+\Delta)^{p-1} < 1$. Now set $p = 1 + \epsilon$, $q = \frac{1+\epsilon}{\epsilon}$ and $C_\epsilon = C_{p,q}$, and we obtain Eq. (19).

We now prove the upper bound of Eq. (20). For $\Delta \geq 0$, the integrand is non-negative and:

$$\left| \int_0^\Delta \frac{x^2}{1+x} dx \right| = \int_0^\Delta \frac{x^2}{1+x} dx.$$

Let $f(\Delta) = \int_0^\Delta \frac{x^2}{1+x} dx$ and $g(\Delta) = \frac{1}{2+\alpha} x^{2+\alpha}$. Then $f(0) = g(0) = 0$, and we claim that $f'(\Delta) \leq g'(\Delta)$ for all $\Delta \geq 0$, which together imply $f(\Delta) \leq g(\Delta)$ for all $\Delta \geq 0$.

To see that $f'(\Delta) \leq g'(\Delta)$, observe that:

$$\frac{g'(\Delta)}{f'(\Delta)} = \frac{\Delta^{1+\alpha}}{\frac{\Delta^2}{1+\Delta}} = \frac{\Delta^{1+\alpha}(1+\Delta)}{\Delta^2} = \frac{\Delta^{1+\alpha} + \Delta^{2+\alpha}}{\Delta^2} = \frac{1}{\Delta^{1-\alpha}} + \Delta^\alpha$$

Both terms on the right-hand side are nonnegative. If $\Delta \in [0, 1]$, then $\frac{1}{\Delta^{1-\alpha}} \geq 1$. If $\Delta \geq 1$ then $\Delta^\alpha \geq 1$. Therefore, the sum is at least one for all $\Delta \geq 0$. \square

A.3.2 Relationship of Decompositions

This section discusses the relationship of our decomposition to that of Le et al. [12, Claim 1].

We first state their Claim 1 in our notation. Define $q_M^{IS}(\mathbf{z}_{1:M}) = \prod_{m=1}^M q(\mathbf{z}_m)$ and define

$$p_M^{IS}(\mathbf{z}_{1:M}, \mathbf{x}) = q_M^{IS}(\mathbf{z}_{1:M}) \frac{1}{M} \sum_{m=1}^M \frac{p(\mathbf{z}_m, \mathbf{x})}{q(\mathbf{z}_m)} = \frac{1}{M} \sum_{m=1}^M p(\mathbf{z}_m, \mathbf{x}) \prod_{m' \neq m} q(\mathbf{z}_{m'}).$$

By construction, the ratio of these two distributions is

$$\frac{p_M^{IS}(\mathbf{z}_{1:M}, \mathbf{x})}{q_M^{IS}(\mathbf{z}_{1:M})} = \frac{1}{M} \sum_{m=1}^M \frac{p(\mathbf{z}_m, \mathbf{x})}{q(\mathbf{z}_m)},$$

and $p_M^{IS}(\mathbf{x}) = p(\mathbf{x})$ and so applying the standard ELBO decomposition (Eq. 1) to p_M^{IS} and q_M^{IS} gives that

$$\log p(\mathbf{x}) = \text{IW-ELBO}_M [q(\mathbf{z}) \| p(\mathbf{z}, \mathbf{x})] + KL[q_M^{IS}(\mathbf{z}_{1:M}) \| p_M^{IS}(\mathbf{z}_{1:M} \mid \mathbf{x})].$$

This is superficially similar to our result because it shows that maximizing the IW-ELBO minimizes the KL-divergence between two augmented distributions. However, it is fundamentally different and does not inform probabilistic inference. In particular, note that the marginals of these two distributions are

$$\begin{aligned} p_M^{IS}(\mathbf{z}_1 \mid \mathbf{x}) &= \frac{1}{M} p(\mathbf{z}_1 \mid \mathbf{x}) + \frac{M-1}{M} q(\mathbf{z}_1), \\ q_M^{IS}(\mathbf{z}_1) &= q(\mathbf{z}_1). \end{aligned}$$

This pair of distributions holds q_M^{IS} "fixed" to be an independent sample of size M from q , and changes p_M^{IS} so that its marginals approach those of q_M^{IS} as $M \rightarrow \infty$. The distribution one can actually sample from, q_M^{IS} , does not approach the desired target.

Contrast this with our approach, where we hold the marginal of p_M fixed so that $p_M(\mathbf{z}_1 \mid \mathbf{x}) = p(\mathbf{z}_1 \mid \mathbf{x})$, and augment q so that $q_M(\mathbf{z}_1)$ gets closer and closer to $p_M(\mathbf{z}_1 \mid \mathbf{x})$ as M increases. Further, since $q_M(\mathbf{z}_1)$ is the distribution resulting from self-normalized importance sampling, it is available for use in a range of inference tasks.