

TO RELIEVE YOUR HEADACHE OF TRAINING AN MRF, TAKE ADVIL

Anonymous authors

Paper under double-blind review

ABSTRACT

We propose a black-box algorithm called *Adversarial Variational Inference and Learning* (AdvIL) to perform inference and learning in a general Markov random field (MRF). AdvIL employs two variational distributions to approximately infer the latent variables and estimate the partition function of an MRF, respectively. The two variational distributions provide an estimate of the negative log-likelihood of the MRF as a minimax optimization problem, which is solved by stochastic gradient descent. AdvIL is proven convergent under certain conditions. On one hand, compared with contrastive divergence, AdvIL requires minimal assumptions about the model structure and can deal with a broader family of MRFs. On the other hand, compared with existing black-box methods, AdvIL provides a tighter estimate of the log partition function and achieves much better empirical results.

1 INTRODUCTION

Undirected graphical models, i.e. Markov random fields (MRFs), find applications in a variety of machine learning areas, including image segmentation (Krähenbühl & Koltun, 2011), generative modelling (Salakhutdinov & Larochelle, 2010) and sequential labelling (Lafferty et al., 2001). However, the inference and learning of general MRFs are challenging due to the presence of loopy structures and a global normalizing factor, i.e. partition function, especially when latent variables are present. Extensive efforts have been devoted to developing approximate methods. On one hand, sample-based methods (Neal, 1993) and variational approaches (Jordan et al., 1999; Welling & Sutton, 2005; Salakhutdinov & Larochelle, 2010) are proposed to infer the latent variables. On the other hand, extensive work (Meng & Wong, 1996; Neal, 2001; Hinton, 2002; Tieleman, 2008; Wainwright et al., 2005; Wainwright & Jordan, 2006) has been done to estimate the partition function. Among these methods, contrastive divergence (Hinton, 2002) is proven effective in certain types of models.

Most of the existing methods highly depend on the model structure and require model-specific analysis in new applications, which makes it important to develop black-box inference and learning methods. Previous work (Ranganath et al., 2014; Schulman et al., 2015) shows the ability to automatically infer the latent variables and obtain gradient estimate in directed models. However, there is no black-box learning method for undirected models except the recent work of NVIL (Kuleshov & Ermon, 2017).

NVIL introduces a variational distribution and derives an upper bound of the partition function in a general MRF, in the same spirit as amortized inference (Kingma & Welling, 2013; Rezende et al., 2014; Mnih & Gregor, 2014) for directed models. NVIL has several advantages over existing methods, including the ability of black-box learning, tracking the partition function during training and getting approximate samples efficiently during testing. However, NVIL also comes with two disadvantages: (1) it leaves the inference problem of MRFs unsolved¹ and only trains simple MRFs with tractable posteriors, and (2) the upper bound of the partition function can be underestimated (Kuleshov & Ermon, 2017), resulting in sub-optimal solutions on high-dimensional data.

We propose *Adversarial Variational Inference and Learning* (AdvIL) to relieve some headache of learning an MRF model. AdvIL is a black-box inference and learning method that partly solves the two problems of NVIL and retains the advantages of NVIL at the same time. First, AdvIL introduces a variational encoder to infer the latent variables, which provides an upper bound of the free energy.

¹NVIL (Kuleshov & Ermon, 2017) presents a hybrid model. The inference in the title refers to directed part but not for an MRF.

Second, AdvIL introduces a variational decoder for the MRF, which provides a lower bound of the log partition function. The two variational distributions provide an estimate of the negative log-likelihood of the MRF. On one hand, the estimate is in an intuitive form of an approximate *contrastive free energy*, which is expressed in terms of the expected energy and the (conditional) entropy of the corresponding variational distribution. On the other hand, similar to GAN (Goodfellow et al., 2014), the estimate is a minimax optimization problem, which is solved by stochastic gradient descent (SGD) in an alternating manner. Theoretically, our algorithm is convergent if the variational encoder approximates the model well. This motivates us to introduce an auxiliary variable to enhance the flexibility of the variational decoder, whose entropy is approximated by the third variational trick.

We evaluate AdvIL in various classical undirected generative models, including restricted Boltzmann machines (RBM) (Ackley et al., 1985), deep Boltzmann machines (DBM) (Salakhutdinov & Hinton, 2009), and Gaussian restricted Boltzmann machines (GRBM) (Hinton & Salakhutdinov, 2006), on several real datasets. We empirically demonstrate that (1) compared with the black-box NVIL (Kuleshov & Ermon, 2017) method, AdvIL provides a tighter estimate of the log partition function and hence achieves much better log-likelihood results; and (2) compared with the standard contrastive divergence based methods (Hinton, 2002; Welling & Hinton, 2002; Welling & Sutton, 2005), AdvIL can deal with a broader family of MRFs without model-specific analysis and obtain better results when the model structure gets complex as in DBM.

2 BACKGROUND

We consider a general case where the model consists of both visible variables v and latent variables h . An MRF defines the joint distribution over v and h as $P(v, h) = \frac{e^{-\mathcal{E}(v, h)}}{\mathcal{Z}}$, where \mathcal{E} denotes the associated energy function that assigns a scalar value for a given configuration of (v, h) and \mathcal{Z} is the partition function such that $\mathcal{Z} = \int_{v, h} e^{-\mathcal{E}(v, h)} dv dh$.

Let $P_{\mathcal{D}}(v)$ denote the empirical distribution of the training data. Minimizing the negative log-likelihood (NLL) of an MRF is a commonly chosen learning criterion and it is given by:

$$\mathcal{L}(\theta) := -\mathbb{E}_{P_{\mathcal{D}}(v)} \left[\log \int_h \frac{e^{-\mathcal{E}(v, h)}}{\mathcal{Z}} dh \right], \quad (1)$$

where θ denotes the trainable parameters in \mathcal{E} . Further, the gradient of θ is:

$$\nabla_{\theta} \mathcal{L}(\theta) = \mathbb{E}_{P_{\mathcal{D}}(v)} [\nabla_{\theta} \mathcal{F}(v)] - \mathbb{E}_{P(v)} [\nabla_{\theta} \mathcal{F}(v)], \quad (2)$$

where $\mathcal{F}(v) = -\log \int_h e^{-\mathcal{E}(v, h)} dh$ denotes the free energy and the gradient in Eqn. (2) is the difference of the free energy in two phases. In the first *positive phase*, the expectation of the free energy under the data distribution is decreased. In the second *negative phase*, the expectation of the free energy under the model distribution is increased.

Unfortunately, both the NLL in Eqn. (1) and its gradient in Eqn. (2) are intractable in general for two reasons. First, the integral of the latent variables in Eqn. (1) or equivalently the computation of the free energy in Eqn. (2) is intractable. Second, the computation of the partition function in Eqn. (1) or equivalently the negative phase in Eqn. (2) is intractable.

Variational inference. Extensive work introduces deterministic approximations for the intractability of inference, including the mean-field approximation (Welling & Hinton, 2002; Salakhutdinov & Hinton, 2009), the Kikuchi and Bethe approximations (Welling & Sutton, 2005) and the recognition model approach (Salakhutdinov & Larochelle, 2010). In this line of work, the intractability of the partition function is addressed using Monte Carlo based methods.

Contrastive free energy. Contrastive divergence (CD) (Hinton, 2002) addresses the intractability of the partition function by approximating the negative phase in Eqn. (2) as follows:

$$\nabla_{\theta} \mathcal{L}(\theta) = \mathbb{E}_{P_{\mathcal{D}}(v)} [\nabla_{\theta} \mathcal{F}(v)] - \mathbb{E}_{P_{CD}(v)} [\nabla_{\theta} \mathcal{F}(v)], \quad (3)$$

where $P_{CD}(v)$ denotes the empirical distribution obtained by starting from a data point and running several steps of Gibbs sampling according to the model distribution and the free energy $\mathcal{F}(v)$ is assumed to be tractable. Existing methods (Welling & Hinton, 2002; Welling & Sutton, 2005) approximate $\mathcal{F}(v)$ using certain function $\mathcal{G}(v)$ and the gradient of θ is:

$$\nabla_{\theta} \mathcal{L}(\theta) \approx \mathbb{E}_{P_{\mathcal{D}}(v)} [\nabla_{\theta} \mathcal{G}(v)] - \mathbb{E}_{P_{CD}(v)} [\nabla_{\theta} \mathcal{G}(v)]. \quad (4)$$




Figure 1: Illustration of the models involved in AdVIL. From left to right: variational encoder $Q(h|v)$, MRF $P(v, h)$, variational decoder $q(v, h)$ with a simple prior and $q(v, h)$ with an expressive prior.

Although these generalized methods exist, it is nontrivial to extend CD-based methods to general MRFs because the Gibbs sampling procedure is highly dependent on the model structure.

Black-box learning. The recent work of NVIL (Kuleshov & Ermon, 2017) addresses the intractability of the partition function in a black-box manner via a variational upper bound of the partition function:

$$\mathbb{E}_{q(v)} \left[\frac{\tilde{P}(v)^2}{q(v)^2} \right] \geq \mathcal{Z}^2, \quad (5)$$

where $\tilde{P}(v) = e^{-\mathcal{F}(v)}$ is the unnormalized marginal distribution on v and $q(v)$ is a neural variational distribution. As a black-box learning method, NVIL potentially allows application to broader model families and improves the capabilities of probabilistic programming systems (Carpenter et al., 2017). Though promising, NVIL leaves the intractability of inference in an MRF unsolved, and the bound in Eqn. (5) is of high variance and is easily underestimated (Kuleshov & Ermon, 2017).

3 METHOD

As stated above, the black-box inference and learning of MRFs are still largely open. In this paper, we make a step towards solving the problems by a new variational approach. For simplicity, we focus on the resulting objective function in this section. See Appendix A for detailed derivation.

3.1 ADVERSARIAL VARIATIONAL INFERENCE AND LEARNING

First, we rewrite the NLL of the MRF (See an illustration in Fig. 1) as follows:

$$\mathcal{L}(\theta) = -\mathbb{E}_{P_{\mathcal{D}}(v)} [-\mathcal{F}(v)] + \log \mathcal{Z}, \quad (6)$$

where the negative free energy and the log partition function are in the form of a logarithm of an integral. Naturally, we can apply the variational trick (Jordan et al., 1999) twice and approximate the two terms individually. Due to the presence of the minus before the first term in Eqn. (6), the two variational tricks bound the two parts of the NLL in the opposite directions, detailed as below.

Formally, on one hand, we introduce an approximate posterior for the latent variables $Q(h|v)$, which is parameterized as a *neural variational encoder* (See an illustration in Fig. 1), to address the intractability of inference as follows:

$$\mathcal{L}(\theta) \leq \mathbb{E}_{P_{\mathcal{D}}(v)Q(h|v)} [\mathcal{E}(v, h) + \log Q(h|v)] + \log \mathcal{Z} := \mathcal{L}_1(\theta, \phi), \quad (7)$$

where ϕ denotes the trainable parameters in $Q(h|v)$. The upper bound is derived via applying the Jensen inequality and the equality holds if and only if $Q(h|v) = P(h|v)$ for all v . In the bound, the first term is the expected energy, which encourages $Q(h|v)$ to infer latent variables that have low values of the energy function $\mathcal{E}(v, h)$, or equivalently high probabilities of $P(v, h)$. The second term corresponds to the negative conditional entropy of $Q(h|v)$, which increases the uncertainty of $Q(h|v)$. In the paper, we denote the conditional entropy of $Q(h|v)$ as $\mathcal{H}(Q) := -\mathbb{E}_{P_{\mathcal{D}}(v)Q(h|v)} [\log Q(h|v)]$.

On the other hand, we introduce an approximate sampler $q(v, h)$, which is parameterized by a *neural variational decoder* (See Fig. 1), to address the intractability of the partition function as follows:

$$\mathcal{L}_1(\theta, \phi) \geq \mathbb{E}_{P_{\mathcal{D}}(v)Q(h|v)} \left[\underbrace{\mathcal{E}(v, h) + \log Q(h|v)}_{\text{Positive Phase}} \right] - \mathbb{E}_{q(v, h)} \left[\underbrace{\mathcal{E}(v, h) + \log q(v, h)}_{\text{Negative Phase}} \right] := \mathcal{L}_2(\theta, \phi, \psi), \quad (8)$$

where ψ denotes the trainable parameters in $q(v, h)$. The lower bound is derived via applying the Jensen inequality as well, and the equality holds if and only if $q(v, h) = P(v, h)$. It can be seen that

the lower bound given by $q(v, h)$ consists of the entropy (denoted as $\mathcal{H}(q)$) and energy terms, which is similar to the upper bound in Eqn. (7), and the overall objective is in the form of approximate *contrastive free energy* (Hinton, 2002; Welling & Sutton, 2005). Because the double variational trick bounds the NLL in opposite directions as above, we have a minimax optimization problem:

$$\min_{\theta} \min_{\phi} \max_{\psi} \mathcal{L}_2(\theta, \phi, \psi). \quad (9)$$

The minimax formulation has been investigated in GAN (Goodfellow et al., 2014), which defines an adversarial game between a generator and a discriminator to train implicit directed models. Therefore, we name our framework *adversarial variational inference and learning* (AdVIL).

Note that $\mathcal{L}_2(\theta, \phi, \psi)$ is neither an upper bound, nor a lower bound of $\mathcal{L}(\theta)$ due to the double variational trick. However, we argue that solving the optimization problem in Eqn. (9) is reasonable because (1) it is equivalent to optimizing $\mathcal{L}(\theta)$ under the nonparametric assumption, which is similar to GAN (Goodfellow et al., 2014); and (2) it converges to a stationary point of $\mathcal{L}_1(\theta, \phi)$, which is an upper bound of $\mathcal{L}(\theta)$, under a weaker assumption, as stated in the following theoretical analysis.

3.2 THEORETICAL ANALYSIS OF ADVIL

In this section, we present our main convergence theorem. See Appendix C for complementary analysis and all proofs. To establish our convergence theorem, we first prove that the angle between $\frac{\partial \mathcal{L}_2(\theta, \phi, \psi)}{\partial \theta}$ and $\frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}$ is positive if $q(v, h)$ approximates $P(v, h)$ well, as stated in Lemma 1.

Lemma 1. *For any (θ, ϕ) , there exists a symmetric positive definite matrix H such that $\frac{\partial \mathcal{L}_2(\theta, \phi, \psi)}{\partial \theta} = H \frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}$ under the assumption: $\|\sum_{v, h} \delta(v, h) \frac{\partial \mathcal{E}(v, h)}{\partial \theta}\|_2 < \|\frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}\|_2$ if $\|\frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}\|_2 > 0$ and $\|\sum_{v, h} \delta(v, h) \frac{\partial \mathcal{E}(v, h)}{\partial \theta}\|_2 = 0$ if $\|\frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}\|_2 = 0$, where $\delta(v, h) = q(v, h) - P(v, h)$.*

Based on Lemma 1 and other commonly used assumptions in the analysis of stochastic optimization (Bottou et al., 2018), AdVIL converges to a stationary point of $\mathcal{L}_1(\theta, \phi)$, as stated in Theorem 1.

Theorem 1. *Solving the optimization problem in Eqn. (9) using stochastic gradient descent, then (θ, ϕ) converges to a stationary point of $\mathcal{L}_1(\theta, \phi)$ under the following assumptions.*

1. $\mathcal{L}_2(\theta, \phi, \psi)$ is twice differentiable with respect to θ, ϕ and ψ .
2. The gradients of $\mathcal{L}_2(\theta, \phi, \psi)$ with respect to θ, ϕ and ψ are Lipschitz.
3. The errors of the stochastic gradients are bounded (Bottou et al., 2018).
4. The stepsize satisfies the diminishing condition (Bottou et al., 2018).
5. The condition of Lemma 1 holds in each step.

Theorem 1 is build upon the analysis of general stochastic optimization (Bottou et al., 2018). See Appendix C.2 for more details. Note that Assumption 5 is much weaker than the nonparametric assumption used in GAN (Goodfellow et al., 2014) because it does not require $q(v, h) = P(v, h)$ in each step. Theorem 1 also provides insights for implementation of AdVIL. Indeed, Assumption 5 motivates us to use a sufficiently powerful $q(v, h)$ with neural networks and auxiliary variables, and update $q(v, h)$ multiple times per update of $P(v, h)$, as detailed in Sec. 3.3 and Sec. 5.1 respectively.

3.3 SPECIFYING THE VARIATIONAL DISTRIBUTIONS

To efficiently get samples, both variational distributions are directed models. We use a directed neural network that maps v to h as the variational encoder $Q(h|v)$ (Kingma & Welling, 2013).

As for the variational decoder, we first factorize it as the product of a prior over h and a conditional distribution, namely $q(v, h) = q(v|h)q(h)$. It is nontrivial to specify the prior $q(h)$ because the marginal distribution of h in the MRF, i.e. $P(h)$, can be correlated across different dimensions. Consequently, a simple $q(h)$ is not flexible enough to track $P(h)$ and can violate the condition of Lemma 1. To this end, we introduce an auxiliary variable z , which can be discrete or continuous, on top of h and define $q(v, h) = \int_z q(z)q(h|z)q(v|h)dz$.² (See an illustration in Fig. 1.) However,

²An alternative way is to use an autoregressive model as $q(h)$. See results and analysis in Appendix E.3.

the entropy term of $q(v, h)$ is intractable because we need to integrate out the auxiliary variable z . Therefore, we introduce the third variational distribution $r(z|h)$ to approximate the entropy of $q(v, h)$. As in Eqn. (7), applying the standard variational trick gives an upper bound:

$$-\mathbb{E}_{q(v,h)} \log q(v, h) \leq -\mathbb{E}_{q(v,h)} \log q(v|h) - \mathbb{E}_{q(h)r(z|h)} \log \left[\frac{q(h, z)}{r(z|h)} \right], \quad (10)$$

which is unsatisfactory because the estimate is minimized with respect to $r(z|h)$ while maximized with respect to $q(v, h)$. Instead, after some transformations (See details in Appendix A) we get a lower bound as follows:

$$-\mathbb{E}_{q(v,h)} \log q(v, h) \geq -\mathbb{E}_{q(v,h)} \log q(v|h) - \mathbb{E}_{q(h,z)} \log \left[\frac{q(h, z)}{r(z|h)} \right]. \quad (11)$$

The equality holds if and only if $r(z|h) = q(z|h)$ for all h . The difference between the two bounds is subtle: the last expectation in Eqn. (10) is over $q(h)r(z|h)$ but that in Eqn. (11) is over $q(h, z)$. Here, a lower bound is preferable because the estimate is maximized with respect to both $r(z|h)$ and $q(v, h)$ and we can train them simultaneously. For simplicity, we absorb the trainable parameters of $r(z|h)$ into ψ . Note that after introducing z and $r(z|h)$, we can still obtain a convergence theorem of AdvIL under the conditions that $r(z|h)$ approximates $q(z|h)$ well and $q(v, h) = \int q(v, h, z) dz$ is sufficiently close to $P(v, h)$ in every step, together with Assumption 1 to Assumption 4 in Theorem 1.

Following GAN (Goodfellow et al., 2014), we optimize θ , ϕ and ψ jointly using stochastic gradient descent (SGD) in an alternating manner. The partial derivatives of ϕ and ψ are estimated via the reparameterization trick (Kingma & Welling, 2013) for the continuous variables and the Gumbel-Softmax trick (Jang et al., 2016; Maddison et al., 2016) for the discrete variables. See Algorithm 1 in Appendix B for the whole training procedure. Note that ψ is updated $K_1 > 1$ times per update of θ .

4 RELATED WORK

In the context of black-box learning in MRFs, AdvIL competes directly with NVIL (Kuleshov & Ermon, 2017). It seems that the upper bound in Eqn. (5) is suitable for optimization because P and q share the same training direction. However, the bound holds only if the support of \tilde{P} is a subset of the support of q . Further, the Monte Carlo estimate of the upper bound is of high variance. Therefore, the bound of NVIL can be easily underestimated, which results in sub-optimal solutions (Kuleshov & Ermon, 2017). In contrast, though AdvIL arrives at a minimax optimization problem, the estimate of Eqn. (8) is tighter and of lower variance. We empirically verify this argument (See Fig. 3) and systematically compare the two methods (See Tab. 1) in Sec.5.4.

As shown in Eqn. (8), AdvIL is closely related to contrastive divergence (CD) based algorithms (Hinton, 2002; Welling & Sutton, 2005). Further, the partial derivative of θ in AdvIL is as follows:

$$\frac{\partial \mathcal{L}_2(\theta, \phi, \psi)}{\partial \theta} = \mathbb{E}_{P_{\mathcal{D}}(v)Q(h|v)} \left[\frac{\partial}{\partial \theta} \mathcal{E}(v, h) \right] - \mathbb{E}_{q(v,h)} \left[\frac{\partial}{\partial \theta} \mathcal{E}(v, h) \right], \quad (12)$$

which also involves a positive phase and a negative phase naturally and is quite similar to Eqn. (3). However, notably, the two phases average over the (v, h) pairs and only require the knowledge of the energy function without any further assumption of the model in AdvIL. Therefore, AdvIL is more suitable to general MRFs than CD (See empirical evidence in Sec. 5.3).

Apart from the work on approximate inference and learning in MRFs as mentioned above, AdvIL is also related to some directed generative models. Kim & Bengio (2016) jointly trains a deep energy model (Ngiam et al., 2011) and a directed generative model by minimizing the KL-divergence between them. Similar ideas have been highlighted in (Finn et al., 2016; Zhai et al., 2016; Dai et al., 2017; Liu & Wang, 2017). In comparison, generally, AdvIL obtains the objective function in a unified perspective on the black-box inference and learning in general MRFs. Note that dealing with latent variables in MRFs is nontrivial (Kim & Bengio, 2016) and therefore existing work focuses on fully observable models. AdvIL uses a sophisticated decoder with auxiliary variables to handle the latent variables and derives a principled variational approximation of the entropy term instead of the heuristics (Kim & Bengio, 2016; Zhai et al., 2016).

Adversarially learned inference (ALI) (Donahue et al., 2016; Dumoulin et al., 2016) is formulated as a similar minimax optimization problem as AdvIL. However, AdvIL involves extra entropy-based regularization and has a different update rule for the variational distributions compared to




Figure 2: Results of AdvIL on the binary digits dataset. (a-c) compare the value of the variational approximations and the corresponding ground truths. All bounds are rather tight after 5,000 iterations. (d) shows that the RBM loss tends to zero and the model converges gradually.

Table 1: Anneal importance sampling (AIS) results in RBM. The results are recorded on the test set according to the best validation performance and averaged over three runs. AdvIL outperforms NVIL consistently and significantly. See the standard deviations in Appendix E.4.

Method	Digits	Adult	Connect4	DNA	Mushrooms	NIPS-0-12	Ocr-letters	RCV1
NVIL-mean	-27.36	-20.05	-24.71	-97.71	-29.28	-290.01	-47.56	-50.47
AdvIL-mean	-26.34	-19.29	-21.95	-97.59	-19.59	-276.42	-45.64	-50.22

ALI. Further, the MRF in AdvIL itself can produce new samples as a generative model while the discriminator in ALI cannot do the same thing.

5 EXPERIMENTS

In this section, we evaluate AdvIL in restricted Boltzmann machines (RBM) (Ackley et al., 1985), deep Boltzmann machines (DBM) (Salakhutdinov & Hinton, 2009) and Gaussian restricted Boltzmann machines (GRBM) (Hinton & Salakhutdinov, 2006) on the Digits dataset, the UCI binary databases (Dheeru & Karra, 2017) and the Frey faces datasets (See detailed settings in Appendix D and the source code³). We compare AdvIL with strong baseline methods systematically and show the promise of AdvIL to learn a broad family of models effectively as a black-box method.

5.1 EMPIRICAL ANALYSIS OF ADVIL

We present a detailed analysis of AdvIL in RBM, whose energy function is defined as $\mathcal{E}(v, h) = -b^\top v - v^\top W h - c^\top h$. The conditional distributions of an RBM are tractable, but we still treat $P(h|v)$ as unknown and train AdvIL in a fully black-box manner. The analysis is performed on the Digits dataset and we augment the data of five times by shifting the digits following the protocol in (Kuleshov & Ermon, 2017). The dimensions of v , h and z are 64, 15 and 10, respectively. Therefore, the log partition function of the RBM and the entropy of the decoder can be computed by brute force.

Firstly, we empirically validate AdvIL in Fig. 2. Specifically, Panel (a) shows that the variational encoder $Q(h|v)$ provides a tight upper bound of the free energy after 2,000 iterations. Panel (b) demonstrates that the variational distribution $r(z|h)$ estimate the entropy of $q(v, h)$ accurately. Panel (c) shows that $q(v, h)$ can successfully track the log partition function after 5,000 iterations. Panel (d) presents that the RBM loss balances well between the negative phase and positive phase, and the model converges gradually. See Appendix E.1 for an empirical test of Assumption 5 in Theorem 1.

Secondly, we empirically show that both P and q can generate data samples in Appendix E.2.

Lastly, we analyze the sensitivity of K_1 . Theoretically, enlarging K_1 will make $q(v, h)$ and $P(v, h)$ to be close and then help the convergence according to Theorem 1. As shown in Fig. 3 (a), a larger K_1 at least won't hurt the convergence, which agrees with Theorem 1. Though $K_1 = 15$ is sufficient on the Digits dataset, we use $K_1 = 100$ as a default setting for AdvIL on larger datasets.

³See the source code in <https://anonymous.4open.science/r/8c779fbc-6394-40c7-8273-e52504814703/>.




Figure 3: (a) Sensitivity analysis of K_1 on the Digits dataset. (b-d) Learning curves of NVIL, AdvIL and PCD on the Mushrooms dataset. Compared with NVIL, AdvIL provides a tighter and lower variance estimate of $\log \mathcal{Z}$ and achieves better performance. Compared with PCD, AdvIL can track the log partition function and achieve comparable results though trained in a black-box manner.

Table 2: AIS results in DBM. The results are recorded according to the best validation performance and averaged by three runs. AdvIL achieves higher averaged AIS results on five out of eight datasets and has a better overall performance than VCD. See the standard deviations in Appendix E.4.

Method	Digits	Adult	Connect4	DNA	Mushrooms	NIPS-0-12	Ocr-letters	RCV1
VCD-mean	-28.49	-22.26	-26.79	-97.59	-23.15	-356.26	-45.77	-50.83
AdvIL-mean	-27.89	-20.29	-26.34	-99.40	-21.21	-287.15	-48.38	-51.02

5.2 RBM RESULTS

To the best of our knowledge, NVIL (Kuleshov & Ermon, 2017) is the only existing black-box learning method for MRFs and hence it is the most direct competitor of AdvIL. In this section, we provide a systematic comparison and analysis of these two methods in terms of the log-likelihood results on the UCI databases (Dheeru & Karra, 2017).

For a fair comparison, we use the widely-adopted and accurate anneal importance sampling (AIS) (Salakhutdinov & Murray, 2008) metric for quantitative evaluation. Besides, we carefully perform grid search over the default settings of NVIL (Kuleshov & Ermon, 2017) and our settings based on their code, and choose the best configuration including $K_1 = 100$ (See details in Appendix D). We directly compare with the best version of NVIL in Tab. 1. It can be seen that AdvIL consistently outperforms NVIL on all datasets, which demonstrate the effectiveness of AdvIL. Besides, the time complexity of AdvIL is comparable to that of NVIL with the same hyperparameters.

We compare the learning curves of NVIL and AdvIL on the Mushroom dataset. As shown in Fig. 3 (b), the upper bound of NVIL is underestimated after 4,000 iterations and then the model can get worse or even diverge. In contrast, as shown in Fig. 3 (c), the lower bound of AdvIL is consistently valid. Besides, the estimate of NVIL is looser and of higher variance than that of AdvIL. The results agree with our analysis in Sec. 4 and explain why AdvIL significantly outperforms NVIL. Further, as shown in Fig. 3 (d), AdvIL is comparable to persistent contrastive divergence (PCD) (Tieleman, 2008), which leverages the tractability of the conditional distributions in an RBM.

5.3 DBM RESULTS

We would like to demonstrate that AdvIL has the ability to deal with highly intractable models such as a DBM conveniently and effectively, compared with standard CD-based methods (Hinton, 2002; Welling & Hinton, 2002; Welling & Sutton, 2005) and NVIL (Kuleshov & Ermon, 2017).

DBM (Salakhutdinov & Hinton, 2009) is a powerful family of deep models that stack multiple RBMs together. The energy function of a two-layer DBM is defined as $\mathcal{E}(v, h_1, h_2) = -b^\top v - v^\top W_1 h_1 - c_1^\top h_1 - h_1^\top W_2 h_2 - c_2^\top h_2$. Learning a DBM is challenging because $P(h_1, h_2|v)$ is not tractable and CD (Hinton, 2002) is not applicable. Inspired by (Welling & Hinton, 2002; Welling & Sutton, 2005), we construct a variational CD (VCD) baseline by employing the same variational encoder $Q(h_1, h_2|v)$ as in AdvIL. The free energy is approximated by the same upper bound as in Eqn. (7), which is minimized with respect to the parameters in $Q(h_1, h_2|v)$. The gradient of the




Figure 4: Filters and samples of a GRBM learned by AdVIL on the Frey faces dataset. (a) presents the training data. (b) presents the first 40 filters of the GRBM. (c) and (d) show random samples from the variational decoder and the GRBM, respectively. We present the mean of v for better visualization.

parameters in the DBM is given by Eqn. (4), where the Gibbs sampling procedure is approximated by $h_1 \sim Q(h_1|v)$ and $v \sim P(v|h_1)$. Note that AdVIL can be directly applied to this case. As for the time complexity, the training speed of AdVIL is around ten times slower than that of VCD in our implementation. However, the approximate inference and sampling procedure of AdVIL is very efficient thanks to the directed variational distributions.

The log-likelihood results on the UCI databases are shown in Tab. 2. It can be seen that AdVIL has a better overall performance even trained in a black-box manner, which shows the promise of AdVIL.

We also extend NVIL by using the same $Q(h_1, h_2|v)$ and $q(v, h_1, h_2)$ as AdVIL. However, NVIL diverges after 300 iterations and gets bad AIS results (e.g., less than -40 on Digits) in our implementation. A potential reason is that the upper bound given by q in NVIL can be underestimated if q is high-dimensional, as analyzed in Sec. 4 and Fig. 3. Note that $q(v, h_1, h_2)$ in DBM involves latent variables and has a higher dimension (e.g. 164 on the Digits dataset) than $q(v)$ in RBM (e.g. 64 on the Digits dataset). The results again demonstrate the advantages of AdVIL over NVIL.

5.4 GRBM RESULTS

We now show the ability of AdVIL to learn a GRBM on the continuous Frey faces dataset. The energy function of a GRBM is $\mathcal{E}(v, h) = \frac{1}{2\sigma^2} \|v - b\|^2 - c^\top h - \frac{1}{\sigma} v^\top W h$, where σ is the standard deviation of the Gaussian likelihood and is set as 1 manually. We standardize the data by subtracting the mean and dividing by the standard deviation. The dimensions of h and z are 200 and 50, respectively.

Though a GRBM is more sensitive to the hyperparameters and hence harder to train than an RBM (Cho et al., 2011; 2013), AdVIL can successfully capture the underlying data distribution using the default hyperparameters (See Appendix D). As shown in Fig. 4, the samples from both the GRBM (via Gibbs sampling after 100,000 burn-in steps) and the decoder are meaningful faces. Besides, the filters of the GRBM outline diverse prototypes of faces, which accords with our expectation.

In summary, the results of the three models together demonstrate that AdVIL can learn a broad family of models conveniently and effectively in a fully black-box manner.

6 CONCLUSION AND DISCUSSION

A novel black-box learning and inference method for undirected graphical models, called adversarial variational inference and learning (AdVIL), is proposed. The key to AdVIL is a double variational trick that approximates the negative free energy and the log partition function separately. A formal convergence theorem, which provides insights for implementation, is established for AdVIL. Empirical results show that AdVIL can deal with a broad family of MRFs in a fully black-box manner and outperforms both the standard contrastive divergence method and the black-box NVIL algorithm.

Though AdVIL shows promising results, we emphasize that the black-box learning and inference of the MRFs are far from completely solved, especially on high-dimensional data. The two intractability problems of MRFs are distinct since the posterior of the latent variables is *local* in terms of v but the partition function is *global* by integrating out v . The additional integral makes estimating the partition function much more challenging. In AdVIL, simply increasing the number of updates of the decoder to obtain a tighter estimate of the partition function on high-dimensional data can be expensive. Future work with tighter bounds on the partition function can benefit the black-box learning in MRFs. Nevertheless, while not perfect, AdVIL can relieve some headache of training an MRF.

REFERENCES

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. TensorFlow: A system for large-scale machine learning. 2016.
- David H Ackley, Geoffrey E Hinton, and Terrence J Sejnowski. A learning algorithm for boltzmann machines. *Cognitive science*, 9(1):147–169, 1985.
- Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *Siam Review*, 60(2):223–311, 2018.
- Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- Bob Carpenter, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: A probabilistic programming language. *Journal of statistical software*, 76(1), 2017.
- Kyung Hyun Cho, Tapani Raiko, and Alexander Ilin. Gaussian-bernoulli deep boltzmann machine. In *Neural Networks (IJCNN), The 2013 International Joint Conference on*, pp. 1–7. IEEE, 2013.
- KyungHyun Cho, Alexander Ilin, and Tapani Raiko. Improved learning of gaussian-bernoulli restricted boltzmann machines. In *International conference on artificial neural networks*, pp. 10–17. Springer, 2011.
- Zihang Dai, Amjad Almahairi, Philip Bachman, Eduard Hovy, and Aaron Courville. Calibrating energy-based generative adversarial networks. *arXiv preprint arXiv:1702.01691*, 2017.
- Dua Dheeru and Taniskidou E Karra. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. Adversarial feature learning. *arXiv preprint arXiv:1605.09782*, 2016.
- Vincent Dumoulin, Ishmael Belghazi, Ben Poole, Olivier Mastropietro, Alex Lamb, Martin Arjovsky, and Aaron Courville. Adversarially learned inference. *arXiv preprint arXiv:1606.00704*, 2016.
- Chelsea Finn, Paul Christiano, Pieter Abbeel, and Sergey Levine. A connection between generative adversarial networks, inverse reinforcement learning, and energy-based models. *arXiv preprint arXiv:1611.03852*, 2016.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, pp. 6626–6637, 2017.
- Geoffrey E Hinton. Training products of experts by minimizing contrastive divergence. *Neural computation*, 14(8):1771–1800, 2002.
- Geoffrey E Hinton and Ruslan R Salakhutdinov. Reducing the dimensionality of data with neural networks. *science*, 313(5786):504–507, 2006.
- Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.
- Taesup Kim and Yoshua Bengio. Deep directed generative models with energy-based probability estimation. *arXiv preprint arXiv:1606.03439*, 2016.

- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *Advances in neural information processing systems*, pp. 109–117, 2011.
- Volodymyr Kuleshov and Stefano Ermon. Neural variational inference and learning in undirected graphical models. In *Advances in Neural Information Processing Systems*, pp. 6734–6743, 2017.
- John Lafferty, Andrew McCallum, and Fernando CN Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. 2001.
- Hugo Larochelle and Iain Murray. The neural autoregressive distribution estimator. In *The Proceedings of the 14th International Conference on Artificial Intelligence and Statistics*, volume 15 of *JMLR: W&CP*, pp. 29–37, 2011.
- Qiang Liu and Dilin Wang. Learning deep energy models: Contrastive divergence vs. amortized mle. *arXiv preprint arXiv:1707.00797*, 2017.
- Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.
- Xiao-Li Meng and Wing Hung Wong. Simulating ratios of normalizing constants via a simple identity: a theoretical exploration. *Statistica Sinica*, pp. 831–860, 1996.
- Andriy Mnih and Karol Gregor. Neural variational inference and learning in belief networks. *arXiv preprint arXiv:1402.0030*, 2014.
- Radford M Neal. Probabilistic inference using markov chain monte carlo methods. 1993.
- Radford M Neal. Annealed importance sampling. *Statistics and computing*, 11(2):125–139, 2001.
- Jiquan Ngiam, Zhenghao Chen, Pang W Koh, and Andrew Y Ng. Learning deep energy models. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pp. 1105–1112, 2011.
- Rajesh Ranganath, Sean Gerrish, and David Blei. Black box variational inference. In *Artificial Intelligence and Statistics*, pp. 814–822, 2014.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- Ruslan Salakhutdinov and Geoffrey Hinton. Deep Boltzmann machines. In *Proceedings of the twelfth international conference on artificial intelligence and statistics*, 2009.
- Ruslan Salakhutdinov and Hugo Larochelle. Efficient learning of deep Boltzmann machines. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 693–700, 2010.
- Ruslan Salakhutdinov and Iain Murray. On the quantitative analysis of deep belief networks. In *Proceedings of the 25th international conference on Machine learning*, pp. 872–879. ACM, 2008.
- John Schulman, Nicolas Heess, Theophane Weber, and Pieter Abbeel. Gradient estimation using stochastic computation graphs. In *Advances in Neural Information Processing Systems*, pp. 3528–3536, 2015.
- Tijmen Tieleman. Training restricted boltzmann machines using approximations to the likelihood gradient. In *Proceedings of the 25th international conference on Machine learning*, pp. 1064–1071. ACM, 2008.
- Martin J Wainwright and Michael I Jordan. Log-determinant relaxation for approximate inference in discrete markov random fields. *IEEE transactions on signal processing*, 54(6):2099–2109, 2006.

Martin J Wainwright, Tommi S Jaakkola, and Alan S Willsky. A new class of upper bounds on the log partition function. *IEEE Transactions on Information Theory*, 51(7):2313–2335, 2005.

Max Welling and Geoffrey E Hinton. A new learning algorithm for mean field boltzmann machines. In *International Conference on Artificial Neural Networks*, pp. 351–357. Springer, 2002.

Max Welling and Charles A Sutton. Learning in markov random fields with contrastive free energies. In *AISTATS*. Citeseer, 2005.

Shuangfei Zhai, Yu Cheng, Rogerio Feris, and Zhongfei Zhang. Generative adversarial networks as variational training of energy based models. *arXiv preprint arXiv:1611.01799*, 2016.

A DERIVATION OF THE OBJECTIVE FUNCTION

Here we derive the objective function of AdvIL in detail. Let θ , ϕ and ψ denote the trainable parameters in the MRF, the variational encoder and the variational decoder, respectively. The first variational trick bounds the free energy as follows:

$$\begin{aligned}\mathcal{L}(\theta) &= -\mathbb{E}_{P_{\mathcal{D}}(v)} \left[\log \left(\int_h e^{-\mathcal{E}(v,h)} dh \right) \right] + \log \mathcal{Z} \\ &= -\mathbb{E}_{P_{\mathcal{D}}(v)} \log \left[\int_h Q(h|v) \frac{e^{-\mathcal{E}(v,h)}}{Q(h|v)} dh \right] + \log \mathcal{Z} \\ &\leq \mathbb{E}_{P_{\mathcal{D}}(v)Q(h|v)} [\mathcal{E}(v, h) + \log Q(h|v)] + \log \mathcal{Z} := \mathcal{L}_1(\theta, \phi),\end{aligned}$$

where the bound is derived via applying the Jensen inequality and the equality holds if and only if $Q(h|v) = P(h|v)$ for all v .

The second variational trick bounds the log partition function as follows:

$$\begin{aligned}\mathcal{L}_1(\theta, \phi) &= \mathbb{E}_{P_{\mathcal{D}}(v)Q(h|v)} [\mathcal{E}(v, h) + \log Q(h|v)] + \log \left(\int_v \int_h e^{-\mathcal{E}(v,h)} dv dh \right) \\ &= \mathbb{E}_{P_{\mathcal{D}}(v)Q(h|v)} [\mathcal{E}(v, h) + \log Q(h|v)] + \log \left(\int_v \int_h q(v, h) \frac{e^{-\mathcal{E}(v,h)}}{q(v, h)} dv dh \right) \\ &\geq \mathbb{E}_{P_{\mathcal{D}}(v)Q(h|v)} [\mathcal{E}(v, h) + \log Q(h|v)] + \mathbb{E}_{q(v,h)} \left[\log \left(\frac{e^{-\mathcal{E}(v,h)}}{q(v, h)} \right) \right] \\ &= \mathbb{E}_{P_{\mathcal{D}}(v)Q(h|v)} \left[\underbrace{\mathcal{E}(v, h)}_{\text{energy term}} + \underbrace{\log Q(h|v)}_{\text{entropy term}} \right] - \mathbb{E}_{q(v,h)} \left[\underbrace{\mathcal{E}(v, h)}_{\text{energy term}} + \underbrace{\log q(v, h)}_{\text{entropy term}} \right] \\ &:= \mathcal{L}_2(\theta, \phi, \psi),\end{aligned}$$

where the bound is also derived via applying the Jensen inequality and the equality holds if and only if $q(v, h) = P(v, h)$.

To enhance the expressive power of the variational decoder, we introduce an auxiliary variable z and define $q(v, h) = \int_z q(z)q(h|z)q(v|h)dz$, which makes the entropy term in the negative phase intractable. To address the problem, we propose the third variational approximation. First, we can decompose the entropy of $q(v, h)$ as $-\mathbb{E}_{q(v,h)} \log q(v, h) = -\mathbb{E}_{q(v,h)} \log q(v|h) - \mathbb{E}_{q(h)} \log q(h)$ and we only need to approximate $-\mathbb{E}_{q(h)} \log q(h)$. However, simply applying the standard variational trick as above, we get an upper bound as follows:

$$\begin{aligned}-\mathbb{E}_{q(h)} \log q(h) &= -\mathbb{E}_{q(h)} \log \int_z q(h, z) dz \\ &= -\mathbb{E}_{q(h)} \log \int_z r(z|h) \frac{q(h, z)}{r(z|h)} dz \\ &\leq -\mathbb{E}_{q(h)r(z|h)} \log \left[\frac{q(h, z)}{r(z|h)} \right],\end{aligned}$$

Algorithm 1 Adversarial variational inference and learning by stochastic gradient descent

```

1: Input: Constants  $K_1$  and  $K_2$ , learning rate schemes  $\alpha$  and  $\gamma$ , randomly initialized  $\theta$ ,  $\phi$  and  $\psi$ 
2: repeat
3:   for  $i = 1, \dots, K_1$  do
4:     Sample a batch of  $(v, h, z) \sim q(v, h, z)$ 
5:     Estimate the objective of  $q$  and  $r$  according to Eqn. (11) and the negative phase in Eqn. (8)
6:     Update  $\psi$  to maximize the objective according to  $\alpha$ 
7:   end for
8:   for  $i = 1, \dots, K_2$  do
9:     Sample a batch of  $(v, h) \sim P_{\mathcal{D}}(v)Q(h|v)$ 
10:    Estimate the objective of  $Q$  according to the positive phase in Eqn. (8)
11:    Update  $\phi$  to minimize the objective according to  $\gamma$ 
12:  end for
13:  Sample a batch of  $(v, h) \sim P_{\mathcal{D}}(v)Q(h|v)$  and another batch of  $(v, h) \sim q(v, h)$ 
14:  Estimate the objective of  $P$  according to Eqn. (8)
15:  Update  $\theta$  to minimize the objective according to  $\gamma$ 
16: until Convergence or reaching certain threshold

```

which is not satisfactory because the optimization problem will be $\min_P \min_Q \max_q \min_r$. Instead, we derive a lower bound as follows:

$$\begin{aligned}
-\mathbb{E}_{q(h)} \log q(h) &= -\mathbb{E}_{q(h)} \log q(h) - \mathbb{E}_{q(h,z)} \log q(z|h) + \mathbb{E}_{q(h,z)} \log q(z|h) \\
&= -\mathbb{E}_{q(h,z)} \log q(h, z) + \mathbb{E}_{q(h,z)} \log q(z|h) \\
&= -\mathbb{E}_{q(h,z)} \log \left[\frac{q(h, z)}{r(z|h)} \right] + \mathbb{D}_{KL}(q(z|h) || r(z|h)) \\
&\geq -\mathbb{E}_{q(h,z)} \log \left[\frac{q(h, z)}{r(z|h)} \right],
\end{aligned}$$

where $\mathbb{D}_{KL}(\cdot || \cdot)$ denotes the KL-divergence and the equality holds if and only if $r(z|h) = q(z|h)$ for all h . The difference between the two bounds is that the expectation is taken over $q(h)r(z|h)$ in the upper bound while over $q(h, z)$ in the lower bound. Using the lower bound, the optimization problem will be $\min_P \min_Q \max_p \max_r$.

B FORMAL TRAINING PROCEDURE

The formal training procedure of AdVIL is presented in Algorithm 1.

C DETAILED THEORETICAL ANALYSIS

For simplicity, we consider discrete v and h (e.g., in an RBM) and the analysis can be extended to the continuous cases. We assume $v \in \{0, 1\}^{d_v}$ and $h \in \{0, 1\}^{d_h}$, where d_v and d_h are the dimensions of the visible and latent variables respectively.

C.1 NONPARAMETRIC CASE

We first analyze the nonparametric case in Proposition 1 as follows.

Proposition 1. *For any $P(v, h) = \exp(-\mathcal{E}(v, h))/\mathcal{Z}$, $\mathcal{L}_2(\theta, \phi, \psi)$ is a tight estimate of the negative log-likelihood of $P(v)$, under the following assumptions*

1. $Q(h|v)$ and $q(v, h)$ are nonparametric.
2. The inner optimization over $Q(h|v)$ and $q(v, h)$ can get their optima.

Proof. Given $P(v, h)$, i.e., $\mathcal{E}(v, h)$, to find $q^*(v, h)$, we optimize \mathcal{L}_2 over $\{q(v, h) | v \in \{0, 1\}^{d_v}, h \in \{0, 1\}^{d_h}\}$ (we will use a shortcut $\{q(v, h)\}$ for simplicity). The optimization problem is equivalent to:

$$\begin{aligned} & \min_{\{q(v, h)\}} \sum_{v, h} q(v, h) [\mathcal{E}(v, h) + \log q(v, h)] \\ & \text{subject to: } \sum_{v, h} q(v, h) = 1, \\ & \quad q(v, h) \geq 0, \forall v, h. \end{aligned}$$

Note that the objective function is convex since its Hessian matrix is positive semi-definite. Besides, the constraints are linear. Therefore, it is a convex optimization problem. Further, we can verify that the Slater's condition (Boyd & Vandenberghe, 2004) holds when q is uniform and then the strong duality holds. Then, we can use the KKT conditions to solve the optimization problem.

The Lagrangian $\mathcal{G}(\{q(v, h)\}, \lambda, \{\mu(v, h)\})$ is:

$$\sum_{v, h} q(v, h) [\mathcal{E}(v, h) + \log q(v, h)] + \lambda (\sum_{v, h} q(v, h) - 1) + \sum_{v, h} \mu(v, h) q(v, h),$$

where λ and $\{\mu(v, h)\}$ are the associated Lagrange multipliers.

To satisfy the stationarity, we take gradients with respect to $q(v, h)$ for all (v, h) and get:

$$[\mathcal{E}(v, h) + \log q^*(v, h) + 1] + \lambda + \mu(v, h) = 0,$$

which implies

$$q^*(v, h) = \exp(-\mathcal{E}(v, h) - (1 + \lambda + \mu(v, h))).$$

According to the complementary slackness, we have

$$\mu(v, h) q^*(v, h) = 0, \forall v, h,$$

which implies $\mu(v, h) = 0, \forall v, h$, since $q^*(v, h) > 0, \forall v, h$.

To satisfy the primal equality constraint, we have

$$\sum_{v, h} q^*(v, h) = \sum_{v, h} \exp(-\mathcal{E}(v, h) - (1 + \lambda)) = 1,$$

which implies

$$q^*(v, h) = \frac{\exp(-\mathcal{E}(v, h))}{\sum_{v', h'} \exp(-\mathcal{E}(v', h'))} = P(v, h), \forall v, h.$$

To find $Q^*(h|v)$, we optimize \mathcal{L}_2 over $\{Q(h|v) | v \in \{0, 1\}^{d_v}, h \in \{0, 1\}^{d_h}\}$ (we will use a shortcut $\{Q(h|v)\}$ for simplicity). The optimization problem is equivalent to:

$$\begin{aligned} & \min_{\{Q(h|v)\}} \sum_v P_{\mathcal{D}}(v) \sum_h Q(h|v) [\mathcal{E}(v, h) + \log Q(h|v)] \\ & \text{subject to: } \sum_h Q(h|v) = 1, \forall v, \\ & \quad Q(h|v) \geq 0, \forall v, h. \end{aligned}$$

Similar to the above procedure, we can get

$$Q^*(h|v) = \frac{\exp(-\mathcal{E}(v, h))}{\sum_{h'} \exp(-\mathcal{E}(v, h'))} = P(h|v), \forall v, h.$$

Under the assumptions that (1) $Q(h|v)$ and $q(v, h)$ are nonparametric, and (2) the inner optimization over ψ and ϕ can get the optimum, the optimal variational distributions $P(v, h)$ and $P(h|v)$ can be obtained. Plugging them back into \mathcal{L}_2 , we get

$$\begin{aligned} \mathcal{L}_2 &= \mathbb{E}_{P_{\mathcal{D}}(v)P(h|v)} [\mathcal{E}(v, h) + \log P(h|v)] - \mathbb{E}_{P(v, h)} [\mathcal{E}(v, h) + \log P(v, h)] \\ &= \mathbb{E}_{P_{\mathcal{D}}(v)P(h|v)} \left[-\log \sum_h e^{-\mathcal{E}(v, h)} \right] + \mathbb{E}_{P(v, h)} [\log \mathcal{Z}] \\ &= \mathbb{E}_{P_{\mathcal{D}}(v)} [\mathcal{F}(v)] + \log \mathcal{Z} = \mathcal{L}. \end{aligned}$$

□

Remark Similar to Theorem 1 in (Goodfellow et al., 2014), Proposition 1 is under the nonparametric assumption, which is relaxed in our following analysis. Namely, we will consider more practical cases where $q(v, h)$ may not be exactly the same as $P(v, h)$ during training.

C.2 MAIN CONVERGENCE THEOREM

For convenience, we summarize the training dynamics of Algorithm 1 with $K_1 = 1, K_2 = 1$ and the exact gradients (not the stochastic ones), as follows:

$$\begin{aligned}\psi_{k+1} &= \psi_k + \alpha_k \frac{\partial \mathcal{L}_2(\theta_k, \phi_k, \psi_k)}{\partial \psi}, \\ \phi_{k+1} &= \phi_k - \gamma_k \frac{\partial \mathcal{L}_2(\theta_k, \phi_k, \psi_{k+1})}{\partial \phi}, \\ \theta_{k+1} &= \theta_k - \gamma_k \frac{\partial \mathcal{L}_2(\theta_k, \phi_k, \psi_{k+1})}{\partial \theta},\end{aligned}\tag{13}$$

where $k = 1, 2, \dots$. We will prove that even though we are optimizing $\mathcal{L}_2(\theta, \phi, \psi)$, (θ_k, ϕ_k) converges to a stationary point of $\mathcal{L}_1(\theta, \phi)$ under certain conditions. To establish this, we first prove that the angle between $\frac{\partial \mathcal{L}_2(\theta, \phi, \psi)}{\partial \theta}$ and $\frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}$ are sufficiently positive if $q(v, h)$ and $P(v, h)$ satisfy certain conditions, as summarized in Lemma 1.

Lemma 1. *For any (θ, ϕ) , there exists a symmetric positive definite matrix H such that $\frac{\partial \mathcal{L}_2(\theta, \phi, \psi)}{\partial \theta} = H \frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}$ under the assumption: $\|\sum_{v,h} \delta(v, h) \frac{\partial \mathcal{E}(v, h)}{\partial \theta}\|_2 < \|\frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}\|_2$ if $\|\frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}\|_2 > 0$ and $\|\sum_{v,h} \delta(v, h) \frac{\partial \mathcal{E}(v, h)}{\partial \theta}\|_2 = 0$ if $\|\frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}\|_2 = 0$, where $\delta(v, h) = q(v, h) - P(v, h)$.*

Proof. According to the Cauchy-Schwarz inequality, we have

$$\left\langle \frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}, \sum_{v,h} \delta(v, h) \frac{\partial \mathcal{E}(v, h)}{\partial \theta} \right\rangle \leq \left\| \frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta} \right\|_2 \left\| \sum_{v,h} \delta(v, h) \frac{\partial \mathcal{E}(v, h)}{\partial \theta} \right\|_2.$$

If $\|\frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}\|_2 > 0$, according to the assumption $\|\sum_{v,h} \delta(v, h) \frac{\partial \mathcal{E}(v, h)}{\partial \theta}\|_2 < \|\frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}\|_2$, we have

$$\left\langle \frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}, \sum_{v,h} \delta(v, h) \frac{\partial \mathcal{E}(v, h)}{\partial \theta} \right\rangle < \left\| \frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta} \right\|_2^2 = \left\langle \frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}, \frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta} \right\rangle,$$

which implies that

$$\left\langle \frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}, \frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta} - \sum_{v,h} \delta(v, h) \frac{\partial \mathcal{E}(v, h)}{\partial \theta} \right\rangle > 0.$$

According to the definitions of $\mathcal{L}_1(\theta, \phi)$ and $\mathcal{L}_2(\theta, \phi, \psi)$, we have

$$\begin{aligned}& \frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta} - \sum_{v,h} \delta(v, h) \frac{\partial \mathcal{E}(v, h)}{\partial \theta} \\ &= \sum_{v,h} [P_{\mathcal{D}}(v)Q(h|v) - P(v, h)] \frac{\partial \mathcal{E}(v, h)}{\partial \theta} - \sum_{v,h} [q(v, h) - P(v, h)] \frac{\partial \mathcal{E}(v, h)}{\partial \theta} \\ &= \sum_{v,h} [P_{\mathcal{D}}(v)Q(h|v) - q(v, h)] \frac{\partial \mathcal{E}(v, h)}{\partial \theta} = \frac{\partial \mathcal{L}_2(\theta, \phi, \psi)}{\partial \theta},\end{aligned}$$

which implies that

$$\left\langle \frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}, \frac{\partial \mathcal{L}_2(\theta, \phi, \psi)}{\partial \theta} \right\rangle > 0.$$

Equivalently, there exists a symmetric positive definite matrix H such that $\frac{\partial \mathcal{L}_2(\theta, \phi, \psi)}{\partial \theta} = H \frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}$. Note that this also holds when $\|\frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}\|_2 = 0$ (i.e., $\frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta} = \vec{0}$) because $\|\frac{\partial \mathcal{L}_2(\theta, \phi, \psi)}{\partial \theta}\|_2 \leq \|\frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \theta}\|_2 + \|\sum_{v,h} \delta(v, h) \frac{\partial \mathcal{E}(v, h)}{\partial \theta}\|_2 = 0$ (i.e., $\frac{\partial \mathcal{L}_2(\theta, \phi, \psi)}{\partial \theta} = \vec{0}$), according to the assumption. \square

Remark Lemma 1 assumes that $q(v, h)$ and $P(v, h)$ are sufficiently close, which is encouraged by choosing a sufficiently powerful family of $q(v, h)$ and updating ψ multiple times per update of θ , i.e. $K_1 > 1$. If Lemma 1 holds, optimizing $\mathcal{L}_2(\theta, \phi, \psi)$ with respect to θ will decrease $\mathcal{L}_1(\theta, \phi)$ with a sufficiently small stepsize. Also note that for any (θ, ψ) , $\frac{\partial \mathcal{L}_2(\theta, \phi, \psi)}{\partial \phi} = \frac{\partial \mathcal{L}_1(\theta, \phi)}{\partial \phi}$ and therefore, optimizing $\mathcal{L}_2(\theta, \phi, \psi)$ with respect to ϕ will decrease $\mathcal{L}_1(\theta, \phi)$ with a sufficiently small stepsize.

Based on Lemma 1 and other commonly used assumptions in the analysis of stochastic gradient descent (Bottou et al., 2018), Algorithm 1 converges to a stationary point of $\mathcal{L}_1(\theta, \phi)$, as stated in Theorem 1.

Theorem 1. *Solving the optimization problem in Eqn. (9) using stochastic gradient descent according to Algorithm 1, then*

$$\lim_{k \rightarrow \infty} \mathbb{E}[\|\frac{\partial \mathcal{L}_1(\theta_k, \phi_k)}{\partial \theta}\|_2^2] = 0,$$

under the following assumptions.

1. The condition of Corollary 4.12 in (Bottou et al., 2018): $\mathcal{L}_2(\theta, \phi, \psi)$ is twice differentiable with respect to θ, ϕ and ψ .
2. Assumption 4.1 in (Bottou et al., 2018): the gradients of $\mathcal{L}_2(\theta, \phi, \psi)$ with respect to θ, ϕ and ψ are Lipschitz.
3. Assumption 4.3 in (Bottou et al., 2018): the first and second moments of the stochastic gradients are bounded by the expected gradients.
4. The stepsize satisfies the diminishing condition (Bottou et al., 2018), i.e., $\alpha_k = \gamma_k$, $\sum_{k=1}^{\infty} \gamma_k = \infty$, $\sum_{k=1}^{\infty} \gamma_k^2 < \infty$.
5. The condition of Lemma 1 holds in each step k . Therefore,

$$\forall k, \exists H_k, \frac{\partial \mathcal{L}_2(\theta_k, \phi_k, \psi_{k+1})}{\partial \theta} = H_k \frac{\partial \mathcal{L}_1(\theta_k, \phi_k)}{\partial \theta}.$$

Proof. See Corollary 4.12 in (Bottou et al., 2018). □

Remark Assumption 1 and Assumption 2 in Theorem 1 are ensured because we use the sigmoid and tanh activation functions. Assumption 3 and Assumption 4 in Theorem 1 are ensured by the sampling and learning rate schemes of the stochastic gradient descent. Assumption 5 in Theorem 1 is much weaker than the nonparametric assumption of Proposition 1 but still requires large K_1 .

C.3 COMPLEMENTARY CONVERGENCE THEOREM

Heusel et al. (2017) propose a two-time scale update rule to train minimax optimization problems with a convergence guarantee even using $K_1 = 1$. AdvIL converges if using the same training method as in (Heusel et al., 2017), which is summarized in Proposition 2.

Proposition 2. *AdvIL trained with a two-time scale update rule (Heusel et al., 2017) converges to a stationary local Nash equilibrium almost surely under the following assumptions.*

1. The gradients with respect to θ, ϕ and ψ are Lipschitz.
2. $\sum_k \alpha_k = \infty$, $\sum_k \alpha_k^2 < \infty$, $\sum_k \gamma_k = \infty$, $\sum_k \gamma_k^2 < \infty$, $\gamma_k = o(\alpha_k)$.
3. The stochastic gradient errors are bounded in expectation.
4. For each θ , the ordinary differentiable equation corresponding to Equation 13 has a local asymptotically stable attractor within a domain of attraction such that the attractor is Lipschitz. Similar assumptions are required for ϕ and ψ .
5. $\sup_k \|\theta_k\| < \infty$, $\sup_k \|\psi_k\| < \infty$, $\sup_k \|\phi_k\| < \infty$.

Proof. See Theorem 1 in (Heusel et al., 2017). □

Table 3: Dimensions of the visible variables and sizes of the train, validation and test splits.

Datasets	# visible	Train	Valid.	Test
Digits	64	1438	359	-
Adult	123	5000	1414	26147
Connect4	126	16000	4000	47557
DNA	180	1400	600	1186
Mushrooms	112	2000	500	5624
NIPS-0-12	500	400	100	1240
OCR-letters	128	32152	10000	10000
RCV1	150	40000	10000	150000
Frey faces	560	1965	-	-

Table 4: The model structures in RBM experiments.

	Digits	Adult	Connect4	DNA	Mushrooms	NIPS-0-12	Ocr-letters	RCV1
dimension of z	15	15	15	15	15	50	15	15
dimension of h	50	50	50	50	50	200	50	50
dimension of v	64	123	126	180	112	500	128	150

Remark Compared to Theorem 1, Proposition 2 ensures the convergence of AdvIL without assuming $q(v, h)$ is sufficiently close to $P(v, h)$ in each step. However, a two time-scale update rule (Heusel et al., 2017) is required to satisfy Assumption 2 and extra weight decay terms are needed to satisfy Assumption 4. Further, the convergence point is not necessarily a stationary point of $\mathcal{L}_1(\theta, \phi)$.

D DATASETS AND EXPERIMENTAL SETTINGS

We evaluate our method on the binary digits dataset⁴, the UCI binary databases and the Frey faces datasets⁵. The information of the datasets is summarized in Tab. 3. We implement our model using the TensorFlow (Abadi et al., 2016) library. In all experiments, q and r are updated 100 times per update of P and Q , i.e. $K_1 = 100$ and $K_2 = 1$. We use the ADAM (Kingma & Ba, 2014) optimizer with the learning rate $\alpha = 0.0003$, the moving average ratios $\beta_1 = 0.5$ and $\beta_2 = 0.999$, and the batch size of 500. We use a continuous z and the sigmoid activation function. All these hyperparameters are set according to the validation performance of an RBM on the binary digits dataset and fixed throughout the paper unless otherwise stated. The sizes of the variational distributions depend on the structure of the MRF and are chosen according to the validation performance. The model structures in RBM and DBM experiments are summarized in Tab. 4 and Tab. 5, respectively.

The authors of NVIL (Kuleshov & Ermon, 2017) propose two variants. The first one employs a mixture of Bernoulli as q . The second one involves auxiliary variables and employs a neural network as q . Both variants scale up to an RBM of at most 64 visible units as reported in their paper (Kuleshov & Ermon, 2017). For a fair comparison, we carefully perform grid search over the default settings of NVIL and our settings based on their code and choose the best configuration. In this setting, the first variant of NVIL still fails to scale up to larger datasets and the best version of the second variant shares the same key hyperparameters as AdvIL, including $K_1 = 100$ and a batch size of 500.

Table 5: The model structures in DBM experiments.

	Digits	Adult	Connect4	DNA	Mushrooms	NIPS-0-12	Ocr-letters	RCV1
dimension of z	15	15	15	15	15	50	15	15
dimension of h_1	50	50	50	50	50	200	50	50
dimension of h_2	50	50	50	50	50	200	50	50
dimension of v	64	123	126	180	112	500	128	150

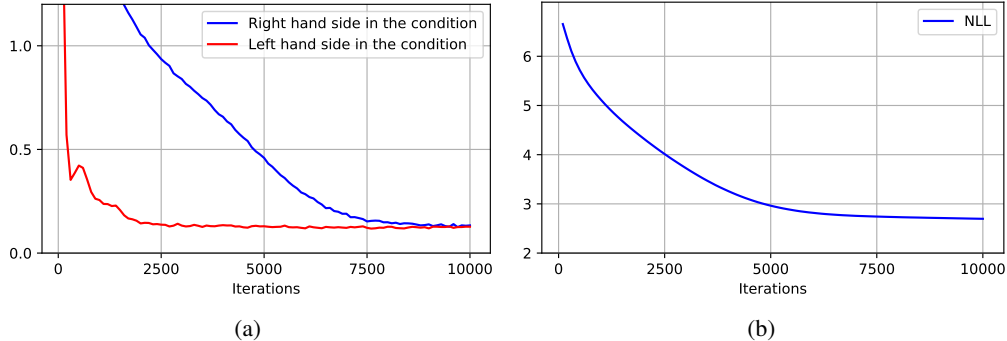


Figure 5: (a) shows that $\|\sum_{v,h} \delta(v,h) \frac{\partial \mathcal{E}(v,h)}{\partial \theta_k}\|_2$ (in red) is less than $\|\frac{\partial \mathcal{L}_1(\theta_k, \phi_k)}{\partial \theta_k}\|_2$ (in blue) during training. (b) shows that the NLL converges gradually.

E MORE RESULTS

E.1 EMPIRICAL VERIFICATION OF ASSUMPTION 5 IN THEOREM 1

We empirically test the assumption $\|\sum_{v,h} \delta(v,h) \frac{\partial \mathcal{E}(v,h)}{\partial \theta_k}\|_2 < \|\frac{\partial \mathcal{L}_1(\theta_k, \phi_k)}{\partial \theta_k}\|_2$ for $k < 10000$, where $\delta(v,h) = q(v,h) - P(v,h)$. Note that computing $\|\frac{\partial \mathcal{L}_1(\theta_k, \phi_k)}{\partial \theta_k}\|_2$ exactly requires summing over v and h and therefore we train a small RBM, where the dimensions of v , h and z are 4, on the Digits dataset. We find $K_1 = 10$ is sufficient in this case.

The results are shown in Fig. 5. It can be seen that a decoder with neural networks and auxiliary variables are sufficiently powerful to track the model distribution and Assumption 5 in Theorem 1 holds. Besides, the model converges gradually, which agrees with our convergence analysis.

E.2 SAMPLES IN RBM

We present the samples from the RBM P and the decoder q in Fig. 6. In this case, we set the number of the hidden units to 50 and other settings remain the same as in Sec. 5.1. The first column demonstrates that the decoder is a good approximate sampler for the RBM. Note that the samples from the decoder are obtained from efficient ancestral sampling but those from the RBM is obtained by Gibbs sampling after 100,000 burn-in steps. The second column shows that if $\mathcal{H}(Q)$ is removed, both models collapse to a certain mode of the data. The third column shows that if $\mathcal{H}(q)$ is removed, both models fail to generate meaningful digits. These results demonstrate the importance of the entropy terms and the necessity of approximating $\mathcal{H}(q)$ in a principled way.

⁴ https://scikit-learn.org/stable/modules/generated/sklearn.datasets.load_digits.html#sklearn.datasets.load_digits

⁵ <http://www.cs.nyu.edu/~roweis/data.html>




Figure 6: (a-c) Samples from the RBM in different settings. (d-f) Samples from the decoder in different settings. We present the mean of v for better visualization in all settings.

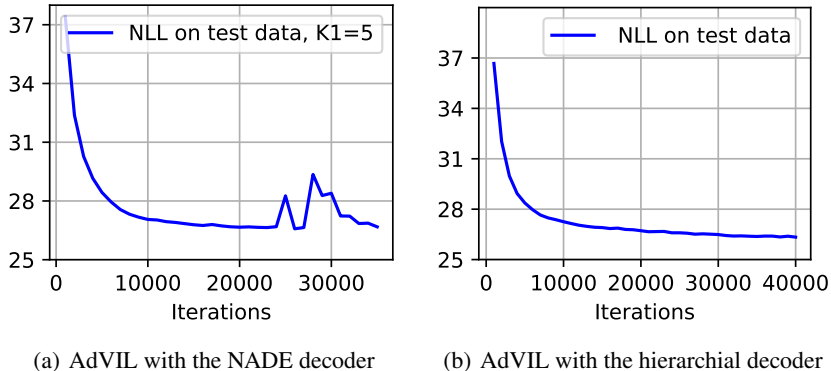


Figure 7: AdvIL with two types of decoders. We set $K_1 = 5$ in the NADE decoder and $K_1 = 100$ in the hierarchical decoder. The two models have a similar model capacity and training time.

E.3 ADVIL WITH AN AUTOREGRESSIVE PRIOR

Here we present the results of AdvIL with a neural autoregressive distribution estimator (NADE) (Larochelle & Murray, 2011) as the prior on the Digits dataset. We use the same RBM as in Sec. 5.1. The dimension of the latent units in NADE is 15, which is the same as the dimension of the auxiliary variables in the hierarchical decoder presented in Sec. 3.3.

Compared with the hierarchical decoder, the NADE decoder has a tractable entropy and hence does not require $r(z|h)$. However, getting samples from NADE is slow while AdvIL requires samples during training. Therefore, $K_1 = 5$ for the NADE decoder has a similar training cost as the hierarchical decoder.

Fig. 7 compares the two decoders. AdvIL with the NADE decoder achieves a slightly worse and unstable result.

E.4 AIS RESULTS WITH STANDARD DEVIATIONS

The AIS results in RBM and RBM with the means and standard deviations are shown in Tab. 6 and Tab. 7 respectively.

Table 6: The AIS results of NVIL and AdvIL in RBM with the means and standard deviations. The results are averaged over three runs with different random seeds.

Method	Digits	Adult	Connect4	DNA	Mushrooms	NIPS-0-12	Ocr-letters	RCV1
NVIL-mean	-27.36	-20.05	-24.71	-97.71	-29.28	-290.01	-47.56	-50.47
NVIL-std	0.13	0.27	0.61	0.12	0.31	2.68	0.14	0.09
AdvIL-mean	-26.34	-19.29	-21.95	-97.59	-19.59	-276.42	-45.64	-50.22
AdvIL-std	0.02	0.07	1.04	0.10	2.01	0.21	0.34	0.06

Table 7: The AIS results of VCD-1 and AdvIL in DBM with the means and standard deviations. The results are averaged over three runs with different random seeds.

Method	Digits	Adult	Connect4	DNA	Mushrooms	NIPS-0-12	Ocr-letters	RCV1
VCD-mean	-28.49	-22.26	-26.79	-97.59	-23.15	-356.26	-45.77	-50.83
VCD-std	0.47	0.51	1.42	0.03	1.42	34.70	1.15	0.62
AdvIL-mean	-27.89	-20.29	-26.34	-99.40	-21.21	-287.15	-48.38	-51.02
AdvIL-std	0.44	0.24	1.50	0.71	0.40	0.63	1.32	0.42