

Bi-level Doubly Variational Learning for Energy-based Latent Variable Models

Ge Kan¹, Jinhu Lü¹, Tian Wang^{1*}, Baochang Zhang¹, Aichun Zhu²
 Lei Huang¹, Guodong Guo³, Hichem Snoussi⁴

¹School of Automation Science and Electrical Engineering, Institute of Artificial Intelligence,
 Beihang University, Beijing China

²School of Computer Science and Technology, Nanjing Tech University, Nanjing China

³Institute of Deep Learning, Baidu Research,
 National Engineering Laboratory for Deep Learning Technology and Application, Beijing China

⁴University of Technology of Troyes, Troyes France

{gekan, wangtian, bczhang}@buaa.edu.cn, jhlu@iss.ac.cn, aichun.zhu@njtech.edu.cn
 huanglei36060520@gmail.com, guogudong01@baidu.com, hichem.snoussi@utt.fr

Abstract

Energy-based latent variable models (EBLVMs) are more expressive than conventional energy-based models. However, its potential on visual tasks are limited by its training process based on maximum likelihood estimate that requires sampling from two intractable distributions. In this paper, we propose Bi-level doubly variational learning (BiDVL), which is based on a new bi-level optimization framework and two tractable variational distributions to facilitate learning EBLVMs. Particularly, we lead a decoupled EBLVM consisting of a marginal energy-based distribution and a structural posterior to handle the difficulties when learning deep EBLVMs on images. By choosing a symmetric KL divergence in the lower level of our framework, a compact BiDVL for visual tasks can be obtained. Our model achieves impressive image generation performance over related works. It also demonstrates the significant capacity of testing image reconstruction and out-of-distribution detection.

1. Introduction

Energy-based models (EBMs) [24] are known as powerful generative models widely studied in the field of machine learning, which define a general distribution explicitly by normalizing exponential negative energy. EBMs have also been successfully applied to solve visual tasks, such as image synthesis [6], classification [12], out-of-distribution (OOD) detection [30] or semi-supervised learning [7]. Further, latent variables are incorporated to define

energy-based latent variable models (EBLVMs), which are more expressive and capable of representation learning [4].

However, the effectiveness of EBLVM deteriorates when applying it to solve computer vision problems, because learning by maximum likelihood estimation (MLE) usually suffers from the *doubly intractable* model problem [42], *i.e.*, sampling from the posterior and the joint distribution of model is nontrivial due to the intractable integrals in normalizing denominator. Recent works [6, 35] leverage gradient-based Markov Chain Monte Carlo (MCMC), such as Langevin dynamics, to approximately sample from EBMs but require lots of steps, since fewer steps may lead to arbitrarily far sampling distribution from the target. Moreover, to handle the *doubly intractable* problem of EBLVMs, double MCMC sampling are needed and as a result making it infeasible on high-dimensional images.

To accomplish EBLVMs efficiently, this paper introduces a **Bi-level Doubly Variational Learning (BiDVL)** framework, which is based on tractable variational posterior and variational joint distribution to estimate the gradients in MLE. We formulate BiDVL as a bi-level optimization (BLO) problem as illustrated in Fig. 1. In specific, the gradient estimate is compelled to fit the real one by exploring variational distributions to achieve the model distributions in the lower level, and the resulting gradient estimate is then used for optimizing the objective in the upper level. Theoretically, BiDVL is equivalent to the original MLE objective under the *nonparametric assumption* [10], while to optimize it practically, we propose an efficient alternative optimization scheme.

Moreover, the drawbacks of doubly variational learning come from its unstableness when learning a deep EBLVM on images. For example, AdvIL [25] that can be approx-

*Corresponding Author

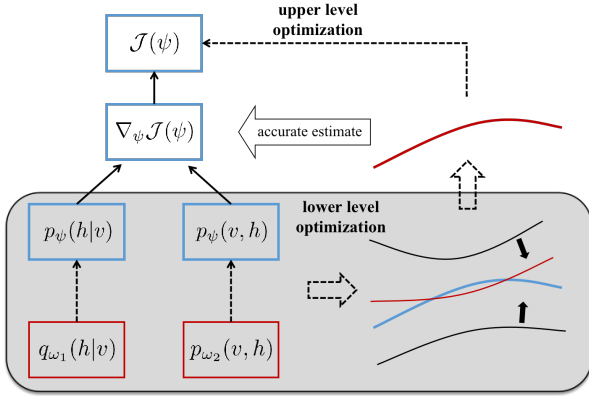


Figure 1. The motivation for BiDVL. Blue items represent the real objective optimization procedure. Red items indicate the approximate one. Real gradient, gradient estimate, and bounds are represented by blue, red and black curves, respectively. The bold red curve is the gradient estimate for upper level optimization. See Sec. 4.1.1 for detailed explanation.

imately degenerated from BiDVL fails to achieve a good performance. Reasons for the failure lie in: 1) A general deep EBLVM has no structural assumption about its posterior and the unconstrained posterior may transfer among diverse mode structures during training, bring about dynamic target for variational learning. 2) Direct interaction between two variational distributions is overlooked, though they all approximate the EBLVM. 3) Both the energy-based posterior and its variational approximation have no effective way of modelling the real one. Such problems may effect other methods [3, 4, 16] for learning deep EBLVMs.

To tackle these issues, we decompose the EBLVM into an energy-based marginal distribution and a structural posterior, where the latter is represented by variational posterior. Then a symmetric KL divergence is chosen in the lower level, doubly linking the variational distributions by sharing posterior in both levels, a compact BiDVL for image tasks is thus obtained. At last, the decoupled EBLVM helps to derive an importance weighted KL divergence, in which samples for Monte Carlo estimator come from dataset. By learning from data through the doubly link, the posterior models the latent space more effectively.

Specifically, we parametrize the variational distributions with an inference model and an implicit generative model (IGM). Recently, [15, 16, 46] propose to jointly train EBM, inference model and IGM. However, their formulations are derived from intuitive combination to some extent, leading to overlap and conflict among models. BiDVL, entirely for learning EBLVMs, interprets the consistency by great improvement. In summary, we make following contributions:

- 1) We introduce a MCMC-free BiDVL framework for learning EBLVMs efficiently. A corresponding alternative optimization scheme is adopted to solve it.
- 2) We notice several problems in learning deep EBLVMs

on images. To overcome these, we define a decoupled EBLVM and choose a symmetric KL divergence in the lower level. The resulting compact objective implies the consistency for learning EBLVMs.

3) We demonstrate the impressive image generation quality among baseline models. Besides, BiDVL also shows the remarkable performance of testing image reconstruction and OOD detection.

2. Related work

Learning EBMs and EBLVMs. As unnormalized probability model, EBMs and EBLVMs are able to fit data distribution more efficiently than other generative models. Nonetheless, they require nontrivial sampling from the model when learned by MLE. Many attempts [2, 6] leverage MCMC to approximate the procedure but suffers from quite poor efficiency. [3, 4, 27, 42] avoid the nontrivial sampling based on Score Matching [18], however need unstable and inefficient computing of score functions. Another approach is Noise Contrastive Estimation [14], which performs worse on high-dimensional images [11]. Furthermore, [13] proposes to learn Stein Discrepancy without sampling, [47] generalizes MLE to a f -divergence variational framework, [8, 11] propose MCMC-free variational approaches, while all of them aim at learning EBMs, our BiDVL accomplishes EBLVMs in a different way.

Recent attempt [25] introduces two variational models to form a min-min-max nested objective. Their method is evaluated on RBM [1] and DBM [40], but is unstable for learning deep EBLVMs on images. We however propose a compact model for stability, see Sec. 4.2 for details.

Joint training with other models. Recent works also notice the compatibility between EBM and other models. [2, 35, 36, 43, 44] utilize an EBM to exponentially tilt a generative distribution, most of them have a two-stage refinement procedure with inefficient MCMC sampling. [5, 41] regard EBM as a discriminator of GAN, but still use MCMC to modify the generative distribution.

[15] proposes a *divergence triangle* loss for joint training EBM, inference model and IGM, [16] extends it to EBLVM and [46] extends it by variational MCMC teaching [45]. These studies integrate KL losses intuitively, by actually assuming the EBM always fits data perfectly, which induces harmful conflict among the models involved. While our compact objective derives an importance weighted KL loss and the variational models in BiDVL are optimized simultaneously rather than their sleep-awake scheme. All the differences originate from our unified framework, showing the consistency for learning EBLVMs.

3. Preliminaries

EBLVM defines a joint scalar energy function over visible variable v and latent variable h . The corresponding joint

distribution is

$$p_\psi(v, h) = \frac{\exp(-\mathcal{E}_\psi(v, h))}{\int \exp(-\mathcal{E}_\psi(v, h)) dh dv}, \quad (1)$$

where $\mathcal{E}_\psi(v, h)$ denotes the energy function over joint space parameterized by ψ , $\mathcal{Z}(\psi) = \int \exp(-\mathcal{E}_\psi(v, h)) dh dv$ is known as the *partition function*. The marginal distribution is given by $p_\psi(v) = \int p_\psi(v, h) dh$, while the marginal energy is $\mathcal{E}_\psi(v) = -\log \int \exp(-\mathcal{E}_\psi(v, h)) dh$.

Training EBLVM is typically based on MLE, or equivalently minimizing the KL divergence $\mathcal{J}(\psi) = D_{\text{KL}}(q(v)||p_\psi(v))$, where $q(v)$ denotes the empirical data distribution. As [24] notes, the gradient of the negative log-likelihood objective w.r.t. ψ can be computed by

$$\begin{aligned} \nabla_\psi \mathbb{E}_{q(v)}[-\log p_\psi(v)] &= \nabla_\psi D_{\text{KL}}(q(v)||p_\psi(v)) = \\ \mathbb{E}_{q(v)p_\psi(h|v)} [\nabla_\psi \mathcal{E}_\psi(v, h)] &- \mathbb{E}_{p_\psi(v, h)} [\nabla_\psi \mathcal{E}_\psi(v, h)], \end{aligned} \quad (2)$$

where $p_\psi(h|v) = \frac{\exp(-\mathcal{E}_\psi(v, h))}{\int \exp(-\mathcal{E}_\psi(v, h)) dh}$ is the posterior distribution over latent variables. In Eq. (2), samples from $p_\psi(h|v)$ and $p_\psi(v, h)$ are required for Monte Carlo gradient estimator, which is *doubly intractable* due to the integrals over the high-dimensional space. Commonly adopted MCMC however suffers from high consumption and is too bulky to be applied to both distributions simultaneously.

4. Learning method

4.1. Proposed Bi-level doubly variational learning

To handle the *doubly intractable* problem efficiently, we propose doubly variational learning, *i.e.*, introduce two variational models to represent the variational posterior $q_{\omega_1}(h|v)$ and the variational joint distribution $p_{\omega_2}(v, h)$, where $\omega \doteq (\omega_1, \omega_2)$ is the trainable parameters. Sampling from variational distribution requires only one forward step, largely improving efficiency.

4.1.1 Framework

In this part, we introduce the doubly variational learning under the bi-level optimization (BLO) framework. Let $\mathcal{E}_\psi \doteq \mathcal{E}_\psi(v, h)$ for clarity. We first rewrite Eq. (2) as follow to show our motivation:

$$\nabla_\psi \mathcal{J}(\psi) = \mathbb{E}_{q(v)p_\psi(h|v)} [\nabla_\psi \mathcal{E}_\psi] - \mathbb{E}_{p_\psi(v, h)} [\nabla_\psi \mathcal{E}_\psi] \quad (3)$$

$$\equiv \mathbb{E}_{q(v)q_{\omega_1}(h|v)} [\nabla_\psi \mathcal{E}_\psi] - \mathbb{E}_{p_{\omega_2}(v, h)} [\nabla_\psi \mathcal{E}_\psi] \quad (3a)$$

$$+ \mathbb{E}_{q(v)p_\psi(h|v)} [\nabla_\psi \mathcal{E}_\psi] - \mathbb{E}_{q(v)q_{\omega_1}(h|v)} [\nabla_\psi \mathcal{E}_\psi] \quad (3b)$$

$$+ \mathbb{E}_{p_{\omega_2}(v, h)} [\nabla_\psi \mathcal{E}_\psi] - \mathbb{E}_{p_\psi(v, h)} [\nabla_\psi \mathcal{E}_\psi], \quad (3c)$$

where \equiv denotes the identity w.r.t. ω , in other word, it holds whatever ω_1, ω_2 are. Equation (3) implies that, for an appropriate ω such that both terms (3b) and (3c) equal zero, then term (3a) is an accurate gradient estimation.

The motivation to formulate a bi-level optimization problem is illustrated in Fig. 1. Variational distributions

$q_{\omega_1}(h|v)$ and $p_{\omega_2}(v, h)$ (red marked items) aim to chase the model distributions $p_\psi(h|v)$ and $p_\psi(v, h)$ (blue marked items) in lower-level optimization, reducing the gap between the upper and lower bound of gradient estimate term (3a) (red curve). The gradient estimate is thus compelled to fit the real one, *i.e.* Eq.(3) (bold blue curve). Then the estimate finally obtained (bold red curve) is used for optimizing the real objective in upper level.

Following the formulas described in [28], we assume the solution of lower-level (LL) subproblem w.r.t. the upper-level (UL) variables is of a singleton, which is known as the *lower-level singleton* assumption [29]. Therefore we define the LL subproblem as follow:

$$\omega^*(\psi) = \arg \min_{\omega \in \Omega} \mathcal{J}_{\text{LL}}(\psi, \omega),$$

$$\begin{aligned} \mathcal{J}_{\text{LL}}(\psi, \omega) &= D(q(v)q_{\omega_1}(h|v), q(v)p_\psi(h|v)) \\ &+ D(p_{\omega_2}(v, h), p_\psi(v, h)), \end{aligned} \quad (4)$$

where Ω is the parameter space of variational distributions. D denotes certain proper metric corresponding to the assumption on $\nabla_\psi \mathcal{E}_\psi$, since the form of terms (3b),(3c) reminds of the general integral probability metrics. See Appendix A.1 for related discussion. We should again emphasize that the solution of LL subproblem is a function w.r.t. the UL variables, and as a result, $\nabla_\psi \mathcal{J}(\psi) = \mathbb{E}_{q(v)q_{\omega_1^*(\psi)}(h|v)} [\nabla_\psi \mathcal{E}_\psi] - \mathbb{E}_{p_{\omega_2^*(\psi)}(v, h)} [\nabla_\psi \mathcal{E}_\psi]$ holds only if $\mathcal{J}_{\text{LL}}(\psi, \omega^*(\psi)) = 0$. Then an equivalent of term (3a) taking $\omega = \omega(\psi)$ into account is given by:

$$\begin{aligned} &\mathbb{E}_{q(v)q_{\omega_1(\psi)}(h|v)} [\nabla_\psi \mathcal{E}_\psi] - \mathbb{E}_{p_{\omega_2(\psi)}(v, h)} [\nabla_\psi \mathcal{E}_\psi] \\ &= \frac{\partial D_{\text{KL}}(q(v)q_{\omega_1}(h|v)||p_\psi(v, h))}{\partial \psi} \Big|_{\omega_1=\omega_1(\psi)} \\ &\quad - \frac{\partial D_{\text{KL}}(p_{\omega_2}(v, h)||p_\psi(v, h))}{\partial \psi} \Big|_{\omega_2=\omega_2(\psi)}. \end{aligned} \quad (5)$$

Inspired by Eq. (5), we define the UL subproblem as follow:

$$\psi^* = \arg \min_{\psi \in \Psi} \mathcal{J}_{\text{UL}}(\psi, \omega^*(\psi)),$$

$$\begin{aligned} \mathcal{J}_{\text{UL}}(\psi, \omega) &= D_{\text{KL}}(q(v)q_{\omega_1}(h|v)||p_\psi(v, h)) \\ &- D_{\text{KL}}(p_{\omega_2}(v, h)||p_\psi(v, h)), \end{aligned} \quad (6)$$

where Ψ is the parameter space of energy function. Further, the gradient of UL objective w.r.t. ψ is denoted by

$$\begin{aligned} \frac{\partial \mathcal{J}_{\text{UL}}(\psi, \omega^*(\psi))}{\partial \psi} &= \frac{\partial \mathcal{J}_{\text{UL}}(\psi, \omega)}{\partial \psi} \Big|_{\omega=\omega^*(\psi)} \\ &+ \left(\frac{\partial \omega^*(\psi)}{\partial \psi^\top} \right)^\top \frac{\partial \mathcal{J}_{\text{UL}}(\psi, \omega)}{\partial \omega} \Big|_{\omega=\omega^*(\psi)}, \end{aligned} \quad (7)$$

where $\frac{\partial \omega^*(\psi)}{\partial \psi}$ is known as the Best-Response (BR) Jacobian [28] which captures the response of LL solution w.r.t. the UL variables.

We argue that, under the *nonparametric assumption* [10], the BLO problem (4,6) in BiDVL is equivalent to optimize the original objective $\mathcal{J}(\psi)$. The formal theorem is pre-

sented as follow:

Theorem 1. *Assuming that $\forall \psi \in \Psi, \exists \omega \in \Omega$ such that $D(q(v)q_{\omega_1}(h|v), q(v)p_{\psi}(h|v)) = 0$, and $D(p_{\omega_2}(v, h), p_{\psi}(v, h)) = 0$, then we have*

$$\mathcal{J}(\psi) = \mathcal{J}_{\text{UL}}(\psi, \omega^*(\psi)), \nabla_{\psi} \mathcal{J}(\psi) = \nabla_{\psi} \mathcal{J}_{\text{UL}}(\psi, \omega^*(\psi)).$$

The *nonparametric assumption* does not always hold in practice, since the variational models have limited capacity. But we can still bound the original objective as $|\mathcal{J}_{\text{UL}}(\psi, \omega^*(\psi)) - \mathcal{J}(\psi)| \leq C' \cdot \mathcal{J}_{\text{LL}}(\psi, \omega)$. See Appendices A.2 and A.3 for detailed derivations and proof.

4.1.2 Alternative optimization

Solving the BLO problem requires to compute the gradient of UL objective (7). As [28] notes, obtaining the gradient suffers from handling the BR Jacobian with a recursive derivation process. Recent attempts [3, 32] leverage *gradient unrolling* technique to estimate the gradient in bi-level optimization problems. Nevertheless, *gradient unrolling* requires inner loop to form a recursive computation graph and more resources are allocated for backtracking. It also incurs inferior performance in our case. For efficiency, we ignore the BR Jacobian term in Eq. (7). Then BiDVL is optimized alternatively, *i.e.*, when updating the parameters involved in current level, we consider the rest as fixed. Notice the missing-term alternative optimization has been widely adopted in reinforcement learning.

4.2. BiDVL for Image Tasks

BiDVL is a general framework for EBLVMs, however, doubly variational learning is usually unstable to scale to learning deep EBLVMs on image datasets. AdvIL [25], which can be approximately derived from BiDVL via choosing KL divergence in the lower level, encounters the same difficulty that unstable training and inferior performance to deal with images. We attribute the failure to three key learning problems which exacerbate the instability on high-dimensional and highly multimodal image datasets:

- 1) The generally defined undirected EBLVM (1) has no explicit structural assumption about its posterior $p_{\psi}(h|v)$ (due to the coupled joint energy represented by network), and resulting in a dynamic target, whose mode structure transfers rapidly, for $q_{\omega_1}(h|v)$ and $p_{\omega_2}(v, h)$ to chase.
- 2) Though both variational distributions are to approximate $p_{\psi}(v, h)$, there is no direct interaction between them, thus misalignment may be induced in variational learning.
- 3) Both $p_{\psi}(h|v)$ and $q_{\omega_1}(h|v)$ cannot model the real posterior effectively, since they learn from each other (4) rather than from data directly.

Next, we introduce our solutions to the problems above with BiDVL as a necessary framework, and then, a practical compact BiDVL for image tasks is consequently conducted.

4.2.1 Decoupled EBLVM

The problems above imply that an unconstrained posterior is confused and superfluous for modeling the latent space, which motivates us to *redefined* $p_{\psi}(v)$ and $p_{\psi}(h|v)$ as an energy-based marginal distribution and a structural posterior. Considering the same parameterization of $q_{\omega_1}(h|v)$ and $p_{\psi}(h|v)$, and by $\min D(q(v)q_{\omega_1}(h|v)||q(v)p_{\psi}(h|v))$ in Eq. (4), we can directly make the energy-based posterior equal to the variational one by sharing parameters ω_1 , and as a consequence, we have $D(q(v)q_{\omega_1}(h|v), q(v)p_{\psi}(h|v)) = 0$, eliminating the confusing chase between posteriors in BiDVL as well as reducing the LL objective. We finally formulate a decoupled version of EBLVM:

$$\begin{aligned} p_{\psi', \omega_1}(v, h) &= p_{\psi'}(v)q_{\omega_1}(h|v), \\ p_{\psi'}(v) &\propto \exp(-\mathcal{E}_{\psi'}(v)), \end{aligned} \quad (8)$$

where $\mathcal{E}_{\psi'}(v)$ is the new marginal energy parametrized with ψ' and then a decoupled joint energy is *redefined* by $\mathcal{E}_{\psi', \omega_1}(v, h) = \mathcal{E}_{\psi'}(v) - \log q_{\omega_1}(h|v)$. Note we are actually optimizing the upper bound of $\mathcal{J}(\psi)$, in other words, if the bi-level objectives are reduced, the original objective is reduced in turn, therefore we can optimize ω_1 in both levels without changing the optimal solution of BiDVL.

4.2.2 Symmetric KL divergence

Since ω_1 is optimized in both levels, a proper metric in the lower level can link variational distributions, where KL divergence is a feasible choice. However as [9, 16] notice, minimizing $D_{\text{KL}}(P||Q)$ w.r.t. Q compels Q to cover the major support of P , while minimizing a reverse version $D_{\text{KL}}(Q||P)$, Q is forced to chase the major modes of P . So for integrating both behavior, we choose a symmetric KL divergence, $S(P||Q) = D_{\text{KL}}(P||Q) + D_{\text{KL}}(Q||P)$, in the lower level. Meanwhile, by taking the alternative optimization from Sec. 4.1.2 into account, a compact BiDVL is derived from Eqs. (4) and (6):

$$\begin{aligned} \min_{\omega} \mathcal{J}_{\text{LL}}(\omega), \\ \mathcal{J}_{\text{LL}}(\omega) &= D_{\text{KL}}(p_{\psi', \omega_1}(v, h)||p_{\omega_2}(v, h)) \\ &+ D_{\text{KL}}(p_{\omega_2}(v, h)||p_{\psi', \omega_1}(v, h)), \end{aligned} \quad (9)$$

$$\begin{aligned} \min_{\psi', \omega_1} \mathcal{J}_{\text{UL}}(\psi', \omega_1) \\ \mathcal{J}_{\text{UL}}(\psi', \omega_1) &= D_{\text{KL}}(q(v)q_{\omega_1}(h|v)||p_{\psi', \omega_1}(v, h)) \\ &- D_{\text{KL}}(p_{\omega_2}(v, h)||p_{\psi', \omega_1}(v, h)), \end{aligned} \quad (10)$$

where the symmetric KL divergence in the lower level sheds light on the connection between variational distributions in two directions. Based on above solutions, we derive a weighted KL loss shown in next section, which is proven to help modelling the data latent space effectively.

4.2.3 Optimizing compact BiDVL

In this part, we present the gradients for optimizing compact BiDVL (9,10) followed by a simplified version. Gradients are computed by:

$$\begin{aligned}
& \nabla_{\omega} \mathcal{J}_{LL}(\omega) \\
&= \nabla_{\omega} D_{KL}(p_{\psi'}(v)q_{\omega_1}(h|v)||p_{\omega_2}(v, h)) \\
&+ \nabla_{\omega_2} \mathbb{E}_{p_{\omega_2}(v, h)}[\mathcal{E}_{\psi'}(v)] + \nabla_{\omega} D_{KL}(p_{\omega_2}(v, h)||q_{\omega_1}(h|v)) \\
&= \nabla_{\omega} \int \frac{p_{\psi'}(v)}{q(v)} q(v)q_{\omega_1}(h|v) \log \frac{q(v)q_{\omega_1}(h|v)}{p_{\omega_2}(v, h)} dh dv \\
&+ \nabla_{\omega_2} \mathbb{E}_{p_{\omega_2}(v, h)}[\mathcal{E}_{\psi'}(v)] + \nabla_{\omega} D_{KL}(p_{\omega_2}(v, h)||q_{\omega_1}(h|v)) \\
&= \nabla_{\omega} D_{r(v)}(q(v)q_{\omega_1}(h|v)||p_{\omega_2}(v, h)) \\
&+ \nabla_{\omega_2} \mathbb{E}_{p_{\omega_2}(v, h)}[\mathcal{E}_{\psi'}(v)] + \nabla_{\omega} D_{KL}(p_{\omega_2}(v, h)||q_{\omega_1}(h|v)), \tag{11}
\end{aligned}$$

$$\begin{aligned}
& \nabla_{\psi', \omega_1} \mathcal{J}_{UL}(\psi', \omega_1) \\
&= \mathbb{E}_{q(v)}[\nabla_{\psi'} \mathcal{E}_{\psi'}(v)] - \mathbb{E}_{p_{\omega_2}(v, h)}[\nabla_{\psi'} \mathcal{E}_{\psi'}(v)] \tag{12} \\
&- \nabla_{\omega_1} D_{KL}(p_{\omega_2}(v, h)||q_{\omega_1}(h|v)),
\end{aligned}$$

where $D_{r(v)}$ denotes the importance weighted KL divergence with $r(v) = \frac{p_{\psi'}(v)}{q(v)}$ as the importance ratio. The principal reason for utilizing importance ratio is the non-trivial sampling from $p_{\psi'}(v)$. Notice computing the ratio is still nontrivial, we present a detailed analysis in Sec. 5.4.

Importance weighted term offers additional benefits. Since $D_{KL}(p_{\psi'}(v)q_{\omega_1}(h|v)||p_{\omega_2}(v, h))$ has a mode-cover behavior, $p_{\psi'}(v)$ defined on the whole space may provide fake mode-cover information for $p_{\omega_2}(v, h)$. However, samples for estimating the weighted term come from dataset, brings $p_{\omega_2}(v, h)$ faithful mode-cover information and further helps the EBLVM faster convergence and stabler training. Furthermore, minimizing $D_{KL}(q(v)q_{\omega_1}(h|v)||p_{\omega_2}(v, h))$ is equivalent to optimize the *evidence lower bound* of VAE, whose reconstructive behavior is proven to model the data latent space effectively.

Minimizing $D_{KL}(p_{\omega_2}(v, h)||q_{\omega_1}(h|v))$ tends to increase the likelihood of generated samples under the posterior, which further enhance the alignment between visible space and latent space. By sharing the variational posterior $q_{\omega_1}(h|v)$, the decoupled EBLVM thus learns a simpler structural posterior easy to model data latent space.

Notice the opposite terms w.r.t. ω_1 contained in Eqs. (11) and (12) can be offset to derive simplified gradients:

$$\begin{aligned}
& \nabla_{\omega} \widehat{\mathcal{J}}_{LL}(\omega) = \nabla_{\omega} D_{r(v)}(q(v)q_{\omega_1}(h|v)||p_{\omega_2}(v, h)) \\
&+ \nabla_{\omega_2} \mathbb{E}_{p_{\omega_2}(v, h)}[\mathcal{E}_{\psi'}(v)] + \nabla_{\omega_2} D_{KL}(p_{\omega_2}(v, h)||q_{\omega_1}(h|v)), \tag{13}
\end{aligned}$$

$$\begin{aligned}
& \nabla_{\psi'} \widehat{\mathcal{J}}_{UL}(\psi') = \mathbb{E}_{q(v)}[\nabla_{\psi'} \mathcal{E}_{\psi'}(v)] - \mathbb{E}_{p_{\omega_2}(v, h)}[\nabla_{\psi'} \mathcal{E}_{\psi'}(v)]. \tag{14}
\end{aligned}$$

The offset gradients are stabler than the original one, as the opposite terms are integrated into one level, however weakens the adversarial learning as we explained in Sec. 5.4.

Algorithm 1 Bi-level doubly variational learning for EBLVM by alternative stochastic gradient descent

Input: Learning rate schemes α and β ; randomly initialized network parameters ω and ψ' ; number of lower-level step N

- 1: **repeat**
 - 2: Sample a batch of data
 - 3: **for** $n = 1, \dots, N$ **do**
 - 4: Optimize the lower-level subproblem (9) by Eq. (11) (or Eq. (13)):
 $\omega \leftarrow \omega - \alpha \nabla_{\omega} \mathcal{J}_{LL}(\omega)$
 - 5: **end for**
 - 6: Optional *gradient unrolling* following [32]
 - 7: Optimize the upper-level subproblem (10) by Eq. (12) (or Eq. (14)):
 $\psi' \leftarrow \psi' - \beta \nabla_{\psi'} \mathcal{J}_{UL}(\psi', \omega_1)$
 $\omega_1 \leftarrow \omega_1 - \alpha \nabla_{\omega_1} \mathcal{J}_{UL}(\psi', \omega_1)$
 - 8: **until** Convergence or reaching certain threshold
-

4.3. Algorithm implementation

The training procedure is summarized in Algorithm 1, and in this section, a specific implementation is introduced. Our model consists of an EBM, an inference model and a generative model, corresponding to the marginal energy-based distribution $p_{\psi'}(v)$, the variational posterior $q_{\omega_1}(h|v)$ and the variational joint distribution $p_{\omega_2}(v, h)$, respectively.

For the generative model, BiDVL has many choices, *e.g.* implicit generative models (IGMs) or flows [39]. In this work, we simply utilize an IGM, from which a sample is obtained via passing a latent variable h through a deterministic generator, *i.e.*, $v = G_{\omega_2}(h)$, where h is sampled from a known distribution $p(h)$. For the inference model, we employ the standard Gaussian encoder of VAE, $\mathcal{N}(\mu_{\omega_1}(v), \Sigma_{\omega_1}(v))$. Sampling from the inference model typically adopts the reparametrization trick, *i.e.* $h = \mu_{\omega_1}(v) + \Sigma_{\omega_1}(v) \cdot \epsilon$, where ϵ denotes the Gaussian noise.

4.4. Discussion

In this section, we present some additional properties of our BiDVL under the specific implementation. By adopting the structure of VAE, Eq. (11) contains an importance weighted VAE loss, where we practically use the importance weighted reconstruction loss and KL regularization in Eq. (11). Meanwhile, $D_{KL}(p_{\omega_2}(v, h)||q_{\omega_1}(h|v))$, which forms another reconstruction over latent space, takes part in a cycle reconstruction, linking the models tightly.

Besides, there are two chase games in Eqs. (9) to (12):

1) The first chase game is demonstrated between the marginal EBM $p_{\psi'}$ and the IGM p_{ω_2} . In lower level, p_{ω_2} tends to chase $p_{\psi'}$, while in upper level, $p_{\psi'}$ turns to data distribution and escapes from p_{ω_2} . The game formulates an adversarial learning similar to GAN, however discriminator guided IGM typically has limited capacity for assigning

density on the learned support [2], while in BiDVL, IGM is guided by a sophisticated energy-based probability model.

2) Opposite terms in Eqs. (11) and (12) contribute to another chase game between variational models. In lower level, q_{ω_1} and p_{ω_2} are drew close over latent space, while in upper level, q_{ω_1} attempts to escape from the other. Optimizing ω_1 to increase $D_{\text{KL}}(p_{\omega_2}(v, h) \| q_{\omega_1}(h|v))$ in the upper level changes the mean and variance of inference model aimlessly, resulting in unstable training. Inspired by the adversarial nature between the reconstruction part and the KL regularization part in VAE loss, we instead minimize $D_{\text{KL}}(q_{\omega_1}(h|v) \| p(h)), v \sim p_{\omega_2}(v)$ in the upper level to interpret the chase game. In fact, it provides q_{ω_1} a fixed target for stable training.

All of the properties such as importance weighted loss, the additional chase game and the bi-level optimization scheme are originated from the unified framework, implying the consistency for learning EBLVMs.

5. Experiments

We evaluate BiDVL on three tasks: image generation, test image reconstruction and out-of-distribution (OOD) detection. Test image reconstruction are conducted to evaluate the ability of modelling latent space, the rest experiments are designed mainly following [16]. In the main paper, experiments are demonstrated on three commonly adopted benchmark datasets including CIFAR-10 [23], SVHN [34] and CelebA [31]. For CIFAR-10 and SVHN, images are resized to 32×32 . For CelebA, we resize images to 32×32 and 64×64 to construct CelebA-32 and CelebA-64. All datasets are scaled to $[-1, 1]$ for pre-processing.

Our model consists of an EBM, an inference model and an IGM. The structure of variational models is following [16], which simply cascades several convolutional layers. For stable training, batch normalization [19] layers are plugged between convolutional layers. The EBM uses convolutional layers to map a sample to real-valued energy with spectral normalization to ensure the constraint on $\nabla_{\psi} \mathcal{E}_{\psi}(v, h)$ discussed in Appendix A.1. Parameters are optimized via Adam [20]. We set the number of lower-level steps N to one and ignore the *gradient unrolling* for faster and stabler training in our experiments. Some major choices are discussed in Sec. 5.4.

In order to evaluate on larger scale images, we also conduct experiments on CelebA resized to 128×128 with stronger Resnet-based structure in Appendix B.

5.1. Image generation

We show the proposed model can generate realistic images with visual similarities to the training images by conducting experiments on CIFAR-10, SVHN and CelebA. Figure 2 shows the result of 32×32 images randomly generated by proposed model and Fig. 3 shows the result of

Model	CIFAR-10	CelebA-64
VAE [22]	109.5	99.09
Divergence triangle [16]	30.1	24.7
IGEBM [6]	40.58	-
GEBM [2]	23.02	-
VERA [11]	27.5	-
SNGAN [33]	21.7	50.4
BiDVL (ours)	20.75	17.24

Table 1. Evaluation of sample quality on CIFAR-10 and CelebA-64 via FID.

Model	CelebA-32
VAE [22]	38.76
DCGAN [38]	12.50
FCE [7]	12.21
GEBM [2]	5.21
BiDVL (ours)	4.47

Table 2. Evaluation of sample quality on CelebA-32 using FID.

64×64 images on CelebA. We find the generated images maintain diversity, which confirms that the learned EBLVM well covers the most modes of data.

We adopt Frechet Inception Distance (FID) [17] to reflect the sample quality and the baseline models are chosen from EBMs, VAEs and GANs to align with the perspectives in Sec. 4.4. Divergence triangle [16] is an energy-guided VAE with a similar model structure to us. IGEBM [6] and GEBM [2] are MCMC-based EBMs while VERA [11] is MCMC-free. SNGAN [33] adopts the structure of Resnet and gets 21.7 FID on CIFAR-10. The recent exponential tilting models which formulate a two-stage refinement procedure with inefficient MCMC sampling, are not considered here. In Tab. 1, we evaluate our best model on CIFAR-10 and CelebA-64. Table 2 reports the FID on CelebA-32. Our model achieves superior performance than baseline models, especially getting 20.75 FID on CIFAR-10, even without elaborate structure such as Resnet.

5.2. Image reconstruction

In this section, we show our model can learn the low-dimensional structure of training samples by assessing the performance on testing image reconstruction. We must note that, EBMs and GANs are incapable of image reconstruction without an inference model, while an inference model is adopted as the variational posterior in our propose model, helping to accomplish the infer-reconstruct procedure. Experimentally, our model achieves a low reconstruction error on both CIFAR-10 and CelebA-64 test datasets, because the reconstruction behavior over latent variables helps to further enhance the alignment between visible space and latent space. We compare our model, which has the best generation in Sec. 5.1, to baseline models in Tab. 3, with rooted

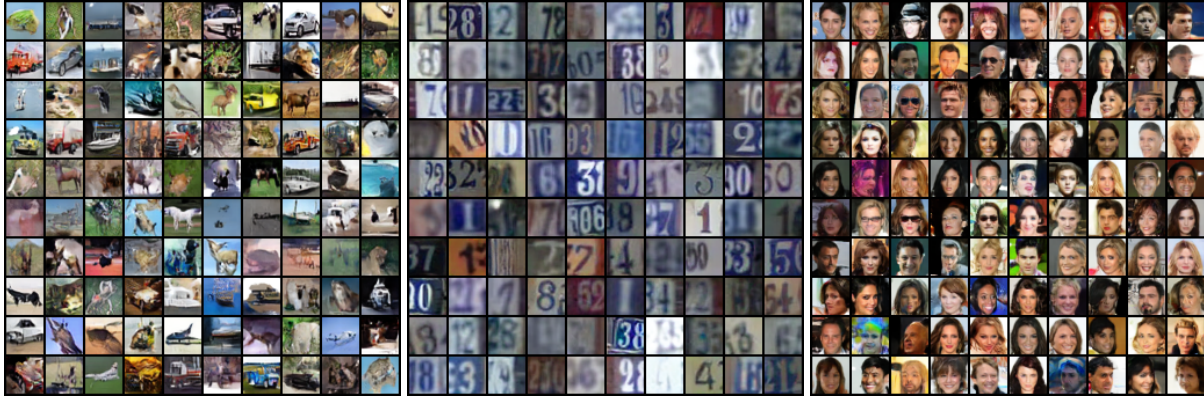


Figure 2. Randomly generated images. Left: CIFAR-10. Middle: SVHN. Right: CelebA-32.



Figure 3. Randomly generated images on CelebA-64.

Model	CIFAR-10	CelebA-64
VAE [22]	0.192	0.197
ALICE [26]	0.185	0.214
SVAE [37]	0.258	0.209
Divergence triangle [16]	0.177	0.190
BiDVL (ours)	0.168	0.187

Table 3. Evaluation of testing image reconstruction on CIFAR-10 and CelebA-64 using RMSE.

mean square error (RMSE) as the metric. Reconstruction images shown in Fig. 4 demonstrate our model can capture the major information in testing images, and the low reconstruction error implies that the information of real data flows to the EBLVM through IGM effectively.

5.3. Out-of-distribution detection

The EBLVM defines an unnormalized density function which can be used for detecting OOD samples. Generative models like GANs are infeasible for OOD because the



Figure 4. Test image reconstruction. Top: CIFAR-10. Bottom: CelebA-64. Left: test images. Right: reconstruction images.

model distribution is implicitly defined, while likelihood models like VAEs and Flows are accused of overestimating OOD regions, and as a result, fails to distinguish OOD samples. However the log unnormalized marginal density of EBLVM, *i.e.*, the negative marginal energy $-\mathcal{E}(v)$, is trained to assign low value on OOD regions and high value on data regions, which is suitable for our case. Since the marginal part of the decoupled EBLVM can model the visible space faithfully, we regard $-\mathcal{E}_{\psi'}(v)$ as the critic of samples. We mainly consider three OOD datasets: uniform noise, SVHN and CelebA. Our proposed model is trained on CIFAR-10 whose testing part is in-distribution dataset. Following [6], area under the ROC curve (AU-ROC) is used as our evaluation metric. Since training EBM is of a high variance process found in related work [16], our best model is chosen from the early stage of training.

Model	Random	SVHN	CelebA
IGEBM [6]	1.0	0.63	0.7
Glow [21]	1.0	0.24	0.57
SVAE [37]	0.29	0.42	0.52
Divergence triangle [16]	1.0	0.68	0.56
JEM [12]	1.0	0.67	0.75
VERA [11]	1.0	0.83	0.33
BiDVL (ours)	1.0	0.76	0.77

Table 4. Out-of-distribution detection on uniform, SVHN, CelebA test datasets, with CIFAR-10 as the in-distribution dataset. We report the AUROC of negative free energy.

Table 4 shows our model achieves compatible performance with most of baseline chosen from recent EBMs, VAEs and flows. Surprisingly, we find our proposed model outperforms JEM slightly, though JEM is an energy-based classifier model supposed to be good at OOD detection. Moreover, VERA [11], a recent MCMC-free method for learning EBMs, performs quite well on SVHN dataset, while our model is better on CelebA.

5.4. Ablations and analysis

We first investigate how the modifications proposed in Sec. 4.2 solve the problems by enhancing the learning stability. We conduct experiments on the original BiDVL with canonical KL, but it is unstable (even diverse) on CIFAR. With a similar phenomenon, only decoupling the EBLVM is still hard to converge. On the other side, the training of undecoupled EBLVM improves a lot with the symmetric KL and gets 32.10 FID on CIFAR, nonetheless much worse than our complete model that gets 20.75 shown in Sec. 5.1.

Next we study the rest choices in BiDVL. Since computing the importance ratio $r(v)$ is nontrivial, we heuristically estimate it with a basic term and a bias term.

$$r(v) = \frac{\mathbb{E}_{p_{\psi'}(v)}[q(v)]}{q(v)} \frac{\exp(-\mathcal{E}_{\psi'}(v))}{\mathbb{E}_{q(v)}[\exp(-\mathcal{E}_{\psi'}(v))]} \quad (15)$$

$$\approx r' \frac{\exp(-\mathcal{E}_{\psi'}(v))}{\mathbb{E}_{q(v)}[\exp(-\mathcal{E}_{\psi'}(v))]},$$

where the basic term r' corresponds to the average ratio, the bias term means samples with higher density under model distribution should be learned more. As training EBM is a high-variance process, which makes the estimate of the bias ratio noisy, we use Sigmoid($\mathbb{E}_{q(v)}[\mathcal{E}_{\psi'}(v)] - \mathcal{E}_{\psi'}(v)$) to replace the bias term but still leads to slightly inferior performance, consequently we ignore the bias term.

Furthermore, we found the proposed model performs quite sensitive to the basic term, as shown in Tab. 5 studying the influences of varying basic ratio on CIFAR. The generation quality degrades as the basic ratio increases and perform pretty worse with $r' = 1.0$. It corresponds to the intuition that the importance ratio is supposed to be small on

r'	0.01	0.05	0.1	0.5	1.0
FID	22.62	20.75	21.90	26.82	29.72

Table 5. The influences of varying basic ratio on FID. Models are trained on CIFAR-10.

	w/o	offset	not offset
FID	20.24	20.75	21.54
RMSE	0.175	0.168	0.170

Table 6. The influences of offsetting on FID and RMSE. Models are trained on CIFAR-10.

average since the energy-based distribution is mild and typically assigns density on unreal samples. Finally, we use 0.05 on 32×32 datasets and 0.1 on 64×64 datasets.

Then we study the effect of offsetting the opposite terms in Eqs. (11) and (12). The implementation details of not offset algorithm are follow Sec. 4.4. For comparison, we evaluate the algorithm without all opposite terms concerned. Not offset algorithm performs unstable on CelebA-64, even not converge sometimes. Table 6 shows the influence. We find the Sigmoid estimation of bias term helps the not offset algorithm stabilizing, however it tends to overfit early and performs slightly worse. Not offset version achieves 0.81 AUROC on SVHN and 0.75 AUROC on CelebA, implying the adversarial learning may help BiDVL better modelling. Algorithm without terms concerned is slightly better on generation, but is inferior on reconstruction and OOD detection. It implies the reconstruction over latent variables enhances the alignment. Our formal experiments are conducted by the offset version.

6. Conclusion

In this paper, we introduce BiDVL, a MCMC-free framework to efficiently learn EBLVMs. For image task, we achieve a decoupled EBLVM, consisting of a marginal energy-based distribution and a structural posterior, and then choose a symmetric KL divergence in the lower level. Experiments show the impressive generative power on images concerned. In this work, simply stacking several convolutional layers largely restricts our model from scaling to a large image dataset. In future work, the deeper and more sophisticated designed networks shall be adopted to evaluate BiDVL on larger-scale images.

Broader impacts. The method is for training EBLVMs, and the trained models have the ability of generating fake content, which may be used with potential malicious.

Acknowledgement This work is partially supported by the National Natural Science Foundation of China (62141604, 61972016, 62032016, 62106012, 62076016, 62101245), Natural Science Foundation of Beijing Municipality (L191007).

References

- [1] David H. Ackley, Geoffrey E. Hinton, and Terrence J. Sejnowski. A learning algorithm for boltzmann machines. *Cogn. Sci.*, 9:147–169, 1985. [2](#)
- [2] Michael Arbel, Liang Zhou, and Arthur Gretton. Generalized energy based models. In *ICLR*, 2021. [2](#), [6](#)
- [3] Fan Bao, Chongxuan Li, Kun Xu, Hang Su, Jun Zhu, and Bo Zhang. Bi-level score matching for learning energy-based latent variable models. In *NeurIPS*, 2020. [2](#), [4](#)
- [4] Fan Bao, Kun Xu, Chongxuan Li, Lanqing Hong, Jun Zhu, and Bo Zhang. Variational (gradient) estimate of the score function in energy-based latent variable models. In *ICML*, 2021. [1](#), [2](#)
- [5] Tong Che, Ruixiang Zhang, Jascha Sohl-Dickstein, H. Larochelle, Liam Paull, Yuan Cao, and Yoshua Bengio. Your gan is secretly an energy-based model and you should use discriminator driven latent sampling. In *NeurIPS*, 2020. [2](#)
- [6] Yilun Du and Igor Mordatch. Implicit generation and modeling with energy based models. In *NeurIPS*, 2019. [1](#), [2](#), [6](#), [7](#), [8](#)
- [7] Ruiqi Gao, Erik Nijkamp, Diederik P. Kingma, Zhen Xu, Andrew M. Dai, and Ying Nian Wu. Flow contrastive estimation of energy-based models. In *CVPR*, pages 7515–7525, 2020. [1](#), [6](#)
- [8] Cong Geng, Jia Wang, Zhiyong Gao, Jes Frellsen, and Søren Hauberg. Bounds all around: training energy-based models with bidirectional bounds. In *NeurIPS*, 2021. [2](#)
- [9] Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT Press, 2016. [4](#)
- [10] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014. [1](#), [3](#)
- [11] Will Grathwohl, Jacob Kelly, Milad Hashemi, Mohammad Norouzi, Kevin Swersky, and David Kristjanson Duvenaud. No mcmc for me: Amortized sampling for fast and stable training of energy-based models. In *ICLR*, 2021. [2](#), [6](#), [8](#)
- [12] Will Grathwohl, Kuan-Chieh Wang, Jörn-Henrik Jacobsen, David Kristjanson Duvenaud, Mohammad Norouzi, and Kevin Swersky. Your classifier is secretly an energy based model and you should treat it like one. In *ICLR*, 2020. [1](#), [8](#)
- [13] Will Grathwohl, Kuan-Chieh Wang, Jörn-Henrik Jacobsen, David Kristjanson Duvenaud, and Richard S. Zemel. Learning the stein discrepancy for training and evaluating energy-based models without sampling. In *ICML*, 2020. [2](#)
- [14] Michael U. Gutmann and Aapo Hyvärinen. Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *J. Mach. Learn. Res.*, 13:307–361, 2012. [2](#)
- [15] Tian Han, Erik Nijkamp, Xiaolin Fang, Mitch Hill, Song-Chun Zhu, and Ying Nian Wu. Divergence triangle for joint training of generator model, energy-based model, and inferential model. In *CVPR*, pages 8662–8671, 2019. [2](#)
- [16] Tian Han, Erik Nijkamp, Linqi Zhou, Bo Pang, Song-Chun Zhu, and Ying Nian Wu. Joint training of variational auto-encoder and latent energy-based model. In *CVPR*, pages 7975–7984, 2020. [2](#), [4](#), [6](#), [7](#), [8](#)
- [17] 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 *NeurIPS*, 2017. [6](#)
- [18] Aapo Hyvärinen. Estimation of non-normalized statistical models by score matching. *J. Mach. Learn. Res.*, 6:695–709, 2005. [2](#)
- [19] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015. [6](#)
- [20] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)
- [21] Diederik P. Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *NeurIPS*, 2018. [8](#)
- [22] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2014. [6](#), [7](#)
- [23] Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009. [6](#)
- [24] Yann LeCun, Sumit Chopra, Raia Hadsell, Aurelio Ranzato, and Fu Jie Huang. A tutorial on energy-based learning. 2006. [1](#), [3](#)
- [25] Chongxuan Li, Chao Du, Kun Xu, Max Welling, Jun Zhu, and Bo Zhang. To relieve your headache of training an mrf, take advil. In *ICLR*, 2020. [1](#), [2](#), [4](#)
- [26] Chunyuan Li, Hanying Liu, Changyou Chen, Yunchen Pu, Liqun Chen, Ricardo Henao, and Lawrence Carin. Alice: Towards understanding adversarial learning for joint distribution matching. In *NeurIPS*, 2017. [7](#)
- [27] Zengyi Li, Yubei Chen, and Friedrich T. Sommer. Learning energy-based models in high-dimensional spaces with multi-scale denoising score matching. *arXiv: Machine Learning*, 2019. [2](#)
- [28] Risheng Liu, Jiabin Gao, Jin Zhang, Deyu Meng, and Zhouchen Lin. Investigating bi-level optimization for learning and vision from a unified perspective: A survey and beyond. *arXiv preprint arXiv:2101.11517*, 2021. [3](#), [4](#)
- [29] Risheng Liu, Pan Mu, Xiaoming Yuan, Shangzhi Zeng, and Jin Zhang. A generic first-order algorithmic framework for bi-level programming beyond lower-level singleton. In *ICML*, 2020. [3](#)
- [30] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. In *NeurIPS*, volume 33, pages 21464–21475, 2020. [1](#)
- [31] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, pages 3730–3738, 2015. [6](#)
- [32] Luke Metz, Ben Poole, David Pfau, and Jascha Sohl-Dickstein. Unrolled generative adversarial networks. In *ICLR*, 2017. [4](#), [5](#)
- [33] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *ICLR*, 2018. [6](#), [13](#)

- [34] Yuval Netzer, Tiejie Wang, Adam Coates, A. Bissacco, Bo Wu, and A. Ng. Reading digits in natural images with unsupervised feature learning. 2011. [6](#)
- [35] Erik Nijkamp, Ruiqi Gao, Pavel Sountsov, Srinivas Vasudevan, Bo Pang, Song-Chun Zhu, and Ying Nian Wu. Learning energy-based model with flow-based backbone by neural transport mcmc. *arXiv preprint arXiv:2006.06897*, 2020. [1](#), [2](#)
- [36] Bo Pang, Tian Han, Erik Nijkamp, Song-Chun Zhu, and Ying Nian Wu. Learning latent space energy-based prior model. *arXiv preprint arXiv:2006.08205*, 2020. [2](#)
- [37] Yunchen Pu, Liqun Chen, Shuyang Dai, Weiyao Wang, Chunyuan Li, and Lawrence Carin. Symmetric variational autoencoder and connections to adversarial learning. In *AISTATS*, 2018. [7](#), [8](#)
- [38] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*, 2015. [6](#)
- [39] Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *ICML*, 2015. [5](#)
- [40] Ruslan Salakhutdinov and Geoffrey E. Hinton. Deep boltzmann machines. In *AISTATS*, 2009. [2](#)
- [41] Yuxuan Song, Q. Ye, Minkai Xu, and Tie-Yan Liu. Discriminator contrastive divergence: Semi-amortized generative modeling by exploring energy of the discriminator. *arXiv preprint arXiv:2004.01704*, 2020. [2](#)
- [42] Eszter Vértés. Learning doubly intractable latent variable models via score matching. In *NeurIPS*, 2016. [1](#), [2](#)
- [43] Zhisheng Xiao, Karsten Kreis, Jan Kautz, and Arash Vahdat. Vaebm: A symbiosis between variational autoencoders and energy-based models. In *ICLR*, 2021. [2](#)
- [44] Zhisheng Xiao, Qing Yan, and Yali Amit. Exponential tilting of generative models: Improving sample quality by training and sampling from latent energy. *arXiv preprint arXiv:2006.08100*, 2020. [2](#)
- [45] Jianwen Xie, Yang Lu, Ruiqi Gao, and Ying Nian Wu. Cooperative learning of energy-based model and latent variable model via mcmc teaching. In *AAAI*, 2018. [2](#)
- [46] Jianwen Xie, Zilong Zheng, and Ping Li. Learning energy-based model with variational auto-encoder as amortized sampler. In *AAAI*, 2021. [2](#)
- [47] Lantao Yu, Yang Song, Jiaming Song, and Stefano Ermon. Training deep energy-based models with f-divergence minimization. In *ICML*, 2020. [2](#)