
Conjugate Energy-Based Models

Hao Wu^{*1} Babak Esmaeili^{*1} Michael Wick² Jean-Baptiste Tristan³ Jan-Willem van de Meent¹

Abstract

In this paper, we propose conjugate energy-based models (CEBMs), a new class of energy-based models that define a joint density over data and latent variables. The joint density of a CEBM decomposes into an intractable distribution over data and a tractable posterior over latent variables. CEBMs have similar use cases as variational autoencoders, in the sense that they learn an unsupervised mapping from data to latent variables. However, these models omit a generator network, which allows them to learn more flexible notions of similarity between data points. Our experiments demonstrate that conjugate EBMs achieve competitive results in terms of image modelling, predictive power of latent space, and out-of-domain detection on a variety of datasets.

1. Introduction

Deep generative models approximate a data distribution by combining a prior over latent variables with a neural generator, which maps latent variables to points on a data manifold. It is common to evaluate these models in terms of their ability to generate realistic examples, or their estimated densities for unseen data. However, an arguably more important use case for these models is unsupervised representation learning. If a generator can faithfully represent the data in terms of a lower-dimensional set of latent variables, then we hope that these variables will encode a set of semantically meaningful factors of variation that will be relevant to a broad range of downstream tasks.

Guiding a model towards a semantically meaningful representation requires some form of inductive bias. A large body of work on variational autoencoders (VAEs, (Kingma & Welling, 2013; Rezende et al., 2014)) has explored the use of priors as inductive biases. Relatively mild biases in

the form of conditional independence are common in the literature on disentangled representations (Higgins et al., 2016; Kim & Mnih, 2018; Chen et al., 2018; Esmaeili et al., 2019). More generally, recent work has shown that defining priors that reflect the structure of the underlying data will lead to representations that are easier to interpret and generalize better. Examples include priors that represent objects in an image (Eslami et al., 2016; Lin et al., 2020b; Engelcke et al., 2019; Crawford & Pineau, 2019b), or moving objects in video (Crawford & Pineau, 2019a; Kosiorek et al., 2018; Wu et al., 2020; Lin et al., 2020a).

Despite steady progress, work on disentangled representations and structured VAEs still predominantly considers synthetic data. VAEs employ a neural generator that is optimized to reconstruct examples in the training set. For complex natural scenes, learning a generator that can produce pixel-perfect reconstructions poses fundamental challenges, given the combinatorial explosion of possible inputs. This is not only a problem for generation, but also from the perspective of the learned representation; a VAE must encode all factors of variation that give rise to large deviations in pixel space, regardless of whether these factors are semantically meaningful (e.g. presence and locations of objects) or not (e.g. shadows of objects in the background of the image).

The motivating question that we consider in this paper is whether it is possible to train latent-variable models without minimizing pixel-level discrepancies between an image and its reconstruction. Instead, we would like to design an objective that minimizes the discrepancy between the encoding of an image and the latent variables, which will in general be in a lower-dimensional space compared to the input. Our hope is that doing so will allow a model to learn more abstract representations, in the sense that it becomes easier to discard factors of variation that give rise to variation in pixel space, but should be considered noise.

In this paper, we consider energy-based models (EBMs) with latent variables as a particular instantiation of this general idea. EBMs with latent variables are by no means new; they have a long history in the context of restricted Boltzmann machines (RBMs) and related models (Smolensky, 1986; Hinton, 2002; Welling et al., 2004). Our motivation in the present work is to design a class of EBMs that retain the desirable features of VAEs, but employ a discriminative

^{*}Equal contribution ¹Khoury College of Computer Sciences, Northeastern University, Boston, MA, USA ²Oracle Labs, MA, USA ³Computer Science department, Boston College, MA, USA. Correspondence to: Hao Wu <wu.hao10@northeastern.edu>, Babak Esmaeili <esmaeili.b@northeastern.edu>.




Figure 1. Comparison between a VAE and a CEBM. A variational autoencoder with a Gaussian or Bernoulli likelihood has an energy that can be expressed in terms of a Bregman divergence in the data space $D_{A^*}(x, \mu_\theta(z))$ between an image x and the reconstruction from the generator network $\mu_\theta(z)$. The energy function in a CEBM can be expressed in terms of a Bregman divergence in the latent space $D_{B^*}(\eta(z), \tilde{\mu}_\theta(x))$ between a vector of natural parameters $\eta(z)$ and the output of an encoder network $\tilde{\mu}_\theta(x)$. See main text for details.

energy function to model data at an intermediate level of representation that does not necessarily encode all features of an image at the pixel level.

Concretely, we propose conjugate EBM (CEBMs), a new family of energy-based latent-variable models in which the energy function defines a neural exponential family. While the normalizer of CEBMs is intractable, we can nonetheless compute the posterior in closed form when we pair the likelihood with an appropriate conjugate bias term. As a result, the neural sufficient statistics in a CEBM fully determine both the marginal likelihood and the encoder, thereby side-stepping the need for a generator (Figure 1).

Our contributions can be summarized as follows:

1. We propose CEBMs, a class of energy-based models for unsupervised representation learning. The density of a CEBM factorizes into a tractable posterior and an energy-based marginal over data. This means that CEBMs can be trained using existing methods for EBMs, whilst inference is tractable at test time.
2. Unlike VAEs, CEBMs model data not at the pixel level, but at the level of the latent representation. We interpret the energy function of CEBMs in terms of a Bregman divergence in the latent space, and show that the density of a VAE can similarly be expressed in terms of a Bregman divergence in the data space.
3. We show that two of the most common inductive biases in VAEs can be incorporated in CEBMs: a spherical Gaussian and a mixture of Gaussians.
4. We evaluate how well CEBMs learned representations agree with class labels (which are not used during training). We show that neighbors are more likely to belong to the same class, which translates to increased performance in downstream classification tasks. Moreover, CEBMs perform competitively in out-of-domain detection. We do also note limitations; in particular we observe that CEBMs suffer from posterior collapse.

2. Background

2.1. Energy-Based Models

An EBM (LeCun et al., 2006) defines a probability density for $x \in \mathbb{R}^D$ via the Gibbs-Boltzmann distribution

$$p_\theta(x) = \frac{\exp\{-E_\theta(x)\}}{Z_\theta}, \quad Z_\theta = \int dx \exp\{-E_\theta(x)\}.$$

The function $E_\theta : \mathbb{R}^D \rightarrow \mathbb{R}$ is called the energy function which maps each configuration to a scalar value, the energy of the configuration. This type of model is widely used in statistical physics, for example in Ising models. The distribution can only be evaluated up to an unknown constant of proportionality, since computing the normalizing constant Z_θ (also known as the partition function) requires an intractable integral with respect to all possible inputs x .

Our goal is to learn a model $p_\theta(x)$ that is close to the true data distribution $p_{\text{data}}(x)$. A common strategy is to minimize the Kullback-Leibler divergence between the data distribution and the model, which is equivalent to maximizing the expected log-likelihood

$$\begin{aligned} \mathcal{L}(\theta) &= \mathbb{E}_{p_{\text{data}}(x)}[\log p_\theta(x)], \\ &= \mathbb{E}_{p_{\text{data}}(x)}[-E_\theta(x)] - \log Z_\theta. \end{aligned} \quad (1)$$

The key difficulty when performing maximum likelihood estimation is that computing the gradient of $\log Z_\theta$ is intractable. This gradient can be expressed as an expectation with respect to $p_\theta(x)$,

$$\nabla \log Z_\theta = \mathbb{E}_{p_\theta(x')} [-\nabla_\theta E_\theta(x')], \quad (2)$$

which means that the gradient of $\mathcal{L}(\theta)$ has the form:

$$\nabla_\theta \mathcal{L}(\theta) = -\mathbb{E}_{p_{\text{data}}(x)}[\nabla_\theta E_\theta(x)] + \mathbb{E}_{p_\theta(x')}[\nabla_\theta E_\theta(x')].$$

This corresponds to *maximizing* the probability of samples $x \sim p_{\text{data}}(x)$ from the data distribution and *minimizing* the probability of samples $x' \sim p_\theta(x')$ from the learned model.

Contrastive divergence methods (Hinton, 2002) compute a Monte Carlo estimate of this gradient, which requires a method for approximate inference to generate samples $x' \sim p_\theta(x')$. A common method for generating samples from EBMs is Stochastic Gradient Langevin Dynamics (SGLD, (Welling & Teh, 2011)), which initializes a sample $x'_0 \sim p_0(x')$ and performs a sequence of gradient updates with additional injected noise ϵ ,

$$x'_{i+1} = x'_i - \frac{\alpha}{2} \frac{\partial E_\theta(x')}{\partial x'} + \epsilon, \quad \epsilon \sim N(0, \alpha). \quad (3)$$

SGLD is motivated as a discretization of a stochastic differential equation whose stationary distribution is equal to the target distribution. It is correct in the limit $i \rightarrow \infty$ and $\alpha \rightarrow 0$, but in practice will have a bias.

The initialization x'_0 is crucial because it determines the number of steps needed to converge to a high-quality sample. For this reason, EBMs are commonly trained using persistent contrastive divergence (PCD, (Du & Mordatch, 2019; Tieleman, 2008)), which initializes some samples from a replay buffer $\mathcal{B}$ of previously generated samples (Nijkamp et al., 2019a; Du & Mordatch, 2019; Xie et al., 2016).

2.2. Energy-Based Latent-Variable Models

Energy-based latent-variable models are a subclass of EBMs where the energy function defines joint density on observed data $x \in \mathbb{R}^D$ and latent variable $z \in \mathbb{R}^K$,

$$p_\theta(x, z) = \frac{\exp\{-E_\theta(x, z)\}}{Z_\theta}. \quad (4)$$

Some of the most well-known examples of this family of models include restricted Boltzmann machines (RBMs, (Smolensky, 1986; Hinton, 2002)), deep belief nets (DBNs, (Hinton et al., 2006)), and deep Boltzmann machines (DBMs, (Salakhutdinov & Hinton, 2009)).

Similar to standard EBMs, energy-based latent-variable models can also be trained using contrastive divergence methods, where the gradient of $\mathcal{L}(\theta)$ can be expressed as:

$$-\mathbb{E}_{p_{\text{data}}(x)p_\theta(z|x)}[\nabla_\theta E_\theta(x, z)] + \mathbb{E}_{p_\theta(x', z')}[\nabla_\theta E_\theta(x', z')].$$

Estimating this gradient has the additional problem of requiring samples from the posterior $p_\theta(z|x)$ which is also intractable in general.

2.3. Conjugate Exponential Families

An exponential family is a set of distributions whose probability density can be expressed in the form

$$p(x | \eta) = h(x) \exp\{\langle t(x), \eta \rangle - A(\eta)\}, \quad (5)$$

where $h : \mathcal{X} \rightarrow \mathbb{R}^+$ is a base measure, $\eta \in \mathcal{H} \subseteq \mathbb{R}^K$ is a vector of natural parameters, $t : \mathcal{X} \rightarrow \mathbb{R}^K$ is a vector of

sufficient statistics, and $A : \mathcal{H} \rightarrow \mathbb{R}$ is the log normalizer (or cumulant function),

$$A(\eta) = \log Z(\eta) = \int dx h(x) \exp\{\langle t(x), \eta \rangle\}. \quad (6)$$

If a likelihood belongs to an exponential family, then there exists a conjugate prior that is itself an exponential family

$$p(\eta | \lambda, \nu) = \exp\{\langle \eta, \lambda \rangle - A(\eta)\nu - B(\lambda, \nu)\}. \quad (7)$$

The convenient property of conjugate exponential families is that both the marginal likelihood $p(x | \lambda, \nu)$ and the posterior $p(\eta | x, \lambda, \nu)$ are tractable. If we define

$$\tilde{\lambda}(x) = \lambda + t(x), \quad \tilde{\nu} = \nu + 1, \quad (8)$$

then the posterior and marginal likelihood are

$$\begin{aligned} p(\eta | x, \lambda, \nu) &= p(\eta | \tilde{\lambda}(x), \tilde{\nu}), \\ p(x | \lambda, \nu) &= h(x) \exp\{B(\tilde{\lambda}(x), \tilde{\nu}) - B(\lambda, \nu)\}. \end{aligned} \quad (9)$$

2.4. Legendre Duality in Exponential Families

Two convex functions $A : \mathcal{H} \rightarrow \mathbb{R}^+$ and $A^* : \mathcal{M} \rightarrow \mathbb{R}^+$ on spaces $\mathcal{H} \subseteq \mathbb{R}^K$ and $\mathcal{M} \subseteq \mathbb{R}^K$ are conjugate duals when

$$A^*(\mu) := \sup_{\eta \in \mathcal{H}} \{\langle \mu, \eta \rangle - A(\eta)\}. \quad (10)$$

When A is a function of Legendre type (see Rockafellar (1970) for details), the gradients of these functions define a bijection between conjugate spaces by mapping points to their corresponding suprema

$$\eta(\mu) = \nabla A^*(\mu), \quad \mu(\eta) = \nabla A(\eta), \quad (11)$$

such that we can express $A^*(\mu)$ at the supremum as

$$A^*(\mu) = \langle \mu, \eta(\mu) \rangle - A(\eta(\mu)) \quad (12)$$

The log normalizer $A(\eta)$ of an exponential family is of Legendre type when the family is regular and minimal ($\mathcal{H}$ is an open set and sufficient statistics $t(x)$ are linearly independent; see Wainwright & Jordan (2008) for details). We refer to $\mathcal{M}$ as the mean parameter space, since we can express any $\mu \in \mathcal{M}$ as the expected value of the sufficient statistics

$$\mu(\eta) = \mathbb{E}_{p(x|\eta)}[t(x)]. \quad (13)$$

2.5. Bregman Divergences and Exponential Families

A Bregman divergence for a function $F : \mathcal{M} \rightarrow \mathbb{R}$ that is continuously-differentiable and strictly convex on a closed set $\mathcal{M}$ has the form

$$D_F(\mu', \mu) = F(\mu') - F(\mu) - \langle \mu' - \mu, \nabla F(\mu) \rangle. \quad (14)$$

Well-known special cases of Bregman divergences include the squared distance ($F(\mu) = \langle \mu, \mu \rangle$) and the Kullback-Leiber (KL) divergence ($F(\mu) = \sum_k \mu_k \log \mu_k$).

Any Bregman divergence can be associated with an exponential family and vice versa, where $F(\mu) = A^*(\mu)$ is the conjugate dual of $A(\eta)$ (see Banerjee et al. (2005)). To see this, we re-express the log density of a (regular and minimal) exponential family using the substitution $\mu = \nabla A(\eta)$ ¹,

$$\begin{aligned} \log p(x | \eta) &= \langle t(x), \eta \rangle - A(\eta), \\ &= (\langle \mu, \eta \rangle - A(\eta)) + \langle t(x) - \mu, \eta \rangle, \\ &= A^*(\mu) + \langle t(x) - \mu, \nabla A^*(\mu) \rangle, \\ &= -D_{A^*}(t(x), \mu) + A^*(t(x)). \end{aligned} \quad (15)$$

In other words, the log density of an exponential family can be expressed in terms of a bias term $A^*(t(x))$ ², and a notion of agreement in the form of a Bregman divergence $D_{A^*}(t(x), \mu)$ between the sufficient statistics $t(x)$ and the mean parameters μ . We will make use of this property of exponential families to provide an interpretation of both CEBMs and VAEs in terms of Bregman divergences.

3. Conjugate Energy-Based Models

We are interested in learning a probabilistic model that defines a joint density $p_{\theta, \lambda}(x, z)$ over high-dimensional data $x \in \mathbb{R}^D$ and a lower-dimensional set of latent variables $z \in \mathbb{R}^K$. The intuition that guides our work is that we would like to measure agreement between latent variables and data at a high level of representation, rather than at the level of individual pixels, where it may be more difficult to distinguish informative features from noise. To this end, we will explore energy-based models as an alternative to VAEs.

Concretely, we propose to consider models of the form

$$p_{\theta, \lambda}(x, z) = \frac{1}{Z_{\theta, \lambda}} \exp \{ -E_{\theta, \lambda}(x, z) \}, \quad (16)$$

where the energy function takes a form that is inspired by exponential family distributions

$$E_{\theta, \lambda}(x, z) = -\langle t_{\theta}(x), \eta(z) \rangle + E_{\lambda}(z). \quad (17)$$

In this energy function, θ are the weights of a network $t_{\theta} : \mathbb{R}^D \rightarrow \mathbb{R}^H$, which plays the role of an encoder by mapping high-dimensional data to a lower-dimensional vector of neural sufficient statistics. The function $\eta : \mathbb{R}^K \rightarrow \mathbb{R}^H$ maps latent variables to a vector of natural parameters in the same space as the neural sufficient statistics. The function $E_{\lambda} : \mathbb{R}^K \rightarrow \mathbb{R}$ serves as an inductive bias, with hyperparameters λ , that plays a role analogous to the prior.

¹We here omit the base measure $h(x)$ for notational simplicity.

²Or $A^*(t(x)) + \log h(x)$ when we include $h(x)$ the density.

We will consider a bias $E_{\lambda}(z)$ in form of a tractable exponential family with sufficient statistics $\eta(z)$

$$E_{\lambda}(z) = -\log p_{\lambda}(z) = -\langle \eta(z), \lambda \rangle + B(\lambda). \quad (18)$$

We can then express the energy function as

$$E_{\theta, \lambda}(x, z) = -\langle \lambda + t_{\theta}(x), \eta(z) \rangle + B(\lambda). \quad (19)$$

This form of the energy function has a convenient property: It corresponds to a model $p_{\theta, \lambda}(x, z)$ in which the posterior $p_{\theta, \lambda}(z | x)$ is tractable. To see this, we make a substitution $\tilde{\lambda}_{\theta}(x) = \lambda + t_{\theta}(x)$ analogous to the one in Equation 8, which allows us to express the energy as

$$E_{\theta, \lambda}(x, z) = -\langle \eta(z), \tilde{\lambda}_{\theta}(x) \rangle + B(\tilde{\lambda}_{\theta}(x)) + E_{\theta, \lambda}(x), \quad (20)$$

$$E_{\theta, \lambda}(x) = -B(\tilde{\lambda}_{\theta}(x)) + B(\lambda). \quad (21)$$

We see that we can factorize the corresponding density

$$p_{\theta, \lambda}(x, z) = p_{\theta, \lambda}(x) p_{\theta, \lambda}(z | x), \quad (22)$$

which yields a posterior and marginal that are analogous the distributions in Equation 9

$$p_{\theta, \lambda}(z | x) = p(z | \tilde{\lambda}_{\theta}(x)), \quad (23)$$

$$\begin{aligned} p_{\theta, \lambda}(x) &= \frac{1}{Z_{\theta, \lambda}} \exp \{ -E_{\theta, \lambda}(x) \}, \\ &= \frac{1}{Z_{\theta, \lambda}} \exp \{ B(\tilde{\lambda}_{\theta}(x)) - B(\lambda) \}. \end{aligned} \quad (24)$$

In other words, the joint density of this model factorizes into a tractable posterior $p_{\theta, \lambda}(z | x)$ and an intractable energy-based marginal likelihood $p_{\theta, \lambda}(x)$. This posterior is conjugate, in the sense that it is in the same exponential family as the bias. For this reason, we refer to this class of models as conjugate energy-based models (CEBMs).

4. Relationship to VAEs

CEBMs differ from VAEs in that they lack a generator network. Instead, the density is fully specified by the encoder network $t_{\theta}(x)$, which defines a notion of agreement $\langle \tilde{\lambda}_{\theta}(x), \eta(z) \rangle$ between data and latent variables in the latent space. As with other exponential families, we can make this notion of agreement explicit by expressing the conjugate posterior in terms of a Bregman divergence using the decomposition in Equation 15

$$\begin{aligned} E_{\theta, \lambda}(x, z) &= D_{B^*}(\eta(z), \tilde{\mu}_{\theta}(x)) \\ &\quad - B^*(\eta(z)) + E_{\theta, \lambda}(x). \end{aligned} \quad (25)$$

Here $B^*(\mu)$ is the conjugate dual of the the log normalizer $B(\lambda)$, and we use $\tilde{\mu}_{\theta}(x) = \mu(\tilde{\lambda}_{\theta}(x))$ as a shorthand for the mean-space posterior parameters. We see that maximizing the density corresponds to minimizing a Bregman divergence in the space of sufficient statistics of the bias.

In Figure 1, we compare CEBMs to VAE in terms of the energy function for the log density of the generative model. In making this comparison, we have to keep in mind that these models are trained using different methods, and that VAEs have a tractable density $p_\theta(x, z)$. That said, the objectives in both models maximize the marginal likelihood, so we believe that it is instructive to write down the corresponding Bregman divergence in the VAE likelihood. This likelihood is typically a Gaussian with known variance, or a Bernoulli distribution (when modeling binarized images). Both distributions have sufficient statistics $t(x) = x$. Once again omitting the base measure $h(x)$ for expediency, we can express the log density of a VAE as an energy

$$\begin{aligned} E_{\theta, \lambda}(x, z) &= -\log p_\theta(x|z) - \log p_\lambda(z), \\ &= -\langle x, \eta_\theta(z) \rangle + A(\eta_\theta(z)) - \log p_\lambda(z). \\ &= D_{A^*}(x, \mu_\theta(z)) - A^*(x) - \log p_\lambda(z) \end{aligned} \quad (26)$$

Here $A^*(x)$ is the conjugate dual of the log normalizer $A(\eta)$, and we use $\eta_\theta(z)$ and $\mu_\theta(z)$ to refer to the output of the generator network in the natural-parameter and the mean-parameter space respectively. To reduce clutter and accommodate the case where a base measure $h(x)$ is needed (e.g. that of a Gaussian likelihood with known variance), we will introduce the additional shorthands

$$E(x) = -A(x) - \log h(x), \quad E_\lambda(z) = -\log p_\lambda(z). \quad (27)$$

We then see that the energy function of a VAE has the form

$$E_{\theta, \lambda}(x, z) = D_{A^*}(x, \mu_\theta(z)) + E(x) + E_\lambda(z). \quad (28)$$

Like that of a CEBM, the energy function of a VAE contains a Bregman divergence, as well as two terms that depend only on x and z . However, whereas the Bregman divergence in CEBM is defined in the mean-parameter space of the latent variables, that of a VAE is computed in the data space.

5. Inductive Biases

CEBMs have a property that is somewhat counter-intuitive. While the posterior $p_{\theta, \lambda}(z | x)$ in this class of models is tractable, the prior is in general not tractable. In particular, although the bias $-E_\lambda(z)$ is the logarithm of a tractable exponential family, it is not the case that $p_{\theta, \lambda}(z) = p_\lambda(z)$. Rather the prior $p_{\theta, \lambda}(z)$ has the form,

$$p_{\theta, \lambda}(z) = \frac{\exp\{-E_\lambda(z)\}}{Z_{\theta, \lambda}} \int dx \exp\{\langle t_\theta(x), \eta(z) \rangle\}.$$

In other words, $E_\lambda(z)$ defines an inductive bias, but this bias is different from the tractable prior in a VAE³, in the

³The bias in a VAE contains the log prior $\log p_\lambda(z)$ and the log normalizer $A(\eta_\theta(z))$ of the likelihood. In a CEBM, by contrast, we omit the term $A_\theta(\eta(z)) = \log \int dx \exp\{\langle t_\theta(x), \eta(z) \rangle\}$, which is intractable, and hereby implicitly absorb it into its prior.

sense that it imposes only a soft constraint on the geometry of the latent space.

In principle, the bias in a CEBM can take the form of any exponential family distribution. Since products of exponential families are also in the exponential family, this covers a broad range of possible biases. For purposes of evaluation in this paper, we will constrain ourselves to two cases:

1. Spherical Gaussian. As a bias that is analogous to the standard prior in VAEs, we consider a spherical Gaussian with fixed hyperparameters $(\mu, \sigma) = (0, 1)$ for each dimension of $z \in \mathbb{R}^K$,

$$E_\lambda(z) = -\sum_k (\langle \eta(z_k), \lambda \rangle - B(\lambda)).$$

Each term has sufficient statistics $\eta(z_k) = (z_k, z_k^2)$, natural parameters λ , and log normalizer $B(\lambda)$ as

$$\begin{aligned} \lambda &= \left(\frac{\mu}{\sigma^2}, -\frac{1}{2\sigma^2} \right) = \left(0, -\frac{1}{2} \right), \\ B(\lambda) &= -\frac{\lambda_1^2}{4\lambda_2} - \frac{1}{2} \log(-2\lambda_2). \end{aligned}$$

The marginal likelihood of the CEBM is then

$$p_{\theta, \lambda}(x) = \frac{1}{Z_{\theta, \lambda}} \exp \left\{ \sum_k (B(\tilde{\lambda}_{\theta, k}(x)) - B(\lambda)) \right\},$$

where $\tilde{\lambda}_{\theta, k}(x) = \lambda + t_{\theta, k}(x)$ and $t_{\theta, k}(x)$ is the sufficient statistics that corresponds to z_k .

2. Mixture of Gaussians. In our experiments, we will consider datasets that are normally used for classification. These datasets, by design, exhibit multimodal structure that we would like to see reflected in the learned representation. In order to design a model that is amenable to uncovering this structure, we will extend the energy function in Equation 17 to contain a mixture component y

$$E_{\theta, \lambda}(x, y, z) = -\langle t_\theta(x), \eta(y, z) \rangle + E_\lambda(y, z).$$

As an inductive bias, we will consider a bias in the form of a mixture of L Gaussians,

$$E_\lambda(y, z) = -\sum_{k, l} I[y = l] (\langle \eta(z_k), \lambda_{l, k} \rangle - B(\lambda_{l, k})).$$

Here $z \in \mathbb{R}^K$ is a vector of features and $y \in \{1, \dots, L\}$ is a categorical assignment variable. The bias for each component l is a spherical Gaussian with hyperparameters $\lambda_{l, k}$ for each dimension k . Again, using the notation $\tilde{\lambda}_{\theta, l, k} = \lambda_{l, k} + t_{\theta, k}(x)$ to refer to the posterior parameters, then we obtain an energy

$$E_{\theta, \lambda}(x, y, z) = -\sum_{k, l} I[y = l] (\langle \eta(z_k), \tilde{\lambda}_{\theta, l, k} \rangle - B(\lambda_{l, k})).$$

We can then define a joint probability over data x and the assignment y in terms the log normalizer $B(\cdot)$,

$$p_{\theta,\lambda}(x, y) = \frac{1}{Z_{\theta,\lambda}} \exp \left\{ \sum_{k,l} I[y = l] (B(\tilde{\lambda}_{\theta,l,k}) - B(\lambda_{l,k})) \right\},$$

which then allows us to compute the marginal $p_{\theta,\lambda}(x)$ by summing over y . We optimize this marginal with respect hyperparameters $\lambda_{l,k}$ as well as the weights θ .

6. Related Work

Energy-Based Latent-Variable Models. The idea of using EBMs to jointly model data and latent variables has a long history in the machine learning literature. Examples of this class of models include restricted Boltzmann machines (RBMs, (Smolensky, 1986; Hinton, 2002)), deep belief nets (DBNs, (Hinton et al., 2006)), and deep Boltzmann machines (DBMs, (Salakhutdinov & Hinton, 2009)). The idea of extending RBMs in exponential families and exploiting conjugacy to yield a tractable posterior is also not new and has been explored in Exponential Family Harmoniums (EFHs; (Welling et al., 2004)). These models differ from CEBMs in that they employ a bilinear interaction term $x^\top Wz$, which ensures that both the likelihood $p(x | z)$ and $p(z | x)$ are tractable. In CEBMs, the corresponding term $t_\theta(x)^\top z$ is nonlinear, which means that the posterior is tractable, but the likelihood is not. We provide a more detailed discussion regarding the connection of our work to this class of models in Appendix A.

EBMs for Image Modelling. Recent work has shown that EBMs with convolutional energy functions can accurately model distributions over images (Xie et al., 2016; Nijkamp et al., 2019a;b; Du & Mordatch, 2019; Xie et al., 2021a). This line of work typically focuses on generation and not on unsupervised representation learning as we do here. A line of work, which is similar to ours in spirit, employs EBMs as priors on the latent space of deep generative models (Pang et al., 2020; Aneja et al., 2020). These approaches, unlike our work, require a generator.

Interpretation of other models as EBMs. Grathwohl et al. (2019); Liu & Abbeel (2020); Xie et al. (2016) have proposed to interpret a classifier as an EBM that defines a joint energy function on the data and labels. CEBMs with a discrete bias can be interpreted as the unsupervised variant of this model class. Che et al. (2020) interpret a GAN as an EBM defined by both the generator and discriminator.

Training EBMs. A commonly used training method is PCD (Tieleman, 2008), where the MCMC is initialized from a replay buffer that stores the previously generated samples (Du & Mordatch, 2019), or from a generator (Xie

et al., 2018; 2020; 2021a). Nijkamp et al. (2019a;b) comprehensively investigate the convergence of PCD based on a variety of factors such as MCMC initialization, network architecture, and the optimizer. They find that the difference between the energy of the data and model samples is a good diagnostic of training stability. Many of these findings were helpful during the training and evaluation in our work.

There is a large literature on alternative training methods. Gao et al. (2020) propose to use the noise contrastive estimation (NCE, (Gutmann & Hyvärinen, 2010)), where they pretrain a flow-based noise model and then train the EBM to discriminate between the real data examples and the ones generated from the noise model. Another popular approach is the score matching (SM, Vértés et al. (2016); Hyvärinen & Dayan (2005); Vincent (2011); Song et al. (2020); Bao et al. (2020)), which learns EBMs by matching the gradient of the log probability density of the model distribution to that of the data distribution. Bao et al. (2020) propose a bi-level version of this method where it is also applicable to latent-variable models. To sidestep the need for MCMC sampling, Han et al. (2019; 2020); Xie et al. (2021b) jointly train an EBM with a VAE in an adversarial manner; Grathwohl et al. (2021) learn a generator by entropy regularization. We refer the readers to Song & Kingma (2021) for a more comprehensive discussion on training methods for EBMs.

7. Experiments

Our experiments evaluate to what extent CEBMs can learn representations that encode meaningful factors of variation, whilst discarding details about the input that we would consider noise. This question is difficult to answer in generality, and in some sense not well-posed; whether a factor of variation should be considered signal or noise can depend on context. For this reason, our experiments primarily focus on the extent to which representations in CEBMs can recover the multimodal structure in datasets that are normally used for classification. While class labels are an imperfect proxy, in the sense that they do not reflect all factors of variation that we may want to encode in a representation, they provide a means of quantifying differences between representations that were learned in an unsupervised manner.

We begin with a qualitative evaluation by visualizing samples and latent representation. We then demonstrate that learned representations align with class structure, in the sense that nearest neighbors in the latent space are more likely to belong to the same class (section 7.2). Next, we evaluate performance on out-of-distribution detection (OOD) tasks which, although not our primary focus in this paper, are a common use case for EBMs (Section 7.3). We then quantify the extent to which the learned representations can improve performance in downstream task, we measure few-label classification accuracy for representations that




Figure 2. Samples generated from a CEBM trained on MNIST, Fashion-MNIST, SVHN and CIFAR-10.




Figure 3. (Left) Samples from CIFAR-10 along with the top 2-nearest-neighbors in pixel space, the latent space of a VAE, and the latent space of a CEBM. (Right) Confusion matrices of 1-nearest-neighbor classification on CIFAR-10 based on L2 distance in the latent space. On average, CEBM representations more closely align with class labels compared to VAE.

were pre-trained without supervision (Section 7.4). Finally, we perform a more in-depth study of the latent space where we investigate to what extent the aggregate posterior distribution is close to the inductive bias as well how vulnerable CEBMs are to posterior collapse (Section 7.5).

7.1. Network Architectures and Training

Architectures & Optimization. The CEBMs in our experiments employ an encoder network $t_\theta(x)$ in the form of 4-layer CNN (as proposed by Nijkamp et al. (2019a)), followed by an MLP output layer. We choose the dimension of latent variables to be 128. We found that the optimization becomes difficult with smaller dimensions. We train our models using 60 SGLD steps, 90k gradient steps, batch size 128, Adam optimizer with learning rate $1e-4$. For training stability, we L2 regularize energy magnitudes (proposed by Du & Mordatch (2019)). See Appendix C for details.

Hyperparameter Sensitivity. As observed in previous work (Du & Mordatch, 2019; Grathwohl et al., 2019), training EBMs is challenging and often requires a thorough hyperparameters search. We found that the choices of activation function, learning rate, number of SGLD steps, and regularization will all affect training stability. Models regularly diverge during training, and it is difficult to perform diagnostics given that $\log p_{\theta,\lambda}(x)$ cannot be computed. As suggested by (Nijkamp et al., 2019a), we found checking the difference in energy between data and model samples

can help to verify training stability. In general we also observed a trade-off between sample quality and the predictive power of latent variables in our experiments. We leave investigation of the source of this trade-off to future work, but we suspect that this is because SGLD has more difficulty converging when the latent space is more disjoint.

7.2. Samples and Latent Space

We begin with a qualitative evaluation by visualizing samples from the model. While generation is not our intended use case in this paper, such samples do serve as a diagnostic that allows us to visually inspect what characteristics of the input data are captured by the learned representation.

Figure 2 shows samples from CEBMs trained on MNIST, Fashion-MNIST, SVHN, and CIFAR-10. We initialize the samples with uniform noise and run 500 SGLD steps. We observe that the distribution over images is diverse and captures the main characteristics of the dataset. Sample quality is roughly on par with samples from other EBMs (Nijkamp et al., 2019a), although it is possible to generate samples with higher visual quality using class-conditional EBMs (Du & Mordatch, 2019; Grathwohl et al., 2019; Liu & Abbeel, 2020) (which assume access to labels).

To assess to whether the representation in CEBMs aligns with classes in each dataset, we look at the agreement between the label of an input and that of its nearest neighbor in the latent space. The latent representations are inferred by

Table 1. AUROC scores in OOD Detection. We use $\log p_\theta(\mathbf{x})$ and $\|\nabla_{\mathbf{x}} \log p_\theta(\mathbf{x})\|$ as score functions. The left block shows results of the models trained on F-MNIST and tested on MNIST, E-MNIST, Constant (C); The right block shows results of the models trained on CIFAR-10 and tested on SVHN, Texture and Constant (C).

	Fashion-MNIST						CIFAR-10					
	$\log p_\theta(\mathbf{x})$			$\ \nabla_{\mathbf{x}} \log p_\theta(\mathbf{x})\ $			$\log p_\theta(\mathbf{x})$			$\ \nabla_{\mathbf{x}} \log p_\theta(\mathbf{x})\ $		
	MNIST	E-MNIST	C	MNIST	E-MNIST	C	SVHN	Texture	C	SVHN	Texture	C
VAE	.50	.39	.09	.61	.57	.01	.42	.58	.41	.38	.51	.37
IGEBM	.35	.36	.90	.78	.82	.96	.45	.31	.64	.33	.17	.62
CEBM	.37	.34	.90	.82	.89	.98	.47	.32	.66	.31	.17	.54
GMM-CEBM	.56	.56	.92	.56	.80	.95	.55	.30	.62	.40	.23	.62

Table 2. Average classification accuracy on the test set. We train a variety of deep generative models on MNIST, Fashion-MNIST, CIFAR-10, and SVHN in an unsupervised way. Then we use the learned latent representations to train logistic classifiers with 1, 10, 100 training examples per class, and the full training set. We train each classifier 10 times on randomly drawn training examples.

Models	MNIST				Fashion-MNIST				CIFAR-10				SVHN			
	1	10	100	full	1	10	100	full	1	10	100	full	1	10	100	full
VAE	42	85	92	95	41	63	72	81	16	22	31	38	13	13	16	36
GMM-VAE	53	86	93	97	49	68	79	84	19	23	33	39	13	14	23	56
BIGAN	33	67	85	91	46	65	75	81	18	30	43	52	11	20	42	56
IGEBM	63	89	95	97	50	70	79	83	16	26	33	42	10	16	35	49
CEBM	67	89	95	97	52	70	77	83	19	30	42	53	12	25	48	70
GMM-CEBM	67	91	97	98	52	70	80	85	16	29	42	52	10	17	39	60

computing the mean of the posterior $p_{\theta,\lambda}(z|x)$. In Figure 3, we show samples from CIFAR-10, along with the images that correspond to the nearest neighbors in pixel space, the latent space of a VAE, and the latent space of a CEBM. The distance in pixel space is a poor measure of similarity in this dataset, whereas proximity in the latent space is more likely to agree with class labels in both VAEs and CEBMs. We additionally show visualization of the latent space with UMAP (McInnes et al., 2018) in Figure 5.

In Figure 3 (right), we quantify this agreement by computing the fraction of neighbors in each class conditioned on the class of the original image. We see a stronger alignment between classes and the latent representation in CEBMs, which is reflected in higher numbers on the diagonal of the matrix. On average, a fraction of 0.38 of the nearest neighbors are in the same class in the VAE, whereas 0.45 of the neighbors are in the same class in the CEBM. This suggest that the representation in CEBMs should lead to higher performance in downstream classification tasks. We will evaluate this performance in Section 7.4.

7.3. Out-of-Distribution Detection

EBMs have formed the basis for encouraging results in out-of-distribution (OOD) detection (Du & Mordatch, 2019; Grathwohl et al., 2019). While not our focus in this paper, OOD detection is a benchmark that helps evaluate whether a

learned model accurately characterizes the data distribution. In Table 1, we report results in terms of two metrics. The first is the area under the receiver-operator curve (AUROC) when thresholding the log marginal $\log p_{\theta,\lambda}(x)$. The second is the gradient-based score function proposed by Grathwohl et al. (2019). We observe that in most cases, CEBM yields a similar score to the VAE and IGEBM baselines.

7.4. Few-label Classification

To evaluate performance in settings where few labels are available, we use pre-trained representations (which were learned without supervision) to train logistic classifiers with 1, 10, 100 training examples per class, as well as the full training set. We evaluate classification performance for a spherical Gaussian bias (CEBM) and the mixture of Gaussians bias (GMM-CEBM). We compare our models against the IGEBM (Du & Mordatch, 2019)⁴, a standard VAE with the spherical Gaussian prior, GMM-VAE (Tomczak & Welling, 2018) where the prior is a mixture of Gaussians (GMM), and BIGAN (Donahue et al., 2016).

We report the classification accuracy on the test set in Table 2. CEBMs overall achieve a higher accuracy compared to VAEs in particular for CIFAR-10 and SVHN where the

⁴Since the IGEBM does not explicitly have latent representations, we extract features from the last layer of the energy function.

Table 3. KL divergence between aggregate posterior and prior and the mutual information between data and latent variables.

	VAE		CEBM		GMM-CEBM	
	KL	MI	KL	MI	KL	MI
MNIST	11.5	9.1	0.9	0.3	18.7	4.7
FMNIST	3.5	9.0	0.6	0.4	8.1	3.9
CIFAR10	21.5	9.2	0.1	0.2	4.5	2.7
SVHN	8.6	10.1	0.1	0.1	5.6	2.2

pixel distance is not good measure for similarity. Moreover, we observe that CEBMs outperform the IGBM. This suggests that the inductive biases in CEBMs can lead to increased performance in downstream tasks. The performance between BIGANs and CEBMs is not as distinguishable which we suspect is due the fact BIGANs, just like CEBMs, do not define a likelihood that measure similarity at the pixel level. We also observe that the CEBM with the GMM inductive bias does not always outperform the one with the Gaussian inductive bias, which we suspect is due to GMM-CEBM having more difficulty to converge.

7.5. Limitations: Posterior Collapse

While our experiments demonstrate that CEBMs are able to reasonably approximate the data distribution and learn latent representations that are in closer agreement with class labels, they do not evaluate the learned notation of posterior uncertainty, and more generally the role of inductive bias. In this subsection, we ask the following two questions: (1) Does the aggregate posterior distribution of the training data live close to the inductive bias $p_\lambda(z)$? (2) What is the mutual information between latent variables and the training data?

To evaluate whether encoded examples are distributed according to the bias $p_\lambda(z)$, we compute the divergence $\text{KL}(\tilde{p}_{\theta,\lambda}(z) \parallel p_\lambda(z))$ between the bias and the aggregate posterior, which is a mixture over training data

$$\hat{p}_{\theta,\lambda}(z) = \frac{1}{N} \sum_n p_{\theta,\lambda}(z|x_n), \quad x_n \sim p_{\text{data}}(x).$$

There are two reasons to consider this distribution, rather than the marginal $p_{\theta,\lambda}(z)$ of the CEBM. The first is computational expedience; it is easier to approximate $\hat{p}_{\theta,\lambda}(z)$ than it is to approximate $p_{\theta,\lambda}(z)$, since the latter requires samples $x \sim p_{\theta,\lambda}(x)$ from the marginal of the CEBM. The second reason is that $\hat{p}_{\theta,\lambda}(z)$ reflects the distribution over features that we might use in a downstream task.

We approximate $\hat{p}_{\theta,\lambda}(z)$ with a Monte Carlo estimate over batches of size 1k (see [Esmacili et al. \(2019\)](#)), which we use to estimate both the KL and the mutual information (see Table 3). Because the marginal KL in CEBMs is significantly lower compared to VAEs across datasets, we conclude that

CEBMs indeed attempted to place the aggregate posterior distribution close to the inductive bias.

Our evaluation of the mutual information proved more surprising: CEBMs learn a representation that has a very low mutual information between x and z . The reason for this is that the posterior parameters $\tilde{\lambda}_\theta(x) = \lambda + t_\theta(x)$ are dominated by the parameters of the bias λ , which means that model essentially ignores the sufficient statistics $t_\theta(x)$, which tend to have a small magnitude relative to λ . This phenomenon could be interpreted as an instance of *posterior collapse* ([Aleml et al., 2017](#)), which has been observed in a variety of contexts when training variational autoencoders by maximizing the marginal likelihood, which in itself is not an objective that guarantees a high mutual information.

8. Discussion

In this paper, we introduced CEBMs, a class of latent-variable models that factorize into an energy-based distribution over data and a tractable posterior over latent variables. CEBMs can be trained using standard methods for EBMs and in this sense have a small “edit distance” relative to existing approaches, whilst also providing a mechanism for incorporating inductive biases for latent variables.

Our experimental results are encouraging but also raise questions. We observe a closer agreement between the unsupervised representation and class labels than in VAEs, which translates into improved performance in downstream classification tasks. At the same time, we observe that CEBMs do not learn a meaningful notion of uncertainty; the CEBM posterior is typically dominated by the inductive bias, which means that there is a very low mutual information between data and latent variables.

This work opens up a number of lines of future research. First and foremost, this work raises the question what objectives would be most suitable for learning energy-based latent-variable models in a manner maximizes agreement with respect to both the data distribution and the inductive bias terms, whilst also ensuring a sufficiently high mutual information between data and latent variables. More generally, we see opportunities to develop CEBMs with structured bias terms as an alternative to models based on VAEs in settings where we are hoping to reason about structured representations with little or no supervision.

Acknowledgements

We would like to thank our reviewers for their thoughtful comments, as well as Heiko Zimmermann, Will Grathwohl and Jacob Kelly for helpful discussions. This work was supported by the Intel Corporation, the 3M Corporation, NSF award 1835309, startup funds from Northeastern University, the Air Force Research Laboratory (AFRL), and DARPA.

References

- Alemi, A. A., Poole, B., Fischer, I., Dillon, J. V., Saurous, R. A., and Murphy, K. Fixing a broken elbow. *arXiv preprint arXiv:1711.00464*, 2017.
- Aneja, J., Schwing, A., Kautz, J., and Vahdat, A. Ncp-vae: Variational autoencoders with noise contrastive priors. *arXiv preprint arXiv:2010.02917*, 2020.
- Banerjee, A., Merugu, S., Dhillon, I. S., and Ghosh, J. Clustering with Bregman Divergences. *Journal of Machine Learning Research*, 6(58):1705–1749, 2005. ISSN 1533-7928.
- Bao, F., Li, C., Xu, T., Su, H., Zhu, J., and Zhang, B. Bi-level Score Matching for Learning Energy-based Latent Variable Models. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Che, T., Zhang, R., Sohl-Dickstein, J., Larochelle, H., Paull, L., Cao, Y., and Bengio, Y. Your GAN is Secretly an Energy-based Model and You Should Use Discriminator Driven Latent Sampling. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Chen, R. T., Li, X., Grosse, R. B., and Duvenaud, D. K. Isolating sources of disentanglement in variational autoencoders. In *Advances in neural information processing systems*, pp. 2610–2620, 2018.
- Crawford, E. and Pineau, J. Exploiting spatial invariance for scalable unsupervised object tracking. *arXiv preprint arXiv:1911.09033*, 2019a.
- Crawford, E. and Pineau, J. Spatially invariant unsupervised object detection with convolutional neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 3412–3420, 2019b.
- Donahue, J., Krähenbühl, P., and Darrell, T. Adversarial feature learning. *arXiv preprint arXiv:1605.09782*, 2016.
- Du, Y. and Mordatch, I. Implicit generation and generalization in energy-based models. *arXiv preprint arXiv:1903.08689*, 2019.
- Engelcke, M., Kosiorek, A. R., Jones, O. P., and Posner, I. Genesis: Generative scene inference and sampling with object-centric latent representations. *arXiv preprint arXiv:1907.13052*, 2019.
- Eslami, S. M. A., Heess, N., Weber, T., Tassa, Y., Szepesvari, D., Kavukcuoglu, K., and Hinton, G. E. Attend, infer, repeat: Fast scene understanding with generative models. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS’16*, pp. 3233–3241, Red Hook, NY, USA, December 2016. Curran Associates Inc. ISBN 978-1-5108-3881-9.
- Esmaeili, B., Wu, H., Jain, S., Bozkurt, A., Siddharth, N., Paige, B., Brooks, D. H., Dy, J., and Meent, J.-W. Structured disentangled representations. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 2525–2534. PMLR, 2019.
- Gao, R., Nijkamp, E., Kingma, D. P., Xu, Z., Dai, A. M., and Wu, Y. N. Flow contrastive estimation of energy-based models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7518–7528, 2020.
- Grathwohl, W., Wang, K.-C., Jacobsen, J.-H., Duvenaud, D., Norouzi, M., and Swersky, K. Your classifier is secretly an energy based model and you should treat it like one. *arXiv preprint arXiv:1912.03263*, 2019.
- Grathwohl, W. S., Kelly, J. J., Hashemi, M., Norouzi, M., Swersky, K., and Duvenaud, D. No {mcmc} for me: Amortized sampling for fast and stable training of energy-based models. In *International Conference on Learning Representations*, 2021.
- Gutmann, M. and Hyvärinen, A. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pp. 297–304. JMLR Workshop and Conference Proceedings, 2010.
- Han, T., Nijkamp, E., Fang, X., Hill, M., Zhu, S.-C., and Wu, Y. N. Divergence triangle for joint training of generator model, energy-based model, and inferential model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8670–8679, 2019.
- Han, T., Nijkamp, E., Zhou, L., Pang, B., Zhu, S.-C., and Wu, Y. N. Joint training of variational auto-encoder and latent energy-based model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. beta-vae: Learning basic visual concepts with a constrained variational framework. 2016.
- Hinton, G. E. Training products of experts by minimizing contrastive divergence. *Neural computation*, 14(8):1771–1800, 2002.
- Hinton, G. E., Osindero, S., and Teh, Y.-W. A Fast Learning Algorithm for Deep Belief Nets. *Neural Computation*, 18(7):1527–1554, May 2006. ISSN 0899-7667. doi: 10.1162/neco.2006.18.7.1527.
- Hyvärinen, A. and Dayan, P. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.

- Kim, H. and Mnih, A. Disentangling by factorising. In *International Conference on Machine Learning*, pp. 2649–2658, 2018.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *International Conference on Learning Representations*, 2013.
- Kosiorrek, A., Kim, H., Teh, Y. W., and Posner, I. Sequential attend, infer, repeat: Generative modelling of moving objects. In *Advances in Neural Information Processing Systems*, pp. 8606–8616, 2018.
- LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M., and Huang, F. A tutorial on energy-based learning. *Predicting structured data*, 1(0), 2006.
- Lin, Z., Wu, Y.-F., Peri, S., Fu, B., Jiang, J., and Ahn, S. Improving generative imagination in object-centric world models. *arXiv preprint arXiv:2010.02054*, 2020a.
- Lin, Z., Wu, Y.-F., Peri, S. V., Sun, W., Singh, G., Deng, F., Jiang, J., and Ahn, S. SPACE: Unsupervised Object-Oriented Scene Representation via Spatial Attention and Decomposition. *arXiv:2001.02407 [cs, eess, stat]*, March 2020b.
- Liu, H. and Abbeel, P. Hybrid discriminative-generative training via contrastive learning. *arXiv preprint arXiv:2007.09070*, 2020.
- McInnes, L., Healy, J., and Melville, J. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- Nijkamp, E., Hill, M., Han, T., Zhu, S.-C., and Nian Wu, Y. On the anatomy of mcmc-based maximum likelihood learning of energy-based models. *arXiv*, pp. arXiv–1903, 2019a.
- Nijkamp, E., Hill, M., Zhu, S.-C., and Wu, Y. N. Learning non-convergent non-persistent short-run mcmc toward energy-based model. In *Advances in Neural Information Processing Systems*, pp. 5232–5242, 2019b.
- Pang, B., Han, T., Nijkamp, E., Zhu, S.-C., and Wu, Y. N. Learning latent space energy-based prior model. *Advances in Neural Information Processing Systems*, 33, 2020.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In Xing, E. P. and Jebara, T. (eds.), *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pp. 1278–1286, Beijing, China, June 2014. PMLR.
- Rockafellar, R. T. *Convex analysis*, volume 36. Princeton university press, 1970.
- Salakhutdinov, R. and Hinton, G. Deep boltzmann machines. In *Artificial intelligence and statistics*, pp. 448–455, 2009.
- Smolensky, P. Information processing in dynamical systems: Foundations of harmony theory. Technical report, Colorado Univ at Boulder Dept of Computer Science, 1986.
- Song, Y. and Kingma, D. P. How to train your energy-based models. *arXiv preprint arXiv:2101.03288*, 2021.
- Song, Y., Garg, S., Shi, J., and Ermon, S. Sliced score matching: A scalable approach to density and score estimation. In *Uncertainty in Artificial Intelligence*, pp. 574–584. PMLR, 2020.
- Tieleman, T. Training restricted boltzmann machines using approximations to the likelihood gradient. In *Proceedings of the 25th international conference on Machine learning*, pp. 1064–1071, 2008.
- Tomczak, J. and Welling, M. Vae with a vampprior. In *International Conference on Artificial Intelligence and Statistics*, pp. 1214–1223, 2018.
- Vértes, E., Unit, U. G., and Sahani, M. Learning doubly intractable latent variable models via score matching, 2016.
- Vincent, P. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- Wainwright, M. J. and Jordan, M. I. Graphical Models, Exponential Families, and Variational Inference. *Foundations and Trends in Machine Learning*, 1(1–2):1–305, 2008. doi: 10/bpnwrm.
- Welling, M. and Teh, Y. W. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pp. 681–688, 2011.
- Welling, M., Rosen-zvi, M., and Hinton, G. E. Exponential family harmoniums with an application to information retrieval. In Saul, L., Weiss, Y., and Bottou, L. (eds.), *Advances in Neural Information Processing Systems*, volume 17, pp. 1481–1488. MIT Press, 2004.
- Wu, H., Zimmermann, H., Sennesh, E., Le, T. A., and van de Meent, J.-W. Amortized population gibbs samplers with neural sufficient statistics. In *Proceedings of the International Conference on Machine Learning*, pp. 10205–10215, 2020.

- Xie, J., Lu, Y., Zhu, S.-C., and Wu, Y. A theory of generative convnet. In *International Conference on Machine Learning*, pp. 2635–2644. PMLR, 2016.
- Xie, J., Lu, Y., Gao, R., and Wu, Y. N. Cooperative learning of energy-based model and latent variable model via mcmc teaching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- Xie, J., Lu, Y., Gao, R., Zhu, S.-C., and Wu, Y. N. Cooperative Training of Descriptor and Generator Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(1):27–45, January 2020. ISSN 1939-3539. doi: 10.1109/TPAMI.2018.2879081.
- Xie, J., Zheng, Z., Fang, X., Zhu, S.-C., and Wu, Y. N. Cooperative training of fast thinking initializer and slow thinking solver for conditional learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021a.
- Xie, J., Zheng, Z., and Li, P. Learning energy-based model with variational auto-encoder as amortized sampler. In *The Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI)*, volume 2, 2021b.