
Energy-Based Processes for Exchangeable Data

Mengjiao Yang^{*1} Bo Dai^{*1} Hanjun Dai¹ Dale Schuurmans^{1,2}

Abstract

Recently there has been growing interest in modeling sets with exchangeability such as point clouds. A shortcoming of current approaches is that they restrict the cardinality of the sets considered or can only express limited forms of distribution over unobserved data. To overcome these limitations, we introduce *Energy-Based Processes (EBPs)*, which extend energy based models to exchangeable data while allowing neural network parameterizations of the energy function. A key advantage of these models is the ability to express more flexible distributions over sets without restricting their cardinality. We develop an efficient training procedure for EBPs that demonstrates state-of-the-art performance on a variety of tasks such as point cloud generation, classification, denoising, and image completion¹.

1. Introduction

Many machine learning problems consider data where each instance is, itself, an unordered set of elements; *i.e.*, such that each observation is a *set*. Data of this kind arises in a variety of applications, ranging from document modeling (Blei et al., 2003; Garnelo et al., 2018a) and multi-task learning (Zaheer et al., 2017; Edwards & Storkey, 2016; Liu et al., 2019) to 3D point cloud modeling (Li et al., 2018; Yang et al., 2019). In unsupervised settings, a dataset typically consists of a set of such sets, while in supervised learning, it consists of a set of (set, label) pairs.

Modeling a distribution over a space of instances, where each instance is, itself, an unordered set of elements involves two key considerations: (1) the elements within a single instance are *exchangeable*, *i.e.*, the elements are order invariant; and (2) the cardinalities of the instances (sets)

vary, *i.e.*, instances need not exhibit the same cardinality. Modeling both *unconditional* and *conditional* distributions over instances (sets) are relevant to consider, since these support unsupervised and supervised tasks respectively.

For unconditional distribution modeling, there has been significant prior work on modeling set distributions, which has sought to balance competing needs to expand model flexibility and preserve tractability on the one hand, with respecting exchangeability and varying instance cardinalities on the other hand. However, managing these trade-offs has proved to be quite difficult, and current approaches remain limited in different respects.

For example, a particularly straightforward strategy for modeling distributions over instances $\mathbf{x} = \{x_1, \dots, x_n\}$ without assuming fixed cardinality is simply to use a recurrent neural network (RNNs) to encode instance probability autoregressively via $p(\mathbf{x}) = \prod_{i=1}^n p(x_i|x_{1:i-1})$ for a permutation of its elements. Such an approach allows the full flexibility of RNNs to be applied, and has been empirically successful (Larochelle & Murray, 2011; Bahdanau et al., 2015), but does not respect exchangeability nor is it clear how to tractably enforce exchangeability with RNNs.

To explicitly ensure exchangeability, a natural idea has been to exploit De Finetti’s theorem, which assures us that for any distribution over an *infinitely* exchangeable sequence, its finite projection distribution on arbitrary finite elements $\mathbf{x} = \{x_1, \dots, x_n\}$ can be decomposed as

$$p(\mathbf{x}) = \int \prod_{i=1}^n p(x_i|\theta) p(\theta) d\theta,^2 \quad (1)$$

for some latent variable θ . In other words, there always exists a latent variable θ such that conditioning on θ renders the instance elements $\{x_i\}_{i=1}^n$ *i.i.d.*. Latent variable models are therefore a natural choice for expressing an exchangeable distribution. Bayesian sets (Ghahramani & Heller, 2005), latent Dirichlet allocation (Blei et al., 2003), and related variants (Blei & Lafferty, 2007; Teh et al., 2006) are classical examples of this kind of approach, where the likelihood and prior in (1) are expressed by simple known distributions. Although the restriction to simple distributions severely limits the expressiveness of these models,

^{*}Equal contribution ¹Google Research, Brain Team ²University of Alberta. Correspondence to: Mengjiao Yang <sherry@google.com>, Bo Dai <bodai@google.com>.

Proceedings of the 37th International Conference on Machine Learning, Vienna, Austria, PMLR 119, 2020. Copyright 2020 by the author(s).

¹The code is available at <https://github.com/google-research/google-research/tree/master/ebp>.

²For simplicity, we consider the distributions with density function exist in this paper.

neural network parameterizations have recently been introduced (Edwards & Storkey, 2016; Korshunova et al., 2018; Yang et al., 2019). These approaches still exhibit limited expressiveness however: Edwards & Storkey (2016) restrict the model to *known* distributions parameterized by neural networks, while Korshunova et al. (2018); Yang et al. (2019) only consider normalizing flow models that require *invertible* neural networks.

If we consider *conditional* rather than *unconditional* distributions over sets, an extensive literature has considered stochastic process representations, which exploits their natural exchangeability and consistency properties. For example, Gaussian processes (\mathcal{GPs}) (Rasmussen & Williams, 2006) and extensions like Student- t processes (\mathcal{TPs}) (Shah et al., 2014), are well known models that, despite their scalability challenges, afford significant modeling flexibility via kernels. Unfortunately, they also restrict the conditional likelihoods to simple known distributions. Damianou & Lawrence (2012); Salimbeni & Deisenroth (2017) enrich the expressiveness of \mathcal{GPs} by stacking \mathcal{GP} -layers, but at the cost of increasing inference intractability with increasing depth. Neural processes (\mathcal{NPs}) (Garnelo et al., 2018b) and subsequent variants (Garnelo et al., 2018a; Kim et al., 2019) attempt to construct neural network to mimic \mathcal{GPs} , but these too rely on *known* distributions for the conditional likelihood, which inherently limits expressiveness.

In this paper, we propose *Energy-Based Processes* (\mathcal{EBPs}), and their extension to *unconditional* distributions, to increase the flexibility of set distribution modeling while retaining *exchangeability* and *varying-cardinality*. After establishing necessary background on energy-based models (EBMs) and stochastic processes in Section 2, we provide a new stochastic process representation theorem in Section 3. This result allows us to then generalize EBMs to *Energy-Based Processes* (\mathcal{EBPs}), which provably obtain the exchangeability and varying-cardinality properties. Interestingly, the stochastic process representation we introduce also covers classical stochastic processes as special cases. We further extend \mathcal{EBP} to the unconditional setting, unifying the previously separate stochastic process and latent variable model perspectives in a common framework. To address the challenge of training \mathcal{EBPs} , we introduce an efficient new *Neural Collapsed Inference* (NCI) in Section 4. Finally, we evaluate the effectiveness of \mathcal{EBPs} with NCI training on a set of *supervised* (e.g., 1D regression and image completion) and *unsupervised* tasks (e.g., point-cloud feature extraction, generation and denoising), demonstrating state-of-the-art performance across a range of scenarios.

2. Background

We provide a brief introduction to energy-based models and stochastic processes, which provide the essential building

blocks for our subsequent development.

2.1. Energy-Based Models

Energy-based models are attractive due to their flexibility (LeCun et al., 2006; Wu et al., 2018) and appealing statistical properties (Brown, 1986). In particular, an EBM over $\Omega \subset \mathbb{R}^d$ with fixed dimension d is defined as

$$p_f(x) = \exp(f(x) - \log Z(f)) \quad (2)$$

for $x \in \Omega$, where $f(x) : \Omega \rightarrow \mathbb{R}$ is the energy function and $Z(f) := \int_{\Omega} \exp(f(x)) dx$ is the partition function. We let $\mathcal{F} := \{f(\cdot) : Z(f) < \infty\}$.

The flexibility of EBMs is well known. For example, classical exponential family distributions can be recovered from (2) by instantiating specific forms for Ω and $f(\cdot)$. Introducing additional structure to the energy function allows both Markov random fields (Kendall & Snell, 1980) and conditional random fields (Lafferty et al., 2001) to be recovered from (2). More recently, the introduction of deep neural energy functions (Xie et al., 2016; Du & Mordatch, 2019; Dai et al., 2019), has led to many successful applications of EBMs to modeling complex distributions in practice.

Although maximum likelihood estimation (MLE) of general EBMs is notoriously difficult, recent techniques such as adversarial dynamics embedding (ADE) appear able to practically train a broader class of such models (Dai et al., 2019). In particular, ADE approximates MLE for EBMs by formulating a saddle-point version of the problem:

$$\max_f \min_{q(x,v) \in \mathcal{P}} \widehat{\mathbb{E}}[f(x)] - H(q(x,v)) - \mathbb{E}_{q(x,v)} \left[f(x) - \frac{\lambda}{2} v^\top v \right], \quad (3)$$

where $p(x, v)$ is parametrized via a learnable Hamiltonian/Langevin sampler. Since we make use of some of the techniques in our main development, we provide some further details of ADE in Appendix A.

Although these recent advances are promising, EBMs remain fundamentally limited for our purposes, in that they are only defined for *fixed-dimensional* data. The question of extending such models to express distributions over *exchangeable* data with *arbitrary cardinality* has not yet been well explored.

2.2. Stochastic Processes

Stochastic processes are usually defined in terms of their finite-dimensional marginal distributions. In particular, consider a stochastic process given by a collection of random variables $\{X_t; t \in \mathcal{T}\}$ indexed by t , where the marginal distribution for any finite set of indices $\{t_1, \dots, t_n\}$ in \mathcal{T} (without order) is specified *i.e.*,

$p(x_{t_1:t_n}) := p(x_{t_1}, \dots, x_{t_n} | \{t_i\}_{i=1}^n)$. For example, Gaussian processes (\mathcal{GP} s) are defined in this way using Gaussians for the marginal distributions (Rasmussen & Williams, 2006), while Student- t processes (\mathcal{TP} s) are similarly defined using multivariate Student- t distributions for the marginals (Shah et al., 2014).

The Kolmogorov extension theorem (Øksendal, 2003) provides the sufficient conditions for designing a valid stochastic processes, namely:

- **Exchangeability** The marginal distribution for any finite set of random variables is *invariant to permutation* order. Formally, for all n and all permutations π , this means

$$p(x_{t_1}, \dots, x_{t_n} | \{t_i\}_{i=1}^n) = p(\pi(x_{t_1:t_n}) | \pi(\{t_i\}_{i=1}^n)),$$

where $p(\pi(x_{t_1:t_n})) := p(x_{\pi(t_1)}, \dots, x_{\pi(t_n)})$.

- **Consistency** The partial marginal distribution, obtained by marginalizing additional variables in the finite sequence, is the same as the one obtained from the original infinite sequence. Formally, if $n \geq m \geq 1$, this means

$$p(x_{t_1:t_m} | \{t_i\}_{i=1}^m) = \int p(x_{t_1:t_n} | \{t_i\}_{i=1}^n) dx_{t_{m+1}:t_n}.$$

Obviously, these conditions also justify stochastic processes as a valid tool for modeling exchangeable data. However, existing classical models, such as \mathcal{GP} s and \mathcal{TP} s, restrict the marginal distributions to simple forms while requiring huge memory and computational cost, which prevents convenient application to large-scale complex data.

3. Energy-Based Processes

We now develop our main modeling approach, which combines a stochastic process representation of exchangeable data with energy-based models. The result is a generalization of Gaussian processes and Student- t processes that exploits EBMs for greater flexibility. We follow this development with an extension to unconditional modeling.

3.1. Representation of Stochastic Processes

Although finite marginal distributions provide a way to parametrize stochastic processes, it is not obvious how to use flexible EBMs to represent marginals while still maintaining exchangeability and consistency. Therefore, instead of such a direct parametrization, we exploit the deeper structure of a stochastic process, based on the following representation theorem.

Theorem 1 *For any stochastic process $(x_{t_1}, x_{t_2}, \dots) \sim \mathcal{SP}$ that can be constructed via Kolmogorov extension theorem, the process can be equivalently represented by a latent variable model*

$$\theta \sim P(\theta), \quad x_{t_i} \sim p(x | \theta, t_i), \quad \forall i \in \{1, \dots, n\} \quad \forall n, \quad (4)$$

where θ can be finite or infinite dimensional and P denotes some measure on θ .

Notice that θ can be either finite or infinite dimensional, and then, $P(d\theta)$ can be either the distribution or stochastic process for the finite or infinite dimensional random variable θ , respectively. Then, Theorem 1 is a straightforward corollary of De Finetti's Theorem.

Proof Since the process \mathcal{SP} is constructed via the Kolmogorov Extension Theorem, it must satisfy exchangeability and consistency, *i.e.*, the sequence $\{x_{t(1)}, \dots, x_{t(n)}\}$ is exchangeable $\forall n$ and following a projective family. This implies, by De Finetti's Theorem, that any marginal distribution can be represented as a mixture of *i.i.d.* processes:

$$p(x_{t_1:t_n} | \{t_i\}_{i=1}^n) = \int \prod_{i=1}^n p(x | \theta, t_i) P(d\theta), \quad (5)$$

which achieves the conclusion. \blacksquare

For simplicity, we assume the density on θ exists as $p(\theta)$. Given such a representation of a stochastic processes, it is now easy to see how to generalize Gaussian, Student- t , and other processes with EBMs.

3.2. \mathcal{EBP} Construction

To enhance the flexibility of a stochastic process representation of exchangeable data, we use EBMs to model the likelihood term in (4), by letting

$$p_w(x | \theta, t) = \frac{\exp(f_w(x, t; \theta))}{Z(f_w, t; \theta)}, \quad (6)$$

where $Z(f_w, t; \theta) = \int \exp(f_w(x, t; \theta)) dx$ and we let w denote the parameters of f , which can be learned. Substituting this into the latent variable representation of stochastic processes (4), leads to the definition of energy-based processes on arbitrary finite marginals as

$$p_w(x_{t_1:t_n} | \{t_i\}_{i=1}^n) = \int \frac{\exp(\sum_{i=1}^n (f_w(x_{t_i}, t_i; \theta)))}{Z^n(f_w, t)} p(\theta) d\theta, \quad (7)$$

given a prior $p(\theta)$ on the finite or infinite latent variable θ . We refer to the resulting process as an energy-based process (\mathcal{EBP}).

Compared to using restricted distributions, such as Gaussian or Student- t , the use of an EBM in an \mathcal{EBP} allows much more flexible energy models f_w , for example in the form of a deep neural network, to represent the complex dependency between x and t . To rigorously verify that the outcome is strictly more general than standard processes, observe that classical process models can be recovered exactly simply by instantiating (7) with specific choices of $f_w(x, t; \theta)$ and $p(\theta)$.

- **Gaussian Processes** Consider the weight-space view of \mathcal{GP} s (Rasmussen & Williams, 2006), which allows the \mathcal{GP} for regression to be re-written as

$$\theta \sim \mathcal{N}(\mathbf{0}, I_d), \quad (8)$$

$$f_w(x, t; \theta) = -\frac{1}{2\sigma^2} \|x - \theta^\top \phi(t)\|^2, \quad (9)$$

where $w = \{\sigma, \phi(\cdot)\}$, with $\phi(\cdot)$ denoting feature mappings that can be finite or infinite dimensional. If we now let $k(t, t') = \phi(t)^\top \phi(t')$ denote the kernel function and $K(t_{1:n}) = [k(t_i, t_j)]_{i,j}^n$, the marginalized distribution can be recovered as

$$p(x_{t_1:t_n} | \{t_i\}_{i=1}^n) = \mathcal{N}(0, K(t_{1:n}) + \sigma^2 I_n),$$

which shows that $X_t \sim \mathcal{GP}(0, K(t_{1:n}) + \sigma^2 I_n)$; see Appendix B.1.

- **Student- t Processes** Denote $\theta = (\alpha, \beta)$ and consider

$$\alpha \sim \mathcal{N}(\mathbf{0}, I_d), \quad \beta^{-1} \sim \Gamma\left(\frac{\nu}{2}, \frac{\gamma}{2}\right), \quad (10)$$

$$f_w(x, t; \theta) = -\frac{\gamma \left\| x - \sqrt{\frac{\beta(\nu-2)}{\gamma}} \alpha^\top \phi(t) \right\|^2}{2\sigma^2(\nu-2)\beta}, \quad (11)$$

where $w = \{\nu, \gamma, \sigma, \phi(\cdot)\}$ with $\nu > 0$ and $\gamma > 0$. These substitutions lead to the marginal distribution

$$p(x_{t_1:t_n} | \{t_i\}_{i=1}^n) = \mathcal{T}(\nu, 0, K(t_{1:n}) + \sigma^2 I_n),$$

which shows that $X_t \sim \mathcal{TP}(\nu, 0, K(t_{1:n}) + \sigma^2 I_n)$; see Appendix B.2.

- **Neural Processes** Neural processes (\mathcal{NP} s) are explicitly defined by a latent variable model in (Garnelo et al., 2018b):³

$$p(x_{t_1:t_n} | \{t_i\}_{i=1}^n) = \int \prod_{i=1}^n \mathcal{N}(x | h_w(t_i; \theta)) p(\theta) d\theta,$$

where $h_w(\cdot; \theta)$ is a neural network. Clearly, \mathcal{NP} s share similarity to \mathcal{EBP} s in that both processes use deep neural networks to enhance modeling flexibility. However, there remain critical differences. In fact, the likelihood function $p(x|t, \theta)$ in \mathcal{NP} s is still restricted to *known* simple distributions, with parameterization given by a neural network. By contrast, \mathcal{EBP} s directly use EBMs with deep neural energy functions to model the likelihood. In this sense, \mathcal{EBP} s are a strict generalization of \mathcal{NP} s: if one fixes the last layer of f_w in \mathcal{EBP} s to be a simple function, such as quadratic, then an \mathcal{EBP} reduces to a \mathcal{NP} .

³The conditional neural processes (Garnelo et al., 2018a) only defines the predictive distribution, hence it is not a proper stochastic processes, as discussed in their paper.

Figure 1 demonstrates the comparison between these process models and an additional variational implicit process (\mathcal{VIP}) model (see Appendix E) in a simple regression setting, highlighting the flexibility of \mathcal{EBP} s in modeling the conditional likelihood.




Figure 1. The ground truth data and learned energy functions of \mathcal{GP} , \mathcal{NP} , \mathcal{VIP} , and \mathcal{EBP} (from left to right). \mathcal{EBP} successfully captures multi-modality of the toy data as \mathcal{GP} and \mathcal{NP} exhibiting only a single mode; see Section 5 for details.

3.3. Unconditional \mathcal{EBP} s Extension

Stochastic processes, such as \mathcal{EBP} s, express the *conditional* distribution over $\{X_t\}$ *conditioned on* an index variable t , which makes this approach naturally applicable to supervised learning tasks on exchangeable data. However, we would also like to tackle unsupervised learning problems given exchangeable observations, so an *unconditional* formulation of the \mathcal{EBP} is required.

To develop an unconditional \mathcal{EBP} , we start with the distribution of an arbitrary finite marginal, $p(x_{t_1:t_n} | \{t_i\}_{i=1}^n)$. Note that when the indices $\{t_i\}_{i=1}^n$ are not observed, we can simply marginalize them out to obtain

$$\begin{aligned} p_w(x_{1:n}) &:= \int p_w(x_{t_1:t_n} | \{t_i\}_{i=1}^n) p(\{t_i\}_{i=1}^n) dt_{1:n} \quad (12) \\ &= \int p_w(x_{t_1:t_n} | \{t_i\}_{i=1}^n, \theta) p(\theta) p(\{t_i\}_{i=1}^n) d\theta dt_{1:n}. \end{aligned}$$

Here we can introduce parameters to the $p(\{t_i\}_{i=1}^n)$, which can also be learned. It can be verified the resulting distribution $p_w(x_{1:n})$ is provably *exchangeable* and *consistent* under mild conditions.

Theorem 2 *If $n \geq m \geq 1$, and the prior is exchangeable and consistent, then the marginal distribution $p(x_{1:n})$ will be exchangeable and consistent.*

The proof can be found in Appendix C.

We refer to this result as the unconditional \mathcal{EBP} . This understanding allows connections to be established with some existing models.

- **\mathcal{GP} -Latent Variable Model** The \mathcal{GP} -latent variable model (GPLVM) (Lawrence, 2004) considers the estimation of the latent index variables by maximizing the

log-marginal likelihood of \mathcal{GP} , *i.e.*,

$$\begin{aligned} \max_{\{t_i\}_{i=1}^n} \log p(x_{t_1:t_n} | \{t_i\}_{i=1}^n) \\ = \log \mathcal{N}(\mathbf{0}, K(t_{1:n}) + \sigma^2 I_n). \end{aligned} \quad (13)$$

This can be understood as using a point estimator with \mathcal{GP} s and an improper uniform prior $p(\{t_i\}_{i=1}^n)$ in (12).

- **Bayesian Recurrent Neural Model** Korshunova et al. (2018) propose a model BRUNO for modeling exchangeable data. This model actually uses degenerate kernels to eliminate $\{t_i\}_{i=1}^n$ in (12). In particular, BRUNO defines a \mathcal{TP} for each latent variable dimension, with the same constant feature mapping $\phi(t) = 1, \forall t$. That is, for the d -th dimension in $x, \forall d \in \{0, \dots, D\}$,

$$p(x_{t_1:t_n}^d | \{t_i\}_{i=1}^n) = p(x_{1:n}^d) \sim \mathcal{T}(\nu_d, \mu_d, K_d), \quad (14)$$

since the kernel is $K_d(t_{1:n}) = \mathbf{1}\mathbf{1}^\top + (\sigma_d)^2 I_n$. The observations are then transformed via an invertible function, *i.e.*, $x' = \psi(x)$ with $\det\left(\frac{\partial \psi(x)}{\partial x}\right)$ invertible.

- **Neural Statistician** Edwards & Storkey (2016) essentially generalize latent Dirichlet allocation (Blei et al., 2003) with neural networks. The model follows (12) with a sophisticated hierarchical prior. However, by comparison with \mathcal{EBPs} , the likelihood function used in neural statistician is still restricted in known simple distributions. Meanwhile, it follows vanilla amortized inference. We will show how \mathcal{EBPs} can work with a more efficient inference scheme in the next section.

We provide more instantiations in Appendix B and the related work in Appendix E.

4. Neural Collapsed Inference for Deep \mathcal{EBPs}

By incorporating EBMs in the latent variable representation of a stochastic process, we obtain a family of flexible models that can capture complex structure in exchangeable data for both *conditional* and *unconditional* distributions. We can exploit deep neural networks in parametrizing the energy function as in Xie et al. (2016); Du & Mordatch (2019); Dai et al. (2019), leading to *deep \mathcal{EBPs}* . However, this raises notorious difficulties in inference and learning as a consequence of flexibility. Therefore, we develop an efficient *Neural Collapsed Inference* (NCI) method for unconditional deep \mathcal{EBPs} . (For the inference and learning of conditional \mathcal{EBPs} , please refer to Appendix D.2.)

4.1. Neural Collapsed Reparametrization

We first carefully analyze the difficulties in inference and learning through the empirical log-marginal distribution of the general \mathcal{EBPs} on given samples $\mathcal{D} = \{x_{1:n}^i\}_{i=1}^N$:

$$\max_w \widehat{\mathbb{E}}_{\mathcal{D}} [\log p_w(x_{1:n})], \quad (15)$$

where $p_w(x_{1:n})$ is defined in (12).

There are several integrations that are not tractable in (15) given a general neural network parameterized $f_w(x, t; \theta)$:

1. The partition function $Z(f_w, t, \theta) = \int \exp(f(x, t; \theta)) dx$ is intractable in $p(x_{t_1:t_n} | \{t_i\}_{i=1}^n, \theta)$;
2. The integration over θ will be intractable for $p(x_{t_1:t_n} | \{t_i\}_{i=1}^n)$;
3. The integration over $\{t_i\}_{i=1}^n$ will be intractable for $p(x_{1:n})$.

One can of course use vanilla amortized inference with the neural network reparameterization trick (Kingma & Welling, 2013; Rezende et al., 2014) for each intractable component, as in (Edwards & Storkey, 2016), but this leads to an optimization over the approximate posteriors $q(x|t, \theta)$ and $q(\theta, \{t_i\}_{i=1}^n)$. The latter distribution requires a complex neural network architecture to capture the dependence in $\{t_i\}_{i=1}^n$, which is usually a significant challenge. Meanwhile, in most unsupervised learning tasks, such as point cloud generation and denoising, one is only interested in $x_{1:n}$, while $\{t_i\}_{i=1}^n$ is not directly used. Since inference over $\{t_i\}_{i=1}^n$ is only an intermediate step, we develop the following *Neural Collapsed Inference* strategy (NCI).

Collapsed inference and sampling strategies have previously been proposed for removing nuisance latent variables that can be tractably eliminated, to reduce computational cost and accelerate inference (Teh et al., 2007; Porteous et al., 2008). Due to the intractability of

$$p_w(x_{1:n}|\theta) = \int p_w(x_{t_1:t_n}|\theta, \{t_i\}_{i=1}^n) p(\{t_i\}_{i=1}^n) dt_{1:n},$$

standard collapsed inference cannot be applied. However, since deep EBMs are very flexible, $p_{w'}(x_{1:n}|\theta)$ can be directly reparameterized with *another* EBM:

$$p_{w'}(x_{1:n}|\theta) \propto \exp(f_{w'}(x_{1:n}; \theta)). \quad (16)$$

Concretely, assume $p(\{t_i\}_{i=1}^n) \propto \exp(\sum_{i=1}^n h_v(t_i))$, so we have

$$\begin{aligned} p_{w'}(x_{1:n}|\theta) &= \prod_{i=1}^n \int p_w(x_{t_i}|\theta, t_i) p(t_i) dt_i \\ &\propto \prod_{i=1}^n \int \exp(f_w(x_{t_i}, t_i; \theta) - Z(f_w, t_i; \theta) + h_v(t_i)) dt_i \\ &\approx \prod_{i=1}^n \frac{1}{Z(f_{w'}; \theta)} \exp(f_{w'}(x_i; \theta)), \end{aligned}$$

where the last step follows because the result of the integration in the second step is a distribution $p(x)$ over Ω , and we are using another learnable EBM to approximate this

distribution. Therefore, we consider the collapsed model:

$$p_{w'}(x_{1:n}|\theta) \propto \exp\left(\sum_{i=1}^n f_{w'}(x_i; \theta)\right), \quad (17)$$

which still satisfies exchangeability and consistency. In fact, with the *i.i.d.* prior on $\{t_i\}_{i=1}^n$, we will obtain a latent variable model based on De Finetti's theorem. With such an approximate collapsed model, the log-marginal distribution can be used as a surrogate:

$$\ell(w) := \log p_{w'}(x_{1:n}) = \log \int p_{w'}(x_{1:n}|\theta) p(\theta) d\theta. \quad (18)$$

We refer to the variational inference in such a task-oriented neural reparameterization model as *Neural Collapsed Inference*, which reduces the computational cost and memory of inferring the posterior compared to using vanilla variational amortized inference.

We can further use the neural collapsing trick for θ ; which will reduce the model to Gibbs point processes (\mathcal{GPP} s) (Dereudre, 2019) and Determinantal point processes (\mathcal{DPP} s) (Lavancier et al., 2015; Kulesza et al., 2012). Therefore, the proposed algorithm can straightforwardly applied for deep \mathcal{GPP} and \mathcal{DPP} estimation. It should be emphasized that by exploiting the proposed primal-dual MLE framework, we automatically obtain a deep neural network parametrized dual sampler with the learned model simultaneously, which can be used in inference and bypass the notorious sampling difficulty in \mathcal{GPP} and \mathcal{DPP} . Please see Appendix D.1 for detailed discussion.

4.2. Amortized Inference

As discussed, $\{t_i\}_{i=1}^n$ can be eliminated by neural collapsed reparameterization. We now discuss variational techniques for integrating over θ and x respectively in the partition function of (18)

ELBO for integration on θ We apply vanilla ELBO to handle the intractability of integration over θ . Specifically, since

$$\begin{aligned} & \log \int p_{w'}(x_{1:n}|\theta) p(\theta) d\theta \\ &= \max_{q(\theta|x_{1:n}) \in \mathcal{P}} \mathbb{E}_{q(\theta|x_{1:n})} [\log p_{w'}(x_{1:n}|\theta)] - KL(q||p), \end{aligned} \quad (19)$$

we can apply the standard reparameterization trick (Kingma & Welling, 2013; Rezende et al., 2014) for $q(\theta|x_{1:n})$.

Primal-Dual form for partition function For the term $\log p_{w'}(x_{1:n}|\theta)$ in (19), which is

$$\log p_{w'}(x_{1:n}|\theta) = f_{w'}(x_{1:n}; \theta) - \log Z(f_{w'}, \theta),$$

Algorithm 1 Neural Collapsed Inference

- 1: Initialize W_1 randomly, set length of steps T .
- 2: **for** iteration $k = 1, \dots, K$ **do**
- 3: Sample mini-batch $\{x_{1:n_j}^j\}_{j=1}^b$ from dataset \mathcal{D} .
- 4: Sample $\theta^j \sim q_\alpha(\theta|x_{1:n})$, $\forall j = 1, \dots, b$.
- 5: Sample $\tilde{x}_{1:n}^j, \tilde{v}^j \sim q_\beta(x_{1:n}, v|\theta)$, $\forall j = 1, \dots, b$.
- 6: $\{\beta_{k+1}\} = \beta_k - \gamma_k \nabla_\beta L(\alpha_k, \beta_k, w'_k)$
- 7: $\{\alpha, w'\}_{k+1} = \{\alpha, w'\}_k + \gamma_k \nabla_{\{\alpha, w'\}} L(\alpha_k, \beta_k; w'_k)$.
- 8: **end for**

we apply an adversarial dynamics embedding technique (Dai et al., 2019) for the log $Z(f_{w'}, \theta)$ as introduced in Section 2. This leads to an equivalent optimization of the form

$$\begin{aligned} & \log p_{w'}(x_{1:n}|\theta) \\ & \propto \min_{q(x_{1:n}, v|\theta) \in \mathcal{P}} f_{w'}(x_{1:n}; \theta) - H(q(x_{1:n}, v|\theta)) \\ & \quad - \mathbb{E}_{q(x_{1:n}, v|\theta)} \left[f_{w'}(x_{1:n}; \theta) - \frac{\lambda}{2} v^\top v \right]. \end{aligned} \quad (20)$$

By combining (19) and (20) into (18), we obtain

$$\max_{w', q(\theta|x_{1:n})} \min_{q(x_{1:n}, v|\theta)} L(q(\theta|x_{1:n}), q(x_{1:n}, v|\theta); w'), \quad (21)$$

where

$$\begin{aligned} & L(q(\theta|x_{1:n}), q(x_{1:n}, v|\theta); w') \\ & := \widehat{\mathbb{E}}_{x_{1:n}} \mathbb{E}_\theta [f_{w'}(x_{1:n}; \theta)] - \widehat{\mathbb{E}}_{x_{1:n}} [KL(q(\theta|x_{1:n})||p(\theta))] \\ & \quad - \widehat{\mathbb{E}}_{x_{1:n}} \mathbb{E}_\theta \left[\mathbb{E}_{x_{1:n}, v|\theta} \left[f_{w'}(x_{1:n}; \theta) - \frac{\lambda}{2} v^\top v \right] \right] \\ & \quad + \widehat{\mathbb{E}}_{x_{1:n}} \mathbb{E}_\theta [H(q(x_{1:n}, v|\theta))]. \end{aligned} \quad (22)$$

where $\widehat{\mathbb{E}}_{x_{1:n}}[\cdot]$, $\mathbb{E}_\theta[\cdot]$ and $\mathbb{E}_{x_{1:n}, v}[\cdot]$ denote the expectation w.r.t. empirical samples, $q(\theta|x_{1:n})$ and $q(x_{1:n}, v|\theta)$, respectively.

Parametrization Finally, we describe some concrete parameterizations for $f_w(x; \theta)$, $q(\theta|x_{1:n})$ and $q(x_{1:n}, v|\theta)$.

The energy function $f_w(x; \theta)$ is parametrized as a MLP that takes input x_{t_i} concatenated with θ . We use the same energy function parameterization for both conditional and unconditional \mathcal{EBPs} .

For $q(\theta|x_{1:n})$ we use a simple Gaussian with mean function parameterized via deepsets (Zaheer et al., 2017):

$$\theta = \text{mlp}_\alpha(x_{1:n}) + \sigma\xi, \quad \xi \sim \mathcal{N}(0, I_d), \quad (23)$$

where $\text{mlp}_\alpha(x_{1:n}) := \sum_{i=1}^n \phi(x_i)$ and α denoting the parameters in $\phi(\cdot)$.

For $q(x_{1:n}, v|\theta)$ we consider dynamics embedding with an RNN or flow-model as the initial distribution; see Appendix A.1 for parameterization and Appendix F for implementation details. We denote the parameters in $q(x_{1:n}, v|\theta)$ as β . We also denote the objective in (21) as $L(\alpha, \beta; w')$. Then, we can use stochastic gradient descent for (21) to optimize $W = (\alpha, \beta, w')$, as illustrated in Algorithm 1.

5. Applications

We test conditional \mathcal{EBPs} on two supervised learning tasks: 1D regression and image completion, and unconditional \mathcal{EBPs} on three unsupervised tasks: point cloud generation, representation learning, and denoising. Details of each experiment can be found in Appendix F.

5.1. Supervised Tasks

1D regression. In order to show that \mathcal{EBPs} are more flexible than \mathcal{GP} s, \mathcal{NP} s and \mathcal{VLP} s in modeling complex distributions, we construct a two-mode synthetic dataset of *i.i.d.* points whose means form two sine waves with a phase offset. In this setting, t_i corresponds to the horizontal axis of the sine wave and x_{t_i} corresponds to the values on the vertical axis. At every training step, we randomly select a subset of the points as observations and estimate the marginal distribution of the observed and unobserved points similar to Garnelo et al. (2018b). We visualize the ground truth and learned energy functions of \mathcal{GP} , \mathcal{NP} , \mathcal{VLP} and \mathcal{EBP} in Figure 1. Clearly, \mathcal{EBP} succeeds as \mathcal{GP} and \mathcal{NP} fail to capture the multi-modality of the underlying data distribution. More comparisons can be found in Appendix G.1.

Image completion. An image can be represented as a set of n pixels $\{(t_i, x_{t_i})\}_{i=1}^n$, where $t_i \in \mathbb{R}^2$ corresponds to the Cartesian coordinates of each pixel and x_{t_i} corresponds to the channel-wise intensity of that pixel ($x_{t_i} \in \mathbb{R}$ for grayscale images and $x_{t_i} \in \mathbb{R}^3$ for RGB images). Conditional \mathcal{EBPs} perform image completion by maximizing $p(x_{t_i:t_n} | \{t_i\}_{i=1}^n)$.

We separately train two conditional \mathcal{EBPs} on the MNIST (LeCun, 1998) and the CelebA dataset (Liu et al., 2015). Examples of completion results are shown in Figure 2 and Figure 3. When a random or consecutive subset of pixels are observed, our method discovers different data modes and generates different MNIST digits, as shown in Figure 2. When a varying number of pixels are observed as in Figure 3, completion with fewer observed pixels (column 2) can lead to a face that is much different from the original face than completion with more observed pixels (column 5), revealing high variance when the number of observations is small (similar to \mathcal{GP} s). More examples of image completion can be found in Appendix G.




Figure 2. Image completion on MNIST. The first row shows the unobserved pixels in gray and observed pixels in black and white. The second and third rows are two different generated samples given the observed pixels from the first row. Generations are based on randomly selected pixels or the top half of an image.



Figure 3. Image completion on CelebA. The first row shows the unobserved pixels in black with an increasing number of observed pixels from left to right (column 1-5). The second row shows the completed image given the observed pixels from the first row.

5.2. Unsupervised Tasks on Point Clouds

Next, we apply unconditional \mathcal{EBPs} to a set of unsupervised learning tasks for point clouds. A point cloud represents a 3D object as the Cartesian coordinates of the set of exchangeable points $\{x_i\}_{i=1}^n \subset \mathbb{R}^3$, where n is the number of points in a point cloud and can therefore be arbitrarily large. Since the point cloud data does not depend on index t_i , they are modeled by unconditional \mathcal{EBPs} which integrate over $\{t_i\}_{i=1}^n$, leading to the unconditional objective $p(x_{1:n}) = \int p(x_{t_1:t_n} | \{t_i\}_{i=1}^n) p(\{t_i\}_{i=1}^n) d\{t_i\}_{i=1}^n$ as first introduced in (12).

Point cloud related work. Earlier work on point cloud generation and representation learning simply treats point clouds as matrices with a fixed dimension (Achlioptas et al., 2017; Gadelha et al., 2018; Zamorski et al., 2018; Sun et al., 2018), leading to suboptimal parameterizations as permutation invariance and arbitrary cardinality of exchangeable data are violated by this representation. Some of the more recent work tries to overcome the cardinality constraint by trading off flexibility of the model. For instance, Yang et al. (2019) uses normalizing flow to transform an arbitrary number of points sampled from the initial distribution, but requires the transformations to be invertible. Yang et al. (2018), as another example, transforms 2D distributions to 3D targets, but assumes that the topology of the generated shape is genus-zero or of a disk topology. Li et al. (2018) demonstrate the straightforward extension of GAN is not




Figure 4. Example point clouds of airplane, chair, and car generated from the learned model.




Figure 5. Energy distributions of the generated samples (fake) and training data (real). x-axis is the energy value and y axis is the count of examples. The energy distributions of the generated and real point clouds show significant overlap.

valid for exchangeable data, and then, provide some ad-hoc strategies to make up such deficiency. Their generator in the proposed PC-GAN is *conditional* on observations, which restricts the usages of the model. Among all generative models for point clouds considered here, \mathcal{EBP} s are the most flexible in handling permutation-invariant data with arbitrary cardinality.

Point cloud generation. We train one unconditional \mathcal{EBP} per category on airplane, chair, and car from the ShapeNet dataset (Wu et al., 2015). Figure 4 shows the accumulative output of the model (see more generated examples in Appendix G). We plot the energy distributions of all real and generated samples for each object category in Figure 5. \mathcal{EBP} s have successfully learned the desired distributions as the energies of real and generated point clouds show significant overlap.

We compare the generation quality of \mathcal{EBP} s with the previous state-of-the-art generative models for point clouds including I-GAN (Achlioptas et al., 2017), PC-GAN (Li et al., 2018), and PointFlow (Yang et al., 2019). Following these prior work, we uniformly sample 2048 points per point

Table 1. Generation results. \uparrow : the higher the better. \downarrow : the lower the better. The best scores are highlighted in bold. JSD is scaled by 10^2 , MMD-CD by 10^3 , and MMD-EMD by 10^2 . Each number for I-GAN is from the model trained using either CD or EMD loss, whichever one is better.

Category	Model	JSD (\downarrow)	MMD (\downarrow)		COV (% \uparrow)	
			CD	EMD	CD	EMD
Airplane	I-GAN	3.61	0.239	3.29	47.90	50.62
	PC-GAN	4.63	0.287	3.57	36.46	40.94
	PointFlow	4.92	0.217	3.24	46.91	48.40
	EBP (ours)	3.92	0.240	3.22	49.38	51.60
Chair	I-GAN	2.27	2.46	7.85	41.39	41.69
	PC-GAN	3.90	2.75	8.20	36.50	38.98
	PointFlow	1.74	2.42	7.87	46.83	46.98
	EBP (ours)	1.53	2.59	7.92	47.73	49.84
Car	I-GAN	2.21	1.48	5.43	39.20	39.77
	PC-GAN	5.85	1.12	5.83	23.56	30.29
	PointFlow	0.87	0.91	5.22	44.03	46.59
	EBP (ours)	0.78	0.95	5.24	51.99	51.70

cloud from the mesh surface of ShapeNet, use both Chamfer distance (CD) and earth mover’s distance (EMD) to measure similarity between point clouds, and use Jensen-Shannon Divergence (JSD), Minimum matching distance (MMD), and Coverage (COV) as evaluation metrics. Table 1 shows that \mathcal{EBP} achieves the best COV for all three categories under both CD and EMD, demonstrating \mathcal{EBP} s advantage in expressing complex distributions and avoiding mode collapse. \mathcal{EBP} also achieves the lowest JSD for two out of three categories. More examples of point cloud generation can be found in Appendix G.

Unsupervised representation learning. Next, we evaluate the representation learning ability of \mathcal{EBP} s. Following the convention of previous work, we first train one \mathcal{EBP} on all 55 object categories of ShapeNet. We then extract the Deep Sets output (θ in our model) for each point cloud in ModelNet40 (Wu et al., 2015) using the pre-trained model, and train a linear SVM using the extracted features. Table 2 shows that our method achieves the second highest classification accuracy among the seven state-of-the-art unsupervised representation learning methods, and is only 0.1% lower in accuracy than the best performing method. Since categories in ShapeNet and ModelNet40 only partially overlap, the representation learning ability of \mathcal{EBP} s can generalize to unseen categories.

Point cloud denoising. Lastly, we apply \mathcal{EBP} s to point cloud denoising by running MCMC sampling using noisy point clouds as initial samples. To create noisy point clouds, we perturb samples from the initial distribution by selecting a random point from the set and add Gaussian perturbations to points within a small radius r of the selected point. We

Table 2. Classification accuracy on ModelNet40. Models are pre-trained on ShapeNet before extracting features on ModelNet40. Linear SVMs are then trained using the learned representations.

Model	Accuracy
VConv-DAE (Sharma et al., 2016)	75.5
3D-GAN (Wu et al., 2016)	83.3
l-GAN (EMD) (Achlioptas et al., 2017)	84.0
l-GAN (CD) (Achlioptas et al., 2017)	84.5
PointGrow (Sun et al., 2018)	85.7
MRTNet-VAE (Gadella et al., 2018)	86.4
PointFlow (Yang et al., 2019)	86.8
PC-GAN (Li et al., 2018)	87.8
FoldingNet (Yang et al., 2018)	88.4
EBP (ours)	88.3

then perform 20 steps of Langevin dynamics with a fixed step size while keeping the unperturbed points fixed. Results in Figure 6 show that the gradient of our learned energy function is capable of guiding the MCMC sampling to recover the original point clouds. More examples of denoising can be found in Appendix G.




Figure 6. Examples of point cloud denoising using MCMC sampling. From left to right: original, perturbed, and denoised point clouds.

6. Conclusion

We have introduced a new energy-based processes representation, \mathcal{EBPs} , that unifies the stochastic process and latent variable modeling perspectives for set distributions. The proposed framework enhances the flexibility of current process and latent variable approaches, with provable exchangeability and consistency, in the conditional and unconditional

settings respectively. We have also introduced a new neural collapsed inference procedure for practical training of \mathcal{EBPs} , which connects the \mathcal{EBPs} to \mathcal{GPPs} , and demonstrated strong empirical results across a range of problems that involve conditional and unconditional set distribution modeling. Extending the approach to distributions over sets of discrete elements remains an interesting direction for future research.

Acknowledgements

We thank Weiyang Liu, Hongge Chen, Adams Wei Yu, and other members of the Google Brain team for helpful discussions.

References

- Achlioptas, P., Diamanti, O., Mitliagkas, I., and Guibas, L. Learning representations and generative models for 3d point clouds. *arXiv preprint arXiv:1707.02392*, 2017.
- Al-Shedivat, M., Wilson, A. G., Saatchi, Y., Hu, Z., and Xing, E. P. Learning scalable deep kernels with recurrent structure. *The Journal of Machine Learning Research*, 18(1):2850–2886, 2017.
- Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.
- Blei, D. and Lafferty, J. A correlated topic model of science. *Annals of Applied Statistics*, 1(1):17–35, 2007.
- Blei, D. and McAuliffe, J. D. Supervised topic models. In *Neural Information Processing Systems*, 2007.
- Blei, D., Ng, A., and Jordan, M. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, January 2003.
- Brown, L. D. *Fundamentals of Statistical Exponential Families*, volume 9 of *Lecture notes-monograph series*. Institute of Mathematical Statistics, Hayward, Calif, 1986.
- Bui, T., Hernández-Lobato, D., Hernandez-Lobato, J., Li, Y., and Turner, R. Deep gaussian processes for regression using approximate expectation propagation. In *International conference on machine learning*, pp. 1472–1481, 2016.
- Chu, W. and Ghahramani, Z. Gaussian processes for ordinal regression. *J. Mach. Learn. Res.*, 6:1019–1041, 2005.
- Dai, B., Liu, Z., Dai, H., He, N., Gretton, A., Song, L., and Schuurmans, D. Exponential family estimation via adversarial dynamics embedding. *arXiv preprint arXiv:1904.12083*, 2019.

- Damianou, A. C. and Lawrence, N. D. Deep gaussian processes. *arXiv preprint arXiv:1211.0358*, 2012.
- Dereudre, D. Introduction to the theory of Gibbs point processes. *Stochastic Geometry*: 181–229. Springer 2019.
- Du, Y. and Mordatch, I. Implicit generation and generalization in energy-based models. *arXiv preprint arXiv:1903.08689*, 2019.
- Duvenaud, D. K., Lloyd, J. R., Grosse, R. B., Tenenbaum, J. B., and Ghahramani, Z. Structure discovery in nonparametric regression through compositional kernel search. In *ICML*, pp. 1166–1174, 2013.
- Edwards, H. and Storkey, A. Towards a neural statistician. *arXiv preprint arXiv:1606.02185*, 2016.
- Gadelha, M., Wang, R., and Maji, S. Multiresolution tree networks for 3d point cloud processing. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 103–118, 2018.
- Garnelo, M., Rosenbaum, D., Maddison, C. J., Ramalho, T., Saxton, D., Shanahan, M., Teh, Y. W., Rezende, D. J., and Eslami, S. Conditional neural processes. *arXiv preprint arXiv:1807.01613*, 2018a.
- Garnelo, M., Schwarz, J., Rosenbaum, D., Viola, F., Rezende, D. J., Eslami, S., and Teh, Y. W. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018b.
- Ghahramani, Z. and Heller, K. Bayesian sets. In *Neural Information Processing Systems*, 2005.
- Ha, D., Dai, A., and Le, Q. V. HyperNetworks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- Hawkes, A. G. *Spectra of some self-exciting and mutually exciting point processes*. *Biometrika*, 58: 83–143, 1971. Oxford University Press
- Hensman, J., Fusi, N., and Lawrence, N. D. Gaussian processes for big data. *CoRR*, abs/1309.6835, 2013. URL <http://arxiv.org/abs/1309.6835>.
- Hinton, G. E. and Salakhutdinov, R. R. Replicated softmax: an undirected topic model. In *Advances in neural information processing systems*, pp. 1607–1614, 2009.
- Hofmann, T. Probabilistic latent semantic indexing. In *ACM Special Interest Group in Information Retrieval (SIGIR)*, 1999.
- Kim, H., Mnih, A., Schwarz, J., Garnelo, M., Eslami, A., Rosenbaum, D., Vinyals, O., and Teh, Y. W. Attentive neural processes. *arXiv preprint arXiv:1901.05761*, 2019.
- Kinderman, R. and Snell, J. L. *Markov Random Fields and their applications*. Amer. Math. Soc., Providence, RI, 1980.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Korshunova, I., Degraeve, J., Huszar, F., Gal, Y., Gretton, A., and Dambre, J. BRUNO: A deep recurrent model for exchangeable data. In *Advances in Neural Information Processing Systems*, pp. 7188–7197, 2018.
- Kulesza, A., and Taskar, B. Determinantal point processes for machine learning *Foundations and Trends® in Machine Learning*, 5–3, 123–286, 2012. Now Publishers, Inc.
- Lafferty, J. D., McCallum, A., and Pereira, F. Conditional random fields: Probabilistic modeling for segmenting and labeling sequence data. In *Proceedings of International Conference on Machine Learning*, volume 18, pp. 282–289, San Francisco, CA, 2001. Morgan Kaufmann.
- Larochelle, H. and Murray, I. The neural autoregressive distribution estimator. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pp. 29–37, 2011.
- Lavancier, F., Møller, J., and Rubak, E. Determinantal point process models and statistical inference *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 77:853–877, 2015. Wiley Online Library.
- Lawrence, N. D. Gaussian process latent variable models for visualisation of high dimensional data. In *Advances in neural information processing systems*, pp. 329–336, 2004.
- LeCun, Y. MNIST handwritten digit database, 1998. URL <http://yann.lecun.com/exdb/mnist/>.
- LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M., and Huang, F. A tutorial on energy-based learning. *Predicting structured data*, 1(0), 2006.
- Li, C.-L., Zaheer, M., Zhang, Y., Póczos, B., and Salakhutdinov, R. Point cloud gan. *arXiv preprint arXiv:1810.05795*, 2018.
- Liu, W., Liu, Z., Rehg, J. M., and Song, L. Neural similarity learning. In *Advances in Neural Information Processing Systems*, pp. 5026–5037, 2019.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*, pp. 3730–3738, 2015.

- Louizos, C., Shi, X., Schutte, K., and Welling, M. The functional neural process. In *Advances in Neural Information Processing Systems*, pp. 8743–8754, 2019.
- Ma, C., Li, Y., and Hernández-Lobato, J. M. Variational implicit processes. *arXiv preprint arXiv:1806.02390*, 2018.
- Øksendal, B. Stochastic differential equations. In *Stochastic differential equations*, pp. 65–84. Springer, 2003.
- Opper, M. and Winther, O. Gaussian processes for classification: Mean field algorithms. *Neural Computation*, 12(11):2655–2684, 2000.
- Porteous, I., Newman, D., Ihler, A., Asuncion, A., Smyth, P., and Welling, M. Fast collapsed gibbs sampling for latent dirichlet allocation. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 569–577. ACM, 2008.
- Quiñonero-Candela, J. and Rasmussen, C. E. A unifying view of sparse approximate gaussian process regression. *The Journal of Machine Learning Research*, 6: 1939–1959, 2005.
- Rasmussen, C. E. and Williams, C. K. I. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, 2006.
- Rezende, D. J. and Mohamed, S. Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*, 2015.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- Salimbeni, H. and Deisenroth, M. Doubly stochastic variational inference for deep gaussian processes. In *Advances in Neural Information Processing Systems*, pp. 4588–4599, 2017.
- Shah, A., Wilson, A., and Ghahramani, Z. Student-t processes as alternatives to gaussian processes. In *Artificial intelligence and statistics*, pp. 877–885, 2014.
- Sharma, A., Grau, O., and Fritz, M. Vconv-dae: Deep volumetric shape learning without object labels. In *European Conference on Computer Vision*, pp. 236–250. Springer, 2016.
- Snelson, E. and Ghahramani, Z. Local and global sparse gaussian process approximations. In *International Conference on Artificial Intelligence and Statistics*, pp. 524–531, 2007.
- Sun, Y., Wang, Y., Liu, Z., Siegel, J. E., and Sarma, S. E. Pointgrow: Autoregressively learned point cloud generation with self-attention. *arXiv preprint arXiv:1810.05591*, 2018.
- Teh, Y., Jordan, M., Beal, M., and Blei, D. Hierarchical dirichlet processes. *Journal of the American Statistical Association*, 101(576):1566–1581, 2006.
- Teh, Y. W., Newman, D., and Welling, M. A collapsed variational Bayesian inference algorithm for latent Dirichlet allocation. In *Advances in Neural Information Processing Systems*, volume 19, pp. 1353–1360, 2007. ISBN 9780262195683.
- Titsias, M. Variational learning of inducing variables in sparse gaussian processes. *Artificial Intelligence and Statistics*, 2009.
- Williams, C. and Rasmussen, C. Gaussian processes for regression. In Touretzky, D. S., Mozer, M. C., and Hasselmo, M. E. (eds.), *Advances in Neural Information Processing Systems 8*, volume 8, pp. 514–520, Cambridge, MA, 1996. MIT Press.
- Williams, C. K. I. and Seeger, M. Using the Nystrom method to speed up kernel machines. In Dietterich, T. G., Becker, S., and Ghahramani, Z. (eds.), *Advances in Neural Information Processing Systems 14*, 2000.
- Wilson, A. G., Hu, Z., Salakhutdinov, R. R., and Xing, E. P. Stochastic variational deep kernel learning. In *Advances in Neural Information Processing Systems*, pp. 2586–2594, 2016.
- Wu, J., Zhang, C., Xue, T., Freeman, B., and Tenenbaum, J. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *Advances in neural information processing systems*, pp. 82–90, 2016.
- Wu, Y. N., Xie, J., Lu, Y., and Zhu, S.-C. Sparse and deep generalizations of the frame model. *Annals of Mathematical Sciences and Applications*, 3(1):211–254, 2018.
- Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., and Xiao, J. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1912–1920, 2015.
- Xie, J., Lu, Y., Zhu, S.-C., and Wu, Y. A theory of generative convnet. In *International Conference on Machine Learning*, pp. 2635–2644, 2016.
- Xie, J., Zheng, Z., Gao, R., Wang, W., and Zhu, S.C., and Nian Wu, Y. Learning descriptor networks for 3D shape synthesis and analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8629–8638, 2018.

- Xie, P., Salakhutdinov, R., Mou, L., and Xing, E. Deep determinantal point process for large-scale multi-label classification. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 473–482, 2017.
- Yang, G., Huang, X., Hao, Z., Liu, M.-Y., Belongie, S., and Hariharan, B. Pointflow: 3d point cloud generation with continuous normalizing flows. *arXiv preprint arXiv:1906.12320*, 2019.
- Yang, Y., Feng, C., Shen, Y., and Tian, D. Foldingnet: Point cloud auto-encoder via deep grid deformation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 206–215, 2018.
- Zaheer, M., Kottur, S., Ravanbakhsh, S., Poczos, B., Salakhutdinov, R. R., and Smola, A. J. Deep sets. In *Advances in neural information processing systems*, pp. 3391–3401, 2017.
- Zamorski, M., Zieba, M., Nowak, R., Stokowiec, W., and Trzcinski, T. Adversarial autoencoders for generating 3d point clouds. *arXiv preprint arXiv:1811.07605*, 2, 2018.

Appendix

A. Adversarial Dynamics Embedding

The details of ADE derivations were originally provided in (Dai et al., 2019). We include relevant details here for completeness, since we make explicit use of these techniques in our training methods.

Given the EBM with $p_f(x) \propto \exp(f(x))$, ADE considers the augmented model

$$p_f(x, v) \propto \exp\left(f(x) - \frac{\lambda}{2}v^\top v\right), \quad (\text{A.1})$$

where v is the auxiliary momentum variable. It has been proved in Dai et al. (2019) that the MLE of (A.1) is the same as the original model, *i.e.*,

$$\operatorname{argmax}_f \widehat{\mathbb{E}}_{\mathcal{D}} \left[\log \int p(x, v) dv \right] = \operatorname{argmax}_f \widehat{\mathbb{E}}_{\mathcal{D}} [\log p_f(x)]. \quad (\text{A.2})$$

We then apply the primal-dual view of the MLE to the augmented model, leading to

$$\max_f \min_{q(x, v)} \widehat{\mathbb{E}} [f(x)] - \mathbb{E}_{q(x, v)} \left[f(x) - \frac{\lambda}{2}v^\top v - \log q(x, v) \right], \quad (\text{A.3})$$

where $q(x, v)$ is the dual sampler in the exponential family of distributions.

To ensure that the dual sampler $q(x, v)$ is flexible and tractable, ADE utilizes the dynamics embedding parametrization. Specifically, we consider the Hamiltonian dynamics embedding as an example. In this setting, the sample first comes from an initial distribution $(x^0, v^0) \sim q_\omega^0(x, v)$, and then moves according to

$$(x', v') = \mathbf{L}_{f, \eta}(x, v) := \begin{pmatrix} v^{\frac{1}{2}} = v + \frac{\eta}{2} \nabla_x f(x) \\ x' = x + \eta v^{\frac{1}{2}} \\ v' = v^{\frac{1}{2}} + \frac{\eta}{2} \nabla_x f(x') \end{pmatrix}, \quad (\text{A.4})$$

where η is defined as the leapfrog stepsize. After T iterations, we obtain

$$(x^T, v^T) = \mathbf{L}_{f, \eta} \circ \mathbf{L}_{f, \eta} \circ \dots \circ \mathbf{L}_{f, \eta}(x^0, v^0), \quad (\text{A.5})$$

where $\mathbf{L}_{f, \eta}$ can be one layer of the neural network. Together with the initial distribution q_ω^0 , we obtain the parametrization of $q(x, v)$ with learnable parameters (η, ω) . As justified in Dai et al. (2019), this parametrization is flexible and has a tractable density,

$$q^T(x^T, v^T) = q_\omega^0(x^0, v^0). \quad (\text{A.6})$$

By plugging this parametrization into (A.3), we obtain the final objective,

$$\max_{f \in \mathcal{F}} \min_{\omega, \eta} \ell(f, \omega) := \widehat{\mathbb{E}}_{\mathcal{D}} [f] - \mathbb{E}_{(x^0, v^0) \sim q_\omega^0(x, v)} \left[f(x^T) - \frac{\lambda}{2} \|v^T\|_2^2 \right] - H(q_\omega^0). \quad (\text{A.7})$$

Dai et al. (2019) also introduces the Langevin and generalized Hamiltonian dynamics embedding for dual density parametrization in ADE. For more details, please refer to Dai et al. (2019).

A.1. Dynamics Embedding Distribution Parametrization

We now present the concrete implementation of the dual sampler $q(x_{1:n}, v|\theta)$ used in our paper. Following the ADE technique introduced above, parametrization of $q(x_{1:n}, v|\theta)$ is separated into parameterizing the initial distribution and the dynamics embedding. In our real-data experiment, we use the block RNN to parametrize $q_\omega^0(x_{1:n}, v|\theta)$, as introduced in Appendix F.1. Block RNN is simply a design choice; other alternatives such as normalizing flows can also parametrize $q_\omega^0(x_{1:n}, v|\theta)$, as shown in our synthetic-data experiment.

The output of the RNN is then treated as the starting sample to which T Hamiltonian/Langevin dynamics updates are applied, *i.e.*,

$$(x_{1:n}^T, v^T) = \mathbf{L} \circ \mathbf{L} \circ \dots \circ \mathbf{L} (x_{1:n}^0, v^0), \quad (\text{A.8})$$

where \mathbf{L} can be either a Hamiltonian layer or a Langevin layer as specialized below,

$$\textbf{Hamiltonian layer: } (x'_{1:n}, v') = \mathbf{L}_{f,\eta} (x_{1:n}, v) := \left(\begin{array}{l} v_i^{\frac{1}{2}} = v_i + \frac{\eta}{2} \nabla_x f(x_i; \theta) \\ x'_i = x_i + \eta v_i^{\frac{1}{2}} \\ v'_i = v_i^{\frac{1}{2}} + \frac{\eta}{2} \nabla_x f(x'_i; \theta) \end{array} \middle| i = 1, \dots, n \right), \quad (\text{A.9})$$

with $v' = \{v'_i\}_{i=1}^n$.

$$\textbf{Langevin layer: } (x'_{1:n}, v') = \mathbf{L}_{f,\eta}^\xi (x_{1:n}) := \left(\begin{array}{l} v'_i = \xi_i + \frac{\eta}{2} \nabla_x f(x_i; \theta) \\ x'_i = x_i + v'_i \end{array} \middle| i = 1, \dots, n \right), \quad (\text{A.10})$$

with $\xi = \{\xi_i\}_{i=1}^n$, $\xi_i \sim q_\omega(\xi)$ and $v' = \{v'_i\}_{i=1}^n$.

Finally, following Theorem 4 in Dai et al. (2019), we obtain the parametrized dual sampler with tractable density as

$$\textbf{Hamiltonian embedding: } q^T(x_{1:n}^T, v^T | \theta) = q^0(x_{1:n}^0, v^0 | \theta), \quad (\text{A.11})$$

$$\textbf{Langevin embedding: } q^T(x_{1:n}^T, \{v^t\}_{t=1}^T | \theta) = q^0(x_{1:n}^0, \xi^0 | \theta) \prod_{t=1}^{T-1} q_{\omega_i}(\xi^t). \quad (\text{A.12})$$

By plugging this into the primal-dual view objective (21), we are able to learn the parameters in \mathcal{EBPs} and dual samplers.

B. Derivation of Special Cases of \mathcal{EBPs}

In this section, we provide the details for instantiating (un)conditional \mathcal{EBPs} to other specific models.

B.1. Latent Variable Representation of Gaussian Processes

We consider the latent variable model specified in (8), *i.e.*,

$$\theta \sim \mathcal{N}(\mathbf{0}, I_d), \quad (\text{B.1})$$

$$f_w(x, t; \theta) = \frac{1}{2\sigma^2} \|x - \theta^\top \phi(t)\|^2. \quad (\text{B.2})$$

To show that the marginal distribution follows

$$p(x_{t_1:t_n} | \{t_i\}_{i=1}^n) = \mathcal{N}(0, K(t_{1:n}) + \sigma^2 I_n),$$

we integrate out of θ , *i.e.*,

$$p(x_{t_1:t_n} | \{t_i\}_{i=1}^n) = \int \prod_{i=1}^n p(x|t_i, \theta) p(\theta) d\theta = \mathcal{N}\left(0, \sigma^2 I_n + \phi(t_{1:n})^\top \phi(t_{1:n})\right),$$

where the last equation come from the integration of Gaussians. This shows that \mathcal{GP} s under the latent variable parametrization are a special case of \mathcal{EBPs} .

B.2. Latent Variable Representation of Student- t Processes

We consider the latent variable model specified in (10), *i.e.*,

$$\alpha \sim \mathcal{N}(\mathbf{0}, I_d), \quad (\text{B.3})$$

$$\beta^{-1} \sim \Gamma\left(\frac{\nu}{2}, \frac{\gamma}{2}\right), \quad (\text{B.4})$$

$$f_w(x, t; \theta) = \frac{\gamma \left\| x - \sqrt{\frac{\beta(\nu-2)}{\gamma}} \alpha^\top \phi(t) \right\|^2}{2\sigma^2 (\nu-2) \beta}, \quad (\text{B.5})$$

To show the marginal distribution follows

$$p(x_{t_1:t_n} | \{t_i\}_{i=1}^n) = \mathcal{T}(\nu, 0, K(t_{1:n}) + \beta I_n),$$

we integrate out of $\theta = (\alpha, \beta)$. By integrating over α , we have

$$\begin{aligned} p(x_{t_1:t_n} | \beta, \{t_i\}_{i=1}^n) &= \int \prod_{i=1}^n p(x_{t_i} | t_i, \alpha, \beta) p(\alpha) d\alpha \\ &= \mathcal{N}\left(\mathbf{0}, \left(\frac{\beta}{\gamma}(\nu - 2)\right) \underbrace{\left(\sigma^2 I_n + \phi(t_{1:n})^\top \phi(t_{1:n})\right)}_{\tilde{K}(t_{1:n})}\right), \end{aligned}$$

and by integrating over β , we have

$$\begin{aligned} p(x_{t_1:t_n} | \{t_i\}_{i=1}^n) &= \int p(x_{t_1:t_n} | \beta, \{t_i\}_{i=1}^n) p(\beta) d\beta \\ &\propto \int \exp\left(-\frac{1 + \frac{\gamma X^\top \tilde{K}^{-1} X}{(\nu-2)}}{2\beta}\right) \beta^{-\frac{\nu}{2}-1} d\beta \\ &\propto \left(1 + \frac{\gamma X^\top \tilde{K}^{-1} X}{(\nu-2)}\right)^{\frac{\nu+n}{2}} = \mathcal{T}(\nu, 0, \tilde{K}) \end{aligned}$$

This shows that \mathcal{TP} s under the latent variable parametrization are a special case of \mathcal{EBP} s.

B.3. Topic Models

The topic models, including Bayesian sets (Ghahramani & Heller, 2005), probabilistic latent semantic index (Hofmann, 1999), latent Dirichlet allocation (LDA) family (Blei et al., 2003; Blei & Lafferty, 2007; Blei & McAuliffe, 2007), and replicated softmax (Hinton & Salakhutdinov, 2009), are also special cases of the unconditional \mathcal{EBP} s.

Here we consider the original LDA and replicated softmax as examples of directed and undirected topic models respectively. Other models follow a similar consideration.

Latent Dirichlet Allocation LDA is a representative of the directed topic model, which treats a document as a set of words. The model defines each component in (12) as

$$\begin{aligned} p(t_i) &= \text{Dir-Multi}(\alpha), \\ p(x_{t_i} | t_i) &= \text{Multi}(\Phi^\top \mathbf{1}_{t_i}), \quad i = 1, \dots, n, \end{aligned}$$

where $t_i \in \{1, \dots, d\}$ denotes the topic of each word and follows the Dirichlet-Multinomial distribution. The x_i is a k -dimensional one-hot encoding for each word in the vocabulary. Φ is a $d \times k$ matrix where each row denotes the distribution of words in one topic.

Replicated Softmax The replicated softmax (Hinton & Salakhutdinov, 2009) is proposed as an undirected topic model, which defines the joint distribution in (12) as a restricted Boltzmann machine with a linear potential function, *i.e.*,

$$p(x_{1:n}, \theta) \propto \exp\left(-\sum_{i=1}^n (\theta^\top W x_i + b^\top x_i) - a^\top \theta\right), \quad (\text{B.6})$$

where $\theta \in \{0, 1\}^d$ can be seen as the latent topic assignment. Each x_i is a k -dimensional indication vector with only one element equals to 1 and the rest equal to 0. Together $\{x_i\}_{i=1}^n$ denote the one-hot encodings for the observed words in the document.

As a result, LDA and replicated softmax are special realizations of unconditional \mathcal{EBP} s, where the index variables are either explicit defined or integrated out.

C. Proof Details of Theorem 2

Theorem 2 *If $n \geq m \geq 1$, and prior is exchangeable and consistent, then the marginal distribution $p(x_{1:n})$ will be exchangeable and consistent.*

Proof We simply verify the consistency of $p(x_{1:n})$ under the consistency of $p(\{t_i\}_{i=1}^m)$.

$$\begin{aligned}
 & \int p(x_{1:n}) dx_{m+1:n} \\
 = & \int \int p(x_{t_1:t_n} | \{t_i\}_{i=1}^n) p(\{t_i\}_{i=1}^n) d\{t_i\}_{i=1}^n dx_{m+1:n} \\
 = & \int \left(\int p(x_{t_1:t_n} | \{t_i\}_{i=1}^n) dx_{m+1:n} \right) p(\{t_i\}_{i=1}^n) d\{t_i\}_{i=1}^n \\
 = & \int p(x_{t_1:t_m} | \{t_i\}_{i=1}^m) \left(\int p(\{t_i\}_{i=1}^n) dt_{m+1:n} \right) dt_{1:m} \\
 = & \int p(x_{t_1:t_m} | \{t_i\}_{i=1}^m) p(\{t_i\}_{i=1}^m) dt_{1:m} = p(x_{1:m}).
 \end{aligned}$$

The exchangeability of $p(x_{1:n})$ directly comes from the exchangeability of $p(x_{t_1:t_n})$ and $p(\{t_i\}_{i=1}^n)$,

$$\begin{aligned}
 p(x_{1:n}) &= \int p(x_{t_1:t_n} | \{t_i\}_{i=1}^n) p(\{t_i\}_{i=1}^n) d\{t_i\}_{i=1}^n \\
 &= \int p(\pi(x_{t_1:t_n}) | \pi(\{t_i\}_{i=1}^n)) p(\pi(\{t_i\}_{i=1}^n)) d\{t_i\}_{i=1}^n \\
 &= p(\pi(x_{t_1:t_n}))
 \end{aligned}$$

■

D. More Details of Inference

D.1. Further Neural Collapsed Inference for Unconditional \mathcal{EBPs}

In the main text, we introduce the neural collapsed inference for \mathcal{EBPs} to eliminate the posterior inference of $\{t_i\}_{i=1}^n$. We can further exploit the neural collapsed inference idea to eliminate θ .

Recall the model with $\{t_i\}_{i=1}^n$ collapsed,

$$\begin{aligned}
 p_{w'}(x_{1:n} | \theta) &= \prod_{i=1}^n \int p_w(x_{t_i} | \theta, t_i) p(t_i) dt_i \\
 &\propto \prod_{i=1}^n \int \exp(f_w(x_{t_i}, t_i; \theta) - Z(f_w, t_i; \theta) + h_v(t_i)) dt_i \\
 &\approx \prod_{i=1}^n \frac{1}{Z(f_{w'}; \theta)} \exp(f_{w'}(x_i; \theta)).
 \end{aligned}$$

We further consider $p(\theta) \propto \exp(g_u(\theta))$, then we have

$$\begin{aligned}
 & \prod_{i=1}^n \int p_{w'}(x_{1:n} | \theta) p(\theta) d\theta \\
 \propto & \int \exp\left(\sum_{i=1}^n f_{w'}(x_i; \theta) + g_u(\theta)\right) d\theta \\
 \approx & \frac{1}{Z(w'')} \exp\left(\sum_{i=1}^n f_{\bar{w}'}\left(x_i; \sum_{i=1}^n \phi_{u'}(x_i)\right)\right) := p_{w''}(x_{1:n}), \tag{D.1}
 \end{aligned}$$

where we denote $w'' = \{\tilde{w}', u'\}$. The last approximation comes from the fact that

- 1) the integration over θ leads to a distribution over $x_{1:n} \in \otimes^n \Omega$;
- 2) since $p(x_{1:n})$ is an integration of latent variable model, it should be exchangeable.

Therefore, we consider deepsets (Zaheer et al., 2017) in the reparametrized EBM as a special choice of (D.1), which satisfies these two conditions in order to approximate the integrated model.

The attention mechanism is another choice for the reparametrized EBM, *i.e.*,

$$p_{w''}(x_{1:n}) = \frac{1}{Z(w'')} \exp\left(\sum_{i=1}^n \text{attn}_{\tilde{w}'}(x_i; x_{1:n})\right),$$

where $\text{attn}_{\tilde{w}'}(x; x_{1:n}) = \sum_{i=1}^n \frac{\exp(\phi(x)^\top \phi(x_i)) h(x_i)}{\sum_{j=1}^n \exp(\phi(x)^\top \phi(x_j))}$ with w'' denoting the parameters in $\phi(\cdot)$ and $h(\cdot)$.

We can apply ADE to such neural collapsed reparametrization if the task only concerns set generation (*i.e.*, does not use θ). We take the deepsets parametrization (D.1) as an example (the attention parametrization follows similarly). Specifically, we learn the parameters in w'' via

$$\begin{aligned} \max_{w''} \min_{q(\{x_i\}_{i=1}^n, v)} \widehat{\mathbb{E}}_{\mathcal{D}} \left[\sum_{i=1}^n f_{\tilde{w}'} \left(x_i; \sum_{i=1}^n \phi_{u'}(x_i) \right) \right] \\ - \mathbb{E}_{q(\{x_i\}_{i=1}^n, v)} \left[\sum_{i=1}^n f_{\tilde{w}'} \left(x_i; \sum_{i=1}^n \phi_{u'}(x_i) \right) - \frac{\lambda}{2} v^\top v \right] - H(q(\{x_i\}_{i=1}^n, v)). \end{aligned} \quad (\text{D.2})$$

Following the ADE technique, we can parametrize the initialization distribution $q_{\omega}^0(\{x_i^0\}_{i=1}^n, v^0)$ using an RNN and refine $q(\{x_i\}_{i=1}^n, v)$ with learnable Hamiltonian/Langevin dynamics as introduced in Section A.

Connection to Gibbs Point Processes (\mathcal{GPP} s) and Determinantal Point Processes (\mathcal{DPP} s): In fact, if we adopt the most general model collapsed unconditional \mathcal{EBP} for arbitrary n , we obtain

$$p_{w''}(x_{1:n}) = \frac{1}{Z_{w''}} \exp(f_{w''}(x_{1:n})), \quad (\text{D.3})$$

which includes the Gibbs point processes (Dereudre, 2019) with $f_{w''}$ satisfying several mild conditions for the regularity of \mathcal{GPP} . Particularly, the determinantal point processes (Lavancier et al., 2015; Kulesza et al., 2012) can be instantiated from \mathcal{GPP} by setting the potential function to be log det of some kernel function, which can also be parametrized by neural network, *e.g.*, Xie et al. (2016).

Therefore, the proposed algorithm can straightforwardly applied for (deep) \mathcal{GPP} s and \mathcal{DPP} s learning. It should be emphasized that by exploiting the proposed primal-dual MLE framework, we automatically obtain a deep neural network parametrized dual sampler with the learned model simultaneously, which can be used in inference and bypass the notorious sampling difficulty in \mathcal{GPP} and \mathcal{DPP} .

D.2. Inference for Conditional \mathcal{EBP}

In the conditional \mathcal{EBP} setting where the index $\{t_i\}_{i=1}^n$ are given for any cardinality n , we are modeling:

$$p(x_{t_1:t_n} | \{t_i\}_{i=1}^n) = \int \prod_{i=1}^n p(x | t_i) p(\theta) d\theta \quad (\text{D.4})$$

For simplicity, we use capital letters to denote the set. We denote T_{train}, X_{train} as the observed sets of points, and T_{test}, X_{test} as the un-observed inputs and targets. For the rest of this section, we denote $T = T_{train} \cup T_{test}$ and $X = X_{train} \cup X_{test}$. Then the predictive model is to infer

$$p(X_{test} | T, X_{train}) = \frac{p(X_{train}, X_{test} | T)}{\sum_{X'} p(X_{train}, X' | T)} = \frac{p(X_{train}, X_{test} | T)}{p(X_{train} | T)} = \frac{p(X_{train}, X_{test} | T)}{p(X_{train} | T_{train})} \quad (\text{D.5})$$

where the last equal sign is due to the consistency condition of stochastic processes. Below we briefly review several existing conditional neural processes.

D.2.1. REVIEW OF CONDITIONAL NEURAL PROCESSES

The learning of processes is generally done via maximizing the marginal likelihood $p(X|T)$. The following neural processes, however, perform learning by maximizing the predictive distribution $p(X|T, X_{train})$ directly.

Conditional Neural Processes (Garnelo et al., 2018a) directly parametrize the conditional distribution $p(X|T, X_{train})$ as:

$$p(X|T, X_{train}) = \prod_{i=1}^{|X|} \mathcal{N}(x_{t_i} | \mu_{t_i}, \sigma_{t_i}) \quad (\text{D.6})$$

where μ_{t_i} and σ_{t_i} are outputs of some neural network $g(t_i, T_{train}, X_{train})$ with permutation invariance.

Neural Processes (Garnelo et al., 2018b) model the distribution of $p(X|T)$ as

$$p(X|T) = \int \prod_{i=1}^{|X|} \mathcal{N}(x_{t_i} | g(t_i, \theta), \sigma) p(\theta) d\theta, \quad (\text{D.7})$$

where $g(t_i, \theta)$ can be learned using ELBO:

$$\log p(X|T) \geq \mathbb{E}_{q(\theta|T, X)} \left[\sum_{i=1}^{|X|} \log \mathcal{N}(x_{t_i} | g(t_i, \theta), \sigma) + \log p(\theta) - \log q(\theta|T, X) \right]. \quad (\text{D.8})$$

During learning, however, Garnelo et al. (2018b) again performs MLE on the predictive model:

$$\log p(X_{pred}|T, X_{context}) \geq \mathbb{E}_{q(\theta|T, X)} \left[\sum_{x_{t_j} \in X_{pred}} \log \mathcal{N}(x_{t_j} | g(t_j, \theta), \sigma) + \log p(\theta|T_{context}, X_{context}) - \log q(\theta|T, X) \right], \quad (\text{D.9})$$

where $X_{context}$ denotes a subset of X_{train} and $X_{pred} = X_{train} \setminus X_{context}$, that are randomly splitted from X_{train} . The $p(\theta|T_{context}, X_{context})$ is the true posterior after observing $T_{context}, X_{context}$. This posterior then serves as the prior of the predictive model according to Bayes' rule. However, since $p(\theta|T_{context}, X_{context})$ is not tractable in general, Garnelo et al. (2018b) uses $q(\theta|T_{context}, X_{context})$ to approximate this term instead.

Attentive Neural Processes (Kim et al., 2019) is an extension to Neural Processes (Garnelo et al., 2018b) by using self-attention to parametrize the variational posterior. Besides θ from the variational posterior, their model also adds the deterministic context embedding r computed using attention, together with location x_{t_i} , into the mean function $g(x_{t_i}, r, \theta)$. Other than these differences, the training procedure is the same as Garnelo et al. (2018b), which still uses Gaussian for observation modeling.

D.2.2. LEARNING CONDITIONAL \mathcal{EBP}

The proposed \mathcal{EBP} s exploits flexible energy-based model, instead of Gaussian distributions in \mathcal{NP} s and their variants:

$$p(X|T, \theta) \propto \exp(f(X, T; \theta)) p(\theta) \propto \exp \left(\sum_{i=1}^{|X|} f(x_{t_i}, t_i; \theta) \right). \quad (\text{D.10})$$

The ELBO then becomes:

$$\log p(X|T) \geq \mathbb{E}_{q(\theta|T, X)} \left[\sum_{i=1}^{|X|} f(x_{t_i}, t_i; \theta) - A(\theta, t_i) + \log p(\theta) - \log q(\theta|T, X) \right], \quad (\text{D.11})$$

where $A(\theta, t) = \log \int \exp(f(x, t; \theta)) dx$ is the log partition function.

Using the ADE technique, we have $A(\theta, t) = \max_{q'(x|\theta, t)} \mathbb{E}_{q'(x|\theta, t)} [f(x, t; \theta)] + H(q')$. By plugging this into the ELBO, we arrive at the learning objective which tries to maximize the marginal likelihood w.r.t.

$$\max_{q, h} \min_{q'} \mathbb{E}_{q(\theta|T, X)} \left[\log p(\theta) - \log q(\theta|T, X) + \sum_{i=1}^{|X|} f(x_{t_i}, t_i; \theta) - \mathbb{E}_{x \sim q'} f(x, t_i; \theta) - H(q') \right]. \quad (\text{D.12})$$

Different from the family of neural processes, we are able to directly optimize for $p(X|T)$ using the above objective. We can easily learn the predictive distribution $p(X|T, X_{train})$ similar to other neural processes by replacing the prior with the variational posterior from the observed set T_{train}, X_{train} .

D.2.3. PREDICTION USING \mathcal{EBPs}

Without loss of generality, we illustrate the prediction of a single point (t, x) given the training set, namely $p(x|t, X_{train}, T_{train})$.

As is pointed out in Eq (D.5), $p(x|t, X_{train}, T_{train}) \propto p(x, X_{train}|t, T_{train})$. We use the variational posterior (denoted as $q(\theta)$ below for simplicity) to approximate the lower bound:

$$\begin{aligned} \log p(x, X_{train}|t, T_{train}) &= \log \left(\int p(\theta) \exp(f(x, t; \theta) - A(\theta, t)) \prod_{i=1}^{|X_{train}|} \exp(f(x_{t_i}, t_i; \theta) - A(\theta, t_i)) d\theta \right) \\ &\geq \max_{q(\theta)} \mathbb{E}_{q(\theta)} [f(x, t; \theta)] + \mathbb{E}_{q(\theta)} [G(\theta, t, X_{train}, T_{train})], \end{aligned} \quad (\text{D.13})$$

where

$$G(\theta, t, X_{train}, T_{train}) = \log p(\theta) - \log q(\theta) - A(\theta, t) + \sum_{i=1}^{|X_{train}|} f(x_{t_i}, t_i; \theta) - A(\theta, t_i) \quad (\text{D.14})$$

In practice, we use $q(\theta|X_{train}|T_{train})$ to replace the above $p(\theta)$, which results in the predictive distribution approximated by:

$$\begin{aligned} p(x|t, T_{train}, X_{train}) &\geq \exp(\mathbb{E}_{q(\theta)} [f(x, t; \theta)]) \cdot \exp(\mathbb{E}_{q(\theta)} G(\theta, t, X_{train}, T_{train})) \\ &\propto \exp(\mathbb{E}_{q(\theta)} [f(x, t; \theta)]) \\ &\simeq \exp(\mathbb{E}_{q(\theta|X_{train}, T_{train})} [f(x, t; \theta)]) \end{aligned} \quad (\text{D.15})$$

E. Additional Related Work

The proposed \mathcal{EBPs} bridge the gap between stochastic processes and models of exchangeable data. We summarize these two related topics below:

Stochastic Processes. Exploiting stochastic processes for conditional distribution modeling has a long line of research, starting from Gaussian processes for regression (Williams & Rasmussen, 1996) to being generalized to classification (Opper & Winther, 2000) and ordinal regression (Chu & Ghahramani, 2005). One of the major bottlenecks of \mathcal{GP} s is the memory and computational costs — $\mathcal{O}(N^2)$ in memory and $\mathcal{O}(N^3)$ in computation respectively, where N denotes the number of training samples. Although low-rank and sparse approximations have been proposed (Williams & Seeger, 2000; Quiñero-Candela & Rasmussen, 2005; Snelson & Ghahramani, 2007; Titsias, 2009; Hensman et al., 2013) to reduce these costs, flexibility of \mathcal{GP} s and its variants remains a critical issue to be addressed.

Student- t processes are derived by adding the inverse-Wishart process prior to the kernel function (Shah et al., 2014), allowing the kernel function to adapt to data. As we discussed in Section 3.2, student- t processes essentially rescales each dimension in the likelihood model, and thus still has restricted modeling flexibility. Another line of research to further improve the flexibility of \mathcal{GP} s is to introduce structures or deep compositions into kernels, such as Duvenaud et al. (2013); Damianou & Lawrence (2012); Bui et al. (2016); Wilson et al. (2016); Al-Shedivat et al. (2017). Even though neural networks can be incorporated in constructing these kernels, the likelihood estimations of these models are still restricted to known distributions, limiting their modeling flexibility.

Neural processes (Garnelo et al., 2018b) and its variants (Garnelo et al., 2018a; Kim et al., 2019; Louizos et al., 2019) introduce neural networks to stochastic processes beyond \mathcal{GP} s. However, their likelihoods are still restricted to known distributions, *e.g.*, Gaussian. Besides the modeling restriction, these models are learned by maximizing *log-predictive* distribution, rather than the *log-marginal* distribution, which may lead to suboptimal solutions.

The most competitive model w.r.t. \mathcal{EBP} s is the implicit processes (Ma et al., 2018), which uses the implicit models as the likelihood in (4). However, due to the intractability of the implicit likelihood, \mathcal{GP} s are introduced for variational inference, which negatively impacts the initially designed flexibility in implicit processes, as demonstrated in Section G.1.

Exchangeable Probabilistic Models. The generative models for exchangeable data is a separate research topic which mostly relies on the De Finetti’s Theorem. Bayesian sets (Ghahramani & Heller, 2005) considers the latent variable model with Bernoulli distribution likelihood and beta prior for a set of binary data. The conjugacy of Bernoulli and beta distributions leads to tractability in Bayesian inference, but limits the flexibility. Topic models, *e.g.*, pLSI (Hofmann, 1999), LDA family (Blei et al., 2003; Blei & Lafferty, 2007; Blei & McAuliffe, 2007), and replicated softmax (Hinton & Salakhutdinov, 2009) generalize Bayesian sets by introducing more complicated local latent variables, but the likelihood is still restricted to simple distributions.

To make distribution modeling more flexible, neural networks have been introduced to likelihood estimation of latent variable models (Edwards & Storkey, 2016). However, this work only exploits the analytic form of parametrization within known distributions. Korshunova et al. (2018); Yang et al. (2019) use normalizing flows to improve modeling flexibility of exchangeable data. Flow-based models, compared to \mathcal{EBP} s, still restrict the underlying distribution (by requiring the transformations to be invertible), and therefore cannot fully utilize the expressiveness of neural networks.

The Gibbs point processes (Dereudre, 2019), including Poisson point processes, Hawkes point processes (Hawkes, 1971) and determinantal point processes (Lavancier et al., 2015; Kulesza et al., 2012) as special cases, is also an alternative flexible model for exchangeable data. As we discussed in Appendix D.1, by the neural collapsed inference technique, we reduce the collapsed \mathcal{EBP} to \mathcal{GPP} , which implies the flexibility of the proposed \mathcal{EBP} . Moreover, this connection also highlights that the proposed primal-dual ADE can be used for \mathcal{GPP} and \mathcal{DPP} estimation.

F. Experiment Details

F.1. Architecture and Training Details

As precluded in Section D, we use deepsets (Zaheer et al., 2017) to parametrize $q(\theta|x_{1:n})$ as a diagonal Gaussian. Specifically, each input x_i undergoes 1D-convolutions with kernel size 1 and filter sizes $\{128, 256\}$ for synthetic-data experiments and filter sizes $\{128, 256, 256, 512\}$ for image completion and point-cloud generation. The 1D-convolution layers are interleaved with ReLU, followed by max-pooling across inputs in the embedding space. We then use the reparametrization trick (Kingma & Welling, 2013) to sample θ from the resulting Gaussian. To compute the energy of a single input $f(x_i; \theta)$, each x_i is concatenated with θ and followed by fully connected layers of $\{128, 64, 1\}$ neurons each interleaved with ReLU. The energy of the set, $f(x_{1:n}; \theta)$, is the average energy of all inputs in the set.

To obtain the ADE initialization distribution $q^0(x_{1:n}^0|\theta)$ on synthetic data, we first use a 2-layer hyper-network (Ha et al., 2017) with 256 hidden neurons each to output the parameters for a 10-layer normalizing planar flow model (Rezende & Mohamed, 2015). The planar flow then takes both θ (learned from T_{train} and X_{train}) and the target indices T_{test} as inputs to produce $q^0(x_{1:t_n}^0|\theta)$. For the synthetic experiments, we directly use $q^0(x_{1:n}^0|\theta)$ as the sampler output without performing additional Hamiltonian/Langevin dynamics, which is sufficient to capture the synthetic data distributions.

To obtain $q(x_{1:n}, v|\theta)$ for image completion and point-cloud generation, we first parametrize the initialization distribution $q^0(x_{1:n}^0|\theta)$ using an RNN where each recurrent LSTM block consists of an MLP with $\{64, 128, 512\}$ hidden neurons interleaved with ReLU. Each LSTM block outputs the mean and variance of a diagonal Gaussian of dimension k times d , where k (block size) is the number of elements generated at once, and d is the dimension of each element. For image completion, k equals the number of pixels in a row of an image (*e.g.*, 28 for MNIST and 32 for CelebA), and d equals the number of channels of a pixel (1 for MNIST and 3 for CelebA). For point-cloud generation, we have $k = 512$ and $d = 3$. We considered a range of k values from 32 to 2048, and selected k based on the best completion/generation performance. The number of recurrent RNN blocks equals the total number of elements in the final generated set (784 for MNIST, 1024 for CelebA, and 2048 for point clouds) divided by k . Using the RNN output as the initialization distribution $q^0(x_{1:n}^0|\theta)$, we then perform $T = 20$ steps of Langevin dynamics similar to Xie et al. (2016) with step size $\eta = 0.1$ and $\xi_i \sim \mathcal{N}(0, 0.05)$

while clipping $\nabla_x f$ to 0.1.

We use spectral normalization in training the energy function and batch normalization in training the sampler. We use Adam with learning rate 10^{-4} to optimize all of our models. We set $\beta_1 = 0.5$ for image completion and $\beta_1 = 0.0$ for the synthetic experiments and for point-cloud generation. The coefficient of $H(q(x_{1:n}, v|\theta))$ in (1) is set to 10^{-5} . All tasks are trained until convergence using batch size 64 on a single NVIDIA V100 GPU with 32 GB memory.

F.2. Point-Cloud Preprocessing

Following Achlioptas et al. (2017), we use shapes from ShapeNet (Wu et al., 2015) that are axis aligned and centered into the unit sphere. For the point-cloud classification task, we apply random rotations along the gravity axis following Achlioptas et al. (2017); Yang et al. (2019), and also normalize shapes from ModelNet40 (Wu et al., 2015) the same way. For generation, the model is trained on the official training split and evaluated on the test split of ShapeNet (where the train-validation-test splits are 70%-20%-10%) similar to Yang et al. (2019).

F.3. Point-Cloud Model Selection

We follow the model selection protocol of Achlioptas et al. (2017), namely select the model with the smallest JSD and reports other measurements of this model. Measurements are created every 100 epochs. We noticed that the measurement of JSD and COV are relatively robust across different evaluation runs, whereas the MMD measurement (due to its small magnitude, e.g., 10^{-4}) is less robust.

F.4. Point-Cloud Generation Metrics

Following Achlioptas et al. (2017); Yang et al. (2019), we use Chamfer distance (CD) and earth mover’s distance (EMD) defined below to measure distance between point clouds.

$$d_{CD}(X, X') = \sum_{x_i \in X} \min_{x'_i \in X'} \|x_i - x'_i\|_2 + \sum_{x'_i \in X'} \min_{x_i \in X} \|x_i - x'_i\|_2, \quad (\text{F.1})$$

$$d_{EMD}(X, X') = \min_{\phi: X \rightarrow X'} \sum_{x_i \in X} \|x_i - \phi(x_i)\|_2. \quad (\text{F.2})$$

Similarly, we use Jensen-Shannon Divergence (JSD), Minimum matching distance (MMD), and Coverage (COV) defined below to evaluate generation quality.

- JSD is computed between the marginal distribution of the entire reference set (p_r) and the marginal distribution of the entire generated set (p_g) of point clouds, specifically,

$$\text{JSD}(p_g, p_r) = \frac{1}{2} D_{KL}(p_r || p_m) + \frac{1}{2} D_{KL}(p_g || p_m), \quad (\text{F.3})$$

where $p_m = \frac{1}{2}(p_r + p_g)$ and D_{KL} stands for the Kullback-Leibler divergence. JSD only measures similarity at the marginal distribution level, and does not provide insights on the generation quality of each individual point cloud.

- MMD measures the average distance between a point-cloud in the reference set S_r and its closest neighbor in the generated set S_g , namely

$$\text{MMD}(S_g, S_r) = \frac{1}{|S_r|} \sum_{X \in S_r} \min_{X' \in S_g} d(X, X'), \quad (\text{F.4})$$

where $d(X, X')$ is the distance between two point clouds according to either CD or EMD.

- COV measures the ratio of point clouds in the reference set S_r that are matched to a distinct closest neighbor in the the generated set S_g :

$$\text{COV}(S_g, S_r) = \frac{1}{|S_r|} |\{\arg \min_{X' \in S_g} d(X, X') | X' \in S_g\}|. \quad (\text{F.5})$$

where $d(X, X')$ can again be based on CD or EMD. COV reflects the generation diversity.

G. More Experimental Results

G.1. Additional Comparisons on Synthetic Data

We compare the proposed \mathcal{EBP} s with Gaussian processes (\mathcal{GP} s), neural processes (\mathcal{NP} s)⁴, and variational implicit processes (\mathcal{VIP} s)⁵.

Figure 7 shows the ground truth data and the normalized probability heatmap for each process. To generate data given an index t_i , we randomly select one of the two modes (e.g., one of the sine waves) as the mean and adds $\epsilon \sim \mathcal{N}(0, 0.1)$ noise to produce x_{t_i} . We plot the heatmap of the learned predictive distribution of each process. For \mathcal{EBP} , we use (D.15) to approximate $p(x_{t_i}|t_i)$.

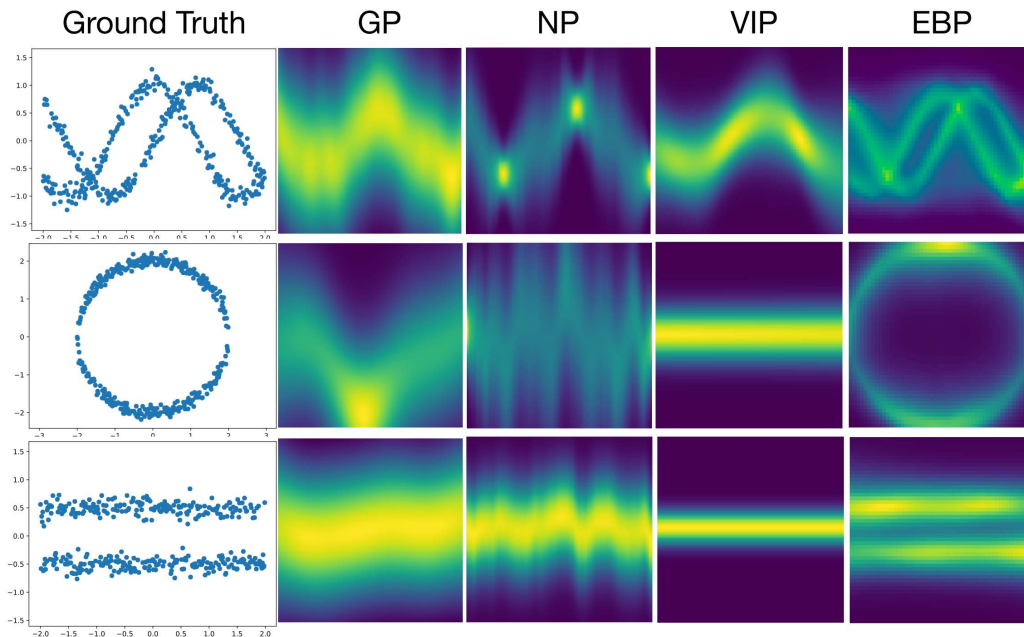


Figure 7. The ground truth data and learned energy functions of \mathcal{GP} , \mathcal{NP} , \mathcal{VIP} , and \mathcal{EBP} (from left to right). \mathcal{EBP} successfully captures multi-modality of the toy data as \mathcal{GP} , \mathcal{NP} , and \mathcal{VIP} exhibiting only a single mode.

⁴<https://github.com/deepmind/neural-processes>

⁵ <https://github.com/LaurantChao/VIP>

G.2. Additional Image Completion Results on MNIST




Figure 8. Additional image completion results on MNIST where the top half of the image serves as context.




Figure 9. Additional image completion results on MNIST where 10, 100, and 1000 (top to bottom) random pixels serve as context.

Figure 8 shows additional image completion results on MNIST when only the upper-half of an image is given. Figure 9 shows MNIST completion results when 10, 100, and 1000 randomly selected pixels are given.

G.3. Additional Image Completion Results on CelebA




Figure 10. Additional image completion results on CelebA for both contiguous (right-most) and random pixels as context.

Figure 10 shows additional completion results on CelebA where 10, 100, 500, 800, 1000 randomly selected pixels are given and when a 16x16 square is removed from the original 32x32 image.

G.4. Additional Point-Cloud Generation Results




Figure 11. Additional examples of airplane, chair, and car point-cloud generation with 4 RNN blocks of block size 512.

G.5. Additional Point-Cloud Denoising Results




Figure 12. Additional examples of airplane, chair, and car point-cloud denoising. Left to right: original, perturbed, denoised point clouds.