

# SetVAE: Learning Hierarchical Composition for Generative Modeling of Set-Structured Data

Jinwoo Kim\* Jaehoon Yoo\* Juho Lee Seunghoon Hong  
 KAIST

{jinwoo-kim, wogns98, juholee, seunghoon.hong}@kaist.ac.kr

## Abstract

Generative modeling of set-structured data, such as point clouds, requires reasoning over local and global structures at various scales. However, adopting multi-scale frameworks for ordinary sequential data to a set-structured data is nontrivial as it should be invariant to the permutation of its elements. In this paper, we propose SetVAE, a hierarchical variational autoencoder for sets. Motivated by recent progress in set encoding, we build SetVAE upon attentive modules that first partition the set and project the partition back to the original cardinality. Exploiting this module, our hierarchical VAE learns latent variables at multiple scales, capturing coarse-to-fine dependency of the set elements while achieving permutation invariance. We evaluate our model on point cloud generation task and achieve competitive performance to the prior arts with substantially smaller model capacity. We qualitatively demonstrate that our model generalizes to unseen set sizes and learns interesting subset relations without supervision. Our implementation is available at <https://github.com/jw9730/setvae>.

## 1. Introduction

There have been increasing demands in machine learning for handling *set-structured* data (*i.e.*, a group of unordered instances). Examples of set-structured data include object bounding boxes [19, 2], point clouds [1, 18], support sets in the meta-learning [6], *etc.* While initial research mainly focused on building neural network architectures to encode sets [32, 17], generative models for sets have recently grown popular [33, 16, 25, 30, 31].

A generative model for set-structured data should verify the two essential requirements: (i) *exchangeability*, meaning that a probability of a set instance is invariant to its elements' ordering, and (ii) handling *variable cardinality*, meaning that a model should flexibly process sets with variable cardinalities. These requirements pose a unique chal-




Figure 1: Color-coded attention learned by SetVAE encoder for three data instances of ShapeNet Airplane [3]. Level 1 shows attention at the most coarse scale. Level 2 and 3 show attention at more fine-scales.

lenge in set generative modeling, as they prevent the adaptation of standard generative models for sequences or images [7, 12, 28, 27]. For instance, typical operations in these models, such as convolution or recurrent operations, exploit implicit ordering of elements (*e.g.*, adjacency), thus breaking the exchangeability. Several works circumvented this issue by imposing heuristic ordering [10, 5, 8, 22]. However, when applied to set-structured data, any ordering assumed by a model imposes an unnecessary inductive bias that might harm the generalization ability.

There are several existing works satisfying these requirements. Edwards et al., [4] proposed a simple generative model encoding sets into latent variables, while other approaches build upon various generative models, such as generative adversarial networks [18, 25], flow-based models [30, 13], and energy-based models [31]. All these works define valid generative models for set-structured data, but with some limitations. To achieve exchangeability, many approaches process set elements independently [18, 30],

\*Equal contribution

Table 1: Summary of several set generative frameworks available to date. Our SetVAE jointly achieves desirable properties, with the advantages of the VAE framework combined with our novel contributions.

Model	Exchangeability	Variable cardinality	Inter-element dependency	Hierarchical latent structure
I-GAN [1]	×	×	○	×
PC-GAN [18]	○	○	×	×
PointFlow [30]	○	○	×	×
EBP [31]	○	○	○	×
SetVAE (ours)	○	○	○	○

limiting the models in reflecting the interactions between the elements during generation. Some approaches take the inter-element dependency into account [25, 31], but have an upper bound on the number of elements [25], or less scalable due to heavy computations [31]. More importantly, existing models are less effective in capturing subset structures in sets presumably because they represent a set with a single-level representation. For sets containing multiple sub-objects or parts, it would be beneficial to allow a model to have structured latent representations such as one obtained via hierarchical latent variables.

In this paper, we propose SetVAE, a novel hierarchical variational autoencoder (VAE) for sets. SetVAE models interaction between set elements by adopting attention-based Set Transformers [17] into the VAE framework, and extends it to a *hierarchy* of latent variables [23, 26] to account for flexible subset structures. By organizing latent variables at each level as a *latent set* of fixed cardinality, SetVAE is able to learn hierarchical multi-scale features that decompose a set data in a coarse-to-fine manner (Figure 1) while achieving *exchangeability* and handling *variable cardinality*. In addition, composing latent variables invariant to input’s cardinality allows our model to generalize to arbitrary cardinality unseen during training.

The contributions of this paper are as follows:

- We propose SetVAE, a novel hierarchical VAE for sets with *exchangeability* and *varying cardinality*. To the best of our knowledge, SetVAE is the first VAE successfully applied for sets with arbitrary cardinality. SetVAE has a number of desirable properties compared to previous works, as summarized in Table 1.
- Equipped with novel Attentive Bottleneck Layers (ABLs), SetVAE is able to model the coarse-to-fine dependency across the arbitrary number of set elements using a hierarchy of latent variables.
- We conduct quantitative and qualitative evaluations of SetVAE on generative modeling of point cloud in various datasets, and demonstrate better or competitive performance in generation quality with less number of parameters than the previous works.

## 2. Preliminaries

### 2.1. Permutation-Equivariant Set Generation

Denote a set as  $\mathbf{x} = \{\mathbf{x}_i\}_{i=1}^n \in \mathcal{X}^n$ , where  $n$  is the cardinality of the set and  $\mathcal{X}$  represents the domain of each element  $\mathbf{x}_i \in \mathbb{R}^d$ . In this paper, we represent  $\mathbf{x}$  as a matrix  $\mathbf{x} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^T \in \mathbb{R}^{n \times d}$ . Note that any operation on a set should be invariant to the elementwise permutation and satisfy the two constraints of *permutation invariance* and *permutation equivariance*.

**Definition 1.** A function  $f : \mathcal{X}^n \rightarrow \mathcal{Y}$  is *permutation invariant* iff for any permutation  $\pi(\cdot)$ ,  $f(\pi(\mathbf{x})) = f(\mathbf{x})$ .

**Definition 2.** A function  $f : \mathcal{X}^n \rightarrow \mathcal{Y}^n$  is *permutation equivariant* iff for any permutation  $\pi(\cdot)$ ,  $f(\pi(\mathbf{x})) = \pi(f(\mathbf{x}))$ .

In the context of generative modeling, the notion of permutation invariance translates into exchangeability, requiring a joint distribution of the elements invariant with respect to the permutation.

**Definition 3.** A distribution for a set of random variables  $\mathbf{x} = \{\mathbf{x}_i\}_{i=1}^n$  is *exchangeable* if for any permutation  $\pi$ ,  $p(\mathbf{x}) = p(\pi(\mathbf{x}))$ .

An easy way to achieve exchangeability is to assume each element to be *i.i.d.* and process a set of initial elements  $\mathbf{z}^{(0)} = \{\mathbf{z}_i^{(0)}\}_{i=1}^n$  independently sampled from  $p(\mathbf{z}_i^{(0)})$  with an elementwise function  $f_{\text{elem}}$  to get the  $\mathbf{x}$ :

$$\mathbf{x} = \{\mathbf{x}_i\}_{i=1}^n \text{ where } \mathbf{x}_i = f_{\text{elem}}(\mathbf{z}_i^{(0)}) \quad (1)$$

However, assuming elementwise independence poses a limit in modeling interactions between set elements. An alternative direction is to process  $\mathbf{z}^{(0)}$  with a permutation-equivariant function  $f_{\text{equiv}}$  to get the  $\mathbf{x}$ :

$$\mathbf{x} = \{\mathbf{x}_i\}_{i=1}^n = f_{\text{equiv}}(\{\mathbf{z}_i^{(0)}\}_{i=1}^n). \quad (2)$$

We refer to this approach as the *permutation-equivariant generative framework*. As the likelihood of  $\mathbf{x}$  does not depend on the order of its elements (because elements of  $\mathbf{z}^{(0)}$  are *i.i.d.*), this approach achieves exchangeability.




Figure 2: Illustration of Multihead Attention Block (MAB) and Induced Set Attention Block (ISAB).

## 2.2. Permutation-Equivariant Set Encoding

To design permutation-equivariant operations over a set, Set Transformer [17] provides attentive modules that model pairwise interaction between set elements while preserving invariance or equivariance. This section introduces two essential modules in the Set Transformer.

First, Multihead Attention Block (MAB) takes the query and value sets,  $Q \in \mathbb{R}^{n_q \times d}$  and  $V \in \mathbb{R}^{n_v \times d}$ , respectively, and performs the following transformation (Figure 2a):

$$\text{MAB}(Q, V) = \text{LN}(\mathbf{a} + \text{FF}(\mathbf{a})) \in \mathbb{R}^{n_q \times d}, \quad (3)$$

$$\text{where } \mathbf{a} = \text{LN}(Q + \text{Multihead}(Q, V, V)) \in \mathbb{R}^{n_q \times d}, \quad (4)$$

where FF denotes elementwise feedforward layer, Multihead denotes multi-head attention [29], and LN denotes layer normalization [17]. Note that the output of Eq. (3) is permutation equivariant to  $Q$  and permutation invariant to  $V$ .

Based on MAB, Induced Set Attention Block (ISAB) processes the input set  $\mathbf{x} \in \mathbb{R}^{n \times d}$  using a smaller set of inducing points  $I \in \mathbb{R}^{m \times d}$  ( $m < n$ ) by (Figure 2b):

$$\text{ISAB}_m(\mathbf{x}) = \text{MAB}(\mathbf{x}, \mathbf{h}) \in \mathbb{R}^{n \times d}, \quad (5)$$

$$\text{where } \mathbf{h} = \text{MAB}(I, \mathbf{x}) \in \mathbb{R}^{m \times d}. \quad (6)$$

The ISAB first transforms the input set  $\mathbf{x}$  into  $\mathbf{h}$  by attending from  $I$ . The resulting  $\mathbf{h}$  is a permutation invariant projection of  $\mathbf{x}$  to a lower cardinality  $m$ . Then,  $\mathbf{x}$  again attends to  $\mathbf{h}$  to produce the output of  $n$  elements. As a result, ISAB is permutation equivariant to  $\mathbf{x}$ .

**Property 1.** In  $\text{ISAB}_m(\mathbf{x})$ ,  $\mathbf{h}$  is permutation invariant to  $\mathbf{x}$ .

**Property 2.**  $\text{ISAB}_m(\mathbf{x})$  is permutation equivariant to  $\mathbf{x}$ .

## 3. Variational Autoencoders for Sets

The previous section suggests that there are two essential requirements for VAE for set-structured data: it should be able to model the likelihood of sets (i) in arbitrary cardinality and (ii) invariant to the permutation (*i.e.*, exchangeable). This section introduces our SetVAE objective that

satisfies the first requirement while achieving the second requirement is discussed in Section 4.

The objective of VAE [15] is to learn a generative model  $p_\theta(\mathbf{x}, \mathbf{z}) = p_\theta(\mathbf{z})p_\theta(\mathbf{x}|\mathbf{z})$  for data  $\mathbf{x}$  and latent variables  $\mathbf{z}$ . Since the true posterior is unknown, we approximate it using the inference model  $q_\phi(\mathbf{z}|\mathbf{x})$  and optimize the variational lower bound (ELBO) of the marginal likelihood  $p(\mathbf{x})$ :

$$\mathcal{L}_{\text{VAE}} = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] - \text{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p_\theta(\mathbf{z})). \quad (7)$$

**Vanilla SetVAE** When our data is a set  $\mathbf{x} = \{\mathbf{x}_i\}_{i=1}^n$ , Eq. (7) should be modified such that it can incorporate the set of arbitrary cardinality  $n$ <sup>1</sup>. To this end, we propose to decompose the latent variable  $\mathbf{z}$  into the two independent variables as  $\mathbf{z} = \{\mathbf{z}^{(0)}, \mathbf{z}^{(1)}\}$ . We define  $\mathbf{z}^{(0)} = \{\mathbf{z}_i^{(0)}\}_{i=1}^n$  to be a set of *initial elements*, whose cardinality is always the same as a data  $\mathbf{x}$ . Then we model the generative process as transforming  $\mathbf{z}^{(0)}$  into a set  $\mathbf{x}$  conditioned on the  $\mathbf{z}^{(1)}$ .

Given the independence assumption, the prior is factorized by  $p(\mathbf{z}) = p(\mathbf{z}^{(0)})p(\mathbf{z}^{(1)})$ . The prior on initial set  $p(\mathbf{z}^{(0)})$  is further decomposed into the cardinality and element-wise distributions as:

$$p(\mathbf{z}^{(0)}) = p(n) \prod_{i=1}^n p(\mathbf{z}_i^{(0)}). \quad (8)$$

We model  $p(n)$  using the empirical distribution of the training data cardinality. We find that the choice of the prior  $p(\mathbf{z}_i^{(0)})$  is critical to the performance, and discuss its implementation in Section 4.2.

Similar to the prior, the approximate posterior is defined as  $q(\mathbf{z}|\mathbf{x}) = q(\mathbf{z}^{(0)}|\mathbf{x})q(\mathbf{z}^{(1)}|\mathbf{x})$  and decomposed into:

$$q(\mathbf{z}^{(0)}|\mathbf{x}) = q(n|\mathbf{x}) \prod_{i=1}^n q(\mathbf{z}_i^{(0)}|\mathbf{x}) \quad (9)$$

We define  $q(n|\mathbf{x}) = \delta(n)$  as a delta function with  $n = |\mathbf{x}|$ , and set  $q(\mathbf{z}_i^{(0)}|\mathbf{x}) = p(\mathbf{z}_i^{(0)})$  similar to [30, 16, 21]. The resulting ELBO can be written as

$$\begin{aligned} \mathcal{L}_{\text{SVAE}} = & \mathbb{E}_{q(\mathbf{z}|\mathbf{x})}[\log p(\mathbf{x}|\mathbf{z})] \\ & - \text{KL}(q(\mathbf{z}^{(0)}|\mathbf{x})||p(\mathbf{z}^{(0)})) \\ & - \text{KL}(q(\mathbf{z}^{(1)}|\mathbf{x})||p(\mathbf{z}^{(1)})). \end{aligned} \quad (10)$$

In the supplementary file, we show that the first KL divergence in Eq. (10) is a constant and can be ignored in the optimization. During inference, we sample  $\mathbf{z}^{(0)}$  by first sampling the cardinality  $n \sim p(n)$  then the  $n$  initial elements independently from the prior  $p(\mathbf{z}_i^{(0)})$ .

<sup>1</sup>Without loss of generality, we use  $n$  to denote the cardinality of a set but assume that the training data is composed of sets in various size.

**Hierarchical SetVAE** To allow our model to learn a more expressive latent structure of the data, we can extend the vanilla SetVAE using hierarchical latent variables.

Specifically, we extend the plain latent variable  $\mathbf{z}^{(1)}$  into  $L$  disjoint groups  $\{\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(L)}\}$ , and introduce a top-down hierarchical dependency between  $\mathbf{z}^{(l)}$  and  $\{\mathbf{z}^{(0)}, \dots, \mathbf{z}^{(l-1)}\}$  for every  $l > 1$ . This leads to the modification in the prior and approximate posterior to

$$p(\mathbf{z}) = p(\mathbf{z}^{(0)})p(\mathbf{z}^{(1)}) \prod_{l>1} p(\mathbf{z}^{(l)}|\mathbf{z}^{(<l)}) \quad (11)$$

$$q(\mathbf{z}|\mathbf{x}) = q(\mathbf{z}^{(0)}|\mathbf{x})q(\mathbf{z}^{(1)}|\mathbf{x}) \prod_{l>1} q(\mathbf{z}^{(l)}|\mathbf{z}^{(<l)}, \mathbf{x}). \quad (12)$$

Applying Eq. (11) and (12) to Eq. (10), we can derive the ELBO as

$$\begin{aligned} \mathcal{L}_{\text{HSVAE}} = & \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} [\log p(\mathbf{x}|\mathbf{z})] \\ & - \text{KL}(q(\mathbf{z}^{(0)}|\mathbf{x})||p(\mathbf{z}^{(0)})) - \text{KL}(q(\mathbf{z}^{(1)}|\mathbf{x})||p(\mathbf{z}^{(1)})) \\ & - \sum_{l=2}^L \mathbb{E}_{q(\mathbf{z}^{(<l)}|\mathbf{x})} \text{KL}(q(\mathbf{z}^{(l)}|\mathbf{z}^{(<l)}, \mathbf{x})||p(\mathbf{z}^{(l)}|\mathbf{z}^{(<l)})). \end{aligned} \quad (13)$$

**Hierarchical prior and posterior** To model the prior and approximate posterior in Eq. (11) and (12) with top-down latent dependency, we employ the bidirectional inference in [23]. We outline the formulations here and elaborate on the computations in Section 4.

Each conditional  $p(\mathbf{z}^{(l)}|\mathbf{z}^{(<l)})$  in the prior is modeled by the factorized Gaussian, whose parameters are dependent on the latent variables of the upper hierarchy  $\mathbf{z}^{(<l)}$ :

$$p(\mathbf{z}^{(l)}|\mathbf{z}^{(<l)}) = \mathcal{N}(\mu_l(\mathbf{z}^{(<l)}), \sigma_l(\mathbf{z}^{(<l)})). \quad (14)$$

Similarly, each conditional in the approximate posterior  $q(\mathbf{z}^{(l)}|\mathbf{z}^{(<l)}, \mathbf{x})$  is also modeled by the factorized Gaussian. We use the residual parameterization in [26] which predicts the parameters of the Gaussian using the displacement and scaling factors  $(\Delta\mu, \Delta\sigma)$  conditioned on  $\mathbf{z}^{(<l)}$  and  $\mathbf{x}$ :

$$q(\mathbf{z}^{(l)}|\mathbf{z}^{(<l)}, \mathbf{x}) = \mathcal{N}(\mu_l(\mathbf{z}^{(<l)}) + \Delta\mu_l(\mathbf{z}^{(<l)}, \mathbf{x}), \sigma_l(\mathbf{z}^{(<l)}) \cdot \Delta\sigma_l(\mathbf{z}^{(<l)}, \mathbf{x})). \quad (15)$$

**Invariance and equivariance** We assume that the decoding distribution  $p(\mathbf{x}|\mathbf{z}^{(0)}, \mathbf{z}^{(1:L)})$  is equivariant to the permutation of  $\mathbf{z}^{(0)}$  and invariant to the permutation of  $\mathbf{z}^{(1:L)}$  since such model induces an exchangeable model:

$$\begin{aligned} p(\pi(\mathbf{x})) &= \int p(\pi(\mathbf{x})|\pi(\mathbf{z}^{(0)}), \mathbf{z}^{(1:L)})p(\pi(\mathbf{z}^{(0)}))p(\mathbf{z}^{(1:L)})d\mathbf{z} \\ &= \int p(\mathbf{x}|\mathbf{z}^{(0)}, \mathbf{z}^{(1:L)})p(\mathbf{z}^{(0)})p(\mathbf{z}^{(1:L)})d\mathbf{z} = p(\mathbf{x}). \end{aligned} \quad (16)$$

We further assume that the approximate posterior distributions  $q(\mathbf{z}^{(l)}|\mathbf{z}^{(<l)}, \mathbf{x})$  are invariant to the permutation of  $\mathbf{x}$ . In the following section, we describe how we implement the encoder and decoder satisfying these criteria.

## 4. SetVAE Framework

We present the overall framework of the proposed SetVAE. Figure 3 illustrates an overview. SetVAE is based on the bidirectional inference [23], which is composed of the bottom-up encoder and top-down generator sharing the same dependency structure. In this framework, the inference network forms the approximate posterior by merging bottom-up information from data with the top-down information from the generative prior. We construct the encoder using a stack of ISABs in Section 2.2, and treat each of the projected set  $\mathbf{h}$  as a deterministic encoding of data.

Our generator is composed of a stack of special layers called Attentive Bottleneck Layer (ABL), which extends the ISAB in Section 2.2 with the stochastic interaction with the latent variable. Specifically, ABL processes a set at each layer of the generator as follows:

$$\text{ABL}_m(\mathbf{x}) = \text{MAB}(\mathbf{x}, \text{FF}(\mathbf{z})) \in \mathbb{R}^{n \times d} \quad (17)$$

$$\text{with } \mathbf{h} = \text{MAB}(I, \mathbf{x}) \in \mathbb{R}^{m \times d}, \quad (18)$$

where FF denotes a feed-forward layer, and the latent variable  $\mathbf{z}$  is derived from the projection  $\mathbf{h}$ . For generation (Figure 3a), we sample  $\mathbf{z}$  from the prior in Eq. (14) by,

$$\mathbf{z} \sim \mathcal{N}(\mu, \sigma) \text{ where } \mu, \sigma = \text{FF}(\mathbf{h}). \quad (19)$$

For inference (Figure 3b), we sample  $\mathbf{z}$  from the posterior in Eq. (15) by,

$$\begin{aligned} \mathbf{z} &\sim \mathcal{N}(\mu + \Delta\mu, \sigma \cdot \Delta\sigma) \\ \text{where } \Delta\mu, \Delta\sigma &= \text{FF}(\mathbf{h} + \mathbf{h}_{\text{enc}}), \end{aligned} \quad (20)$$

where  $\mathbf{h}_{\text{enc}}$  is obtained from the corresponding ISAB layer of the bottom-up encoder. Following [23], we share the parameters between the generative and inference networks. A detailed illustration of ABL is in the supplementary file.

To generate a set, we first sample the initial elements  $\mathbf{z}^{(0)}$  and the latent variable  $\mathbf{z}^{(1)}$  from the prior  $p(\mathbf{z}^{(0)})$  and  $p(\mathbf{z}^{(1)})$ , respectively. Given these inputs, the generator iteratively samples the subsequent latent variables  $\mathbf{z}^{(l)}$  from the prior  $p(\mathbf{z}^{(l)}|\mathbf{z}^{(<l)})$  one by one at each layer of ABL, while processing the set conditioned on the sampled latent variable via Eq. (17). The data  $\mathbf{x}$  is then decoded elementwise from the final output.

### 4.1. Analysis

**Modeling Exchangeable Likelihood** The architecture of SetVAE satisfies the invariance and equivariance criteria in Section 3. This is, in part, achieved by producing latent variables from the projected sets of bottom-up ISAB and top-down ABL. As the projected sets are permutation invariant to input (Section 2), the latent variables  $\mathbf{z}^{(1:L)}$  provide an invariant representation of the data. Furthermore, due to

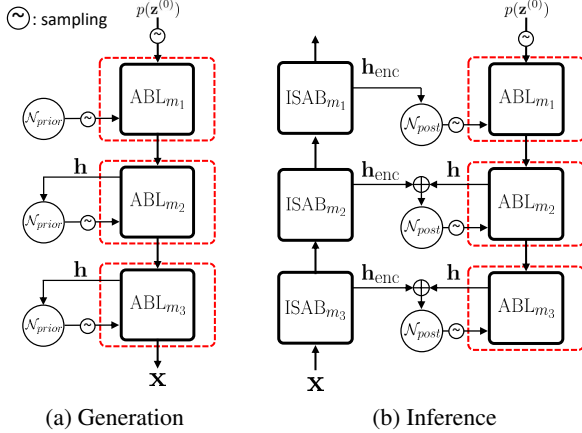


Figure 3: The hierarchical SetVAE.  $\mathcal{N}_{prior}$  denotes the prior (Eq. (19)) and  $\mathcal{N}_{post}$  denotes the posterior (Eq. (20)).

permutation equivariance of ISAB, the top-down stack of ABLs produce an output equivariant to the initial set  $\mathbf{z}^{(0)}$ . This renders the decoding distribution  $p(\mathbf{x}|\mathbf{z}^{(0)}, \mathbf{z}^{(1:L)})$  permutation equivariant to  $\mathbf{z}^{(0)}$ . Consequently, the decoder of SetVAE induces an exchangeable model.

**Learning Coarse-to-Fine Dependency** In SetVAE, both the ISAB and ABL project the input set  $\mathbf{x}$  of cardinality  $n$  to the projected set  $\mathbf{h}$  of cardinality  $m$  via multi-head attention (Eq. (6) and (18)). In the case of  $m < n$ , this projection functions as a bottleneck to the cardinality. This allows the model to encode some features of  $\mathbf{x}$  into the  $\mathbf{h}$  and discover interesting *subset* dependencies across the set elements. Denoting  $m_l$  as the bottleneck cardinality at layer  $l$  (Figure 3), we set  $m_l < m_{l+1}$  to induce the model to discover coarse-to-fine dependency of the set, such as object parts. Such bottleneck also effectively reduces network size, allowing our model to perform competitive or better than the prior arts with less than 50% of their parameters. This coarse-to-fine structure is a unique feature of SetVAE.

## 4.2. Implementation Details

This section discusses the implementation of SetVAE. We leave comprehensive details on the supplementary file.

**Multi-Modal Prior.** Although a unimodal Gaussian is a typical choice for the initial element distribution  $p(\mathbf{z}_i^{(0)})$  [16, 30], we find that the model converges significantly faster when we employ the multi-modal prior. We use a mixture of Gaussians (MoG) with  $K$  components:

$$p(\mathbf{z}_i^{(0)}) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{z}_i^{(0)}; \mu_k^{(0)}, \sigma_k^{(0)}). \quad (21)$$

**Likelihood.** For the likelihood  $p_\theta(\mathbf{x}|\mathbf{z})$ , we may consider a Gaussian distribution centered at the reconstruction. In

the case of point sets, we design the likelihood by

$$\begin{aligned} \mathcal{L}_{recon}(\mathbf{x}) &= -\log p_\theta(\mathbf{x}|\mathbf{z}) \\ &= \frac{1}{2} d(\mathbf{x}, \hat{\mathbf{x}}) + \text{const}, \end{aligned} \quad (22)$$

where  $d(\mathbf{x}, \hat{\mathbf{x}})$  is the optimal matching distance defined as

$$d(\mathbf{x}, \hat{\mathbf{x}}) = \min_{\pi} \sum_i \|\mathbf{x}_i - \hat{\mathbf{x}}_{\pi(i)}\|_2^2. \quad (23)$$

In other words, we measure the likelihood with the Gaussian at optimally permuted  $\mathbf{x}$ , and thus maximizing this likelihood is equivalent to minimizing the optimal matching distance between the data and the reconstruction. Unfortunately, directly maximizing this likelihood requires  $O(n^3)$  computation due to the matching. Instead, we choose the Chamfer Distance (CD) as a proxy reconstruction loss,

$$\begin{aligned} \mathcal{L}_{recon}(\mathbf{x}) &= \text{CD}(\mathbf{x}, \hat{\mathbf{x}}) \\ &= \sum_i \min_j \|\mathbf{x}_i - \hat{\mathbf{x}}_j\|_2^2 + \sum_j \min_i \|\mathbf{x}_i - \hat{\mathbf{x}}_j\|_2^2. \end{aligned} \quad (24)$$

The CD may not admit a direct interpretation as a negative log-likelihood of  $p_\theta(\mathbf{x}|\mathbf{z})$ , but shares the optimum with the matching distance having a proper interpretation. By employing the CD for the reconstruction loss, we learn the VAE with a surrogate for the likelihood  $p_\theta(\mathbf{x}|\mathbf{z})$ . CD requires  $O(n^2)$  computation time, so is scalable to the moderately large sets. Note also that the CD should be scaled appropriately to match the likelihood induced by optimal matching distance. We implicitly account for this by applying weights to KL divergence in our final objective function:

$$\mathcal{L}_{HSVAE}(\mathbf{x}) = \mathcal{L}_{recon}(\mathbf{x}) + \beta \mathcal{L}_{KL}(\mathbf{x}), \quad (25)$$

where  $\mathcal{L}_{KL}(\mathbf{x})$  is the KL divergence in Eq. (13).

## 5. Related Work

**Set generative modeling.** SetVAE is closely related to recent works on permutation-equivariant set prediction [33, 16, 2, 21, 20]. Closest to our approach is the autoencoding TSPN [16] that uses a stack of ISABs [17] to predict a set from randomly initialized elements. However, TSPN does not allow sampling, as it uses a pooling-based deterministic set encoding (FSPool) [34] for reconstruction. SetVAE instead discards FSPool and access projected sets in ISAB directly, which allows an efficient variational inference and a direct extension to hierarchical multi-scale latent.

Our approach differs from previous generative models treating each element *i.i.d.* and processing a random initial set with an elementwise function [4, 30, 13]. Notably, PointFlow [30] uses a continuous normalizing flow (CNF) to process a 3D Gaussian point cloud into an object. However, assuming elementwise independence could pose a



Table 2: Comparison against the state-of-the-art generative models.  $\uparrow$ : the higher the better.  $\downarrow$ : the lower the better. The best scores are highlighted in bold. MMD-CD is scaled by  $10^3$ , and MMD-EMD by  $10^2$ .

Category	Model	# Parameters (M)		MMD( $\downarrow$ )		COV( $\%, \uparrow$ )		1-NNA( $\%, \downarrow$ )	
		Full	Gen	CD	EMD	CD	EMD	CD	EMD
Airplane	I-GAN (CD)[30]	1.97	1.71	0.239	4.27	43.21	21.23	86.30	97.28
	I-GAN (EMD)[30]	1.97	1.71	0.269	3.29	<b>47.90</b>	<b>50.62</b>	87.65	85.68
	PC-GAN[30]	9.14	1.52	0.287	3.57	36.46	40.94	94.35	92.32
	PointFlow[30]	1.61	1.06	0.217	3.24	46.91	48.40	75.68	<b>75.06</b>
	SetVAE (Ours)	<b>0.75</b>	<b>0.39</b>	<b>0.199</b>	<b>3.07</b>	43.45	44.93	<b>75.31</b>	77.65
	Training set	-	-	0.226	3.08	42.72	49.14	70.62	67.53
Chair	I-GAN (CD)[30]	1.97	1.71	2.46	8.91	41.39	25.68	64.43	85.27
	I-GAN (EMD)[30]	1.97	1.71	2.61	7.85	40.79	41.69	64.73	65.56
	PC-GAN[30]	9.14	1.52	2.75	8.20	36.50	38.98	76.03	78.37
	PointFlow[30]	1.61	1.06	<b>2.42</b>	7.87	46.83	<b>46.98</b>	60.88	<b>59.89</b>
	SetVAE (Ours)	<b>0.75</b>	<b>0.39</b>	2.55	<b>7.82</b>	<b>46.98</b>	45.01	<b>58.76</b>	61.48
	Training set	-	-	1.92	7.38	57.25	55.44	59.67	58.46
Car	I-GAN (CD)[30]	1.97	1.71	1.55	6.25	38.64	18.47	63.07	88.07
	I-GAN (EMD)[30]	1.97	1.71	1.48	5.43	39.20	39.77	69.74	68.32
	PC-GAN[30]	9.14	1.52	1.12	5.83	23.56	30.29	92.19	90.87
	PointFlow[30]	1.61	1.06	0.91	5.22	44.03	<b>46.59</b>	60.65	<b>62.36</b>
	SetVAE (Ours)	<b>0.75</b>	<b>0.39</b>	<b>0.88</b>	<b>5.05</b>	<b>48.58</b>	44.60	<b>59.66</b>	63.35
	Training set	-	-	1.03	5.33	48.30	51.42	57.39	53.27

limit in modeling complex element interactions. Also, CNF requires the invertibility of the generative model, which could further limit its expressiveness. SetVAE resolves this problem by adopting permutation equivariant ISAB that models inter-element interactions via attention, and a hierarchical VAE framework with flexible latent dependency.

Contrary to previous works specifically designed for a certain type of set-structured data (e.g., point cloud [1, 18, 30]), we emphasize that SetVAE can be trivially applied to arbitrary set-structured data. We demonstrate this by applying SetVAE to the generation of a scene layout represented by a set of object bounding boxes.

**Hierarchical VAE.** Our model is built upon the prior works on hierarchical VAEs for images [11], such as Ladder-VAE [23], IAF-VAE [14], and NVAE [26]. To model long-range pixel correlations in images, these models organize latent variables at each hierarchy as images while gradually increasing their resolution via upsampling. However, the requirement for permutation equivariance has prevented applying multi-scale approaches to sets. ABLs in SetVAE solve this problem by defining latent variables in the projected scales of each hierarchy.

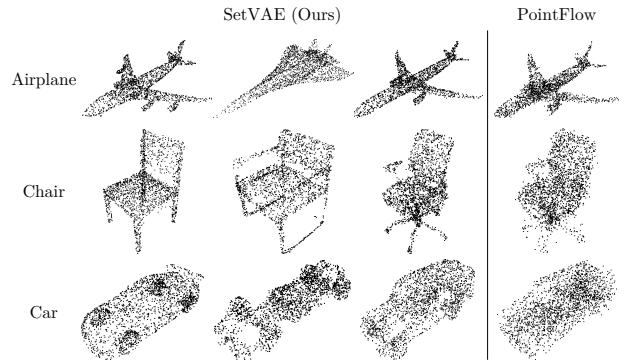


Figure 4: Examples of randomly generated point sets from SetVAE (ours) and PointFlow in ShapeNet.

## 6. Experiments

### 6.1. Experimental Setup

**Dataset** We examine SetVAE using ShapeNet [3], Set-MNIST [33], and Set-MultiMNIST [5] datasets. For ShapeNet, we follow the prior work using 2048 points sampled uniformly from the mesh surface [30]. For Set-MNIST, we binarized the images in MNIST and scaled the coordinates to  $[0, 1]$  [16]. Similarly, we build Set-MultiMNIST using  $64 \times 64$  images of MultiMNIST [5] with two digits randomly located without overlap.

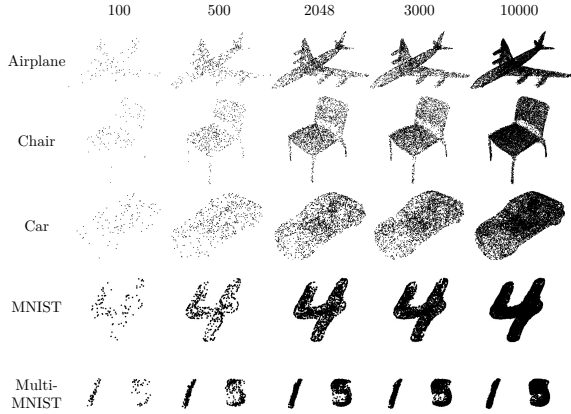


Figure 5: Samples from SetVAE for different cardinalities. At each row, the hierarchical latent variables are fixed and the initial set is re-sampled with different cardinality.



Figure 6: Samples from SetVAE and PointFlow in a high cardinality setting. Only the initial sets are re-sampled.

**Evaluation Metrics** For evaluation in ShapeNet, we compare the standard metrics including Minimum Matching Distance (MMD), Coverage (COV), and 1-Nearest Neighbor Accuracy (1-NNA), where the similarity between point clouds are computed with Chamfer Distance (CD) (Eq. (24)), and Earth Mover’s Distance (EMD) based on optimal matching. The details are in the supplementary file.

## 6.2. Comparison to Other Methods

We compare SetVAE with the state-of-the-art generative models for point clouds including l-GAN [1], PC-GAN [18], and PointFlow [30]. Following these works, we train our model for each category of airplane, chair, and car.

Table 2 summarizes the evaluation result. SetVAE achieves better or competitive performance to the prior arts using a *much smaller* number of parameters (8% to 45% of competitors). Notably, SetVAE often outperforms PointFlow with a substantial margin in terms of minimum matching distance (MMD) and has better or comparable coverage (COV) and 1-NNA. Lower MMD indicates that SetVAE generates high-fidelity samples, and high COV and low 1-NNA indicate that SetVAE generates diverse samples covering various modes in data. Together, the results indicate that SetVAE generates realistic, high-quality point sets. Notably, we find that SetVAE trained with CD (Eq. (24)) generalizes well to EMD-based metrics than l-GAN.

We also observe that SetVAE is significantly faster than

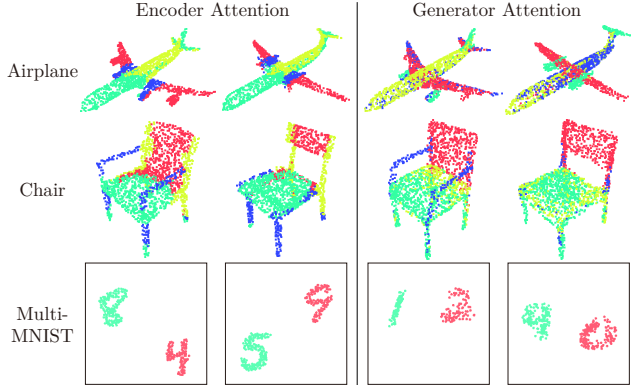


Figure 7: Attention visualization at a selected layer. Each point is color-coded by its assignment based on attention.

PointFlow in both training ( $56\times$  speedup; 0.20s vs. 11.2s) and testing ( $68\times$  speedup; 0.052s vs. 3.52s)<sup>2</sup>. It is because PointFlow requires a costly ODE solver for both training and inference, and has much more parameters.

Figure 4 illustrates qualitative comparisons. Compared to PointFlow, we observe that SetVAE generates sharper details, especially in small object parts such as wings and engines of an airplane or wheels of a car. We conjecture that this is because our model generates samples considering inter-element dependency while capturing shapes in various granularities via a hierarchical latent structure.

## 6.3. Internal Analysis

**Cardinality disentanglement** Ideally, a generative model for sets should be able to disentangle the cardinality of a set from the rest of the generative factors (e.g., structure, style). SetVAE partly achieves this by decomposing the latent variables into the initial set  $\mathbf{z}^{(0)}$  and the hierarchical latent variables  $\mathbf{z}^{(1:L)}$ . To validate this, we generate samples by changing the initial set’s cardinality while fixing the rest. Figure 5 illustrates the result. We observe that SetVAE generates samples having consistent global structure with a varying number of elements. Surprisingly, it even generalizes well to the cardinalities not seen during training. For instance, the model generalizes to any cardinality between 100 and 10,000, although it is trained with only 2048 points in ShapeNet and less than 250 points in Set-MNIST. It shows that SetVAE disentangles the cardinality from the factors characterizing an object.

We compare SetVAE to PointFlow in extremely high cardinality setting (100k points) in Figure 6. Although PointFlow innately disentangles cardinality by modeling each element independently, we observe that it tends to generate noisy, blurry boundaries in large cardinality settings. In contrast, SetVAE retains the sharpness of the structure even for extreme cardinality, presumably because it considers inter-element dependency in the generation process.

<sup>2</sup>Measured on a single GTX 1080ti with a batch size of 16.

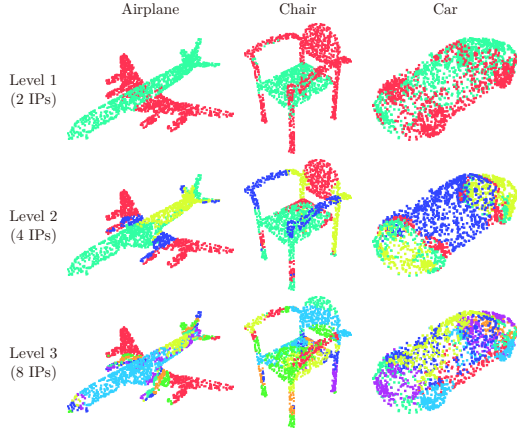


Figure 8: Visualization of encoder attention across multiple layers. IP notes the number of inducing points at each level. See the supplementary file for more results.




Figure 9: Layer-wise classification results using the hierarchical latent variables in Set-MultiMNIST dataset.

**Discovering coarse-to-fine dependency** SetVAE discovers interesting subset structures via hierarchical bottlenecks in ISAB and ABL. To demonstrate this, we visualize the encoder attention (Eq. (6)) and generator attention (Eq. (18)) in Figure 7, where each point is color-coded based on its hard assignment to one of  $m$  inducing points<sup>3</sup>. We observe that SetVAE attends to semantically interesting parts consistently across samples, such as wings and engines of an airplane, legs and backs of a chair, and even different instances of multiple digits.

Figure 8 illustrates the point-wise attention across levels. We observe that the top-level tends to capture the coarse and symmetric structures, such as wings and wheels, which are further decomposed into much finer granularity in the subsequent levels. We conjecture that this coarse-to-fine subset dependency helps the model to generate accurate structure in various granularity from global structure to local details.

Interestingly, we find that the hierarchical structure of SetVAE sometimes leads to the disentanglement of generative factors across layers. To demonstrate this, we train two classifiers in Set-MultiMNIST, one for digit class and the other for their positions. We train the classifiers using latent variables at each generator layer as an input, and measure the accuracy at each layer. In Figure 9, the latent variables at the lower layers tend to contribute more in locating the digits, while higher layers contribute to generating shape.

<sup>3</sup>For illustrative purposes, we present results from a selected head.

Table 3: Ablation study performed on Set-MultiMNIST dataset using FID scores for  $64 \times 64$  rendered images.

Model	FID( $\downarrow$ )
SetVAE (Ours)	<b>1047</b>
Non-hierarchical	1470
Unimodal prior	1252




Figure 10: Generation results from SetVAE trained on SUN-RGBD dataset. **Zoom-in** for a better view.

**Ablation study** Table 3 summarizes the ablation study of SetVAE (see the supplementary file for qualitative results and evaluation detail). We consider two baselines: Vanilla SetVAE using a global latent variable  $z^{(1)}$  (Section 3), and hierarchical SetVAE with unimodal prior (Section 4.2).

The Vanilla SetVAE performs much worse than our full model. We conjecture that a single latent variable is not expressive enough to encode complex variations in MultiMNIST, such as identity and position of multiple digits. We also find that multi-modal prior stabilizes the training of attention and guides the model to better local optima.

**Extension to categorical bounding boxes** SetVAE provides a solid exchangeability guarantee over sets, thus applicable to any set-structured data. To demonstrate this, we trained SetVAE on categorical bounding boxes in indoor scenes of the SUN-RGBD dataset [24]. As shown in Figure 10, SetVAE generates plausible layouts, modeling a complicated distribution of discrete semantic categories and continuous spatial instantiation of objects.

## 7. Conclusion

We introduced SetVAE, a novel hierarchical VAE for sets of varying cardinality. Introducing a novel bottleneck equivariant layer that learns subset representations, SetVAE performs hierarchical subset reasoning to encode and generate sets in a coarse-to-fine manner. As a result, SetVAE generates high-quality, diverse sets with reduced parameters. We also showed that SetVAE achieves cardinality disentanglement and generalization.

**Acknowledgement** This work was supported in part by Institute of Information & communications Technology Planning & Evaluation (IITP) grant (2020-0-00153, 2016-0-00464, 2019-0-00075), Samsung Electronics, HPC support by MSIT & NIPA.



## References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas J. Guibas. Learning representations and generative models for 3d point clouds. In *ICLR*, 2018.
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020.
- [3] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical report, Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015.
- [4] Harrison Edwards and Amos J. Storkey. Towards a neural statistician. In *ICLR*, 2017.
- [5] S. M. Ali Eslami, Nicolas Heess, Theophane Weber, Yuval Tassa, David Szepesvari, Koray Kavukcuoglu, and Geoffrey E. Hinton. Attend, infer, repeat: Fast scene understanding with generative models. In *NeurIPS*, 2016.
- [6] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.
- [7] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014.
- [8] Klaus Greff, Raphaël Lopez Kaufman, Rishabh Kabra, Nick Watters, Christopher Burgess, Daniel Zoran, Loic Matthey, Matthew Botvinick, and Alexander Lerchner. Multi-object representation learning with iterative variational inference. In *ICML*, 2019.
- [9] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017.
- [10] Seunghoon Hong, Dingdong Yang, Jongwook Choi, and Honglak Lee. Inferring semantic layout for hierarchical text-to-image synthesis. In *CVPR*, 2018.
- [11] Matthew J. Johnson, David Duvenaud, Alexander B. Wiltschko, Sandeep R. Datta, and Ryan P. Adams. Structured vae: Composing probabilistic graphical models and variational autoencoders. *CoRR*, 2018.
- [12] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019.
- [13] Hyeonju Kim, Hyeonseung Lee, Woo Hyun Kang, Joun Yeop Lee, and Nam Soo Kim. Softflow: Probabilistic framework for normalizing flow on manifolds. *CoRR*, 2020.
- [14] Diederik P. Kingma, Tim Salimans, and Max Welling. Improving variational inference with inverse autoregressive flow. In *NeurIPS*, 2016.
- [15] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014.
- [16] Adam R. Kosiorek, Hyunjik Kim, and Danilo J. Rezende. Conditional set generation with transformers. *CoRR*, 2020.
- [17] Juho Lee, Yoonho Lee, Jungtaek Kim, Adam R. Kosiorek, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In *ICML*, 2019.
- [18] Chun-Liang Li, Manzil Zaheer, Yang Zhang, Barnabás Póczos, and Ruslan Salakhutdinov. Point cloud GAN. In *ICLR*, 2019.
- [19] Manyi Li, Akshay Gadi Patil, Kai Xu, Siddhartha Chaudhuri, Owais Khan, Ariel Shamir, Changhe Tu, Baoquan Chen, Daniel Cohen-Or, and Hao (Richard) Zhang. GRAINS: generative recursive autoencoders for indoor scenes. In *ACM Trans. Graph.*, 2019.
- [20] Yang Li, Haidong Yi, Christopher M. Bender, Siyuan Shan, and Junier B. Oliva. Exchangeable neural ODE for set modeling. *NeurIPS*, 2020.
- [21] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. In *NeurIPS*, 2020.
- [22] Daniel Ritchie, Kai Wang, and Yu-An Lin. Fast and flexible indoor scene synthesis via deep convolutional generative models. In *CVPR*, 2019.
- [23] Casper Kaae Sønderby, Tapani Raiko, Lars Maaløe, Søren Kaae Sønderby, and Ole Winther. Ladder variational autoencoders. In *NeurIPS*, 2016.
- [24] Shuran Song, Samuel P. Lichtenberg, and Jianxiong Xiao. SUN RGB-D: A RGB-D scene understanding benchmark suite. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, 2015.
- [25] Karl Stelzner, Kristian Kersting, and Adam R. Kosiorek 2. Generative adversarial set transformers, 2020.
- [26] Arash Vahdat and Jan Kautz. NVAE: A deep hierarchical variational autoencoder. In *NeurIPS*, 2020.
- [27] Aäron van den Oord, Nal Kalchbrenner, Lasse Espeholt, Koray Kavukcuoglu, Oriol Vinyals, and Alex Graves. Conditional image generation with pixelcnn decoders. In *NeurIPS*, 2016.
- [28] Aäron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *ICML*, 2016.
- [29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [30] Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge J. Belongie, and Bharath Hariharan. Pointflow: 3d point cloud generation with continuous normalizing flows. In *ICCV*, 2019.
- [31] Mengjiao Yang, Bo Dai, Hanjun Dai, and Dale Schuurmans. Energy-based processes for exchangeable data. In *ICML*, 2020.
- [32] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabás Póczos, Ruslan Salakhutdinov, and Alexander J. Smola. Deep sets. In *NeurIPS*, 2017.
- [33] Yan Zhang, Jonathon S. Hare, and Adam Prügel-Bennett. Deep set prediction networks. In *NeurIPS*, 2019.

- [34] Yan Zhang, Jonathon S. Hare, and Adam Prügel-Bennett. Fspool: Learning set representations with featurewise sort pooling. In *ICLR*, 2020.