

Deep Message Passing on Sets

Yifeng Shi, Junier Oliva, Marc Niethammer

Department of Computer Science, UNC-Chapel Hill, USA
 {yifengs, joliva, mn}@cs.unc.edu

Abstract

Modern methods for learning over graph input data have shown the fruitfulness of accounting for relationships among elements in a collection. However, most methods that learn over set input data use only rudimentary approaches to exploit intra-collection relationships. In this work we introduce *Deep Message Passing on Sets* (DMPS), a novel method that incorporates relational learning for sets. DMPS not only connects learning on graphs with learning on sets via *deep kernel learning*, but it also bridges message passing on sets and traditional diffusion dynamics commonly used in denoising models. Based on these connections, we develop two new blocks for relational learning on sets: the *set-denoising block* and the *set-residual block*. The former is motivated by the connection between message passing on general graphs and diffusion-based denoising models, whereas the latter is inspired by the well-known residual network. In addition to demonstrating the interpretability of our model by learning the true underlying relational structure experimentally, we also show the effectiveness of our approach on both synthetic and real-world datasets by achieving results that are competitive with or outperform the state-of-the-art. For readers who are interested in the detailed derivations of several results that we present in this work, please see the supplementary material at: <https://arxiv.org/abs/1909.09877>.

Introduction

Significant effort in machine learning has been devoted to methods that operate over fixed-length, finite vectors. These methods ultimately perform some variant of a classic functional estimation task where one maps one fixed input vector $x \in \mathbb{R}^d$ to another fixed output vector $y \in \mathbb{R}^p$ via an estimated function $\hat{f} : \mathbb{R}^d \mapsto \mathbb{R}^p$. Notwithstanding the impressive progress of these approaches, the world we live in is filled with data that does not come neatly pre-packaged into fixed finite vectors. Instead, often data is observed and reasoned over in collections such as sets. For instance, when performing object detection on point clouds, one assigns a label (the object type) based on the underlying shape that is inferred collectively using all observed points (as opposed to labelling any one individual 3d point). In these, and many

other tasks, one seeks to assign an output response to an *entire* set of elements in which the elements can themselves be related to each other in a complex manner. Based on this need for analysis approaches that can operate on sets relationally, we develop a machine learning (ML) estimation technique that incorporates relational learning for set-structured data.

Set-structured data in its general form presents fundamental challenges to many existing machine learning methods. First, an appropriate method should be invariant to the order of the elements in the input set; i.e., different permutations should not influence the final response since the underlying instance is an *unordered* set. Second, a method should allow input sets of variable cardinalities; i.e., we should be able to associate a single label with a variable number of set elements. Both of these challenges render traditional approaches based on ordered inputs of fixed dimension (e.g., vectors or images of given sizes) like standard multilayer perceptrons (MLP) inapplicable. Some rectifications have been proposed to cope with these challenges without directly addressing them. For example, one can train a recurrent neural network (RNN) to adapt to variable-sized input sets. However, there is no guarantee that a network can learn to be permutation invariant (Vinyals, Bengio, and Kudlur 2015), especially in the context of input sets with large cardinality and small sample sizes. Hence, in this work we instead consider deep learning (DL) architectures that are specifically constructed to handle set-structured data.

Architectures that handle set-structured data exist. These techniques use, for instance, global-pooling operations (Qi et al. 2016) and intermediate equivariant mappings (Zaheer et al. 2017) to produce estimators that are invariant to permutations. While these architectures are asymptotically universal, they are notably limited in the intra-set dependencies they model. For instance, processing each set element independently using an MLP and aggregating features via max pooling to produce a set-level feature representation, as proposed in (Qi et al. 2016), may have difficulty capturing pairwise relations among the set elements. This indicates that the effectiveness of such methods may be limited for complex real-world data with finite samples. Hence, we propose to explicitly incorporate relational learning on elements of

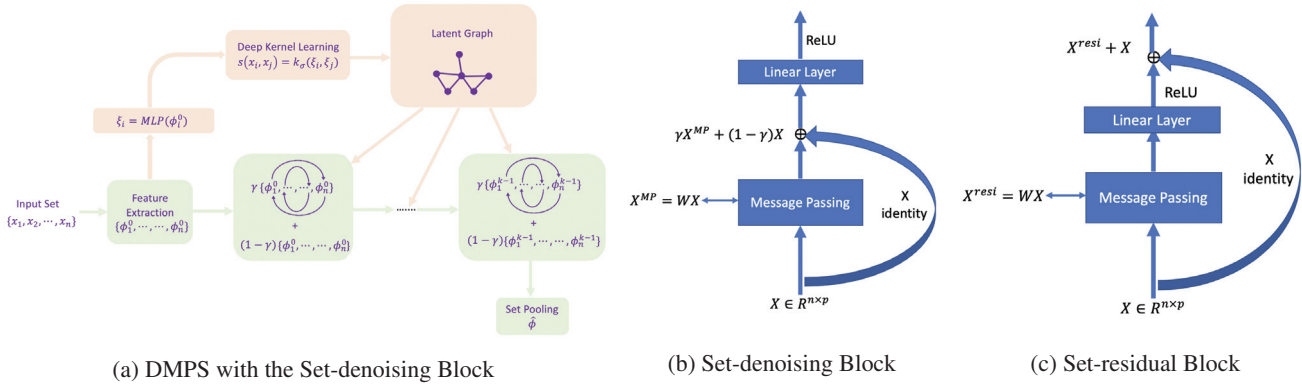


Figure 1: An overview of DMPS coupled with the set-denoising block as an illustrative example, and magnified views of the set-denoising block and the set-residual block. We briefly explain each figure here. (a): Given a set of objects, $\{x_1, x_2, \dots, x_n\}$, we first extract feature representations of the set elements, $\{\phi_1^0, \phi_2^0, \dots, \phi_n^0\}$, through, for example, a convolution neural network if the set elements are images. We estimate the underlying latent graph of the set using the extracted features via deep kernel learning (Wilson et al. 2015). We then apply a predefined number (k , in this case) of set-denoising blocks, which can be replaced by the message passing step or the set-residual block, to the extracted features in order to produce meaningful final feature representations that encode complex interactions among set elements. Lastly, we use a set pooling operation to generate a set-level feature representation for downstream tasks; (b): We first process the feature matrix, X , with a message passing step to produce X^{MP} , and then compute a weighted sum between X and X^{MP} with a learnable diffusion coefficient γ . Lastly we process the added matrix through a linear layer followed by a non-linear operator; (c): We still process the feature matrix, X , with a message passing step to produce X^{resi} . Instead of adding X^{resi} immediately back to X , we process X^{resi} with a linear layer followed by a non-linear operator first, and then add the resulting feature matrix to X .

an input set, a proven strategy for better learning (Santoro et al. 2017), by first constructing a latent graph that appropriately reflects the underlying relational structure of the input set, and then applying the message passing scheme on the constructed graph to leverage such relational information. Although not the first work that can be interpreted as considering relational information for modeling set-structured data through graphs (e.g. an alternative interpretation of Vaswani et al. (2017) given in Battaglia et al. (2018)), we explicitly represent the relational structure of input sets via learned latent graphs, and more importantly, place sensible resections on such graphs that are more consistent with the graph learning literature, facilitate the interpretability of our approach, and allow for the natural development and interpretation of our approach from a diffusion point of view.

Main Contributions In this paper, we further relational learning on sets with our framework *Deep Message Passing on Sets* (DMPS). Our main contributions are: 1) we unite learning on graphs with learning on sets through deep kernel learning, allowing for flexible relational learning on sets; 2) we develop two novel blocks, the *set-denoising block* and the *set-residual block*, to further facilitate learning interactions among set elements; 3) in addition to demonstrating the interpretability of our approach by successfully learning the true underlying relational structures, we show the effectiveness of our approach on both synthetic and real-world datasets by achieving results that are competitive with or outperform the state-of-the-art.

Background

This section introduces background material relevant for the development of our DMPS approach.

General Formulation of Valid Set Functions

Given an input set $X = \{x_1, x_2, \dots, x_n\}$, most permutation invariant functions that operate on sets studied in recent literature (Zaheer et al. 2017; Qi et al. 2016; Ilse, Tomczak, and Welling 2018) belong to the following class of functions:

$$\mathbb{F}(f) = \{f : f(X) = \kappa(\text{pool}\{\phi_X(S_1), \dots, \phi_X(S_m)\})\}, \quad (1)$$

where S_i for $1 \leq i \leq m$ is a subset of the power set of X , m is the number of such subsets being modeled, ϕ_X is a function that acts on the power set of X , pool is a permutation invariant pooling operation that produces a set-level representation for downstream tasks, and κ is another function that transforms the set-level representation into the output space of the model. Most methods in the literature differ in the choice of ϕ_X . To explicitly encode interactions among set elements, we will construct a function $\hat{\phi}_X$ that acts on a nontrivial subset of X using a message passing scheme.

Message Passing on Graphs

Representation learning on graphs is an active research area (Hamilton, Ying, and Leskovec 2017). We focus on the message passing scheme. Consider a graph $G = \{V, E\}$, where V is the set of vertices and E is the set of edges. We assume that each node v has a node feature h_v and each edge between two vertices, say v and w , has an edge feature e_{vw} . Message passing on graphs can be summarized in terms of

a *message passing phase* and a *readout phase* (Gilmer et al. 2017). In the message passing phase we update each h_v^t as

$$m_v^{t+1} = \sum_{w \in N(v)} M_t(h_v^t, h_w^t, e_{vw}) \text{ (propagate neighborhood)}$$

$$h_v^{t+1} = U_t(m_v^{t+1}, h_v^t) \text{ (update the feature vector)}$$

where M_t and U_t are the feature aggregation and update functions, respectively, $N(v)$ denotes the neighborhood of v , and t enumerates the message passing step. In the *readout phase*, we produce a graph-level feature representation from the node features for downstream tasks with \hat{f} as

$$\hat{f} = R(\{h(v)|v \in V\})$$

where R is the readout function. To ensure the message passing scheme is permutation invariant with respect to graph isomorphisms, an example of a trio of M_t , U_t , and R can be a function that computes a weighted average of the neighborhood features based on the edge weights, a concatenation operator, and a sum operator, respectively.

Deep Message Passing on Sets

We bridge learning on graphs and learning on sets with the message passing scheme, by first learning an underlying latent graph that represents the connectivity of the set elements, and then applying message passing to this latent graph to incorporate relational learning into learning on sets (see Fig. 1a for an illustration). In this way, DMPS leverages relational information to encode input sets in contrast to more traditional approaches (Zaheer et al. 2017; Qi et al. 2016), where each set element is either processed through some rudimentary equivariant transformations or in an independent manner.

Latent Graph Learning

Message passing on graphs is based on neighborhood structures and edge weights of the graphs. Our goal is to leverage relational information to encode elements in an input set via message passing, a natural way to capture the intra-dependencies among the elements if a graph that appropriately underpins the input set exists. However, unlike graph data, we generally do not know a-priori what neighbors a particular set element has and how strongly it is connected to these neighbors. Instead we need to infer such a graph structure from the set itself, ideally in an end-to-end fashion that is optimized jointly with the message passing scheme and a downstream task objective. To this end, we propose to learn edge weights among set elements, say x_i and x_j , via a similarity function, $e_{i,j} = s(x_i, x_j)$, where the similarity function itself is learned through the *deep kernel learning* scheme (Wilson et al. 2015). Specifically, deep kernel learning uses a shared multilayer perceptron network (MLP) to transform set elements into a feature space in which a kernel function is applied to the resulting feature representations, i.e. $s(x_i, x_j) = k_\sigma(\text{MLP}(x_i), \text{MLP}(x_j))$ where k_σ is a valid kernel function with an adaptive hyperparameter σ . Following this kernel strategy, we effectively use an infinite number of adaptive basis functions to estimate the similarities among set elements.

Message Passing on Sets

We simplify message passing on graphs to adapt to our setting:

Definition 1. Given a set $X = \{x_1, x_2, \dots, x_n\}$ where $x_i \in R^p$ for all i , we define message passing on sets as an iterative updating procedure that updates each set element as a weighted sum of the entire set $x_i \leftarrow \sum_{j=1}^n w_{i,j} x_j$, where $\sum_j w_{i,j} = 1, \forall i : 1 \leq i \leq n$, and $w_{i,j} \geq 0, \forall i, j : 1 \leq i, j \leq n$.

More compactly, we have $X^{t+1} = WX^t$ where t denotes the time step, $X^{t+1}, X^t \in R^{n \times p}$, and $W \in R_+^{n \times n}$ is an adaptive, row-normalized stochastic matrix constructed using the deep kernel learning scheme. More specifically, given an input set, we first construct the kernel matrix K where $K_{i,j} = s(x_i, x_j)$ using the learned similarity function. We then obtain W by applying the `softmax` operator to K to ensure the rows of K sum up to one, allowing us to interpret W as the weighted, row-normalized adjacency matrix of the underlying (fully-connected) latent graph. If one possesses prior knowledge about the set elements, choosing an appropriate threshold, say δ , and setting the entries of W that are smaller than δ to zero would result in a sparser graph. Furthermore, if one were to stack, for example, k message passing steps (i.e. $1 \leq t \leq k+1$), the estimated weight matrix, W , of the learned latent graph is shared across the stacked k steps. This is to say, although W is jointly estimated with other parameters, it is not step-specific, i.e., the underlying latent graph is assumed to be static. Stacking multiple message passing steps thus can be interpreted as propagating information from each set element’s multi-hop neighbors to encode higher-order information. As shown in our experiments, such learned W is capable of representing the relational structure of each input set in an intuitive way, allowing for a deeper understanding of the data in hand while leading to more effective learning results. Next, we develop some concepts that are needed to introduce the set-denoising block.

Diffusion on General Graphs

We present the update equation that explains message passing and motivates the set-denoising block from a diffusion point-of-view. We refer interested readers to the supplementary material for a more detailed treatment. We first define the discretized Dirichlet energy for graphs as

$$E(\{x_i\}) = \frac{C}{2} \sum_{(i,j) \in E} w_{i,j} \|x_i - x_j\|_2^2, \quad (2)$$

where x_i denotes the feature vector of node i , $w_{i,j}$ is the weight of the edge connecting nodes i and j , and C is a constant. It is crucial to note that, intuitively, performing message passing on graph moves the feature vectors of the nodes “closer” to each other, which in turn is equivalent to minimizing Eq. (2) (i.e. a weighted sum of feature vector differences measured by the L_2 norm). Differentiating Eq. (2) with respect to x_i , discretizing time with an Euler-forward

approximation, and rearranging the terms, we obtain

$$x_i^{t+1} = (1 - \delta t C \sum_j w_{i,j}) x_i^t + \delta t C \sum_j w_{i,j} x_j^t, \quad (3)$$

$$= (1 - \delta t C) x_i^t + \delta t C \sum_j w_{i,j} x_j^t. \quad (4)$$

where δt denotes the time step.

Hence, the update rule is a convex combination of the current node feature x_i^t and the update feature obtained by the message passing step. Subject to time-step constraints on δt , larger δt will result in a smoother solution. If we choose the constant C and the time step δt such that $C\delta t = 1$, we recover the message passing step. This observation offers us another interpretation of DMPS: message passing with the weight matrix W is equivalent to diffusing each set element based on the entire set, i.e. updating each set element by a weighted average of the entire set.

Set-denoising and Set-residual Blocks

Recall that beyond first-order relations, stacking a desired number of message passing steps allows us to take higher-order interactions among set elements into account. While enhancing the model’s capability of capturing complex relational signals, such stacking results in deeper networks, which can potentially cause problems in terms of over-smoothing the feature representations and other common difficulties in training deep networks. Based on the previously derived update equation (4) and the residual network (He et al. 2016), we propose to address both concerns by introducing the set-denoising block and the set-residual block. Before detailing the architectures of those two blocks, we elaborate on the intuitions behind them:

- It is well-known that a discretized diffusion process with Neumann boundary conditions on a graph converges to a steady state (i.e. all node features being the same) with enough time steps. In our setting, by stacking n message passing steps we effectively are “running” the discretized diffusion process on the latent graph over n time-steps. Although accounting for interactions among set elements is crucial, it is also indispensable to retain meaningful distinctions among them (e.g. all set elements sharing the same feature representation when the diffusion process has converged is obviously not ideal for learning). It is thus important to avoid over-smoothing when stacking message passing steps, a critical observation that is attested by one of our experimental studies.
- Ideally, the architectures of either the set-denoising block or the set-residual block should alleviate the concerns of vanishing gradient or difficulty of learning the identity mapping when building a deep network.

Now we formally introduce the two proposed set blocks

- **Set-denoising Block:** As depicted in Fig. 1b, we first apply message passing to the feature matrix X , resulting in X^{MP} . The original feature matrix X is then combined with X^{MP} via a convex combination. This convex combination is based on a learnable diffusion coefficient

Algorithm 1: Deep Message Passing on Sets with the Set-denoising Block

Result: Final label prediction \hat{y} ;

Input: $\{x_1, x_2, \dots, x_n\}$: a set of objects; k : the number of set-denoising blocks desired;

Learnable Parameters: $H_t, \forall t : 1 \leq t \leq k$: parameters of the linear layers; γ : the diffusion coefficient;

Initialization: $t = 1$;

—Extract p -dimensional numerical feature of each set element x_i and form a feature matrix $X \in R^{n \times p}$;

—Construct the weight matrix W from the feature matrix X using deep kernel learning;

while $t \leq k$ **do**

$$X^{t+1} = (1 - \gamma)X^t + \gamma W X^t;$$

$$X^{t+1} = \tau(X^{t+1} H_t);$$

$$t = t + 1;$$

end

$$x = \text{pool}(X^k);$$

$$\hat{y} = \kappa(x)$$

$\gamma \in (0, 1)$, which corresponds to $\delta t C$ in the discretized anisotropic heat equation (4). In other words, we do not explicitly choose the constant C and the step size δt in (4); instead we either fix γ upfront or jointly learn it with other model parameters. It is worth noting that if γ is learnable, we effectively allow our model to adaptively adjust the optimal degree of smoothing. We choose γ to be the same for each block, though making γ block-specific is possible. The linear layer followed by a non-linear operator at the end serve to further increase the expressiveness of the learned features.

- **Set-residual Block:** Fig. 1c is directly motivated by the residual network (He et al. 2016): assuming there exists an optimal feature matrix for learning, X^{optim} , it might be easier for the network to learn the difference, $X^{\text{optim}} - X$, through message passing as opposed to learn the optimal feature matrix X^{optim} from scratch. Moreover, the architecture of the set-residual block alleviates some common problems that come with training a deep network (He et al. 2016).

Deep Message Passing on Sets (DMPS)

As a concrete example, Algorithm 1 outlines *Deep Message Passing on Sets* (DMPS) coupled with the set-denoising block, where τ is an element-wise non-linear operator, pool is a set pooling operator, and κ is a function (e.g a fully-connected layer) that transforms the set-level feature representation into the output space.

Analysis

Notice that with or without the set-denoising/set-residual blocks, the message passing step in DMPS is permutation equivariant with respect to the set elements. Since the composition of a permutation equivariant operation and a valid set pooling operation is permutation invariant, we have the following proposition

Proposition 2. *DMPS is permutation invariant to the order of the elements in the input set.*

Additionally, being able to approximate any valid set function is desirable for a model. Building upon the results in Zaheer et al. (2017) and Qi et al. (2016), we have

Proposition 3. *DMPS is an universal approximator for any permutation invariant function.*

Proof. See supplementary material. \square

Related Work

Relational and Non-local Reasoning

- **Relational Reasoning:** Our proposed approach fits and generalizes the relation network (Santoro et al. 2017). To elaborate, if we do not perform any feature transformation after message passing and define the set pooling operator as the sum operator, a model comprised of one message passing step would result in $f(X) = \kappa \left(\sum_{i,j} w_{i,j} x_j^T H_1 \right)$ where x_j^T is the j -th row of the feature matrix X . This agrees with the form suggested in Santoro et al. (2017) if we let $g_\theta(x_i, x_j) = w_{i,j} x_j^T H_1$. Therefore, our proposed network can be seen as a generalization of the relation network in the sense that stacking multiple message passing steps or set blocks enables us to learn high-order relations among the set elements.
- **Non-local Networks:** While non-local relational reasoning, whose main objective is to extend the frameworks of local learning schemes like the convolution operator or the fully-connected layer to allow for non-local learning, has been proposed in earlier work (Wang et al. 2017), DMPS is an extension of that to set-structured data where our goal is to leverage non-trivial relations among set elements. We also proposed a new way of learning the similarity function, $s(x_i, x_j)$, that infers pairwise relations, namely *deep kernel learning*. Last but not least, in order to further strengthen the flexibility of our model, we coupled DMPS with the set-denoising and set-residual blocks, motivated by the diffusion dynamics and the residual network, respectively, whereas Wang et al. (2017) only considered residual connections.

Learning Deep Networks

The network proposed by He et al. (2016) allows for sensible training of deep networks, an advantage that is inherited by the set-residual block. While the set-denoising block shares certain structural similarities with the highway network (Srivastava, Greff, and Schmidhuber 2015), it was introduced in a different context, namely through the lens of diffusion on graphs and the avoidance of over-smoothing. In retrospect, the connection between diffusion dynamics and the set-denoising block established in this work provides another useful interpretation of the highway network in terms of relational learning. Last but not least, the set-denoising and set-residual blocks are introduced specifically to handle set-structured data, an attribute that is beyond the scope of the residual network or the highway network.

Learning on Latent Graphs of Sets

We now refer back to Eqn. 1. A special case of ϕ_X , as it can be seen in Qi et al. (2016) for example, acts on set elements independently. Notwithstanding, this can still be considered a trivial form of message passing on a latent graph whose edge set is empty, including the approaches in Qi et al. (2016) and Zaheer et al. (2017) as special cases. Franceschi et al. (2019) considers the challenge of jointly learning a probabilistic graph structure (to represent the relational structure of the given data points) and the parameters of the model in a semi-supervised setting in which a prior relational structure is corrupted or completely missing, by casting the two estimation problems in terms of a bilevel programming optimization task. The set transformer (Lee et al. 2019), by directly applying the transformer (Vaswani et al. 2017) on set-structured data, essentially uses a different W for every message passing step, and there are effectively many different W 's at each step, since each attention head uses different keys/values. Unlike the transformer, DMPS places sensible restrictions on W that are well-motivated: they are based on physics (diffusion dynamics) and graph learning (static latent graph). Furthermore, it is worth emphasizing that DMPS only learns one latent graph (i.e. a single W as opposed to multiple W 's). Such design choice is critical because not only it allows for more intuitive and less ambiguous analysis/visualization, and leads to fewer free parameters and potentially better generalization given limited data, but it also suits data for which one latent graph suffices for the purpose of capturing the underlying relational structure better, explaining the experimental advantage of DMPS over the set transformer in the experiments that we conduct.

Experiments

We apply DMPS and its extensions to a range of synthetic and real-world datasets. For each experiment, we compare our methods against, to the best of our knowledge, the state-of-the-art results for that dataset. Unless otherwise specified, three message passing steps, set-denoising blocks, or set-residual blocks are stacked to form the final model. Furthermore, we adopt the following abbreviations in this section: ET—equivariant transformation (Zaheer et al. 2017); V-DMPS—vanilla DMPS, i.e the block component being the message passing step; R-DMPS—the block component being the set-residual block; D-DMPS w/(FDC or LDC)—the block component being the set-denoising block with (fixed or learnable, respectively) diffusion coefficient; and UG—uniform graph, i.e a fully-connected, undirected graph with equal weights for all edges.

Classifying Gaussian Sets

In this experiment we investigate a traditional problem of classifying random samples drawn from two different multivariate Gaussian distributions with the same mean and different covariance matrices. We create sets of real numbers by drawing the set elements as vector random samples from one of the two Gaussian distributions. The latent graph underlying each set is thus determined by the covariance matrix of its corresponding Gaussian distribution. To use the




Figure 2: Fig. (a) and (b) show the learned kernel matrices with input sampled from the respective two normal distributions at test time, and Fig. (c) depicts how the testing accuracy varies with ρ .

covariance matrix as the ground truth for the latent graph, our main goal is to test DMPS’s ability to capture and recover the true relational structure that underpins each set. Given a set $\mathbf{X} = [X_1, X_2, \dots, X_p]^T$ where $\mathbf{X} \sim N(\mu, \Sigma)$, it is worth emphasizing that the order among the elements in \mathbf{X} does not matter, as \mathbf{X} is permutation equivariant with respect to its mean and covariance. The canonical order is used here for convenience, although the result would be the same if one were to permute the order upfront. We next describe the experiment in more detail.

We sample input sets from two 5-dimensional Gaussian distributions $N(\mathbf{0}, \mathbf{I})$ and $N(\mathbf{0}, \Sigma)$. To further test DMPS’s capability to apprehend sparse relational signals, we choose Σ to be the same as the identity matrix except at a (randomly) chosen pair of indices $((2, 4),$ in this case) in which $\Sigma_{2,4} = \Sigma_{4,2} = \rho$ and $\rho \in [0, 1)$. This is to say, the only, yet subtle difference between sets drawn from those two distributions is that the values of the second and the fourth elements are positively correlated for sets drawn from $N(\mathbf{0}, \Sigma)$, a relational information that can only be captured if the elements in the sets are modeled interactively. Fig. 2a and 2b showcase the latent graphs recovered by DMPS at test time with input sets sampled from the two chosen distributions, respectively, and with $\rho = 0.95$, while Fig. 2c conveys how the test results vary with different choices of ρ .

We have shown the following advantages of DMPS through this experiment: a). DMPS is able to take advantage of non-trivial covariance relations, thus outperforming methods that do not explicitly take such information into account. Furthermore, the trend of the curve in Fig. 2c confirms the intuition that DMPS performs better when the underlying relational signal gets stronger; b). in contrast to other relational learning methods that focus on sets like Lee et al. (2019), DMPS is intuitively interpretable in that the kernel matrix learned through the *latent graph learning* recovers the covariance structure of the underlying Gaussian distribution. This property of DMPS is highly desirable, and is also consistent with theoretical probability as functional transformations of random variables preserve the independence and covariance relations among those variables (the network acts on each set element, i.e each dimension of the Gaussian random sample, independently before the message passing step).

Counting Unique Characters




Figure 3: Learned latent graph for a test input set with nine images categorized by three unique characters. The color transparency of the edges is proportional to their learned weights, with heavier-colored edges carrying larger weights.

To test the model’s ability to model set-structured data relationally, Lee et al. (2019) proposed the task of counting unique characters using the characters dataset (Lake, Salakhutdinov, and Tenenbaum 2015), where the goal is to predict the number of unique characters in an input set of character images. Please refer to the supplementary material for detailed experimental setup. Tab. 1 shows the testing results. We emphasize that we align as much architectural choices, such as learning rate, number of training batches, batch size, etc., as we can with Lee et al. (2019) for fair comparison. We make some additional comments below.

Table 1: Counting unique characters

Architecture	Test Accuracy
DeepSets w/ Mean ET	0.4617 ± 0.0076
DeepSets w/ Max ET	0.4359 ± 0.0077
Set Transformer	0.6037 ± 0.0075
V-DMPS w/ UG	0.1357 ± 0.0000
D-DMPS w/ LDC & UG	0.4661 ± 0.0085
V-DMPS	0.6446 ± 0.0174
R-DMPS	0.6600 ± 0.0103
D-DMPS w/ FDC	0.6748 ± 0.0120
D-DMPS w/ LDC	0.6674 ± 0.0080

Firstly, DMPS and its variants outperform other methods by significant margins, showing the effectiveness of our proposed model. Secondly, since relational learning is the

centerpiece of our model, we perform two ablation studies (V-DMPS w/ UG and D-DMPS w/ LDC & UG) where we perform message passing on the uniform graph instead of the learned latent graph. The idea is that regular DMPS and its variants should outperform models with a fixed uniform graph if the learned latent graph indeed captures useful relational information. As shown in Tab. 1, V-DMPS w/ UG performs poorly, while D-DMPS w/ LDC & UG does better (perhaps because of the less erroneous smoothing induced by the set-denoising block) it still does not perform as well. This shows the significance of proper relational learning provided by regular DMPS and its variants. To further demonstrate the interpretability of our approach, Fig. 3 shows the learned latent graph for a test input set with nine images categorized by three unique characters. We see that the learned latent graph in which three clusters emerge delineates the relations among the set elements in a reasonable manner, underscoring a straightforward intuition that images corresponding to the same character are more closely related.

Point Cloud Classification

Table 2: ModelNet 40 Classification Task

Architecture	100 points	1000 points
DeepSets w/ Max ET	0.82 ± 0.02	0.87 ± 0.01
Set Transformer	0.8454 ± 0.0144	0.8915 ± 0.0144
PointNet++	—	$0.907 \pm —$
V-DMPS	0.8367 ± 0.0047	0.8751 ± 0.0029
R-DMPS	0.8475 ± 0.0036	0.8935 ± 0.0016
D-DMPS w/ FDC	0.8564 ± 0.0031	0.8783 ± 0.0032
D-DMPS w/ LDC	0.8571 ± 0.0062	0.8798 ± 0.0020

We apply DMPS and its variants to the ModelNet40 dataset (Chang et al. 2015), which contains objects that are represented as sets of 3d points (point clouds) in 40 different categories. Our model allows us to directly model sets of 3d points. For this experiment, we construct input sets with sizes 100 and 1,000 points per set by uniformly sampling from the mesh representations of the objects. Tab. 2 shows the test performances of our approaches compared with other state-of-the-art methods that directly operate on raw point clouds. We make some additional comments below.




Figure 4: Depiction of how the diffusion coefficient γ affects the test accuracy in the case of 100 points per set.

Firstly, we point out that the task is harder when there

are fewer points in the input set, thus requiring more efficient relational learning. We observe that D-DMPS w/ LDC outperforms other methods by significant margins in the case of 100 points in the input set. As for the case of 1000 points, PointNet++ achieved the state-of-the-art among methods that directly process raw point clouds. Our method performs on par with the set transformer (Lee et al. 2019), and outperforms deep sets (Zaheer et al. 2017) by leveraging relational information. Secondly, to investigate the importance of balancing between appropriate message passing and over-smoothing, we perform a study in which we fix the diffusion coefficient to various values and see how the test accuracy varies. Fig. 4 shows the result. As γ ranges from 0 to 1, the model effectively ranges from DeepSets w/o ET to V-DMPS, with anywhere in-between being D-DMPS w/ FDC at that particular γ . The test accuracy peaks when γ approaches 0.5, and decreases when γ becomes either too large or too small. This shows the significance of controlling the degree of smoothing. Along with the update equation, the novelty of the set-denoising block is affirmed theoretically and empirically.

Histopathology Dataset

Table 3: Breast Cancer

Architecture	Test Accuracy
Attention	0.745 ± 0.018
Gated Attention	0.755 ± 0.016
V-DMPS	0.800 ± 0.023
R-DMPS	0.818 ± 0.029
D-DMPS w/ FDC	0.846 ± 0.019
D-DMPS w/ LDC	0.836 ± 0.023

The concept of learning on sets also applies well to *weakly-labeled data*. In this section we perform experiment on classifying weakly-labeled real-life histopathology images provided in the breast cancer dataset (Gelasca et al. 2008). A common approach is to divide an image into smaller patches and think of the patches as a set of "small images" with a single label for the set.

The breast cancer dataset introduced in Gelasca et al. (2008) consists of 58 weakly-labeled 896×768 H&E images. An image is labeled malignant if it contains breast cancer cells; otherwise it is labeled benign. We follow a similar procedure to pre-process the images as in Ilse, Tomczak, and Welling (2018). We divide the images into 32×32 patches, which results in 672 patches per set (i.e., per image). Furthermore, because of the small number of available images, we perform data augmentation at the training stage by randomly rotating and mirroring the patches. We point out that Ilse, Tomczak, and Welling (2018) also randomly adjusted the amount of H&E by decomposing the RGB color of the tissue into the H&E color space. We compare the performances of DMPS and its variants to the *attention* and *gated attention* models introduced in Ilse, Tomczak, and Welling (2018). The testing results are shown in Tab. 3. Despite the framework of multiple instance learning, and

thus the attention scheme, being particularly suitable to computational histopathology (Kandemir and Hamprecht 2015; Ilse, Tomczak, and Welling 2018), we see that DMPS and its variants perform uniformly better by significant margins.

Conclusion and Future Work

In this paper, we introduced DMPS, a set-learning scheme that explicitly takes interactions among set elements into account when modeling set-structured data. We also proposed two variants of message passing on sets, the set-denoising block and the set-residual block. Although this is a step towards relational learning on sets, there are many possible extensions. For example, the message passing scheme can be interpreted as a gradient descent step based on the weighted Dirichlet integral with the functional two-norm. Would one, for example, discretize an energy of the form $\int w(x) \|\nabla u(x)\|_2 dx$ instead, we would obtain a form of weighted total-variation message passing. Hence, one interesting future work would be to derive a family of message passing algorithms by changing the functional two-norm to the more general functional p-norm and to explore the behavior of the resulting message passing schemes.

References

- Battaglia, P. W.; Hamrick, J. B.; Bapst, V.; Sanchez-Gonzalez, A.; Zambaldi, V. F.; Malinowski, M.; Tacchetti, A.; Raposo, D.; Santoro, A.; Faulkner, R.; Çaglar Gülçehre; Song, H. F.; Ballard, A. J.; Gilmer, J.; Dahl, G. E.; Vaswani, A.; Allen, K. R.; Nash, C.; Langston, V.; Dyer, C.; Heess, N. M. O.; Wierstra, D.; Kohli, P.; Botvinick, M. M.; Vinyals, O.; Li, Y.; and Pascanu, R. 2018. Relational inductive biases, deep learning, and graph networks. *ArXiv abs/1806.01261*.
- Chang, A. X.; Funkhouser, T. A.; Guibas, L. J.; Hanrahan, P.; Huang, Q.-X.; Li, Z.; Savarese, S.; Savva, M.; Song, S.; Su, H.; Xiao, J.; Yi, L.; and Yu, F. 2015. Shapenet: An information-rich 3d model repository. *ArXiv abs/1512.03012*.
- Franceschi, L.; Niepert, M.; Pontil, M.; and He, X. 2019. Learning discrete structures for graph neural networks. In *ICML*.
- Gelasca, E. D.; Byun, J.; Obara, B.; and Manjunath, B. S. 2008. Evaluation and benchmark for biological image segmentation. *2008 15th IEEE International Conference on Image Processing* 1816–1819.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural message passing for quantum chemistry. In *ICML*.
- Hamilton, W. L.; Ying, R.; and Leskovec, J. 2017. Representation learning on graphs: Methods and applications. *IEEE Data Eng. Bull.* 40:52–74.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* 770–778.
- Ilse, M.; Tomczak, J. M.; and Welling, M. 2018. Attention-based deep multiple instance learning. In *ICML*.
- Kandemir, M., and Hamprecht, F. A. 2015. Computer-aided diagnosis from weak supervision: A benchmarking study. *Computerized medical imaging and graphics : the official journal of the Computerized Medical Imaging Society* 42:44–50.
- Lake, B. M.; Salakhutdinov, R.; and Tenenbaum, J. B. 2015. Human-level concept learning through probabilistic program induction. *Science* 350:1332–1338.
- Lee, J.; Lee, Y.; Kim, J.; Kosiorek, A. R.; Choi, S.; and Teh, Y. W. 2019. Set transformer: A framework for attention-based permutation-invariant neural networks. In *ICML*.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2016. Pointnet: Deep learning on point sets for 3d classification and segmentation. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* 77–85.
- Santoro, A.; Raposo, D.; Barrett, D. G. T.; Malinowski, M.; Pascanu, R.; Battaglia, P. W.; and Lillicrap, T. P. 2017. A simple neural network module for relational reasoning. In *NIPS*.
- Srivastava, R. K.; Greff, K.; and Schmidhuber, J. 2015. Highway networks. *ArXiv abs/1505.00387*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need. In *NIPS*.
- Vinyals, O.; Bengio, S.; and Kudlur, M. 2015. Order matters: Sequence to sequence for sets. *CoRR abs/1511.06391*.
- Wang, X.; Girshick, R. B.; Gupta, A.; and He, K. 2017. Non-local neural networks. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* 7794–7803.
- Wilson, A. G.; Hu, Z.; Salakhutdinov, R.; and Xing, E. P. 2015. Deep kernel learning. In *AISTATS*.
- Zaheer, M.; Kottur, S.; Ravanbakhsh, S.; Póczos, B.; Salakhutdinov, R.; and Smola, A. J. 2017. Deep sets. In *NIPS*.