
Scalable Normalizing Flows for Permutation Invariant Densities

Marin Bilos¹ Stephan Günnemann¹

Abstract

Modeling sets is an important problem in machine learning since this type of data can be found in many domains. A promising approach defines a family of permutation invariant densities with continuous normalizing flows. This allows us to maximize the likelihood directly and sample new realizations with ease. In this work, we demonstrate how calculating the trace, a crucial step in this method, raises issues that occur both during training and inference, limiting its practicality. We propose an alternative way of defining permutation equivariant transformations that give closed form trace. This leads not only to improvements while training, but also to better final performance. We demonstrate the benefits of our approach on point processes and general set modeling.

1. Introduction

We say that data is exchangeable if it lacks ordering, i.e., any permutation of random variables has the same probability. Examples include point clouds, items in a shopping cart, tracking household electricity consumption in a city etc. Often, the number of elements is also a random variable, like in locations of cellular stations or trees in a forest. We wish to find a generative model that best describes such data. To do that, we need to specify a symmetric density — one that is invariant to data permutations.

A powerful framework to model densities are normalizing flows (Kobyzev et al., 2020). They parametrize a transformation from a simple base distribution into a complex one. To define a permutation invariant density, it is enough to have a permutation invariant base distribution, and a permutation equivariant transformation (Papamakarios et al., 2019). The latter is impossible to obtain with most of the traditional architectures (Dinh et al., 2017; Kingma et al., 2016), unless we assume the points are independent of each other.

¹Technical University of Munich, Germany. Correspondence to: <bilos@in.tum.de>.




Figure 1. Illustration of our approach. Transforming random points from the permutation invariant base distribution with an ODE gives a realization in the target distribution. The opposite direction calculates the likelihood of the observed sample. The dynamics f has to be a permutation equivariant neural network.

However, the i.i.d. assumption is a strong limitation because the presence of one object can influence the distribution of the others. For example, short trees grow near each other, but are inhibited by taller trees (Ogata & Tanemura, 1985). We want to find a way to model invariant densities that capture this behavior. A special class of flows, called continuous normalizing flows (Chen et al., 2018), support unrestricted transformations that allow us to directly use permutation equivariant layers (Zaheer et al., 2017; Lee et al., 2019). The general approach is illustrated in Figure 1.

The main disadvantage of continuous flows is the requirement to calculate the trace of the Jacobian of the transformation. In practice, this is done using an unbiased estimator. In this work, we show that this introduces additional noise during training, making it unstable. Further, for inference and in some downstream tasks (e.g., out-of-distribution detection), we would still have to use costly exact trace computation.

To combat this, we propose a straightforward decoupling of permutation equivariant functions to obtain a constant trace, i.e., we construct a volume preserving flow. We further extend this to have exact trace computation while keeping the original expressiveness. In the experiments we show that our method speeds-up the training, makes it more stable, provides better results, and scales with big datasets.

Additionally, we introduce novel models for point processes that account for complex interactions between the points and have closed form likelihood with easy sampling. Previous work resorted to less expressive models or using pseudo-likelihood (Besag, 1975). We tackle other use cases as well, such as modeling point clouds.

2. Background

The realizations of a finite point process on a bounded region $\mathcal{B} \subset \mathbb{R}^d$ are finite sets of points $X = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, $\mathbf{x}_i \in \mathcal{B}$. To fully specify a point process we need two distributions: the distribution of the number of points $p(n)$, and the permutation invariant probability density $\tilde{p}(\mathbf{x}_1, \dots, \mathbf{x}_n)$ of their locations. The probability of seeing n points at locations X is (Daley & Vere-Jones, 2007):

$$p(X) = n!p(n)\tilde{p}(\mathbf{x}_1, \dots, \mathbf{x}_n). \quad (1)$$

Traditionally, point processes are defined with an intensity function $\lambda(\cdot)$ that measures how many points we expect to see on some subset. This is equivalent to what we described in Eq. 1, therefore, we can safely use our parametrization without losing generality.

Perhaps the simplest non-trivial point process model is an inhomogeneous Poisson process. The number of points n follows a Poisson distribution and the locations of the points are assumed to be generated i.i.d. from some density $\tilde{p}(X) = \prod_i p(\mathbf{x}_i)$ (Chiu et al., 2013). To generate a new realization, we first sample $n \sim \text{Poisson}(\lambda)$, and then sample n points $\mathbf{x} \sim p(\mathbf{x})$. However, IHP alone cannot capture more complex permutation invariant densities, with interactions between the points.

A solution candidate are normalizing flows since they define complex distributions with invertible smooth transformations of the initial random variable. That means, if we apply a function $f : \mathcal{B}^d \rightarrow \mathcal{B}^d$ to a random variable $\mathbf{z} \sim q(\mathbf{z})$, where f is invertible and differentiable, we can get the density of $\mathbf{x} = f(\mathbf{z})$ by calculating the change of variables formula: $p(\mathbf{x}) = q(\mathbf{z}) |\det J_f(\mathbf{z})|^{-1}$, where J is the Jacobian of f . The main challenge is defining a function that can be efficiently inverted and whose determinant is not prohibitively expensive to calculate. Thus, most approaches design transformations that have a triangular Jacobian, using coupling (Dinh et al., 2017) or autoregressive transformations (Kingma et al., 2016). This, however, permits us from using permutation equivariant functions since having rich interactions between the elements requires a *full* Jacobian.

Another way to transform an initial sample \mathbf{z} is with an ordinary differential equation $f(\mathbf{z}(t), t) = \partial \mathbf{z}(t) / \partial t$. The instantaneous change in log-density is (Chen et al., 2018):

$$\frac{\partial \log p(\mathbf{z}(t))}{\partial t} = -\text{Tr} \left(\frac{\partial f}{\partial \mathbf{z}(t)} \right), \quad (2)$$

and the final log-density is obtained by integrating across time, e.g., using a black-box ODE solver:

$$\log p(\mathbf{x}) = \log q(\mathbf{z}(t_0)) - \int_{t_0}^{t_1} \text{Tr} \left(\frac{\partial f}{\partial \mathbf{z}(t)} \right) dt. \quad (3)$$

It is easy to show that if f is permutation equivariant and the base distribution q is permutation invariant, the final

density p is also permutation invariant (see Supplementary Material). The only other constraint is that f , along with its derivatives, needs to be Lipschitz continuous which is satisfied by using smooth activations in a neural network. The benefit of this approach is that we are restricting f only mildly, meaning the dimensions can interact with each other arbitrarily. We utilize this to model interactions between the points in a point process.

3. Model

A general way to define an equivariant map $f : \mathcal{B}^n \rightarrow \mathcal{B}^n$ is to define an invariant function $g : \mathcal{B}^n \rightarrow \mathbb{R}$, and let $f(X) = \nabla_X g(X)$ (Papamakarios et al., 2019). However, this is computationally costly and numerically unstable (Köhler et al., 2020). We can instead construct an equivariant (deep set) layer directly (Zaheer et al., 2017; Maron et al., 2020):

$$f(X)_i = g(\mathbf{x}_i) + \sum_{j \neq i} h(\mathbf{x}_j), \quad (4)$$

where $g, h : \mathcal{B} \rightarrow \mathcal{B}$ are neural networks.

An alternative equivariant mapping is a multihead self-attention (Vaswani et al., 2017). In general, the attention layer is defined for matrices $Q \in \mathbb{R}^{n \times d_k}$, $K \in \mathbb{R}^{n \times d_k}$ and $V \in \mathbb{R}^{n \times d_v}$ as follows:

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V. \quad (5)$$

Our transformation is then:

$$f(X)_i = \text{Attention}(f_q(X), f_k(X), f_v(X))_i \quad (6)$$

where we treat X as a stacked matrix of \mathbf{x}_i set elements. Functions f_q, f_k and f_v are permutation invariant neural networks, i.e., transform each point independently.

3.1. Efficient trace calculation

Computing the trace in Eq. 2 during training is expensive since it scales quadratically $\mathcal{O}(n^2 d^2)$ with the number of points and the dimension of \mathcal{B} . Grathwohl et al. (2019) propose using Hutchinson’s trace estimator which reduces the cost to $\mathcal{O}(nd)$. However, this introduces new problems: 1) the estimator brings noise into training making it unstable; and 2) during inference, we still need to evaluate the exact trace to get the correct density. In the following we introduce a way to have closed form trace in linear time $\mathcal{O}(nd)$, while retaining the expressiveness of layers in Eq. 4 and 6. We do that by decoupling the original transformation into different nested parts, that give us exact trace computation (Figure 2).

Interactions between points. If we take a closer look at Eq. 4, we can see that it constitutes of two decoupled functions, one is acting on the current element: $g(\mathbf{x}_i)$, and the

second corresponds to the contribution from all the other elements: $\sum_{j \neq i} h(\mathbf{x}_j)$. If we ignore the first part, the trace will be always equal to 0, meaning we do not have to evaluate Eq. 2 since it is constant. This gives us a volume-preserving normalizing flow, similar to (Dinh et al., 2015; Toth et al., 2020). Here, every point \mathbf{x}_i is transformed based on all the other points $\mathbf{x}_{j \neq i}$. This is illustrated in Figure 2 as a *between points* transformation.

In practice, we efficiently implement $\mathbf{h}_i = \sum_{j \neq i} h(\mathbf{x}_j)$ for the i th element by calculating the sum of all $h(\mathbf{x}_j)$ and then subtract $h(\mathbf{x}_i)$. If instead of the sum we use mean aggregation, we divide everything by $n - 1$. We get a constant trace for the operation of finding the maximum value by returning the largest value for all elements that are not the actual maximum, i.e., with the usual max-pooling. At the true maximum we return the second largest value.

Interactions within point. The next question is if we can keep the function g inside Eq. 4 while retaining the fixed trace property. To do so, we apply the same trick as before — separate the part that produces the constant trace from the rest. This is visualized in Figure 2 as a *within point* transformation. Each dimension in one element \mathbf{x}_i should be transformed using other dimensions but should not depend on itself. Since g does not carry any invariance constraints, we can use arbitrary neural network layers. We implement the transformation g by masking the weights, in particular, we use the existing MADE architecture (Germain et al., 2015). Such masked networks ensure that there is no computation path from the input dimension \mathbf{x}_{ij} to the output $g(\mathbf{x}_i)_j$. The resulting Jacobian of g always has zero values on the diagonal. The output dimension can be larger than the input, with groups of dimensions that are independent of \mathbf{x}_{ij} . This allows us to build multilayer networks and keeps the constant trace requirement fulfilled.

Per dimension transformation. So far we got a transformation with a Jacobian matrix that has zeros on the diagonal — a volume preserving flow. Finally, to get the full Jacobian, the last remaining part is the *per dimension* transformation. In the simplest case, this is an $\mathbb{R} \rightarrow \mathbb{R}$ mapping, and can be easily calculated in closed form. A more expressive function, one that recovers g fully, is $g(\mathbf{x}_{ij}) = \tau(\mathbf{x}_{ij}, \mathbf{g}_i)$, where $\tau: \mathbb{R}^{d_g+1} \rightarrow \mathbb{R}$ and \mathbf{g}_i does not depend on \mathbf{x}_{ij} , i.e., it is a result of masked layers from above. This kind of decoupling was first introduced by Chen & Duvenaud (2019). By exploiting the properties of modern deep learning frameworks, the trace can now be calculated cheaply. We can extend this idea by also including the result of the between point interactions \mathbf{h}_i .

The complete transformation of a single dimension of one set element \mathbf{x}_i is $f(\mathbf{x}_{ij}) = \tau(\mathbf{x}_{ij}, \mathbf{g}_i, \mathbf{h}_i)$. The function τ transforms the j th dimension of the i th element, conditioned on the within point interactions \mathbf{g}_i , and the between points

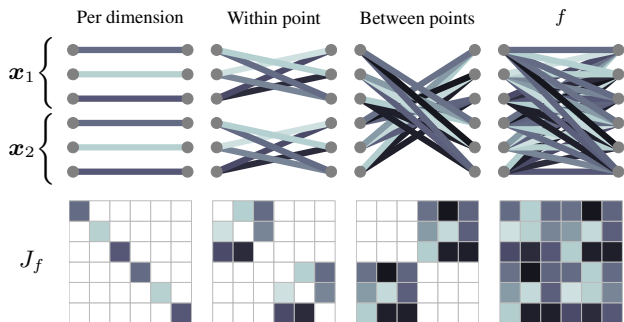


Figure 2. Illustration of different interactions and the corresponding Jacobian decoupling for equivariant f and inputs $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^3$.

interactions \mathbf{h}_i . To capture the original function g exactly, the vector \mathbf{g}_i has to have a dimension of $d - 1$; however, in practice, smaller dimensions can be used with good results and better efficiency. Function τ can be any neural network. We can detach the zero trace networks that produce \mathbf{g}_i and \mathbf{h}_i completely from the computation graph and just calculate $\partial \tau_{ij} / \partial x_{ij}$ in parallel to obtain the exact trace. When backpropagating we need to attach the networks back to allow learning of the parameters. More details can be found in Chen & Duvenaud (2019).

Alternatively, we can keep the original function g , and use stochastic trace estimation. That is, we only use the constant trace for between points interaction, while g is unrestricted. The Jacobian produced by g is a block diagonal matrix (see Figure 2). Therefore, we can estimate the trace as a sum of traces in these $d \times d$ blocks. This results in the reduced variance of the estimator. Chen et al. (2021) use a similar approach for the attention-based models, in which they detach the gradient connections between certain elements. We can even resort to the exact calculation with the $\mathcal{O}(nd^2)$ cost, which can be acceptable for $d \ll n$.

Attention. In the previous part we described how to implement the exact trace for Eq. 4. The attention layer (Eq. 6) presents more challenges since it introduces the similarity measure in form of a dot product. The final attention weights for the i th element is the i th row of the $\text{softmax}(QK^T)$ matrix. If we want a zero diagonal Jacobian, we need to obtain the elements of Q with a MADE network $f_q(\mathbf{x}_i)$. We also need to mask the diagonal of QK^T with $-\infty$ to get zeros there after the softmax operation. This ensures that we never take the dot product of an element with itself. Therefore, the function $f_k(\cdot)$ can be an unrestricted network. At last, the output of an attention layer is a weighted sum of \mathbf{v}_i . Since we have already zeroed-out diagonal entries in the similarity matrix, the function $f_v(\cdot)$ can be unrestricted too. This gives constant trace for the *within* and *between points* interaction. The *per dimension* mapping can be implemented as discussed above. We experimentally show that our implementation does not hinder performance.

Discussion. The key to building an exact trace model is designing functions g and h — interaction of one dimension with all the others, and interaction of points between themselves — with zero trace Jacobian. The novelty of our approach is utilizing the properties of permutation equivariant functions. We propose exact trace models that can be either: 1) used as volume preserving flows, 2) combined with trace estimation or exact trace calculation, 3) extended within the framework of Jacobian decoupling (Chen & Duvenaud, 2019). In the experiments, we mainly use the third option. In the ablation studies we also explore the capabilities of the volume preserving flow.

3.2. Architecture

Our model is a continuous normalizing flow (CNF) that constitutes of a permutation invariant base distribution and a permutation equivariant function f that defines the ordinary differential equation dynamics. A CNF can be seen as a single transformation: we can combine it with other layers (fixed or learnable) or stack multiple CNFs. We use a normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$ for our base distribution. If we require a distribution on a bounded domain, e.g., uniform $\mathcal{U}(\mathbf{0}, \mathbf{1})$, we can go to the $\mathbb{R}^{n \times d}$ space with logit function, so that we can transform points with our CNF. We then map everything back to the original space with a sigmoid function.

The model is trained by maximizing the log-likelihood $p(X)$. Following the previous works, we set the integral bounds in Eq. 3 to: $t_0 = 0$ and $t_1 = 1$. Instead of back-propagating through solver steps, we use a memory efficient adjoint method to obtain the gradients (Farrell et al., 2013; Chen et al., 2018). We use an adaptive ODE solver (Dormand & Prince, 1980), whose error is bounded, in contrast to fixed-step solvers that have bounded number of steps. In the transformation f , we use permutation equivariant layers based on Eq. 4 or Eq. 6. One important measure of scalability of ODE models is the number of evaluations that the solver performs. In the experiments we show our exact trace model requires fewer steps.

To sample new realizations, we first sample n points from the base distribution $Z \sim q(Z)$. We then solve an ODE in the opposite direction (from 1 to 0) to get samples X from the learned distribution $p(X)$.

4. Related work

4.1. Models for exchangeable data and sets

Zaheer et al. (2017); Qi et al. (2017) propose permutation equivariant layers (Eq. 4) with universal approximation. Lee et al. (2019) demonstrate the same property holds for the attention based network and propose an extension for large sets that uses a small number of learnable inducing points. Wagstaff et al. (2019) study the theoretical limitations of

invariant functions with sum aggregation. Alternative aggregation schemes include max-pooling, Janossy pooling (Murphy et al., 2019), featurewise sort pooling (Zhang et al., 2020), and using graph methods (Grover & Leskovec, 2016).

Korshunova et al. (2018) treat points independently and use an exchangeable Student-t base distribution. Bender et al. (2020) use an autoregressive function on ordered sequences and optimize the following likelihood:

$$p(X) = \frac{1}{n!} p(\mathbf{x}_{(1)}, \dots, \mathbf{x}_{(n)}). \quad (7)$$

An obvious disadvantage is the requirement to impose a canonical order, which may not be possible and is often arbitrary. Other autoregressive models for *set-like* data include graph generation on breath-first search ordering (You et al., 2018); and ordering on z -axis for meshes and point clouds (Nash et al., 2020; Sun et al., 2020). In temporal processes the order of points is known so existing approaches mostly use autoregressive models (Du et al., 2016; Mei & Eisner, 2017; Shchur et al., 2020a;b).

In contrast, Yang et al. (2019) model point clouds with a variational autoencoder, having one CNF for posterior and another CNF for point locations. Yuan et al. (2020) apply the variational autoencoder approach in spatial point processes by defining a nonparametric kernel. VAE approach relies on de Finetti’s theorem (De Finetti, 1937; O’Neill, 2009):

$$p(X) = \int p(\mathbf{z}) \prod_{\mathbf{x}_i \in X} p(\mathbf{x}_i | \mathbf{z}) d\mathbf{z}, \quad (8)$$

and maximizes the evidence lower bound:

$$\log p(X) \geq \mathbb{E}_q [\log p(X | \mathbf{z})] - D_{\text{KL}} [q(\mathbf{z} | X) || p(\mathbf{z} | X)].$$

Using amortized inference, parameters of the posterior are defined as a permutation invariant function of X . Taking the gradients w.r.t. the samples is enabled with the reparametrization trick (Kingma & Welling, 2014). Yuan et al. (2020) demonstrate an application of their approach for recommender systems by embedding the non-spatial data with a graph neural network. Our model supports this too, without any changes to the architecture.

4.2. Equivariant continuous normalizing flows

The idea of using CNFs together with equivariant functions has been proposed by Köhler et al. (2020) in the context of multi-body physical systems by modeling the Gibbs distribution with a simple equivariant function, based on a Gaussian kernel. In Section 3 we offer ways to use strictly more expressive equivariant functions with closed-form trace. In the Supplementary Material we demonstrate how our method outperforms Gaussian kernel by a large margin. Toth et al. (2020) propose modeling physics systems by transforming

the initial samples with the Hamiltonian dynamics. This method is similar to CNFs in general, but requires variational approximation for training. Li et al. (2020) propose modeling sets similar to our baseline model with stochastic trace. We demonstrate how this is not scalable and propose an efficient alternative. We also show how this can be applied for spatial point process models with exact likelihood — an important, previously unsolved problem.

4.3. Improvements to Neural ODEs

Neural ODEs (Lu et al., 2018; Chen et al., 2018) were recently proposed as a continuous equivalent of the widely adopted residual architecture (He et al., 2016). By taking advantage of their properties we can construct a continuous version of the normalizing flow. Grathwohl et al. (2019) demonstrate the benefits of this when using a stochastic trace estimator. Chen & Duvenaud (2019) first used general network decoupling to train CNFs which led to improvements in performance. We use their approach to extend our base models with fixed trace. A lot of recent work focused on improving the stability of ODEs, including better representation power (Dupont et al., 2019), improved training (Gholami et al., 2019; Zhuang et al., 2020), regularization of the solver dynamics (Finlay et al., 2020; Kelly et al., 2020), etc. Our work is orthogonal to these methods since it focuses on the special case of modeling invariant data but can also be easily combined with them.

4.4. Point processes

Cox process (Cox, 1955) is a doubly stochastic Poisson process that defines intensity as a realization of a random field, e.g. a Gaussian process. For example, a Neyman-Scott process (Neyman & Scott, 1958) generates initial (parent) points from a Poisson process, and final (children) points around parents. A more specific example is a Matérn cluster process (Matérn, 2013) that places children in balls centered around parents. A variation where children follow a normal distribution is called a modified Thomas process (Thomas, 1949). Even though we do not specify the random field explicitly, the randomness from the base density is allowing us to model these processes with equivariant transformations.

Explicit interactions can be defined with a Gibbs process by assigning energy values to point configurations, originally defined in the context of statistical physics (Ruelle, 1969). However, this comes at a cost of calculating the normalization constant to find exact densities. Besides pseudolikelihood (Besag, 1975), other works used logistic composite likelihood (Clyde & Strauss, 1991) and Monte Carlo methods (Huang & Ogata, 1999). We avoid this by designing a model with exact likelihood at no flexibility trade-off.

5. Experiments

In the experiments we aim to show how having an exact trace model does not limit the expressiveness — on the contrary — we get the benefits of the faster and more stable training. First, we show that our versions of the models are equivalent to their original implementation, then demonstrate the modeling capacity for the point processes, and finally, show how we can scale to bigger datasets.

The detailed explanation of the data processing, hyperparameter tuning, and additional results can be found in the Supplementary Material.¹ We train all of our models on a single GPU (12GB).

5.1. Matching performance with the exact trace

We use established tasks (Lee et al., 2019; Li et al., 2020) to compare our novel models with their original versions. By obtaining similar performance, we demonstrate their equivalence while additionally having exact trace computation.

Maximum value regression. We construct a toy task where the goal is to retrieve the maximum value from a set. We use sets of real numbers, and compare the deep set and attention original architectures with their zero trace and exact trace versions. We run each experiment 5 times and report the mean and the standard deviation of the test set mean absolute error. Table 1 shows that there is no significant difference in performance and that our models sometimes perform better.

Counting the number of unique digits. We construct sets of digits from MNIST dataset where the target is the number of unique classes in a set. The model needs to learn to classify images and then count the number of unique classes. Our end-to-end model processes the digits with a multilayer convolutional neural network, aggregates the embeddings using deep set or attention, and outputs the λ parameter of a Poisson distribution. We train by maximizing the log-likelihood and evaluate the test performance with accuracy, defined as the frequency of predicting the Poisson mode that matches the true value, as in (Lee et al., 2019).

We run each experiment 5 times and report the mean and the standard deviation of the test set accuracy. The results in Table 1 again indicate that our architecture matches the performance of the original models. Note that we set the latent dimension d_g and the hidden dimensions such that the total number of parameters in the exact trace models and the original models is approximately the same.

Point cloud classification. Similar to (Li et al., 2020), we use labeled point clouds (Wu et al., 2015) in a supervised classification task. We feed in a set of 512 points to a two-layer convolutional network followed by an ODE with the

¹<https://www.daml.in.tum.de/scalable-nf>

	Deep set	Attention
Max value regression		
Unchanged	0.032 ± 0.003	0.029 ± 0.002
Zero trace	0.032 ± 0.003	0.022 ± 0.002
Exact trace	0.027 ± 0.003	0.036 ± 0.011
Counting digits		
Unchanged	0.838 ± 0.010	0.854 ± 0.012
Zero trace	0.836 ± 0.017	0.872 ± 0.017
Exact trace	0.843 ± 0.013	0.842 ± 0.005
Point cloud		
Unchanged	0.7685 ± 0.017	0.7495 ± 0.007
Exact trace	0.7523 ± 0.008	0.7714 ± 0.012

Table 1. Mean absolute error on maximum value regression (top), accuracy of counting the unique digits in a set (middle), and accuracy of point cloud classification (bottom).

deep set or attention network, apply max-pooling over all the elements and output the logits with a two-layer feedforward network. The hidden dimensions range in (32, 128), and there are 40 classes in total. The results in Table 1 show that both the exact trace and original transformation perform similarly.

Note that the models in these experiments were not tuned to achieve the best performance on these tasks but rather to compare the performance between two different implementations of the deep set and attention architecture, the original versus the exact trace version. We can conclude that there is no significant difference in these two implementations, as intended. x

5.2. Modeling point processes

Synthetic data. For all datasets we simulate 1000 realizations on $(0, 1) \times (0, 1)$ area. Thomas dataset is simulated by first sampling $m \sim \text{Poisson}(3)$ parents uniformly, and then for each, we sample $n_i \sim \text{Poisson}(5)$ children from the normal distribution centered on the parents ($\sigma = 0.01$). In Matérn the children are uniformly sampled inside discs centered on parents ($R = 0.1$). Mixture is an inhomogeneous process that generates $n \sim \text{Poisson}(64)$ points from a mixture of 3 normal distributions.

Real-world data. Check-ins NY (Cho et al., 2011) is a collection of locations of social network users. We consider points in New York that have 1938 unique users. The single realization corresponds to all the recorded locations for one user. We also construct a smaller dataset for a different city (Check-ins Paris) that has 286 users. Crimes² dataset contains daily records of locations and types of crimes that

²<https://nij.ojp.gov/funding/real-time-crime-forecasting-challenge>




Figure 3. Relative difference in log-likelihood (higher is better) centered at the approximate trace model scores. We compare two different architectures for our approach.

occurred in Portland. Each day contains between 298 and 736 points with 480 on average.

Baselines. Inhomogeneous Poisson process (IHP) models $p(X) = \prod_i p(x_i)$, i.e., it assumes independence between the points. We use a normalizing flow with spline coupling layers (Durkan et al., 2019). IHP cannot capture interactions between the points but models a fixed density on \mathcal{B} . Additionally, we use an autoregressive model (Bender et al., 2020), where the input points are sorted on the first dimension and we maximize Eq. 7. We also include the model without exact likelihood, an importance weighted autoencoder (IWAE). For the permutation invariant encoder we either use deep set or attention layers with layer normalization (Ba et al., 2016) and max-pooling to output parameters of the posterior $p(z|X)$. The likelihood term $p(X|z)$ is an IHP conditioned on z . To estimate the likelihood (Eq. 8) on the held-out data we use 5000 samples (Burda et al., 2015). The detailed description of data processing steps and the hyperparameters we used is in the Supplementary Material.

Results. Datasets are split into training, validation and test sets (60%-20%-20%). We train with early stopping, use mini-batches of size 64 and Adam optimizer with the learning rate of 10^{-3} (Kingma & Ba, 2015). The loss we use is per-point negative log-likelihood, i.e. $\mathcal{L}(X) = -\log p(X)/n$. Table 2 reports the mean and the variance of the loss on a held-out test set averaged over 5 runs for models selected based on their validation set performance. Our method beats others in most of the cases, and is within the margin of error for the remaining. In Figure 3, we compare our method to an approximate trace approach. We see that the relative difference to the log-likelihood of an approximate model is mostly in favour of our method.

We additionally quantitatively test the sampling quality by comparing the distributions of in-between point distances for real data and new model samples. Again, our model outperforms the competitors. We conclude that CNFs are well suited for modeling exchangeable data. Further, samples in Figure 4 (Left) qualitatively demonstrate that we capture the true underlying process, in this case, the clustering behavior. We can see that the individual points group around, forming small clusters, as expected. We additionally compute

Scalable Normalizing Flows for Permutation Invariant Densities

	Check-ins NY	Check-ins Paris	Crimes	Matérn	Mixture	Thomas
IHP	-1.61 ± 0.02 (0.035)	-2.87 ± 0.04 (0.026)	-2.34 ± 0.01 (0.002)	-0.21 ± 0.00	-2.06 ± 0.00	-0.01 ± 0.00
Autoregr.	-1.69 ± 0.02 (0.016)	-3.39 ± 0.05 (0.022)	-1.61 ± 0.06 (0.006)	-0.38 ± 0.02	-1.55 ± 0.01	0.13 ± 0.02
IWAE	-2.03 ± 0.07 (0.035)	-3.57 ± 0.02 (0.033)	-2.35 ± 0.00 (0.002)	-0.59 ± 0.01	-2.05 ± 0.01	-0.23 ± 0.00
CNF (Our)	-2.30 ± 0.09 (0.012)	-3.58 ± 0.03 (0.014)	-2.34 ± 0.01 (0.002)	-0.77 ± 0.08	-2.07 ± 0.00	-0.55 ± 0.00

Table 2. Per-point negative log-likelihood (in brackets: Wasserstein distance of the between points distance distributions) on point process data. Lower values are better for both metrics.




Figure 4. Left: Thomas data (top) and samples (bottom). Middle: Density realizations on NY data. Right: Comparison to GP-PP model.

the Ripley’s K function that measures the homogeneity of the process, where the higher number denotes a more clustered process. The real data has $K = 1.92$ and our model achieves $K = 1.93$, almost a perfect match. For reference, the baseline inhomogeneous model has $K = 1.41$.

Finally, in the Supplementary Material we show how our method can be orders of magnitude faster during test time evaluation. This is because approximate trace models have to compute the trace exactly in this setting, thus, even if the training steps take a small amount of time, the model in the production can be prohibitively slow. We also demonstrate that removing variance imposed by stochastic estimators can lead to much better convergence in the early stages of training. This means our method achieves better results faster, and the results stay better according to Figure 3.



Figure 5. Setting different set sizes n results in different types of samples for data containing discretized digits.

Set cardinality. Above, we followed previous works and used the loss that ignores the cardinality of the set captured in the $p(n)$ term. The empirical distribution of the number of points in our data is almost always simple and unimodal, meaning that the burden of the modeling is on the locations of the points and their interactions. It is easy to have a full point process model by adding $p(n) = \text{Poisson}(\lambda)$, where λ is learnable.

To show this we construct a toy dataset consisting of points that form shapes of two digits: digit 0 that has $n = 40$ points and digit 1 with $n = 20$. Learning $p(n)$ is trivial in this case, the hard part is learning the distribution of point locations given n , $p(X|n)$. After fitting the data with our attention-based model, we show in Figure 5 how sampling with different n gives different shapes, as expected.

Stochastic density. Since the density $p(X)$ is influenced by interactions between the points, it is not stationary but depends on X (unlike inhomogeneous model). We can draw a new *density realization* on \mathcal{B} by sampling $X \sim p(X)$. To be precise, we first draw $X \sim p(X)$, which defines the conditional density at non-observed points $x \in \mathcal{B}$, $x \notin X$. We can sample from this distribution or use it to detect the out of distribution data. Figure 4 (Middle) shows different conditional densities based on the real data samples for a model trained on Check-ins NY data. Having different density realizations connects to the notion of the doubly stochastic processes (Cox, 1955).

A popular way to define a Cox process is by defining the random field with a Gaussian process (Rasmussen & Williams, 2006), and then drawing an intensity function as a random realization. Figure 4 (Right) shows how a Gaussian process modulated Poisson process (GP-PP) (Lloyd et al., 2015) fits the coal mining disaster data, with 191 points collected from 1851 to 1962. We additionally plot the empirical intensity measure estimated with a Gaussian kernel. Finally, we plot the mean and the standard deviation of 100 *intensity function* realizations from our model. It closely matches the underlying intensity and provides the uncertainty measure. We get the intensity function from our model by scaling the density function with the learned parameter λ .

The above two examples demonstrate how the stochasticity of the base distribution samples leads to different final conditional distributions, without Gaussian process modulation.




Figure 6. Real data (top), and our model samples (bottom) for airplane class in point clouds data and discretized digits from MNIST.




Figure 7. Comparing the time (log scale) to evaluate trace for different estimators across different dimension sizes of data.

5.3. Scaling to large sets

In the previous section we demonstrated that CNFs are a good choice for modeling exchangeable data. Now we test how different trace calculation approaches affect the performance on bigger datasets.

First, we compare the speed of different trace estimation methods, namely: our exact trace model, stochastic trace estimation with Hutchinson’s estimator, and computing the trace exactly without using any specialized architectures. We use the same type of the model and the same number of parameters. We record the time to calculate the trace on dimension sizes in the range (50, 1000). Figure 7 shows that our method is faster than both competing approaches. Finding the trace exactly on arbitrary architectures (not using methods from Section 3) becomes untractable very soon. That means that even if we used relatively fast stochastic trace estimation during training, our model would be of limited use in production when we want to know the exact density.

Further, we test the performance on real-world large sets of points. For that we model point clouds, in particular, we use airplane and chair classes (Wu et al., 2015) where each sample contains 512 points. We additionally use discretized MNIST digits dataset, i.e., sample the coordinates of 40 black pixels and add a small noise to them to avoid degenerate solutions (placing infinite mass on a single point). The model consists of 12 stacked continuous normalizing flows with the deep set or the attention transformations, and a similar hyperparameter setup as in Li et al. (2020). We observe that both versions of the model perform similarly.

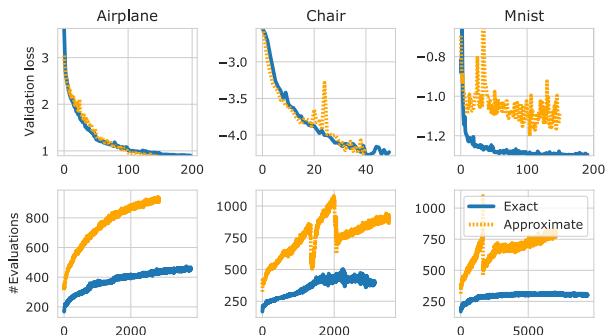


Figure 8. Comparison of the number of solver evaluations and the validation loss during training for exact and approximate trace.

Figure 6 shows samples of airplane point clouds and discretized digits from our model. We can see that the model captures the complex data well, and the samples are visually similar to previous works (Li et al., 2020). Next, we test an important metric that shows scalability — number of ODE solver evaluations. In Figure 8 we can see that having an exact trace in a model allows us to obtain the lower number of total evaluations, meaning that both the training and inference will be faster compared to the stochastic trace estimation. At the same time we do not sacrifice performance, as is seen in the evolution of validation loss over training, where our method matches or outperforms the competing approach. Note that we also observe noisier training and validation curves for approximate trace meaning the training is less stable. An interesting observation is that both models adapt the complexity by requiring more evaluations later in the training, when they learn the underlying distribution better.




Figure 9. How different latent size d_g influences expressiveness.

5.4. Ablation studies

Latent dimension d_g . We do an ablation study for different latent dimensions d_g for the network g , as described in Section 3. We test the sizes of 1, 2 and 8. Figure 9 shows that higher d_g leads to better performance as expected, but in the case of the digits dataset, we reach the point of diminishing returns very fast.

Volume preserving flow. The version of our model which has the constant zero trace only accounts for the interactions between the points, but not the full transformation of the points themselves. This means that if the base process is homogeneous, the volume preserving flow cannot learn the inhomogeneous part of the process, i.e., the fixed density over \mathcal{B} . In the Supplementary Material we show how it performs against IHP model and demonstrate that on some datasets it is enough to use the model with zero trace transformations alone. In the end, it is the best to use the full model with exact trace since it captures both the inhomogeneous part and the interaction between the points.

ODE regularization. The goal of our method is to reduce the variance during training, which leads to having a better performance both during training and inference. Most of the recent works on improving neural ODE models (see Section 4) focused on learning functions f that reduce the *stiffness*, meaning we get better behavior in the ODE solver. In the point process experiments we tried Frobenius norm regularization (Finlay et al., 2020), but did not observe significant changes in performance.

Induced attention. Lee et al. (2019) propose an alternative to the attention layer that uses a set of inducing points m , aiming to reduce the computation cost when $m \ll n$. We compare it with the usual self attention on the point process data, and notice that it usually performs worse. The details can be found in the Supplementary Material. Having an extra hyperparameter is another drawback of this method. However, the lower memory footprint of induced attention can help train models on sets with very large cardinality.

6. Conclusion

In this work we presented a novel model for point processes and random sets that uses continuous normalizing flows with equivariant functions that has the exact trace of the Jacobian available. The exact trace, in particular, allows us to scale to bigger datasets since this approach is more efficient but at the same time, it achieves better performance.

Throughout the experimental evaluation we demonstrated how we can successfully model complex point processes with interactions between the points while having tractable likelihood and straightforward sampling. This is the first spatial point process model with these properties. We further

show that our model outperforms all other competitors on this task. At the same time, it does not require any special assumptions on the data, like independence between the points, nor does it need canonical ordering.

Our method relies in particular on two types of deep learning models: continuous normalizing flows and equivariant neural networks. By utilizing new techniques developed in these two subfields (new types of equivariant layers and ODE regularizations) we hope it will perform even better in the future.

Acknowledgments

This research was supported by BMW AG and by the German Federal Ministry of Education and Research (BMBF), grant no. 01IS18036B. The authors of this work take full responsibilities for its content.

References

- Ba, J. L., Kiros, J. R., and Hinton, G. E. Layer normalization. *arXiv:1607.06450*, 2016.
- Bender, C., Garcia, J. J., O’Connor, K., and Oliva, J. Exchangeable generative models with flow scans. In *AAAI*, 2020.
- Besag, J. Statistical analysis of non-lattice data. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 24(3):179–195, 1975.
- Burda, Y., Grosse, R., and Salakhutdinov, R. Importance weighted autoencoders. In *ICLR*, 2015.
- Chen, R. T. and Duvenaud, D. K. Neural networks with cheap differential operators. In *NeurIPS*, 2019.
- Chen, R. T., Amos, B., and Nickel, M. Neural spatio-temporal point processes. In *ICLR*, 2021.
- Chen, T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K. Neural ordinary differential equations. In *NeurIPS*, 2018.
- Chiu, S. N., Stoyan, D., Kendall, W. S., and Mecke, J. *Stochastic geometry and its applications*. John Wiley & Sons, 2013.
- Cho, E., Myers, S. A., and Leskovec, J. Friendship and mobility: user movement in location-based social networks. In *ACM SIGKDD*, 2011.
- Clyde, M. and Strauss, D. Logistic regression for spatial pair-potential models. *Lecture Notes-Monograph Series*, pp. 14–30, 1991.

- Cox, D. R. Some statistical methods connected with series of events. *Journal of the Royal Statistical Society: Series B (Methodological)*, 17(2):129–157, 1955.
- Daley, D. and Vere-Jones, D. An introduction to the theory of point processes: Volume I: Elementary theory and methods. 2007.
- De Finetti, B. La prévision: ses lois logiques, ses sources subjectives. In *Annales de l’institut Henri Poincaré*, volume 7, pp. 1–68, 1937.
- Dinh, L., Krueger, D., and Bengio, Y. Nice: Non-linear independent components estimation. *ICLR Workshop*, 2015.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using Real NVP. In *ICLR*, 2017.
- Dormand, J. R. and Prince, P. J. A family of embedded runge-kutta formulae. *Journal of computational and applied mathematics*, 6(1):19–26, 1980.
- Du, N., Dai, H., Trivedi, R., Upadhyay, U., Gomez-Rodriguez, M., and Song, L. RMTTP: Embedding event history to vector. In *KDD*, 2016.
- Dupont, E., Doucet, A., and Teh, Y. W. Augmented neural odes. In *NeurIPS*, 2019.
- Durkan, C., Bekasov, A., Murray, I., and Papamakarios, G. Neural spline flows. In *NeurIPS*, 2019.
- Farrell, P. E., Ham, D. A., Funke, S. W., and Rognes, M. E. Automated derivation of the adjoint of high-level transient finite element programs. *SIAM Journal on Scientific Computing*, 35(4):C369–C393, 2013.
- Finlay, C., Jacobsen, J.-H., Nurbekyan, L., and Oberman, A. M. How to train your neural ode. In *ICML*, 2020.
- Germain, M., Gregor, K., Murray, I., and Larochelle, H. Made: Masked autoencoder for distribution estimation. In *ICML*, 2015.
- Gholami, A., Keutzer, K., and Biro, G. Anode: Unconditionally accurate memory-efficient gradients for neural odes. In *IJCAI*, 2019.
- Grathwohl, W., Chen, R. T., Bettencourt, J., Sutskever, I., and Duvenaud, D. FFJORD: Free-form continuous dynamics for scalable reversible generative models. In *ICLR*, 2019.
- Grover, A. and Leskovec, J. node2vec: Scalable feature learning for networks. In *KDD*, 2016.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *CVPR*, 2016.
- Huang, F. and Ogata, Y. Improvements of the maximum pseudo-likelihood estimators in various spatial statistical models. *Journal of Computational and Graphical Statistics*, 8(3):510–530, 1999.
- Kelly, J., Bettencourt, J., Johnson, M. J., and Duvenaud, D. Learning differential equations that are easy to solve. In *NeurIPS*, 2020.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *ICLR*, 2014.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. Improved variational inference with inverse autoregressive flow. In *NeurIPS*, 2016.
- Kobyzev, I., Prince, S., and Brubaker, M. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- Köhler, J., Klein, L., and Noé, F. Equivariant flows: exact likelihood generative learning for symmetric densities. In *ICML*, 2020.
- Korshunova, I., Degraeve, J., Huszár, F., Gal, Y., Gretton, A., and Dambre, J. Bruno: A deep recurrent model for exchangeable data. In *NeurIPS*, 2018.
- Lee, J., Lee, Y., Kim, J., Kosiorek, A., Choi, S., and Teh, Y. W. Set transformer: A framework for attention-based permutation-invariant neural networks. In *ICML*, 2019.
- Li, Y., Yi, H., Bender, C. M., Shan, S., and Oliva, J. B. Exchangeable neural ode for set modeling. In *NeurIPS*, 2020.
- Lloyd, C., Gunter, T., Osborne, M., and Roberts, S. Variational inference for gaussian process modulated poisson processes. In *ICML*, 2015.
- Lu, Y., Zhong, A., Li, Q., and Dong, B. Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations. In *ICML*, 2018.
- Maron, H., Litany, O., Chechik, G., and Fetaya, E. On learning sets of symmetric elements. In *ICML*, 2020.
- Matérn, B. *Spatial variation*, volume 36. Springer Science & Business Media, 2013.
- Mei, H. and Eisner, J. M. The neural hawkes process: A neurally self-modulating multivariate point process. In *NeurIPS*, pp. 6754–6764, 2017.

- Murphy, R. L., Srinivasan, B., Rao, V., and Ribeiro, B. Janossy pooling: Learning deep permutation-invariant functions for variable-size inputs. In *ICLR*, 2019.
- Nash, C., Ganin, Y., Eslami, S., and Battaglia, P. W. Polygen: An autoregressive generative model of 3d meshes. In *ICML*, 2020.
- Neyman, J. and Scott, E. L. Statistical approach to problems of cosmology. *Journal of the Royal Statistical Society: Series B (Methodological)*, 20(1):1–29, 1958.
- Ogata, Y. and Tanemura, M. Estimation of interaction potentials of marked spatial point patterns through the maximum likelihood method. *Biometrics*, pp. 421–433, 1985.
- O’Neill, B. Exchangeability, correlation, and bayes’ effect. *International statistical review*, 77(2):241–250, 2009.
- Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., and Lakshminarayanan, B. Normalizing flows for probabilistic modeling and inference. *arXiv:1912.02762*, 2019.
- Qi, C. R., Su, H., Mo, K., and Guibas, L. J. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017.
- Rasmussen, C. E. and Williams, C. K. I. *Gaussian processes for machine learning*. Adaptive computation and machine learning. MIT Press, 2006.
- Ruelle, D. *Statistical mechanics: Rigorous results*. W. A. Benjamin, Inc., New York, 1969.
- Shchur, O., Biloš, M., and Günnemann, S. Intensity-free learning of temporal point processes. In *ICLR*, 2020a.
- Shchur, O., Gao, N., Biloš, M., and Günnemann, S. Fast and flexible temporal point processes with triangular maps. In *NeurIPS*, 2020b.
- Sun, Y., Wang, Y., Liu, Z., Siegel, J., and Sarma, S. Point-grow: Autoregressively learned point cloud generation with self-attention. In *IEEE WACV*, 2020.
- Thomas, M. A generalization of poisson’s binomial limit for use in ecology. *Biometrika*, 36(1/2):18–25, 1949.
- Toth, P., Rezende, D. J., Jaegle, A., Racanière, S., Botev, A., and Higgins, I. Hamiltonian generative networks. In *ICLR*, 2020.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *NeurIPS*, 2017.
- Wagstaff, E., Fuchs, F. B., Engelcke, M., Posner, I., and Osborne, M. On the limitations of representing functions on sets. In *ICML*, 2019.
- Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., and Xiao, J. 3d shapenets: A deep representation for volumetric shapes. In *CVPR*, 2015.
- Yang, G., Huang, X., Hao, Z., Liu, M.-Y., Belongie, S., and Hariharan, B. Pointflow: 3d point cloud generation with continuous normalizing flows. In *ICCV*, 2019.
- You, J., Ying, R., Ren, X., Hamilton, W. L., and Leskovec, J. Graphrnn: Generating realistic graphs with deep autoregressive models. In *ICML*, 2018.
- Yuan, B., Wang, X., Ma, J., Zhou, C., Bertozzi, A. L., and Yang, H. Variational autoencoders for highly multivariate spatial point processes intensities. In *ICLR*, 2020.
- Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R. R., and Smola, A. J. Deep sets. In *NeurIPS*, 2017.
- Zhang, Y., Hare, J., and Prügel-Bennett, A. FSPool: Learning set representations with featurewise sort pooling. In *ICLR*, 2020.
- Zhuang, J., Dvornik, N., Li, X., Tatikonda, S., Papademetris, X., and Duncan, J. Adaptive checkpoint adjoint method for gradient estimation in neural ode. In *ICML*, 2020.