

# Probabilistic Forecasting of Irregularly Sampled Time Series with Missing Values via Conditional Normalizing Flows

Vijaya Krishna Yalavarthi<sup>\*1</sup>, Randolf Scholz<sup>\*1</sup>, Stefan Born<sup>2</sup>, Lars Schmidt-Thieme<sup>1</sup>

<sup>1</sup> Information Systems and Machine Learning Lab, University of Hildesheim, Germany

<sup>2</sup> Institute of Mathematics, TU Berlin, Germany  
{yalavarthi, scholz, schmidt-thieme}@ismll.de,  
born@math.tu-berlin.de

## Abstract

Probabilistic forecasting of irregularly sampled multivariate time series with missing values is crucial for decision making in various domains, including health care, astronomy, and climate. State-of-the-art methods estimate only marginal distributions of observations in single channels and at single timepoints, assuming a Gaussian distribution for the data. In this work, we propose a novel model, ProFITi using conditional normalizing flows to learn multivariate conditional distribution: joint distribution of the future values of the time series conditioned on past observations and specific channels and timepoints, without assuming any fixed shape of the underlying distribution. As model components, we introduce a novel invertible triangular attention layer and an invertible non-linear activation function on and onto the whole real line. Through extensive experiments on 4 real-world datasets, ProFITi demonstrates significant improvement, achieving an average log-likelihood gain of 2.0 compared to the previous state-of-the-art method.

## 1 Introduction

Irregularly sampled multivariate time series with missing values (IMTS) are common in various real-world scenarios such as health, astronomy and climate. Accurate forecasting of IMTS is important for decision-making, but estimating uncertainty is crucial to avoid overconfidence. State-of-the-art models applied to this task are Ordinary Differential Equations (ODE) based models (Schirmer et al. 2022; De Brouwer et al. 2019; Biloš et al. 2021) which are 1) computationally inefficient, and 2) offer only marginal likelihoods. In practice, joint or multivariate distributions are desired to capture dependencies and study forecasting scenarios. With joint distributions one can estimate the likelihood of specific combinations of future variables ex. likelihood of having rain and strong winds, which marginal or point forecast models cannot deliver.

For this, we propose a novel conditional normalizing flow model called **ProFITi**, for **Probabilistic Forecasting of Irregularly sampled Multivariate Time series**. ProFITi is designed to learn conditional joint distributions. We also propose two novel model components that can be used in flow

models: a sorted invertible triangular attention layer, **SITA**, parametrized by conditioning input to learn joint distributions, and an invertible non-linear activation function designed for flows, **Shiesh**, that is on and onto whole real line. ProFITi consists of several invertible blocks build using SITA and Shiesh functions. Being a flow-based model, ProFITi can learn any random conditional joint distribution, while existing models (Schirmer et al. 2022; De Brouwer et al. 2019; Biloš et al. 2021) learn only Gaussians.

Experiments on 4 real-world IMTS datasets, attest the superior performance of ProFITi. Our contributions are:

1. Introduced **ProFITi**, a novel and, to the best of our knowledge, first normalizing flow based probabilistic forecasting model for predicting multivariate conditional distributions of irregularly sampled time series with missing values (Section 6).
2. Proposed a novel invertible triangular attention layer, named **SITA**, which enables target variables to interact and capture dependencies within a conditional normalizing flow framework (Section 4).
3. Proposed a novel non-linear, invertible, differentiable activation function on and onto the whole real line, **Shiesh** (Section 5). This activation function can be used in normalizing flows.
4. Conducted experiments on 4 IMTS datasets for normalized joint negative loglikelihood. On average, ProFITi provides a loglikelihood gain of 2.0 over the previously best model (Section 7).

A detailed version of this work, along with additional experiments and proofs to lemmas, is available in the technical report (Yalavarthi et al. 2024b). Implementation: [github.com/yalavarthivk/ProFITi](https://github.com/yalavarthivk/ProFITi)

## 2 Literature Review

There have been multiple works dealing point forecasting of irregular time series (Yalavarthi et al. 2024a; Ansari et al. 2023; Chen et al. 2024). Very few models provide uncertainty quantification for such forecasts.

**Probabilistic Forecasting Models for IMTS.** Probabilistic IMTS forecasting often relies on variational inference or predicting distribution parameters. Neural ODE models (Chen et al. 2018) combine probabilistic latent states

<sup>\*</sup>These authors contributed equally.

with deterministic networks. Other approaches like latent-ODE (Rubanova, Chen, and Duvenaud 2019), GRU-ODE-Bayes (De Brouwer et al. 2019), Neural-Flows (Biloš et al. 2021), and Continuous Recurrent Units (Schirmer et al. 2022) provide only marginal distributions, no joint distributions. Similarly, probabilistic interpolation models, such as HETVAE (Shukla and Marlin 2022) and Triplet-Former (Yalavarthi, Burchert, and Schmidt-Thieme 2023), also provide only Gaussian marginal distributions. In contrast, Gaussian Process Regression models (GPR; Dürichen et al. 2015; Li and Marlin 2015, 2016; Bonilla, Chai, and Williams 2007) offer full joint posterior distributions for forecasts, but struggle with the computational demands on long time series due to the dense matrix inversion operations. All the models assume the data distribution is to be Gaussian and fail if the true distribution is different.

**Normalizing Flows for variable input size.** We deal with predicting distributions for variable many targets. This utilizes equivariant transformations, as shown in (Biloš and Günnemann 2021; Satorras, Hoogeboom, and Welling 2021; Liu et al. 2019). All the models apply continuous normalizing flows which require solving an ODE driven by a neural network using a slow numerical integration process. Additionally, they cannot incorporate conditioning inputs.

**Conditioning Normalizing Flows.** Learning conditional densities has been largely explored within computer vision (Khorashadizadeh et al. 2023; Winkler et al. 2019; Anantha Padmanabha and Zabarar 2021). They apply normalizing flow blocks such as affine transformations (Dinh, Sohl-Dickstein, and Bengio 2017), autoregressive transformations (Kingma and Welling 2013) or Sylvester flow blocks (van den Berg et al. 2018). Often the conditioning input is appended to the target while passing through the flow layers as shown in (Winkler et al. 2019). For continuous data representations only a few works exist (Kumar et al. 2020; de Bézenac et al. 2020; Rasul et al. 2021; Si, Kuleshov, and Bishop 2022). However, methods that deal with regular multivariate time series (such as Rasul et al. 2021) cannot handle IMTS due to its missing values. We solve this by using invertible attention that allows flexible size.

**Flows with Invertible Attention.** To the best of our knowledge, there have been only two works that develop invertible attention for Normalizing Flows. Sukthanker et al. (Sukthanker et al. 2022) proposed an invertible attention by adding the identity matrix to a softmax attention. However, softmax yields only positive values in the attention matrix and does not learn negative covariances. Zha et al. (Zha et al. 2021) introduced residual attention similar to residual flows (Behrmann et al. 2019) that suffer from similar problems as residual flows such as the lack of an explicit inverse making inference slow. Additionally, computing determinants of dense attention matrices has cubic complexity which is not desired.

### 3 Problem Setting & Analysis

**The IMTS Forecasting Problem.** An **irregularly sampled multivariate times series with missing values** (called

briefly just **IMTS** in the following), is a sequence  $x^{\text{OBS}} = ((t_\tau, v_\tau))_{\tau=1:\mathcal{T}}$  where  $v_\tau \in \{\mathbb{R} \cup \text{NaN}\}^C$  is an observation event at timepoint  $t_\tau \in \mathbb{R}$ .  $v_{\tau,c} \in \mathbb{R}$  indicates observed value and NaN indicates a missing value. Horn et al. (Horn et al. 2020) introduced set notation where only observed values are considered and missing values are ignored. For notational convenience, we use sequences whose order does not matter to represent sets.

Now,  $x^{\text{OBS}} = ((t_i^{\text{OBS}}, c_i^{\text{OBS}}, o_i^{\text{OBS}}))_{i=1:I}$  is a sequence of unique triples, where  $t_i^{\text{OBS}} \in \mathbb{R}$  denotes the time,  $c_i^{\text{OBS}} \in \{1, \dots, C\}$  the channel and  $o_i^{\text{OBS}} \in \mathbb{R}$  the value of an observation,  $I \in \mathbb{N}$  the total number of observations across all channels and  $C \in \mathbb{N}$  the number of channels.

An **IMTS query** is a sequence  $x^{\text{QRY}} = ((t_k^{\text{QRY}}, c_k^{\text{QRY}}))_{k=1:K}$  of just timepoints and channels (also unique), a sequence  $y \in \mathbb{R}^K$  we call an **answer**. It is understood that  $y_k$  is the answer to the query  $(t_k^{\text{QRY}}, c_k^{\text{QRY}})$ .

The **IMTS probabilistic forecasting problem** then is, given a dataset  $\mathcal{D}^{\text{train}} := ((x^{\text{OBS}^n}, x^{\text{QRY}^n}, y^n))_{n=1:N}$  of triples of time series, queries and answers from an unknown distribution  $p$  (with earliest query timepoint is beyond the latest observed timepoint for series  $n$ ,  $\min_k t_k^{\text{QRY}} > \max_i t_i^{\text{OBS}}$ ), to find a model  $\hat{p}$  that maps each observation/query pair  $(x^{\text{OBS}}, x^{\text{QRY}})$  to a joint density over answers,  $\hat{p}(y_1, \dots, y_k | x^{\text{OBS}}, x^{\text{QRY}})$ , such that the expected joint negative log likelihood is minimal:

$$\ell^{\text{NLL}}(\hat{p}; p) := -\mathbb{E}_{(x^{\text{OBS}}, x^{\text{QRY}}, y) \sim p} \log \hat{p}(y | x^{\text{OBS}}, x^{\text{QRY}})$$

Please note, that the number  $C$  of channels is fixed, but the number  $I$  of past observations and the number  $K$  of future observations queried may vary over instances  $(x^{\text{OBS}}, x^{\text{QRY}}, y)$ . If query sizes  $K$  vary, instead of (joint) negative log likelihood one also can normalize by query size to make numbers comparable over different series and limit the influence of large queries, the **normalized joint negative log likelihood** njNLL:

$$\ell^{\text{njNLL}}(\hat{p}; p) := - \mathbb{E}_{(x^{\text{OBS}}, x^{\text{QRY}}, y) \sim p} \frac{1}{|y|} \log \hat{p}(y | x^{\text{OBS}}, x^{\text{QRY}}) \quad (1)$$

**Problem Analysis and Characteristics.** As the problem is not just an (unconditioned) density estimation problem, but the distribution of the outputs depends on both, the past observations and the queries, a **conditional density model** is required (**requirement 1**).

A crucial difference from many settings addressed in the related work (Schirmer et al. 2022; Biloš et al. 2021; De Brouwer et al. 2019) is that we look for probabilistic models of the **joint distribution** of all queried observation values  $(y_1, \dots, y_K)$ , not just at the **single variable marginal distributions**  $p(y_k | x^{\text{OBS}}, x_k^{\text{QRY}})$  (for  $k = 1:K$ ). The problem of marginal distributions is a special case of our formulation where all queries happen to have just one element (always  $K = 1$ ). So for joint probabilistic forecasting of IMTS, models need to output densities on a **variable number of variables (requirement 2)**.

Furthermore, since we deal with the set representation of IMTS, whenever two query elements get swapped, a generative model should swap its output accordingly, a density

model should yield the same density value, i.e., the model should be **permutation invariant (requirement 3)**. For any permutation  $\pi$ :

$$\hat{p}(y_1, \dots, y_K \mid x_1^{\text{OBS}}, x_1^{\text{QRY}}, \dots, x_K^{\text{QRY}}) = \hat{p}(y_{\pi(1)}, \dots, y_{\pi(K)} \mid x_{\pi(1)}^{\text{OBS}}, x_{\pi(1)}^{\text{QRY}}, \dots, x_{\pi(K)}^{\text{QRY}}) \quad (2)$$

## 4 Invariant Conditional Normalizing Flow Models

**Normalizing flows.** Parametrizing a specific distribution such as the Gaussian is a simple and robust approach to probabilistic forecasting. It can be added on top of any point forecasting model (for marginal distributions or fixed-size queries at least). However, such models are less suited for targets having a more complex distribution. Then typically normalizing flows are used (Rippel and Adams 2013; Papamakarios et al. 2021). A normalizing flow is an (unconditional) density model for variables  $y \in \mathbb{R}^K$  consisting of a simple base distribution, typically a standard normal  $p_Z(z) := \mathcal{N}(z; 0_K, \mathbb{I}_{K \times K})$ , and an invertible, differentiable, parametrized map  $f(z; \theta) : \mathbb{R}^K \rightarrow \mathbb{R}^K$ ; then

$$\hat{p}(y; \theta) := p_Z(f^{-1}(y; \theta)) \left| \det \left( \frac{\partial f^{-1}(y; \theta)}{\partial y} \right) \right| \quad (3)$$

is a proper density, i.e., integrates to 1, and can be fitted to data minimizing negative log likelihood via gradient descent algorithms. A normalizing flow can be conditioned on predictor variables  $x \in \mathbb{R}^M$  by simply making  $f$  dependent on predictors  $x$ , too:  $f(z; x, \theta)$  (satisfying **requirement 1**).  $f$  then has to be invertible w.r.t.  $z$  for any  $x$  and  $\theta$  (Tripe and Turner 2018).

**Invariant conditional normalizing flows.** A conditional normalizing flow represents an invariant conditional distribution in the sense of eq. 2 (**requirement 3**), if i) its predictors  $x$  also can be grouped into  $K$  elements  $x_1, \dots, x_K$  and possibly common elements  $x^{\text{com}}$ :  $x = (x_1, \dots, x_K, x^{\text{com}})$ , and ii) its transformation  $f$  is equivariant in stacked  $x_{1:K}$  and  $z_{1:k}$ :

$$f(z^\pi; x_{1:K}^\pi, x^{\text{com}}, \theta)^\pi = f(z; x_{1:K}, x^{\text{com}}, \theta) \quad \forall \text{permutations } \pi \quad (4)$$

where  $z^\pi := (z_{\pi(1)}, \dots, z_{\pi(K)})$  denotes a permuted vector. We call this an **invariant conditional normalizing flow model**. If  $K$  is fixed, we call it **fixed size**, otherwise **dynamic size**. In our work, we consider  $x_{1:K}$  as the embedding of  $x_1^{\text{QRY}}, \dots, x_K^{\text{QRY}}$  and  $x^{\text{OBS}}$  and ignore  $x^{\text{com}}$  (see eq. 16).

**Invariant conditional normalizing flows via invertible attention.** The primary choice for a dynamic size (**requirement 2**), equivariant, parametrized function is attention (Attn; Vaswani et al. 2017):

$$\begin{aligned} A(X_Q, X_K) &:= X_Q W_Q (X_K W_K)^T, \\ A^{\text{softmax}}(X_Q, X_K) &:= \text{softmax}(A(X_Q, X_K)) \\ \text{Attn}(X_Q, X_K, X_V) &:= A^{\text{softmax}}(X_Q, X_K) \cdot X_V W_V \end{aligned}$$

where  $X_Q, X_K, X_V$  are query, key and value matrices,  $W_Q, W_K, W_V$  are parameter matrices (not depending on the

number of rows of  $X_Q, X_K, X_V$ ) and the softmax is taken rowwise.

Self attention mechanism ( $X_Q = X_K = X_V$ ) has been used in the literature as is for unconditional vector fields (Köhler, Klein, and Noe 2020; Li et al. 2020; Biloš and Günemann 2021). To be used in a conditional vector field,  $X_Q = X_K = X$  will have to contain the condition elements  $x_{1:K}$ ,  $X_V = Z$  contains the base samples  $z_{1:K}$  and  $W_V = 1$ .

$$X := \begin{bmatrix} x_1^T \\ \vdots \\ x_K^T \end{bmatrix}, \quad Z := \begin{bmatrix} z_1 \\ \vdots \\ z_K \end{bmatrix} \quad (5)$$

Now, we make attention matrix itself invertible. To get **invertible attention (iAttn)**, we regularize the attention matrix  $A$  sufficiently to become invertible (see Lemma 1 in Yalavarthi et al. 2024b for proof)

$$A^{\text{reg}}(X) := \frac{1}{\|A(X, X)\|_2 + \epsilon} A(X, X) + \mathbb{I} \quad (6)$$

$$\text{iAttn}(Z, X) := A^{\text{reg}}(X) \cdot Z \quad (7)$$

where  $\epsilon > 0$  is a hyperparameter. We note that,  $A^{\text{reg}}(X)$  is not a parameter of the model, but computed from the conditioners  $x$ . Our approach is different from iTrans attention (Sukthanker et al. 2022, fig. 17) that makes attention invertible more easily via  $A^{\text{iTrans}}(X) := A^{\text{softmax}}(X, X) + \mathbb{I}$  using the fact that the spectral radius  $\sigma(A^{\text{softmax}}(X, X)) \leq 1$ , but therefore is restricted to non-negative interaction weights.

The attention matrix  $A^{\text{reg}}(X)$  will be dense in general and thus slow to invert, taking  $\mathcal{O}(K^3)$  operations. Following ideas for autoregressive flows and coupling layers, a triangular matrix would allow a much more efficient inverse pass, as its determinant can be computed in  $\mathcal{O}(K)$  and linear systems can be solved in  $\mathcal{O}(K^2)$ . This does not restrict the expressivity of the model, as due to the Knothe–Rosenblatt rearrangement (Villani 2009) from optimal transport theory, any two probability distributions on  $\mathbb{R}^K$  can be transformed into each other by flows with a locally triangular Jacobian. Unfortunately, just masking the upper triangular part of the matrix will destroy the equivariance of the model. We resort to the simplest way to make a function equivariant: we sort the inputs before passing them into the layer and revert the outputs to the original ordering. We call this approach **sorted invertible triangular attention (SITA)**:

$$\pi := \text{argsort}(x_1 S, \dots, x_K S) \quad (8)$$

$$A^{\text{tri}}(X) := \text{softplus-diag}(\text{lower-triang}(A(X, X))) + \epsilon \mathbb{I} \quad (9)$$

$$\text{SITA}(Z, X) := (A^{\text{tri}}(X^\pi) \cdot Z^\pi)^{\pi^{-1}} \quad (10)$$

where  $\pi$  operates on the rows of  $X$  and  $Z$ . Softplus activation is applied to diagonal elements making them positive. Sorting is a simple lexicographic sort along the dimensions of vector  $x_k S$ . The matrix  $S$  allows to specify a sorting criterion, e.g., a permutation matrix. Note that sorting is unique only when  $x$  has unique elements. *In practice, we compute  $\pi$  from  $x^{\text{QRY}}$  instead of  $x$ , first sort by timepoint, and then by channel.*

Activation	E1	E2	E3
ReLU	×	×	×
Leaky-ReLU	✓	✓	✓
P-ReLU	✓	✓	✓
ELU	✓	×	✓
SELU	✓	×	✓
GELU	×	×	✓
Tanh	✓	×	✓
Sigmoid	✓	×	✓
Tanh-shrink	✓	✓	×
Shiesh	✓	✓	✓

Table 1: Properties of existing activation functions.

**Example 1** (Demonstration of sorting in SITA). Given  $x^{\text{QRY}} = ((1, 2), (0, 2), (2, 1), (3, 1), (0, 1), (3, 3))$  where first and second elements in  $x_k^{\text{QRY}}$  indicate queried time and channel respectively. Assume  $S = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ . Then

$$\begin{aligned} \pi &= \text{argsort}(x_1^{\text{QRY}} S, \dots, x_5^{\text{QRY}} S) \\ &= \text{argsort}((1, 2), (0, 2), (2, 1), (3, 1), (0, 1), (3, 3)) \\ &= (5, 2, 1, 3, 4, 6) \end{aligned}$$

## 5 Shiesh: A New Activation Function for Normalizing Flows

The transformation function  $f$  of a normalizing flow usually is realized as a stack of several simple functions. As in any other neural network, elementwise applications of a function, called activation functions, is one of those layers that allows for non-linear transformations. However, most common activation functions used in deep learning are not diffeomorphic and do not have both their domain and codomain on the entire real line, making them unsuitable for normalizing flows. For instance:

- **ReLU** is not invertible (E1)
- **ELU** cannot be used consistently throughout the layer stack because its output domain,  $\mathbb{R}^+$ , does not span the entire real number line (E2)
- **Tanh-shrink** ( $\text{tanhshrink}(u) := u - \tanh(u)$ ) is invertible and covers the entire real line, but it has a zero gradient at some points (e.g., at  $u = 0$ ). This zero gradient makes it impossible to compute the inverse of the normalizing factor,  $\left| \det \left( \frac{\partial f(u)}{\partial u} \right) \right|$ , required for normalizing flows (E3)

To serve as a standalone layer in a normalizing flow, an activation function must fulfill these three requirements: **E1**. be invertible, **E2**. cover the whole real line and **E3**. have no zero gradients. Out of all activation functions in the pytorch library (V 2.2) only Leaky-ReLU and P-ReLU meet all three requirements (see Table ??). Leaky-ReLU and P-ReLU usually are used with a slope on their negative branch being well less than 1, so that stacking many of them might lead to small gradients also causing problems for the normalizing constant of a normalizing flow.

To address these challenges, we propose a new activation function derived from unconstrained monotonic neural networks (UMNN; Wehenkel and Louppe 2019). UMNN have




Figure 1: (left) Shiesh function, (right) partial derivative.

been proposed as versatile, learnable activation functions for normalizing flows, being basically a continuous flow for each scalar variable  $u$  separately and a scalar field  $g$  implemented by a neural network:

$$a(u) := v(1) \text{ with } v : [0, 1] \rightarrow \mathbb{R}$$

$$\text{being the solution of } \frac{\partial v}{\partial \tau} = g(\tau, v(\tau)), \quad v(0) := u \quad (11)$$

In consequence, they suffer from the same issues as any continuous normalizing flow: they are slow as they require explicit integration of the underlying ODE. Besides requirements E1–E3, activation functions will profit from further desired properties: **D1**. having an explicit inverse, **D2**. having an explicit Jacobian and **D3**. having a bounded gradient. UMNN do not have desired property D1 and provide no guarantees for property D3.

Instead of parameterizing the scalar field  $g$  and learn it from data, we make an educated guess and choose a specific function with few parameters for which eq. 11 becomes explicitly solvable and requires no numerics at runtime: for the scalar field  $g(\tau, a; b) := \tanh(b \cdot a(\tau))$  the resulting ODE

$$\frac{\partial v}{\partial \tau} = \tanh(b \cdot v(\tau)), \quad v(0) := u$$

has an explicit solution (Sec. G; Yalavarthi et al. 2024b)

$$v(\tau; u, b) = \frac{1}{b} \sinh^{-1} (e^{b\tau} \cdot \sinh(b \cdot u))$$

yielding our activation function Shiesh:

$$\begin{aligned} \text{Shiesh}(u; b) &:= a(u) := v(1; u, b) \\ &= \frac{1}{b} \sinh^{-1} (e^b \sinh(b \cdot u)) \quad (12) \end{aligned}$$

being invertible, covering the whole real line and having no zero gradients (E1–E3) and additionally with analytical inverse and gradient (D1 and D2)

$$\begin{aligned} \text{Shiesh}^{-1}(u; b) &= \frac{1}{b} \sinh^{-1} (e^{-b} \cdot \sinh(b \cdot u)) \\ \frac{\partial}{\partial u} \text{Shiesh}(u; b) &= \frac{e^b \cosh(b \cdot u)}{\sqrt{1 + (e^b \sinh(b \cdot u))^2}} \quad (13) \end{aligned}$$

and bounded gradient (D3) in the range  $(1, e^b]$  (G.4, Yalavarthi et al. 2024b). Figure 1 shows a function plot. In our experiments we fixed its parameter  $b = 1$ .

Since the Shiesh is applied element-wise, it does not violate the Requirements 1, 2 and 3 in Section 3.

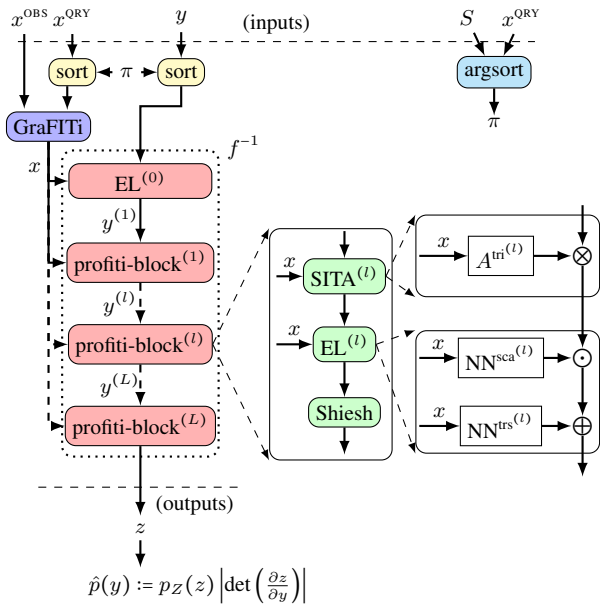


Figure 2: ProFITi architecture;  $\otimes$ : dot product,  $\odot$ : Hadamard product,  $\oplus$ : addition. Functions referred to their equation numbers: sort, argsort (eq. 8), GraFITi (eq. 16), SITA (eq. 10), EL (eq. 14), Shiesh (eq. 12). For efficiency, we perform, sorting only once directly on  $x^{\text{ORY}}$  and  $y$ .

## 6 Overall ProFITi Model Architecture

Invertible attention and the Shiesh activation function systematically model inter-dependencies between variables and non-linearity respectively, but do not move the zero point. To accomplish the latter, we use a third layer called **elementwise linear transformation layer (EL)**:

$$\text{EL}(y_k; x_k) := y_k \cdot \text{NN}^{\text{sca}}(x_k) + \text{NN}^{\text{trs}}(x_k) \quad (14)$$

where  $\text{NN}^{\text{sca}}$  and  $\text{NN}^{\text{trs}}$  are neural networks for scaling and translation.  $\text{NN}^{\text{sca}}$  is equipped with a  $\exp(\tanh)$  output function to make it positive and bounded, guaranteeing the inverse. We combine all three layers from eq. 7, 12, and 14 to a block

$$\text{profiti-block}(y; x) := \text{Shiesh}(\text{EL}(\text{SITA}(y; x); x)) \quad (15)$$

and stack  $L$  of those blocks to build the inverse transformation  $f^{-1}$  of our conditional invertible flow model ProFITi. We add a transformation layer with slope fixed to 1 as initial encoding on the  $y$ -side of the model. See Figure 2 for an overview of its architecture. As shown in the figure, for efficiency reasons we perform sorting (eq. 10) only once directly on the queries  $x^{\text{ORY}}$  and answers  $y$ .

**Encoder for Query embedding.** As discussed in Section 4, for probabilistic time series forecasting we have to condition on both, the past observations  $x^{\text{OBS}}$  and the queried time point/channel pairs  $x^{\text{ORY}}$  of interest. While in principle any equivariant encoder could be used, an encoder that leverages the relationships between those two pieces of the conditioner is crucial. We use GraFITi (Yalavarthi et al. 2024a),

a graph based equivariant point forecasting model for IMTS that provides state-of-the-art performance (in terms of accuracy and efficiency) as encoder

$$(x_1, \dots, x_K) := \text{GraFITi}(x_1^{\text{ORY}}, \dots, x_K^{\text{ORY}}, x^{\text{OBS}}) \quad (16)$$

The GraFITi encoder is trained end-to-end within the Profiti model, we did not pretrain it.

Note that for each query, other IMTS forecasting models yield a scalar, the predicted value, not an embedding vector. While it would be possible to use IMTS forecasting models as (scalar) encoders, due to their limitations to a single dimension we did not follow up on this idea.

**Training.** We train the ProFITi model  $\hat{p}$  for the normalized joint negative log-likelihood loss (njNLL; eq. 1) which written in terms of the transformation  $f^{-1}(\cdot; \cdot; \theta)$  of the normalizing flow and its parameters  $\theta$  yields:

$$\begin{aligned} \ell^{\text{njNLL}}(\theta) &:= \ell^{\text{njNLL}}(\hat{p}; p) \\ &= \mathbb{E}_{(x^{\text{OBS}}, x^{\text{ORY}}, y) \sim p} \frac{1}{|y|} \log p_Z(f^{-1}(y; x^{\text{OBS}}, x^{\text{ORY}}; \theta)) \\ &\quad \left| \det \left( \frac{\partial f^{-1}(y; x^{\text{OBS}}, x^{\text{ORY}}; \theta)}{\partial y} \right) \right| \end{aligned} \quad (17)$$

## 7 Experiments

### 7.1 Experiment for Joint Likelihoods

**Datasets.** We use 3 publicly available real-world medical IMTS datasets: **MIMIC-III** (Johnson et al. 2016), **MIMIC-IV** (Johnson et al. 2021), and **Physionet’12** (Silva et al. 2012). Datasets contain ICU patient records collected over 48 hours. The preprocessing procedures outlined in (Yalavarthi et al. 2024a; Biloš et al. 2021; De Brouwer et al. 2019) were applied. Observations in Physionet’12, MIMIC-III and MIMIC-IV were rounded to intervals of 1 hr, 30 min and 1 min respectively. We also evaluated on publicly available climate dataset **USHCN** (Menne, Williams Jr, and Vose 2015). It consists of climate data observed for 150 years from 1218 weather stations in USA.

**Baseline Models.** ProFITi is compared to 3 probabilistic IMTS forecasting models: **CRU** (Schirmer et al. 2022), **Neural-Flows** (Biloš et al. 2021), and **GRU-ODE-Bayes** (De Brouwer et al. 2019). To disentangle lifts originating from GraFITi (encoder) and those originating from ProFITi, we add **GraFITi+** as a baseline. GraFITi+ predicts an elementwise mean and variance of a normal distribution. As often interpolation models can be used seamlessly for forecasting, too, we include **HETVAE** (Shukla and Marlin 2022), a state-of-the-art probabilistic interpolation model, for comparison. Furthermore, we include Multi-task Gaussian Process Regression (**GPR**; Dürichen et al. 2015) as a baseline able to provide joint densities.

**Protocol.** We split the dataset into Train, Validation and Test in ratio 70:10:20, respectively. We select the hyperparameters from 10 random hyperparameter configurations based on their validation performance. We perform 5 fold cross validation with the chosen hyperparameters. Following (Biloš et al. 2021) and (Yalavarthi et al. 2024a), we use

	USHCN	$\frac{\text{time}}{\text{epoch}}$	Physionet' 12	$\frac{\text{time}}{\text{epoch}}$	MIMIC-III	$\frac{\text{time}}{\text{epoch}}$	MIMIC-IV	$\frac{\text{time}}{\text{epoch}}$
GPR	2.011±1.376	2s	1.367±0.074	35s	3.146±0.359	71s	2.789±0.057	227s
HETVAE	198.9±397.3	1s	0.561±0.012	8s	0.794±0.032	8s	OOM	–
GRU-ODE	0.766±0.159	100s	0.501±0.001	155s	0.961±0.064	511s	0.823±0.318	1052s
Neural-Flows	0.775±0.152	21s	0.496±0.001	34s	0.998±0.177	272s	0.689±0.087	515s
CRU	0.761±0.191	35s	0.741±0.001	40s	1.234±0.076	131s	OOM	–
GraFITi+	0.489±0.173	3s	0.367±0.021	32s	0.721±0.053	80s	0.287±0.040	84s
ProFITi (ours)	<b>-3.226±0.225</b>	6s	<b>-0.647±0.078</b>	59s	<b>-0.377±0.032</b>	97s	<b>-1.777±0.066</b>	123s
↑	3.5		1.0		1.1		2.1	

Table 2: Normalized Joint Negative Log-likelihood (njNLL), lower the better, best in bold, OOM indicates out of memory error, ↑ shows improvment in njNLL w.r.t. next best model.

	USHCN	Physionet' 12	MIMIC-III	MIMIC-IV
HETVAE	168.1±335.5	0.519±0.018	0.947±0.071	OOM
GRU-ODE	0.776±0.172	0.504±0.061	0.839±0.030	0.876±0.589
Neural-Flows	0.775±0.180	0.492±0.029	0.866±0.097	0.796±0.053
CRU	0.762±0.180	0.931±0.019	1.209±0.044	OOM
GraFITi+	<i>0.462±0.122</i>	<i>0.505±0.015</i>	<i>0.657±0.040</i>	<i>0.351±0.045</i>
ProFITi_marg (ours)	<b>-2.575±1.336</b>	<b>-0.368±0.033</b>	<b>0.092±0.036</b>	<b>-0.782±0.023</b>

Table 3: Results for Marginal Negative Log-likelihood (mNLL), lower the better. Best in bold and second best in italics. ProFITi\_marg is ProFITi trained for marginals.

the first 36 hours as observation range and forecast the next 3 time steps for medical datasets and first 3 years as observation range and forecast the next 3 time steps for climate dataset. Note that 3 time steps meaning do not mean 3 observations. For example, in Physionet' 12,  $K$  varies between 3 and 49. All models are implemented in PyTorch and run on GeForce RTX-3090 and 1080i GPUs. We compare the models for Normalized Joint Negative Log-likelihood (njNLL) loss (eq. 1). Except for GPR, and ProFITi, we take the average of the marginal negative log-likelihoods of all the observations in a series to compute njNLL for that series.

Sampling-based metrics like the Energy Score (Gneiting and Raftery 2007) not only suffer from the curse of dimensionality but also evaluate multivariate distributions improperly (Marcotte et al. 2023). Similarly, Continuous Ranked Probability Score sum (CRPS-sum; Rasul et al. 2021) for multivariate probabilistic forecasting can be misled by simple noise models, where random forecasts may outperform state-of-the-art methods (Koochali et al. 2022).

**Results.** Table 2 demonstrates the Normalized Joint Negative Log-likelihood (njNLL, lower the better) and run time per epoch for all the datasets. Best results are presented in bold. ProFITi outperforms all the prior approaches with significant margin on all the four datasets. While GraFITi+ is the second-best performing model, ProFITi outperforms it by an average gain of 2.0 in njNLL. We note that although GPR is predicting joint likelihoods, it performs poorly, likely because of having very few parameters. We note that njNLL of HETVAE is quite high for the USHCN dataset. The reason is HETVAE predicted a very small variance ( $10^{-4}$ ) for 1

sample whose predicted mean is farther from the target. We do not provide results for CRU on MIMIC-IV as our GPU (48GB VRAM) gives out of memory errors. The reason for such high likelihoods compared to the baseline models is i.) not assuming a Gaussian underlying distribution and ii.) directly predicting joint distributions (see Section 7.3).

## 7.2 Auxiliary Experiments for Marginals

Existing models in the related work (De Brouwer et al. 2019) and (Biloš et al. 2021) cannot predict multivariate distributions, hence their evaluation was restricted to Marginal Negative Log-likelihood (mNLL):

$$\ell^{\text{mNLL}}(\hat{p}; \mathcal{D}^{\text{test}}) := -\frac{\sum_{(x^{\text{OBS}}, x^{\text{QRY}}, y) \in \mathcal{D}^{\text{test}}} \sum_{k=1}^{|y|} \log \hat{p}(y_k | x^{\text{OBS}}, x_k^{\text{QRY}})}{\sum_{(x^{\text{OBS}}, x^{\text{QRY}}, y) \in \mathcal{D}^{\text{test}}} |y|} \quad (18)$$

We trained ProFITi only for marginals by removing SITA from the architecture and called ProFITi\_marg. Results presented in Table 3 shows that ProFITi\_marg outperforms baseline models again.

In addition to mNLL, we also compare the models in terms of CRPS score, a metric for marginal probabilistic forecasting in Table 4; and MSE, a metric for point forecasting in Table 5. While all the baseline models can compute CRPS and MSE explicitly from Gaussian parameters, ProFITi requires sampling. For this, we randomly sampled 100 samples. We computed MSE from robustious mean of samples which is the mean of the samples with outliers removed.

	USHCN	Physionet'12	MIMIC-III	MIMIC-IV
HETVAE	0.229±0.017	0.278±0.001	0.359±0.009	OOM
GRU-ODE	0.313±0.012	0.278±0.001	0.308±0.005	0.281±0.004
Neural-flows	0.306±0.028	0.277±0.003	0.308±0.004	0.281±0.004
CRU	0.247±0.010	0.363±0.002	0.410±0.005	OOM
GraFITi+	<i>0.222±0.011</i>	<i>0.256±0.001</i>	<i>0.279±0.006</i>	<i>0.217±.005</i>
ProFITi_marg (ours)	<b>0.192±0.019</b>	<b>0.253±0.001</b>	<b>0.276±0.001</b>	<b>0.206±0.001</b>

Table 4: Results for CRPS score, lower the better. Best in bold and second best in italics

	USHCN	Physionet'12	MIMIC-III	MIMIC-IV
HETVAE	0.298±0.073	0.304±0.001	0.523±0.055	OOM
GRU-ODE	0.410±0.106	0.329±0.004	0.479±0.044	0.365±0.012
Neural-Flows	0.424±0.110	0.331±0.006	0.479±0.045	0.374±0.017
CRU	<i>0.290±0.060</i>	0.475±0.015	0.725±0.037	OOM
GraFITi	<b>0.256±0.027</b>	<b>0.286±0.001</b>	<b>0.401±0.028</b>	<b>0.233±0.005</b>
ProFITi_marg (ours)	0.300±0.053	<i>0.295±0.002</i>	<i>0.443±0.028</i>	<i>0.246±0.004</i>

Table 5: Results for MSE, lower the better. Best in bold and second best in italics.

ProFITi\_marg outperforms all the baseline models in terms of CRPS score. On the other hand, for point forecasting, ProFITi\_marg is the second best model in terms of MSE, and GraFITi remains the best. While we leverage GraFITi as the encoder for ProFITi, ProFITi incorporates various components specifically designed to predict distributions, even if this sacrifices some point forecast accuracy. Also, GraFITi is trained to predict a Gaussian distribution, and the Mean Squared Error (MSE) is probabilistically equivalent to the negative log-likelihood of a Gaussian distribution with a fixed variance. This probabilistic interpretation can lead to better performance compared to models that are not specifically trained for Gaussian distributions. Additionally, models for uncertainty quantification often suffer from slightly worse point forecasts, as observed in several studies (Lakshminarayanan, Pritzel, and Blundell 2017; Seitzer et al. 2021; Shukla and Marlin 2022). In the domain of probabilistic forecasting, the primary metric of interest is negative log-likelihood, and ProFITi demonstrates superior performance.

### 7.3 Ablation studies: Varying model components

Model	Physionet2012
ProFITi	-0.647±0.078
ProFITi-SITA	-0.470±0.017
ProFITi-Shiesh	0.285±0.061
ProFITi-SITA-Shiesh	0.372±0.021
ProFITi-Shiesh+PReLU	0.384±0.060
ProFITi-Shiesh+LReLU	NaN Error
ProFITi- $A^{\text{tri}}$ + $A^{\text{iTrans}}$	-0.199±0.141
ProFITi- $A^{\text{tri}}$ + $A^{\text{reg}}$	-0.778±0.016

Table 6: Varying model components; Metric: njNLL; ProFITi-A+B: component A is removed and B is added.

To analyze ProFITi’s superior performance, we conduct an ablation study on Physionet’12 (Table 6). The Shiesh activation significantly improves performance by enabling learning of non-Gaussian distributions (compare ProFITi and ProFITi-Shiesh). Similarly, learning joint distributions (ProFITi) outperforms ProFITi-SITA in njNLL. ProFITi-SITA-Shiesh, which learns only Gaussian marginals, performs worse than ProFITi. Replacing Shiesh with PReLU (ProFITi-Shiesh+PReLU) degrades performance, and using Leaky-ReLU (LReLU) results in small Jacobians and vanishing gradients. The  $A^{\text{iTrans}}$  (ProFITi- $A^{\text{tri}}$  +  $A^{\text{iTrans}}$ ) variant performs poorly due to its limitation to positive covariances. While ProFITi with  $A^{\text{reg}}$  or  $A^{\text{tri}}$  shows similar results,  $A^{\text{reg}}$  faces scalability challenges: computing the determinant of a full attention matrix has  $\mathcal{O}(K^3)$  complexity, compared to  $\mathcal{O}(K)$  for triangular matrices. Additionally,  $A^{\text{reg}}$  underperforms for longer forecast horizons (see appendix H.2; Yalavarthi et al. 2024b). Further experiments on varying observation and forecast horizons and sparsity levels are provided in appendix H of Yalavarthi et al. (2024b).

## Conclusions

In this work, we propose a novel model ProFITi for probabilistic forecasting of irregularly sampled multivariate time series with missing values using conditional normalizing flows. ProFITi is designed to learn multivariate conditional distribution of varying length sequences. To the best of our knowledge, ProFITi is the first normalizing flow based model that can handle irregularly sampled time series with missing values. We propose two novel model components, sorted invertible triangular attention and Shiesh activation function in order to learn any random target distribution. Our experiments on four real-world datasets demonstrate that ProFITi provides significantly better likelihoods than existing models.

## Acknowledgments

This work was supported by the Federal Ministry for Economic Affairs and Climate Action (BMWK), Germany, within the framework of the IIP-Ecosphere project (project number: 01MK20006D); co-funded by the Lower Saxony Ministry of Science and Culture under grant number ZN3492 within the Lower Saxony “Vorab” of the Volkswagen Foundation and supported by the Center for Digital Innovations (ZDIN); and also by the German Federal Ministry of Education and Research (BMBF) through the Program “International Future Labs for Artificial Intelligence” (Grant IDD20002A – KIWI-BioLab)”

## References

- Anantha Padmanabha, G.; and Zabarab, N. 2021. Solving Inverse Problems Using Conditional Invertible Neural Networks. *Journal of Computational Physics*, 433: 110194.
- Ansari, A. F.; Heng, A.; Lim, A.; and Soh, H. 2023. Neural Continuous-Discrete State Space Models for Irregularly-Sampled Time Series. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, 926–951. PMLR.
- Behrmann, J.; Grathwohl, W.; Chen, R. T. Q.; Duvenaud, D.; and Jacobsen, J.-H. 2019. Invertible Residual Networks. In *Proceedings of the 36th International Conference on Machine Learning*, 573–582. PMLR.
- Biloš, M.; and Günnemann, S. 2021. Normalizing Flows for Permutation Invariant Densities. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 957–967. PMLR.
- Biloš, M.; Sommer, J.; Rangapuram, S. S.; Januschowski, T.; and Günnemann, S. 2021. Neural Flows: Efficient Alternative to Neural ODEs. *Advances in Neural Information Processing Systems*, 34: 21325–21337.
- Bonilla, E. V.; Chai, K.; and Williams, C. 2007. Multi-Task Gaussian Process Prediction. In *Advances in Neural Information Processing Systems*, volume 20.
- Chen, R. T.; Rubanova, Y.; Bettencourt, J.; and Duvenaud, D. K. 2018. Neural Ordinary Differential Equations. *Advances in neural information processing systems*, 31.
- Chen, Y.; Ren, K.; Wang, Y.; Fang, Y.; Sun, W.; and Li, D. 2024. ContiFormer: Continuous-time Transformer for Irregular Time Series Modeling. *Advances in Neural Information Processing Systems*, 36.
- de Bézenac, E.; Rangapuram, S. S.; Benidis, K.; Bohlke-Schneider, M.; Kurle, R.; Stella, L.; Hasson, H.; Gallinari, P.; and Januschowski, T. 2020. Normalizing Kalman Filters for Multivariate Time Series Analysis. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 2995–3007. Curran Associates, Inc.
- De Brouwer, E.; Simm, J.; Arany, A.; and Moreau, Y. 2019. GRU-ODE-Bayes: Continuous Modeling of Sporadically Observed Time Series. *Advances in neural information processing systems*, 32.
- Dinh, L.; Sohl-Dickstein, J.; and Bengio, S. 2017. Density Estimation Using Real NVP. In *International Conference on Learning Representations*.
- Dürichen, R.; Pimentel, M. A. F.; Clifton, L.; Schweikard, A.; and Clifton, D. A. 2015. Multitask Gaussian Processes for Multivariate Physiological Time-Series Analysis. *IEEE Transactions on Biomedical Engineering*, 62(1): 314–322.
- Gneiting, T.; and Raftery, A. E. 2007. Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, 102(477): 359–378.
- Horn, M.; Moor, M.; Bock, C.; Rieck, B.; and Borgwardt, K. 2020. Set Functions for Time Series. In *International Conference on Machine Learning*, 4353–4363. PMLR.
- Johnson, A.; Bulgarelli, L.; Pollard, T.; Horng, S.; and Celi, L.A. 2021. Mark. R. *MIMIC-IV (version 1.0)*. *PhysioNet*.
- Johnson, A. E.; Pollard, T. J.; Shen, L.; Lehman, L.-w. H.; Feng, M.; Ghassemi, M.; Moody, B.; Szolovits, P.; Anthony Celi, L.; and Mark, R. G. 2016. MIMIC-III, a Freely Accessible Critical Care Database. *Scientific data*, 3(1): 1–9.
- Khorashadizadeh, A.; Kothari, K.; Salsi, L.; Harandi, A. A.; de Hoop, M.; and Dokmanić, I. 2023. Conditional Injective Flows for Bayesian Imaging. *IEEE Transactions on Computational Imaging*, 9: 224–237.
- Kingma, D. P.; and Welling, M. 2013. Auto-Encoding Variational Bayes.
- Köhler, J.; Klein, L.; and Noe, F. 2020. Equivariant Flows: Exact Likelihood Generative Learning for Symmetric Densities. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 5361–5370. PMLR.
- Koochali, A.; Schichtel, P.; Dengel, A.; and Ahmed, S. 2022. Random Noise vs. State-of-the-Art Probabilistic Forecasting Methods: A Case Study on CRPS-Sum Discrimination Ability. *Applied Sciences*, 12(10): 5104.
- Kumar, M.; Babaeizadeh, M.; Erhan, D.; Finn, C.; Levine, S.; Dinh, L.; and Kingma, D. 2020. VideoFlow: A Conditional Flow-Based Model for Stochastic Video Generation. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and Scalable Predictive Uncertainty Estimation Using Deep Ensembles. *Advances in neural information processing systems*, 30.
- Li, S. C.-X.; and Marlin, B. 2016. A Scalable End-to-End Gaussian Process Adapter for Irregularly Sampled Time Series Classification. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS’16*, 1812–1820. Red Hook, NY, USA: Curran Associates Inc. ISBN 978-1-5108-3881-9.
- Li, S. C.-X.; and Marlin, B. M. 2015. Classification of Sparse and Irregularly Sampled Time Series with Mixtures of Expected Gaussian Kernels and Random Features. In *UAI*, 484–493.

- Li, Y.; Yi, H.; Bender, C.; Shan, S.; and Oliva, J. B. 2020. Exchangeable Neural ODE for Set Modeling. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 6936–6946. Curran Associates, Inc.
- Liu, J.; Kumar, A.; Ba, J.; Kiros, J.; and Swersky, K. 2019. Graph Normalizing Flows. *Advances in Neural Information Processing Systems*, 32.
- Marcotte, É.; Zantedeschi, V.; Drouin, A.; and Chapados, N. 2023. Regions of Reliability in the Evaluation of Multivariate Probabilistic Forecasts. In *International Conference on Machine Learning*, 23958–24004. PMLR.
- Menne, M. J.; Williams Jr, C.N.; and Vose, R. S. 2015. United States Historical Climatology Network Daily Temperature, Precipitation, and Snow Data. *Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, Oak Ridge, Tennessee*.
- Papamakarios, G.; Nalisnick, E.; Rezende, D. J.; Mohamed, S.; and Lakshminarayanan, B. 2021. Normalizing Flows for Probabilistic Modeling and Inference. *Journal of Machine Learning Research*, 22(1).
- Rasul, K.; Sheikh, A.-S.; Schuster, I.; Bergmann, U. M.; and Vollgraf, R. 2021. Multivariate Probabilistic Time Series Forecasting via Conditioned Normalizing Flows. In *International Conference on Learning Representations*.
- Rippel, O.; and Adams, R. P. 2013. High-Dimensional Probability Estimation with Deep Density Models. *arXiv preprint arXiv:1302.5125*.
- Rubanova, Y.; Chen, R. T.; and Duvenaud, D. K. 2019. Latent Ordinary Differential Equations for Irregularly-Sampled Time Series. *Advances in neural information processing systems*, 32.
- Satorras, V. G.; Hogeboom, E.; and Welling, M. 2021. E(n) Equivariant Graph Neural Networks. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 9323–9332. PMLR.
- Schirmer, M.; Eltayeb, M.; Lessmann, S.; and Rudolph, M. 2022. Modeling Irregular Time Series with Continuous Recurrent Units. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162, 19388–19405. PMLR.
- Seitzer, M.; Tavakoli, A.; Antic, D.; and Martius, G. 2021. On the Pitfalls of Heteroscedastic Uncertainty Estimation with Probabilistic Neural Networks. In *International Conference on Learning Representations*.
- Shukla, S. N.; and Marlin, B. 2022. Heteroscedastic Temporal Variational Autoencoder for Irregularly Sampled Time Series. In *International Conference on Learning Representations*.
- Si, P.; Kuleshov, V.; and Bishop, A. 2022. Autoregressive Quantile Flows for Predictive Uncertainty Estimation. In *International Conference on Learning Representations*.
- Silva, I.; Moody, G.; Scott, D. J.; Celi, L. A.; and Mark, R. G. 2012. Predicting In-Hospital Mortality of Icu Patients: The Physionet/Computing in Cardiology Challenge 2012. In *2012 Computing in Cardiology*, 245–248. IEEE.
- Sukthanker, R. S.; Huang, Z.; Kumar, S.; Timofte, R.; and Van Gool, L. 2022. Generative Flows with Invertible Attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11234–11243.
- Trippe, B. L.; and Turner, R. E. 2018. Conditional Density Estimation with Bayesian Normalising Flows. arxiv:1802.04908.
- van den Berg, R.; Hasenclever, L.; Tomczak, J. M.; and Welling, M. 2018. Sylvester Normalizing Flows for Variational Inference. In Globerson, A.; and Silva, R., eds., *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence, UAI 2018, Monterey, California, USA, August 6-10, 2018*, 393–402. AUAI Press.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention Is All You Need. *Advances in neural information processing systems*, 30.
- Villani, C. 2009. *Optimal Transport*, volume 338 of *Grundlehren Der Mathematischen Wissenschaften*. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-540-71049-3.
- Wehenkel, A.; and Louppe, G. 2019. Unconstrained Monotonic Neural Networks. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; dAlché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Winkler, C.; Worrall, D.; Hogeboom, E.; and Welling, M. 2019. Learning Likelihoods with Conditional Normalizing Flows. arxiv:1912.00042.
- Yalavarthi, V. K.; Burchert, J.; and Schmidt-Thieme, L. 2023. Tripletformer for Probabilistic Interpolation of Irregularly Sampled Time Series. *2023 IEEE International Conference on Big Data (Big Data)*.
- Yalavarthi, V. K.; Madhusudhanan, K.; Scholz, R.; Ahmed, N.; Burchert, J.; Jawed, S.; Born, S.; and Schmidt-Thieme, L. 2024a. GraFITi: Graphs for Forecasting Irregularly Sampled Time Series. In Wooldridge, M. J.; Dy, J. G.; and Natarajan, S., eds., *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada*, 16255–16263. AAAI Press.
- Yalavarthi, V. K.; Scholz, R.; Born, S.; and Schmidt-Thieme, L. 2024b. Probabilistic Forecasting of Irregular Time Series via Conditional Flows. arxiv:2402.06293.
- Zha, J.; Zhong, Y.; Zhang, J.; Hartley, R.; and Zheng, L. 2021. Invertible Attention. *CoRR*, abs/2106.09003.