RELIABLE PROBABILISTIC FORECASTING OF IRREG-ULAR TIME SERIES THROUGH MARGINALIZATION-CONSISTENT FLOWS

Anonymous authors

Paper under double-blind review

ABSTRACT

Probabilistic forecasting of joint distributions for irregular time series with missing values is an underexplored area in machine learning. Existing models, such as Gaussian Process Regression and ProFITi, are limited: while ProFITi is highly expressive due to its use of normalizing flows, it often produces unrealistic predictions because it lacks marginalization consistency—marginal distributions of subsets of variables may not match those predicted directly, leading to inaccurate marginal forecasts when trained on joints. We propose MOSES (Mixtures of Separable Flows), a novel model that parametrizes a stochastic process via a mixture of normalizing flows, where each component combines a latent multivariate Gaussian with separable univariate transformations. This design allows MOSES to be analytically marginalized, enabling accurate and reliable predictions for various probabilistic queries. Experiments on four datasets show that MOSES achieves highly accurate joint and marginal predictions. Thanks to its inherent marginalization consistency, it matches ProFITi in joint prediction performance while significantly outperforming all baselines on marginal predictions.

1 Introduction

Probabilistic forecasting of irregular time series requires models that can make predictions at *arbitrary* time points and for an arbitrary subset of variables (channels). Unlike regular time series forecasting, where predictions are confined to a fixed grid, irregular settings require modeling a **stochastic process** as the number of targets can vary.

A core requirement of any stochastic process model is **marginalization consistency**: the marginal distribution of a subset of variables must agree whether computed directly or by integrating out the joint. This property, guaranteed by **Kolmogorov's extension theorem**, ensures that the model defines a mathematically coherent stochastic process. Consistency is not only necessary for mathematical rigor but also highly desirable in practice: by the **data processing inequality**, accurate joint predictions together with consistency ensure reliable marginals, providing both correctness and performance guarantees.

Many recent methods prioritize flexibility over consistency. For example, ProFITi (Yalavarthi et al., 2025) achieves strong empirical performance using normalizing flows but violates marginalization consistency, producing contradictory marginal and joint predictions (Figure 1).

Example 1.1. In an ICU monitoring system, such a model might predict that blood pressure remains stable with 90% probability when queried individually, but only 60% when inferred from the joint distribution of vitals. These contradictions undermine reliability and could mislead clinical decisions.

Due to inconsistency, ProFITi suffers significantly on marginal distributions despite having good joints. On the other hand, existing consistent approaches, such as Gaussian Process Regression (GPR) (Dürichen et al., 2015), avoid this problem but are limited to Gaussian distributions, creating a false contradiction between correctness and flexibility.

We address this problem with **MOSES** (**Mixtures of Separable Flows**), a model based on mixtures of normalizing flows. Each component of MOSES uses a multivariate Gaussian, induced by a Gaussian process, as the base distribution, combined with univariate invertible transformations to modulate

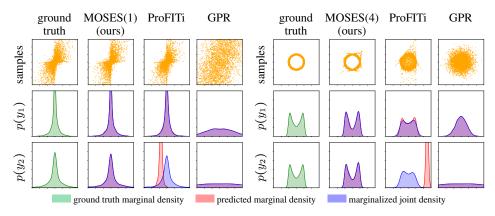


Figure 1: Demonstration of marginal consistency for MOSES (ours), ProFITi Yalavarthi et al. (2025), and Gaussian Process Regression Bonilla et al. (2007) on two toy datasets: blast (left) and circle (right). ProFITi is inconsistent w.r.t. the marginals of the second variable y_2 , while MOSES is consistent with the marginals of both y_1 and y_2 . MOSES(D) indicates D mixture components. Gaussian Process Regression (GPR) is marginalization consistent but predicts incorrect distributions.

the marginals—similar in spirit to Gaussian Copula Processes (Wilson & Ghahramani, 2010). Since mixtures, Gaussian processes, and univariate transformations are each marginalization consistent, MOSES is consistent **by design**. At the same time, it retains the expressiveness of flow-based methods.

Our contributions are as follows:

- We formalize marginalization consistency as a fundamental requirement for probabilistic forecasting models and show that violating it renders a model incoherent.
- 2. We propose MOSES which achieves the expressiveness of modern flow approaches while maintaining strict marginalization consistency through a principled architectural design.
- 3. We show that consistency does not come at high performance cost: MOSES matches inconsistent models on joint prediction tasks while significantly outperforming them on marginal predictions, providing a principled solution to irregular time series forecasting. PyTorch implementation is provided as a supplementary material.

2 PRELIMINARIES

We represent an irregular time series X as a sequence of N triplets (Horn et al., 2020; Yalavarthi et al., 2025):

$$X = \left(\left(t_n^{\text{OBS}}, c_n^{\text{OBS}}, v_n^{\text{OBS}} \right) \right)_{n=1:N} \in \text{Seq}(\mathcal{X}), \quad \mathcal{X} = \mathbb{R} \times \{1, \dots, C\} \times \mathbb{R}, \tag{1}$$

where $t_n^{\rm OBS}$ is the observation time, $c_n^{\rm OBS}$ the channel, and $v_n^{\rm OBS}$ the observed value.

A **time series query** Q is a sequence of K pairs:

$$Q = \left(\left(t_k^{\text{QRY}}, c_k^{\text{QRY}} \right) \right)_{k=1:K} \in \text{Seq}(\mathcal{Q}), \quad \mathcal{Q} = \mathbb{R} \times \{1, \dots, C\},$$
 (2)

where t_k^{QRY} and c_k^{QRY} specify the future time and queried channel.

A forecasting answer is $y = (y_1, \dots, y_K)$ with $y_k \in \mathbb{R}$ predicting the value in channel c_k^{QRY} at t_k^{QRY} . All query points occur after the observations: $\min_k t_k^{QRY} > \max_n t_n^{OBS}$. Here, $Seq(\mathcal{X})$ denotes the space of finite sequences over \mathcal{X} .

Requirements. A marginalization consistent probabilistic irregularly sampled time series forecasting model must satisfy the following requirements:

R1 Joint Multivariate Prediction. The goal of probabilistic irregular time series forecasting is to model the joint distribution $\hat{p}(y \mid Q, X)$ over responses y, given query points Q and observed

series X, allowing both context length N = |X| and query length K = |Q| to vary dynamically:

$$\hat{p}: \operatorname{Seq}(\mathbb{R} \times \mathcal{Q}) \times \operatorname{Seq}(\mathcal{X}) \to \mathbb{R}_{\geq 0},$$

$$(y, Q, X) \mapsto \hat{p}(y_1, \dots, y_K \mid Q_1, \dots, Q_K, X_1, \dots, X_N)$$
(3)

For a given (Q, X), the mapping $(y_1, \dots, y_K) \mapsto \hat{p}(y_1, \dots, y_K \mid Q, X)$ defines a probability density on $\mathbb{R}^{|Q|}$.

Unlike standard multivariate time series forecasting, where future time points are fixed and typically ignored, irregular time series require conditioning on future time points and channels, as these can vary across instances.

R2 Permutation Invariance. As the time stamp and channel-ID are included in each sample, the order of the samples does not matter, and hence any model prediction should be independent of the order of both the query or context:

$$\hat{p}(y \mid Q, X) = \hat{p}(y^{\pi} \mid Q^{\pi}, X^{\tau}), \ \forall \pi \in S_{|Q|}, \tau \in S_{|X|}$$
(4)

R3 Marginalization Consistency/Projection Invariance. Predicting the joint density for the subquery Q_{-k} given by removing the k-th item from Q should yield the same result as marginalizing the k-th variable from the complete query Q.

$$\hat{p}(y_{-k} \mid Q_{-k}, X) = \int_{\mathbb{R}} \hat{p}(y \mid Q, X) \, dy_k \tag{5}$$

This generalizes to any subset $K_S \subseteq \{1, \dots, K\}$.

For a model satisfying **R1-R3**, we will only have to marginalize if we try to validate the marginalization consistency. For this validation we added requirement **R3**. Yalavarthi et al. (2025) discussed **R1** and **R2**, but did not consider **R3**. We argue that irregularly sampled time series are realizations of a stochastic process and **R3** is a fundamental property of any model that mimics it.

Theorem 2.1. Any model that satisfies **R1-R3** realizes an \mathbb{R} -valued stochastic process over the index set $T = \mathbb{R} \times \{1, \dots, C\}$.

Proof. This is a direct application of Kolmogorov's extension theorem (Øksendal, 2003)

WHY DO WE NEED MARGINALIZATION CONSISTENCY?

- 1. Without marginalization consistency, probabilistic models can be unreliable. Two test cases with the same context X but different queries Q may share overlapping targets. An inconsistent model can assign different distributions to these targets, despite identical conditioning information, which is unrealistic. Such models are inapplicable in many real-world applications.
- 2. Marginalization consistency provides performance guarantees. For a consistent model, if the joint prediction over K points is accurate, predictions for any subset of size < K are also guaranteed to be accurate. This follows from the **data processing inequality (DPI; Murphy, 2022)**, which ensures marginalizing cannot increase KL divergence:

$$D_{KL}(p(y_1, ..., y_K \mid Q_1, ..., Q_K, X) \parallel \hat{p}(y_1, ..., y_K \mid Q_1, ..., Q_K, X))$$

$$\geq D_{KL}(p(y_1, ..., y_{K-1} \mid Q_1, ..., Q_{K-1}, X) \parallel \hat{p}(y_1, ..., y_{K-1} \mid Q_1, ..., Q_{K-1}, X))$$

Intuitively, marginalization "averages out" errors, so accuracy cannot worsen on subsets. Consistent models thus remain reliable for marginal predictions, as confirmed in Tables 1 and 2, where inconsistent model, ProFITi, perform poorly on marginals despite good joint predictions.

3 CONSTRUCTING MARGINALIZATION CONSISTENT CONDITIONAL DISTRIBUTIONS

Our goal is to build a model for the conditional joint distribution $p(y_1, \ldots, y_K \mid Q_1, \ldots, Q_k, X)$, as in (3). Since the model should satisfy **R3**, it follows that the marginal distribution of y_k must only depend on Q_k and X.

ground truth

MOSES (1) MOSES (2) MOSES (3) MOSES (4)

GMM (1) GMM (5) GMM (10) GMM (15)

Figure 2: (Top) Importance of multiple flow components: MOSES(1) cannot represent the correct distribution, but MOSES(4) can. (Bottom) Limitation of Gaussian Mixture Models: GMM needs 15 components to match the distribution of MOSES(4).

Separably Parametrized Gaussians. A simple way to model a permutation-invariant conditional distribution for a variable number of targets is a multivariate Normal distribution $\mathcal{N}(y \mid \mu(X,Q),\Sigma(X,Q))$, where the mean and covariance are **separably parametrized**:

$$\mu_k = \tilde{\mu}(Q_k, X), \quad \Sigma_{k,\ell} = \widetilde{\Sigma}(Q_k, Q_\ell, X),$$
 (6)

with functions $\tilde{\mu}$ and $\widetilde{\Sigma}$ as in Gaussian Processes. Such a setup is **marginalization consistent** by design, since marginalizing a Gaussian simply involves selecting the relevant rows and columns of Σ and elements of μ . However, Gaussian Processes are limited because they can only model joint Gaussian distributions. For more flexible distributions, **normalizing flows** are a common alternative (Rezende & Mohamed, 2015).

Separable Normalizing Flows. Normalizing flows model a distribution by applying an invertible transformation $f: \mathbb{R}^K \to \mathbb{R}^K$ to a source distribution p_Z on \mathbb{R}^K . The resulting target distribution is given by the change-of-variable formula:

$$p_Y(y) = p_Z(f^{-1}(y;\theta)) \cdot \left| \det \left(\frac{\partial f^{-1}(y;\theta)}{\partial y} \right) \right|. \tag{7}$$

Most existing flows use simple source distributions, typically a standard multivariate normal, and model interactions between variables through the transformation (Rezende & Mohamed, 2015; Papamakarios et al., 2021). Conditional flows for a variable number of targets follow the same idea, designing expressive transformations for vectors of arbitrary size (Liu et al., 2019; Biloš & Günnemann, 2021; Yalavarthi et al., 2025). However, these models generally **do not guarantee marginalization consistency**, and no simple condition on the transform alone can provide it.

We propose a different approach: use **simple, separable transformations** combined with a **richer source distribution**, namely a Gaussian Process with full covariance. In this setup, dependencies between variables are captured by the source distribution rather than the transformation, enabling marginalization consistency. This is similar in spirit to copula-based architectures, where dependencies between variables can be separated from their univariate marginal distributions.

Lemma 3.1. A conditional flow model over \mathbb{R}^K or $Seq(\mathbb{R})$ is separable, if it is expressed in the form

$$f(z \mid Q, X) = (\phi(z_1 \mid Q_1, X), \dots, \phi(z_K \mid Q_K, X))$$
(8)

for some univariate function $\phi: \mathbb{R} \times \mathcal{Q} \times \operatorname{Seq}(\mathcal{X}) \to \mathbb{R}$, that is invertible in the first argument. Any model that consists of such a separable flow transformation, combined with a marginalization consistent model for the source distribution, is itself marginalization consistent. (Proof: Appendix A.1)

Conditional Mixtures of Flows. Using Gaussians as base distributions, we can only linear dependencies between variables. To increase expressiveness, we combine multiple separable flows into a mixture. Even a few components can significantly improve the model, reaching performance comparable to a simple GMM without flow transformations (Figure 2, Appendix C.7).

Figure 3: Illustration of MOSES. D-many flows (fixed). K-many variables (variable). Encoder (enc) takes X,Q (observed series and query timepoint-channel ids.) as input, and outputs an embedding \mathbf{h} (depends on both X, and Q) and w (depends on X only). μ,Σ of p_{Z_d} are parametrized by \mathbf{h}_d . Flow transformation of p_{Z_d} : parametrized by \mathbf{h}_d . Transformation layer: K-many univariate transformations ϕ that transforms z_k of $z \sim p_{Z_d}(z \mid \mathbf{h}_D)$ to y_k of $y \sim p_d^{\mathsf{FLOW}}(y \mid \mathbf{h}_d)$.

Lemma 3.2. Given probabilistic models $(\hat{p}_d)_{d=1:D}$ that satisfy **R1-R3**, then a mixture model

$$\hat{p}(y \mid Q, X) = \sum_{d=1}^{D} w_d(X) \, \hat{p}_d(y \mid Q, X) \tag{9}$$

with permutation invariant weight function $w: \operatorname{Seq}(\mathcal{X}) \to \Delta^D$, were Δ^D denotes probability simplex in D variables: $\Delta^D := \{ w \in \mathbb{R}^D \mid w_d \geq 0, \sum_d w_d = 1 \}$, also satisfies **R1-R3**. (Proof: Appendix A.2)

4 MIXTURES OF SEPARABLE FLOWS (MOSES)

Based on the constructions from the previous section, we build a marginalization-consistent model for forecasting irregular time series using four components (Figure 3):

1. Separable encoder:

- A shared encoding of the observations, $\mathbf{h}^{\text{OBS}} := \text{enc}^{\text{OBS}}(X; \theta^{\text{OBS}})$, used for all queries.
- D query-specific encodings, $\mathbf{h}_{d,k} \coloneqq \mathrm{enc}^{\mathrm{QRY}}(Q_k, X; \theta_d^{\mathrm{QRY}})$, capturing each query and the entire context.
- 2. **Separable Gaussians:** D multivariate Gaussians, $p_{Z_d}(z \mid \mu_d, \Sigma_d)$, parametrized separately by the query encodings \mathbf{h}_d .
- 3. Separable transformations: D many separable transformations f_d are parametrized using the encoded queries \mathbf{h}_d . Each f_d is applied on top of Gaussian base p_{Z_d} providing a separable normalizing flow \hat{p}_d^{FLOW} .
- 4. **Mixture of flows:** The D many separable normalizing flows are combined using mixing weights $w := w(\mathbf{h}^{\text{OBS}})$, which depend only on the shared observation encoding \mathbf{h}^{OBS} , not on the queries.
- **1. Separable Encoder.** To encode both the observations $X = ((t_n^{\text{OBS}}, c_n^{\text{OBS}}, v_n^{\text{OBS}}))_{n=1:N}$ and queries $Q = ((t_k^{\text{QRY}}, c_k^{\text{QRY}}))_{k=1:K}$, we apply a positional embedding with learnable parameters $(a_f, b_f)_{f=1:F}$ to the time component (Kazemi et al., 2019).

$$pos_embed(t)_f := \begin{cases} a_f t + b_f & \text{if } f = 1\\ \sin(a_f t + b_f) & \text{else} \end{cases}$$
 (10)

And one-hot encodings for the channel component. The value is simply passed through.

$$\mathbf{x} \coloneqq [\operatorname{pos_embed}(t_n^{\operatorname{OBS}}), \operatorname{one-hot}(c_n^{\operatorname{OBS}}), v_n^{\operatorname{OBS}}]_{n=1:N}$$
 (11a)

$$\mathbf{q} \coloneqq [\mathsf{pos_embed}(t_k^{\mathsf{QRY}}), \mathsf{one\text{-}hot}(c_k^{\mathsf{QRY}})]_{k=1:K} \tag{11b}$$

The observations $\mathbf{x} \in \mathbb{R}^{N \times (F+C+1)}$ are further encoded via self-attention and the queries $\mathbf{q} \in \mathbb{R}^{K \times (F+C)}$ via cross-attention w.r.t. the encoded observations:

$$\mathbf{h}^{\text{OBS}} := \text{MHA}(\mathbf{x}, \mathbf{x}, \mathbf{x}; \theta^{\text{OBS}}) \qquad (\in \mathbb{R}^{N \times M})$$
 (12a)

$$\widetilde{\mathbf{h}} := \mathrm{MHA}(\mathbf{q}, \mathbf{h}^{\mathrm{OBS}}, \mathbf{h}^{\mathrm{OBS}}; \theta^{\mathrm{QRY}}) \quad (\in \mathbb{R}^{K \times D \cdot M})$$
 (12b)

$$\mathbf{h} := \operatorname{reshape}(\widetilde{\mathbf{h}}) \qquad (\in \mathbb{R}^{D \times K \times M})$$
 (12c)

where MHA denotes multihead attention. For the encoding of the queries we use an encoding dimension $D \cdot M$ and reshape each \mathbf{h}_k into D encodings $\mathbf{h}_{d,k}$ of dimension M.

2. Separable Gaussians $p_{Z_d}(z \mid \mu_d, \Sigma_d)$ **.** We model the means and covariances of the source multivariate Gaussian using simple linear and quadratic functions of the encoded queries \mathbf{h}_d :

$$\mu(\mathbf{h}_d) = \mathbf{h}_d \theta^{\text{MEAN}} \qquad \Longrightarrow \mu(\mathbf{h}_d)_k = \mathbf{h}_{d,k} \theta^{\text{MEAN}},$$
 (13a)

$$\Sigma(\mathbf{h}_d) = \mathbb{I}_K + \frac{(\mathbf{h}_d \theta^{\text{cov}})(\mathbf{h}_d \theta^{\text{cov}})^T}{\sqrt{M'}} \implies \Sigma(\mathbf{h}_d)_{k,l} = \delta_{kl} + \frac{(\mathbf{h}_{d,k} \theta^{\text{cov}})(\mathbf{h}_{d,l} \theta^{\text{cov}})^T}{\sqrt{M'}}, \quad (13b)$$

where $\theta^{\text{MEAN}} \in \mathbb{R}^{M \times 1}$ and $\theta^{\text{COV}} \in \mathbb{R}^{M \times M'}$ are trainable weights shared across all D mixture components, \mathbb{I}_K is the $K \times K$ identity matrix, and δ_{kl} is the Kronecker delta.

The scaling by $\sqrt{M'}$ in (13b) ensures stable learning, following (Vaswani et al., 2017). Since $\Sigma(\mathbf{h}_d)$ is the sum of a positive semi-definite and a positive definite matrix, it remains positive definite. The encoding \mathbf{h}_d incorporates both the context X and the queries Q, consistent with the separable Gaussian setup in (6). Note that, while the source Gaussian is defined via Gaussian Processes (GP), we do not perform GP-style inference.

3. Separable transformations f_d . To obtain separable invertible transformations, we apply a univariate bijective function to each variable independently. Spline-based functions are particularly popular due to their expressiveness and good generalization (Durkan et al., 2019; Dolatabadi et al., 2020). We use computationally efficient Linear Rational Spline (LRS) transformations (Dolatabadi et al., 2020). For a conditional LRS $\phi(z_k; \mathbf{h}_{d,k}, \theta^{\text{FLOW}})$, parameters such as bin widths and heights, derivatives at the knots, and λ are computed from the conditioning input $\mathbf{h}_{d,k}$ and shared model parameters θ^{FLOW} . The same θ^{FLOW} is used for all variables $z_{1:K}$, allowing the transformation to handle varying numbers of variables, and is also shared across all D mixture components. For details, see Appendix A.4. In summary, the conditional flow model is separable across the query size $f = f_1 \times \cdots \times f_K$ with

$$f_d(y) := f(y \mid \mathbf{h}_d) = (\phi(y_1, \mathbf{h}_{d,1}), \dots, \phi(y_K, \mathbf{h}_{d,K}))$$
 (14)

4. Mixture Model. We model the mixture weights via cross attention, using trainable parameters $\beta \in \mathbb{R}^{D \times M}$ as attention queries, and a softmax to ensure the weights to sum to 1:

$$w := \operatorname{softmax}(MHA(\beta, \mathbf{h}^{OBS}, \mathbf{h}^{OBS}; \theta^{MIX}))$$
(15)

Theorem 4.1. Our model, MOSES, satisfies R1-R3 and hence realizes a stochastic process via Kolmogorov's Extension Theorem (see Theorem 2.1). Proof. See Appendix A.3.

Computational Complexities. The D-separable flows are computationally efficient: since they are separable, their Jacobian matrix is diagonal and computing determinant requires $\mathcal{O}(K)$ operations. The main computational cost lies in evaluating Σ_d^{-1} and $\det \Sigma_d^{-1}$ for the base distribution, which typically requires $\mathcal{O}(K^3)$ operations. However, for large K, our low-rank modification $\Sigma_d = \mathbb{I}_K + UU^T$ (see (13b)) reduces their computation to $\mathcal{O}(M'^2K)$ using the Woodbury and Weinstein–Aronszajn identities. This approach scales well for large values of $K \gg M'$, as M' is independent of K.

Training. Given a batch \mathcal{B} of training instances (Q, X, y), we minimize the normalized joint negative log-likelihood (njNLL) (Yalavarthi et al., 2025):

$$\mathcal{L}^{\text{njNLL}}(\theta) = \frac{1}{|\mathcal{B}|} \sum_{(Q,X,y)\in\mathcal{B}} -\frac{1}{|y|} \log \hat{p}(y \mid Q, X)$$
 (16)

where θ includes all the model parameters. njNLL generalizes NLL to varying target sizes.

Relationship to Copulas. Each component of MOSES is conceptually similar to a Gaussian Copula Process Wilson & Ghahramani (2010), in that it separates the linear dependency structure from the marginal distributions. The covariance matrix of base Gaussian can be decomposed into a correlation matrix and diagonal variances. The diagonal variances, together with the base mean and the univariate spline transformations, determine contribution of each component to the marginal distribution of each variable. On the other hand, since MOSES is a mixture of D such components—each with its own transformations—the overall model does not have a simple separation between dependency structure and marginal distributions.

Unlike many copula-based architectures that train marginals and dependencies separately to avoid identifiability issues, MOSES is trained end-to-end, and fully identifiable due to the restricted domain of the spline transformations. The shift in mean of the base distribution cannot be compensated by the spline transformation.

To the best of our knowledge, copula-based architectures have not been applied to probabilistic forecasting of irregular time series. Models such as TACTiS and TACTiS-2 (Drouin et al., 2022; Ashok et al., 2024) are designed for regularly sampled, fully observed data. While their principles could, in principle, extend to irregular time series, their implementations are not readily adaptable and they lack marginalization consistency due to their non-separable encoder and copula structure.

5 RELATED WORK

There have been multiple works that deal with point forecasting of irregular time series (Ansari et al., 2023; Che et al., 2018; Chen et al., 2024; Yalavarthi et al., 2024; Zhang et al., 2024). In this work we deal with probabilistic forecasting of irregular time series. Models such as NeuralFlows (Biloš et al., 2021), GRU-ODE (De Brouwer et al., 2019), and CRU (Schirmer et al., 2022) predict only the marginal distribution for a single time stamp. Additionally, interpolation models like HetVAE (Shukla & Marlin, 2022) and Tripletformer (Yalavarthi et al., 2023) can also be applied for probabilistic forecasting. However, they also produce only marginal distributions. All the above models assume underlying distribution is Gaussian which is not the case for many real-world datasets. On the other hand, Gaussian Process Regression (GPR; Dürichen et al., 2015), and ProFITi (Yalavarthi et al., 2025) can predict proper joint distributions. ProFITi is not marginalization consistent because of non-separable encoder and probabilistic component.

There have been works on models for tractable and consistent marginals for fixed size variables (e.g. tabular data). Probabilistic Circuits (Choi et al., 2020) create a sum-prod network on the marginal distributions. Later, Sidheekh et al. (2023) added univariate normalizing flows to the leaf nodes of the circuit for better expressivity. However, it is not trivial to extend such circuits to deal with sequential data of variable size. Gaussian Mixture Models (GMMs) (Duda & Hart, 1974) are often used only for unconditional density estimation, but can be extended to conditional density estimation. They can provide tractable and consistent marginal distributions. However, GMMs are not expressive enough and often require a very large number of components to approximate even simple distributions, see Figure 2. Note that normalizing flow models such as Dinh et al. (2017); Papamakarios et al. (2017; 2021) neither provide tractable marginals nor are applicable to varying number of variables.

Prior work on mixtures of normalizing flows has focused on fixed-length sequences. Pires & Figueiredo (2020) and Ciobanu (2021) used affine coupling and masked autoregressive flows for density estimation, while Postels et al. (2021) applied them to reconstruction tasks. These models face challenges with variable-length sequences and intractable marginals. Sendera et al. (2021) introduced non-Gaussian Gaussian Processes for few-shot learning but only support single-variable prediction. In contrast, MOSES handles multiple variables and missing values.

6 EXPERIMENTS

6.1 TOY EXPERIMENT

We illustrate marginalization consistency using two synthetic bivariate distributions (Blast and Circle; see Figure 1, Appendix B). The task is to estimate the unconditional joint distribution. MOSES accurately models both joint and marginal distributions while preserving consistency. In contrast,

Table 1: Comparing njNLL. Lower the better, best results in bold, second best in italics.

	Model	USHCN	Physionet'12	MIMIC-III	MIMIC-IV
inconsistent	ProFITi	-3.226 ± 0.225	-0.647 ± 0.078	-0.377 ± 0.032	-1.777 ± 0.066
consistent univariate	GRU-ODE NeuralFlows CRU Tripletformer+	0.766 ± 0.159 0.775 ± 0.152 0.761 ± 0.191 4.632 ± 8.179	0.501 ± 0.001 0.496 ± 0.001 1.057 ± 0.007 0.519 ± 0.112	0.961 ± 0.064 0.998 ± 0.177 1.234 ± 0.076 1.051 ± 0.141	0.823 ± 0.318 0.689 ± 0.087 OOM 0.686 ± 0.115
consistent multivariate	GPR GMM	2.011 ± 1.376 1.050 ± 0.031	1.367 ± 0.074 1.063 ± 0.002	3.146 ± 0.359 1.160 ± 0.020	2.789 ± 0.057 1.076 ± 0.003
	MOSES (ours)	-3.357 ± 0.176	-0.491 ± 0.041	-0.305 ± 0.027	-1.668 ± 0.097

Table 2: Trained for njNLL and evaluate for mNLL, lower the better.

	Model	USHCN	Physionet'12	MIMIC-III	MIMIC-IV
inconsistent	ProFITi	-3.324 ± 0.206	-0.016 ± 0.085	0.408 ± 0.030	0.500 ± 0.322
	GRU-ODE	0.776 ± 0.172	0.504 ± 0.061	0.839 ± 0.030	0.876 ± 0.589
consistent	Neural-Flows	0.775 ± 0.180	0.492 ± 0.029	0.866 ± 0.097	0.796 ± 0.053
univariate	CRU	0.762 ± 0.180	0.931 ± 0.019	1.209 ± 0.044	OOM
	Tripletformer+	0.411 ± 7.506	0.524 ± 0.110	0.894 ± 0.083	0.751 ± 0.063
consistent	GPR	1.235 ± 0.096	1.161 ± 0.065	1.341 ± 0.009	1.161 ± 0.010
multivariate	GMM	1.042 ± 0.021	1.069 ± 0.002	1.124 ± 0.007	1.075 ± 0.007
inani ranace	MOSES (ours)	-3.355 ± 0.156	-0.271 ± 0.028	0.163 ± 0.026	-0.634 ± 0.017

ProFITi captures the joint well—especially for Blast—but fails on marginals due to its triangular attention mechanism, which enforces a fixed dependency order. Gaussian Process Regression maintains consistency but lacks predictive accuracy. To predict the marginalization inconsistency we use **2-Wasserstein distance (WD)**. For each variable y_k , we compare:

- i) $\hat{p}(y_k \mid Q_k, X)$: the predicted marginal,
- ii) $\hat{p}^{MAR}(y_k \mid Q_k, X)$: the marginal obtained by integrating the joint $\hat{p}(y \mid Q, X)$.

Since sampling directly from \hat{p}^{MAR} is difficult, we sample from the joint and extract the k-th component. The marginalization inconsistency is defined as the average WD across all K variables:

$$MI = \frac{1}{K} \sum_{k=1}^{K} WD(\hat{p}(y_k \mid Q_k, X), \hat{p}^{MAR}(y_k \mid Q_k, X)).$$
(17)

We use 1000 samples to compute MI. Currently, we compute only univariate marginals; multivariate marginals are possible in principle but computationally prohibitive.

6.2 Main experiment

We evaluate our model on four real-world datasets: one climate dataset (USHCN) and three medical datasets (Physionet'12, MIMIC-III, and MIMIC-IV). Following prior work (Yalavarthi et al., 2025; Biloš et al., 2021), we observe the first 36h and predict the next 3 time steps for medical datasets, and observe 3 years and predict 3 time steps for USHCN. Both the number of observations (N) and queries (K) vary across samples (see Table 4). We split each dataset into training, validation, and test sets using a 70:10:20 ratio. We train MOSES using the Adam optimizer with a learning rate of 0.001 and batch size of 64. Hyperparameter search is over mixture components $D \in \{1, 2, 5, 7, 10\}$, attention heads $\in \{1, 2, 4\}$, and latent sizes $M, F \in \{16, 32, 64, 128\}$. All models are implemented in PyTorch and trained on NVIDIA RTX 3090 and GTX 1080 Ti GPUs.

Baselines. We use NeuralFlows (Biloš et al., 2021), GRU-ODE (De Brouwer et al., 2019), CRU (Schirmer et al., 2022), GPR (Dürichen et al., 2015), and ProFITi (Yalavarthi et al., 2025). Our encoder is similar to Tripletformer (Yalavarthi et al., 2023) that predict marginal distributions

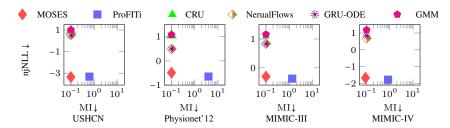


Figure 4: njNLL vs. MI. MOSES is marginalization consistent within sampling error.

Table 3: Comparing njNLL and mNLL across datasets to verify the contribution of probabilistic component of ProFITi. "ProFITi-TF" denotes ProFITi-Transformer using same encoder as MOSES

Dataset	njN	ILL	mN	ILL
Danaser	ProFITi-TF	MOSES	ProFITi-TF	MOSES
USHCN Physionet'12 MIMIC-III MIMIC-IV	-3.415±0.271 -0.657±0.034 0.516±0.111 -1.405±0.220	-3.357±0.176 -0.491±0.041 -0.305±0.027 -1.668±0.097	-3.440±0.243 0.017±0.042 1.279±0.057 0.345±0.325	-3.355±0.156 -0.271±0.028 0.163±0.026 -0.634±0.017

for interpolation. We used it for the forecasting and called the model Tripletformer+. NeuralFlows, GRU-ODE, CRU, and Tripletformer+ predict only marginals and are marginalization consistent, as their joint distribution is the product of marginals. GPR is also marginalization consistent. We also compare with Gaussian Mixture Model (GMM) which is MOSES without flows attached to highlight the advantage of flows in MOSES.

Results. To highlight the importance of Marginalization Consistency in probabilistic forecasting models we train the model for njNLL in (17) and evaluate for two metrics: 1. Normalized Joint Negative Log-Likelihood (njNLL; Table 1) and 2. Marginal Negative Log-Likelihood (mNLL; Table 2). While njNLL measures joint density, mNLL assesses univariate marginal density (Biloš et al., 2021; Schirmer et al., 2022). A good model must perform well on both the metrics.

MOSES outperforms all marginalization-consistent models across both metrics. Although it performs similarly or slightly worse than ProFITi on njNLL, it significantly surpasses ProFITi on mNLL. For USHCN, both models perform comparably within standard deviation. From Figure 4, our model achieves similar likelihoods to ProFITi, with MI values close to zero, while ProFITi has up to an order of magnitude larger due to sampling (rounded to 0.1). This difference arises because ProFITi prioritizes joint distributions but neglects marginalization consistency, leading to performance degradation when queried on single instances (Figure 1). In contrast, our model maintains consistency. Similar results are observed when comparing Energy Score (for multivariate distributions) and CRPS (for univariate distributions) as seen in Tables 8 and 9. Additionally, ProFITi's gains primarily stem from its encoder. When using the same encoder (ProFITi-TF), our model achieves superior accuracy in MIMIC-III and MIMIC-IV (Table 3).

CONCLUSIONS

We introduced MOSES, a marginalization-consistent mixture of separable flows for probabilistic forecasting of irregular time series with missing values. By carefully parametrizing its components, we ensured both decomposability and marginalization consistency. Experiments on four real-world irregularly sampled datasets show that MOSES achieves substantial gains in marginal predictions while preserving competitive accuracy in joint predictions compared to more flexible but inconsistent models. This work provides an initial step toward addressing marginalization inconsistency in probabilistic forecasting. Future research will likely focus on enhancing performance while preserving consistency. A key limitation of MOSES is that mixture weights are independent of the query Q. Although adapting weights to query time points could improve joint predictions, requirements R1–R3 enforce query-independence as a necessary condition for marginalization consistency.

REFERENCES

- Abdul Fatir Ansari, Alvin Heng, Andre Lim, and Harold Soh. Neural continuous-discrete state space models for irregularly-sampled time series. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 926–951. PMLR, 2023.
- Arjun Ashok, Étienne Marcotte, Valentina Zantedeschi, Nicolas Chapados, and Alexandre Drouin. TACTiS-2: Better, faster, simpler attentional copulas for multivariate time series. In *The Twelfth International Conference on Learning Representations*, 2024.
- Marin Biloš and Stephan Günnemann. Normalizing flows for permutation invariant densities. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 957–967. PMLR, 2021.
- Marin Biloš, Johanna Sommer, Syama Sundar Rangapuram, Tim Januschowski, and Stephan Günnemann. Neural flows: Efficient alternative to neural ODEs. *Advances in Neural Information Processing Systems*, 34:21325–21337, 2021.
- Edwin V Bonilla, Kian Chai, and Christopher Williams. Multi-task gaussian process prediction. In *Advances in Neural Information Processing Systems*, volume 20, 2007.
- Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. Recurrent neural networks for multivariate time series with missing values. *Scientific reports*, 8(1):1–12, 2018. doi: 10.1038/s41598-018-24271-9.
- Yuqi Chen, Kan Ren, Yansen Wang, Yuchen Fang, Weiwei Sun, and Dongsheng Li. ContiFormer: Continuous-time transformer for irregular time series modeling. *Advances in Neural Information Processing Systems*, 36, 2024.
- YooJung Choi, Antonio Vergari, and Guy Van den Broeck. Probabilistic Circuits: A Unifying Framework for Tractable Probabilistic Models*. 2020.
- Sebastian Ciobanu. Mixtures of normalizing flows. In *Proceedings of ISCA 34th International Conference on Computer Applications in Industry and Engineering, EPiC Series in Computing*, volume 79, pp. 82–90, 2021.
- Edward De Brouwer, Jaak Simm, Adam Arany, and Yves Moreau. GRU-ODE-Bayes: Continuous modeling of sporadically-observed time series. *Advances in Neural Information Processing Systems*, 32, 2019.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. In *International Conference on Learning Representations*, 2017.
- Hadi Mohaghegh Dolatabadi, Sarah Erfani, and Christopher Leckie. Invertible generative modeling using linear rational splines. In *International Conference on Artificial Intelligence and Statistics*, pp. 4236–4246. PMLR, 2020.
- Alexandre Drouin, Étienne Marcotte, and Nicolas Chapados. Tactis: Transformer-attentional copulas for time series. In *International Conference on Machine Learning*, pp. 5447–5493. PMLR, 2022.
- Richard O. Duda and Peter E. Hart. Pattern classification and scene analysis. In *A Wiley-Interscience Publication*, 1974.
- Robert Dürichen, Marco A. F. Pimentel, Lei Clifton, Achim Schweikard, and David A. Clifton. Multitask gaussian processes for multivariate physiological time-series analysis. *IEEE Transactions on Biomedical Engineering*, 62(1):314–322, 2015. doi: 10.1109/TBME.2014.2351376.
- Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural spline flows. *Advances in neural information processing systems*, 32, 2019.
- Max Horn, Michael Moor, Christian Bock, Bastian Rieck, and Karsten Borgwardt. Set functions for time series. In *International Conference on Machine Learning*, pp. 4353–4363. PMLR, 2020.

542

543

544

546

547

548

549

550

551

552

553 554

555

556

558

559

561

562

563

564

565

566 567

568

569 570

571

572

573 574

575

576

577

578

579

580 581

582

583

584

585

586

587

588 589

590

591

592

- 540 A Johnson, L Bulgarelli, T Pollard, S Horng, and LA Celi. Mark. R. MIMIC-IV (version 1.0). PhysioNet, 2021.
 - Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. MIMIC-III, a freely accessible critical care database. Scientific data, 3(1):1–9, 2016.
 - Seyed Mehran Kazemi, Rishab Goel, Sepehr Eghbali, Janahan Ramanan, Jaspreet Sahota, Sanjay Thakur, Stella Wu, Cathal Smyth, Pascal Poupart, and Marcus Brubaker. Time2Vec: Learning a Vector Representation of Time, July 2019. URL http://arxiv.org/abs/1907.05321. arXiv:1907.05321 [cs].
 - Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. Advances in Neural Information Processing Systems (NIPS), 30, 2017.
 - Jenny Liu, Aviral Kumar, Jimmy Ba, Jamie Kiros, and Kevin Swersky. Graph normalizing flows. Advances in Neural Information Processing Systems, 32, 2019.
 - Étienne Marcotte, Valentina Zantedeschi, Alexandre Drouin, and Nicolas Chapados. Regions of reliability in the evaluation of multivariate probabilistic forecasts. In *International Conference on* Machine Learning, pp. 23958–24004. PMLR, 2023.
 - Matthew J Menne, CN Williams Jr, and Russell S Vose. United States historical climatology network daily temperature, precipitation, and snow data. Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, Oak Ridge, Tennessee, 2015.
 - Kevin P. Murphy. Probabilistic Machine Learning: An Introduction. Adaptive Computation and Machine Learning Series. MIT Press, Cambridge, MA, USA, March 2022. ISBN 978-0-262-04682-4.
 - Bernt Øksendal. Stochastic Differential Equations. Universitext. Springer, Berlin, Heidelberg, 2003. ISBN 978-3-540-04758-2 978-3-642-14394-6. doi: 10.1007/978-3-642-14394-6.
 - George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017.
 - George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. Journal of Machine Learning Research, 22(1), January 2021. ISSN 1532-4435.
 - Guilherme G. P. Freitas Pires and Mário A. T. Figueiredo. Variational mixture of normalizing flows. In 28th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning, ESANN 2020, Bruges, Belgium, October 2-4, 2020, pp. 205-210, 2020.
 - Janis Postels, Mengya Liu, Riccardo Spezialetti, Luc Van Gool, and Federico Tombari. Go with the flows: Mixtures of normalizing flows for point cloud generation and reconstruction. In 2021 International Conference on 3D Vision (3DV), pp. 1249–1258. IEEE, 2021. doi: 10.1109/ 3DV53792.2021.00132.
 - Kashif Rasul, Abdul-Saboor Sheikh, Ingmar Schuster, Urs M Bergmann, and Roland Vollgraf. Multivariate probabilistic time series forecasting via conditioned normalizing flows. In *International* Conference on Learning Representations, 2021.
 - Danilo Rezende and Shakir Mohamed. Variational Inference with Normalizing Flows. In International Conference on Machine Learning, pp. 1530–1538. PMLR, June 2015.
 - Mona Schirmer, Mazin Eltayeb, Stefan Lessmann, and Maja Rudolph. Modeling irregular time series with continuous recurrent units. In Proceedings of the 39th International Conference on Machine Learning, volume 162, pp. 19388–19405. PMLR, 2022.

- Maximilian Seitzer, Arash Tavakoli, Dimitrije Antic, and Georg Martius. On the Pitfalls of Heteroscedastic Uncertainty Estimation with Probabilistic Neural Networks. In *International Conference on Learning Representations*, 2021.
- Marcin Sendera, Jacek Tabor, Aleksandra Nowak, Andrzej Bedychaj, Massimiliano Patacchiola, Tomasz Trzcinski, Przemysław Spurek, and Maciej Zieba. Non-gaussian gaussian processes for few-shot regression. *Advances in Neural Information Processing Systems*, 34:10285–10298, 2021.
- Satya Narayan Shukla and Benjamin Marlin. Heteroscedastic temporal variational autoencoder for irregularly sampled time series. In *International Conference on Learning Representations*, 2022.
- Sahil Sidheekh, Kristian Kersting, and Sriraam Natarajan. Probabilistic Flow Circuits: Towards Unified Deep Models for Tractable Probabilistic Inference. In *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*, pp. 1964–1973. PMLR, July 2023.
- Ikaro Silva, George Moody, Daniel J Scott, Leo A Celi, and Roger G Mark. Predicting in-hospital mortality of icu patients: The physionet/computing in cardiology challenge 2012. In 2012 Computing in Cardiology, pp. 245–248. IEEE, 2012.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 30*, pp. 5998–6008. Curran Associates, Inc., 2017.
- Andrew G Wilson and Zoubin Ghahramani. Copula processes. *Advances in Neural Information Processing Systems*, 23, 2010.
- Vijaya Krishna Yalavarthi, Johannes Burchert, and Lars Schmidt-Thieme. Tripletformer for Probabilistic Interpolation of Irregularly sampled Time Series. In 2023 IEEE International Conference on Big Data (BigData), pp. 986–995, December 2023. doi: 10.1109/BigData59044.2023.10386325.
- Vijaya Krishna Yalavarthi, Kiran Madhusudhanan, Randolf Scholz, Nourhan Ahmed, Johannes Burchert, Shayan Jawed, Stefan Born, and Lars Schmidt-Thieme. GraFITi: Graphs for Forecasting Irregularly Sampled Time Series. In Michael J. Wooldridge, Jennifer G. Dy, and Sriraam Natarajan (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, pp. 16255–16263. AAAI Press, 2024. doi: 10.1609/AAAI.V38I15.29560.
- Vijaya Krishna Yalavarthi, Randolf Scholz, Stefan Born, and Lars Schmidt-Thieme. Probabilistic Forecasting of Irregularly Sampled Time Series with Missing Values via Conditional Normalizing Flows. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(20):21877–21885, April 2025. ISSN 2374-3468. doi: 10.1609/aaai.v39i20.35494.
- Weijia Zhang, Chenlong Yin, Hao Liu, Xiaofang Zhou, and Hui Xiong. Irregular Multivariate Time Series Forecasting: A Transformable Patching Graph Neural Networks Approach. In *Forty-First International Conference on Machine Learning*, June 2024.

A THEORY

A.1 PROOF OF LEMMA 3.1

Proof. Since X is a common conditional to all the marginals, we can ignore it. So, assume that f is a separable transformation:

$$f(z \mid Q) = (\phi(z_1 \mid Q_1), \dots, \phi(z_K \mid Q_K))$$
(18)

and that $\hat{p}_Z(z \mid Q)$ is marginalization consistent model. Then, the predictive distribution is

$$\hat{p}(y \mid Q) = \hat{p}_Z(f^{-1}(y \mid Q) \mid Q) \cdot \left| \det \frac{\mathrm{d}f^{-1}(y \mid Q)}{\mathrm{d}y} \right|$$
 (19)

Since f is separable, it follows that the Jacobian is diagonal:

$$\frac{\mathrm{d}f^{-1}(y \mid Q)}{\mathrm{d}y} = \frac{\mathrm{d}\left(\phi^{-1}(y_1 \mid Q_1), \dots \phi^{-1}(y_1 \mid Q_1)\right)}{\mathrm{d}(y_1, \dots, y_K)}$$

$$= \mathrm{diag}\left(\frac{\mathrm{d}\phi^{-1}(y_1 \mid Q_1)}{\mathrm{d}y_1}, \dots, \frac{\mathrm{d}\phi^{-1}(y_K \mid Q_K)}{\mathrm{d}y_K}\right) \tag{20}$$

Hence, the determinant of the Jacobian is the product of the diagonal elements:

$$\left| \det \frac{\mathrm{d}f^{-1}(y \mid Q)}{\mathrm{d}y} \right| = \prod_{k=1:K} \left| \det \frac{\mathrm{d}\phi^{-1}(y_k \mid Q_k)}{\mathrm{d}y_k} \right| = \prod_{k=1:K} \left| \frac{\mathrm{d}\phi^{-1}(y_k \mid Q_k)}{\mathrm{d}y_k} \right| \tag{21}$$

Using this fact, we can integrate the joint density over y_k to get the marginal density:

$$\int \hat{p}(y \mid Q) \, \mathrm{d}y_{k} \\
= \int \hat{p}_{Z}(f^{-1}(y \mid Q) \mid Q) \quad \cdot \left| \det \frac{\mathrm{d}f^{-1}(y \mid Q)}{\mathrm{d}y} \right| \mathrm{d}y_{k} \qquad \triangleright (7) \\
= \int \hat{p}_{Z}(f^{-1}(y \mid Q) \mid Q) \quad \cdot \prod_{k=1:K} \left| \frac{\mathrm{d}\phi^{-1}(y_{k} \mid Q_{k})}{\mathrm{d}y_{k}} \right| \mathrm{d}y_{k} \qquad \triangleright (21) \\
= \left(\prod_{l \neq k} \left| \frac{\mathrm{d}\phi^{-1}(y_{k} \mid Q_{k})}{\mathrm{d}y_{k}} \right| \right) \quad \cdot \int \hat{p}_{Z}(f^{-1}(y \mid Q) \mid Q) \cdot \left| \frac{\mathrm{d}\phi^{-1}(y_{k} \mid Q_{k})}{\mathrm{d}y_{k}} \right| \mathrm{d}y_{k} \\
= \left(\prod_{l \neq k} \left| \frac{\mathrm{d}\phi^{-1}(y_{k} \mid Q_{k})}{\mathrm{d}y_{k}} \right| \right) \quad \cdot \int \hat{p}_{Z}(z \mid Q) \, \mathrm{d}z_{k} \qquad \triangleright \text{transf.-thm} \\
= \left(\prod_{l \neq k} \left| \frac{\mathrm{d}\phi^{-1}(y_{k} \mid Q_{k})}{\mathrm{d}y_{k}} \right| \right) \hat{p}_{Z}(z_{-k} \mid Q_{-k}) \qquad \triangleright (5) \\
= \hat{p}_{Z}(z_{-k} \mid Q_{-k}) \left| \det \frac{\mathrm{d}f^{-1}(y_{-k} \mid Q_{-k})}{\mathrm{d}y_{-k}} \right| \qquad \triangleright (7)$$

A.2 PROOF OF LEMMA 3.2

Proof. Consider a mixture model of the form

$$\hat{p}(y \mid Q, X) := \sum_{l=1}^{D} w_d(X) \hat{p}_d(y \mid Q, X)$$
(22)

satisfying the conditions from Lemma 3.2, i.e. the component models \hat{p}_d satisfy the requirements **R1**-**R3** and the weight function $w: \operatorname{Seq}(\mathcal{X}) \to \Delta^D$ is permutation invariant with respect to X.

- 1. \hat{p} satisfies **R1**: By construction of the mixture model, it has the same domain and codomain as the component models.
- 2. \hat{p} satisfies **R2**: Let $\pi \in S_{|Q|}$ and $\tau \in S_{|X|}$, then

$$\hat{p}(y \mid Q^{\pi}, X^{\tau}) = \sum_{d=1}^{D} w_d(X^{\tau}) \hat{p}_d(y \mid Q^{\pi}, X^{\tau})$$

$$= \sum_{d=1}^{D} w_d(X) \hat{p}_d(y \mid Q, X)$$

$$\Rightarrow \text{ permutation invariance of } w \text{ and } \hat{p}_d$$

$$= \hat{p}(y \mid Q, X)$$

3. \hat{p} satisfies **R3**:

$$\int \hat{p}(y \mid Q, X) \, \mathrm{d}y_k = \int \sum_{d=1}^D w_d(X) \hat{p}_d(y \mid Q, X) \, \mathrm{d}y_k$$

$$= \sum_{d=1}^D w_d(X) \int \hat{p}_d(y \mid Q, X) \, \mathrm{d}y_k$$

$$= \sum_{d=1}^D w_d(X) \hat{p}_d(y_{-k} \mid Q_{-k}, X) \quad \triangleright \hat{p}_d \text{ is marginalization consistent}$$

$$= \hat{p}_Y(y_{-k} \mid Q_{-k}, X)$$

A.3 Proof of Theorem 4.1

Proof. Due to Lemma 1, it is sufficient to show that all the component models satisfy the requirements **R1-R3**. Since we use Gaussian Processes as the base distribution, Lemma 3.1 ensures that each component model is marginalization consistent, establishing **R3**. Requirement **R1** is by construction. Finally, permutation invariance **R2** can be seen as follows:

First, note that, by Equation (12), it follows that if $\mathbf{h}^{\mathrm{OBS}}$ is permutation equivariant with respect to X, and $\widetilde{\mathbf{h}}$ and \mathbf{h} are both permutation equivariant with respect to Q and permutation invariant with respect to X. Now, let $\pi \in S_{|Q|}$ and $\tau \in S_{|X|}$, then, for the d-th component model $\hat{p}_{Y_d}(y \mid Q, X)$. In particular, the flow satisfies $f_d^{-1}(y^\pi, Q^\pi, X^\tau) = f^{-1}(y^\pi, \mathbf{h}_d^\pi) = z^\pi$. Therefore:

$$\begin{split} &\hat{p}_{Y_d}(y^{\pi} \mid Q^{\pi}, X^{\tau}) \\ &= \hat{p}_{Z_d}(f^{-1}(y^{\pi}, Q^{\pi}, X^{\tau}) \mid Q^{\pi}, X^{\tau}) \cdot \left| \det \frac{\mathrm{d}f^{-1}(y^{\pi}, Q^{\pi}, X^{\tau})}{\mathrm{d}y^{\pi}} \right| \\ &= \mathcal{N}(f^{-1}(y^{\pi}, \mathbf{h}_{d}^{\pi}) \mid \mu(\mathbf{h}_{d}^{\pi}), \Sigma(\mathbf{h}_{d}^{\pi})) \cdot \left| \det \frac{\mathrm{d}f^{-1}(y^{\pi}, \mathbf{h}^{\pi})}{\mathrm{d}y^{\pi}} \right| \quad \triangleright \text{ by remark above} \\ &= \mathcal{N}(z^{\pi} \mid \mu(\mathbf{h}_{d}^{\pi}), \Sigma(\mathbf{h}_{d}^{\pi})) \cdot \left| \det \frac{\mathrm{d}f^{-1}(y^{\pi}, \mathbf{h}^{\pi})}{\mathrm{d}y^{\pi}} \right| \\ &= \mathcal{N}(z \mid \mu, \Sigma) \cdot \left| \det \frac{\mathrm{d}f^{-1}(y^{\pi}, \mathbf{h}^{\pi})}{\mathrm{d}y^{\pi}} \right| \quad \triangleright \text{ permutation invariance of GP} \\ &= \mathcal{N}(z \mid \mu, \Sigma) \cdot \left| \det \frac{\mathrm{d}f^{-1}(y, \mathbf{h})}{\mathrm{d}y} \right| \quad \triangleright \text{ by (21)} \\ &= \hat{p}_{Y_d}(y \mid Q, X) \end{split}$$

A.4 LINEAR RATIONAL SPLINES

Linear Rational Splines (LRS) are computationally efficient spline functions Dolatabadi et al. (2020). Formally, given a set of monotonically increasing points $\{(u_m, v_m)\}_{m=1:M}$ called knots, that is $u_m < u_{m+1}$ and $v_m < v_{m+1}$, along with their corresponding derivatives $\{\Delta_m > 0\}_{m=1:M}$, then the LRS transformation $\phi(u)$ within a bin $u \in [u_m, u_{m+1}]$ is:

$$\phi(u) = \begin{cases} \frac{\alpha_m v_m(\lambda_m - \tilde{u}) + \bar{\alpha}_m \bar{v}_m \tilde{u}}{\alpha_m(\lambda_m - \tilde{u}) + \bar{\alpha}_m \tilde{u}} & : \quad 0 \le \tilde{u} \le \lambda_m \\ \frac{\bar{\alpha}_m \bar{v}_m (1 - \tilde{u}) + \alpha_{m+1} v_{m+1} (\tilde{u} - \lambda_m)}{\bar{\alpha}_m (1 - \tilde{u}) + \alpha_{m+1} (\tilde{u} - \lambda_m)} : \quad \lambda_m \le \tilde{u} \le 1 \end{cases}$$

$$\text{where} \quad \tilde{u} = \frac{u - u_m}{u_{m+1} - u_m} \in [0, 1]$$

$$(23)$$

Here, $\lambda_m \in (0,1)$ signifies the location of automatically inserted virtual knot between u_m and u_{m+1} with value \bar{v}_m . The values of λ_m , α_m , $\bar{\alpha}_m$ and \bar{v}_m are all automatically derived from the original

Table 4: Statistics of the datasets used in our experiments. Sparsity means the percentage of missing observations in the time series. N is the total number of observations and K is the number of queries in our experiments in Section 6.

Name	#Samples	#Channels	Sparsity	N	K
USHCN	1100	5	77.9%	8 - 322	3 - 6
PhysioNet'12	12,000	37	85.7%	3 - 519	1 - 53
MIMIC-III	21,000	96	94.2%	4 - 709	1 - 85
MIMIC-IV	18,000	102	97.8%	1 - 1382	1 - 79

Table 5: We compare models with respect to MSE, where lower values indicate better performance. The best-performing model is shown in bold, the second-best in italics, and models with performance within one standard deviation of the best are underlined

	Model	USHCN	PhysioNet'12	MIMIC-III	MIMIC-IV
inconsistent	ProFITi GRU-ODE	$\frac{0.308 \pm 0.061}{0.410 \pm 0.106}$	0.305 ± 0.007 0.329 ± 0.004	0.548 ± 0.063 0.479 ± 0.044	0.389 ± 0.015 0.365 ± 0.012
consistent univariate	Neural-Flows CRU Tripletformer+	0.424 ± 0.110 0.290 ± 0.060 0.349 ± 0.131	0.331 ± 0.006 0.475 ± 0.015 0.293 ± 0.018	0.479 ± 0.045 0.725 ± 0.037 0.547 ± 0.068	0.374 ± 0.017 OOM 0.369 ± 0.030
consistent multivariate	GPR GMM	0.597 ± 0.110 0.294 ± 0.083	0.575 ± 0.059 0.293 ± 0.005	0.862 ± 0.016 0.535 ± 0.064	0.609 ± 0.014 0.332 ± 0.015
manifemen	MOSES (ours)	0.411 ± 0.099	0.307 ± 0.006	0.517 ± 0.057	0.342 ± 0.028

knots and their derivatives Dolatabadi et al. (2020). For a conditional LRS $\phi(z_k; \mathbf{h}_{d,k}, \theta)$, the function parameters such as width and height of each bin, the derivatives at the knots, and λ are computed from the conditioning input $\mathbf{h}_{d,k}$ and some model parameters θ . θ helps to project $\mathbf{h}_{d,k}$ to the function parameters, and is common to all the variables $z_{1:K}$ so that the transformation ϕ can be applied for varying number of variables K. Additionally, we set θ common to all the components as well. Since, each component has separate embedding for a variable z_k ($\mathbf{h}_{d,k}$), we achieve different transformations in different components for same variable.

In summary, the conditional flow model is separable across the query size $f = f_1 \times \cdots \times f_K$ with

$$f_d(y) := f(y \mid \mathbf{h}_d) = (\phi(y_1, \mathbf{h}_{d,1}), \dots, \phi(y_K, \mathbf{h}_{d,K}))$$
 (24)

Table 6: Experiment on varying observation and forecast horizons. Evaluation metric-njNLL, Lower the better. Best results are in bold.

	36/12	24/24	12/36
NeuralFlows ProFITi	0.708 ± 0.048 -0.768 ± 0.041	1.097±0.044 - 0.355±0.243	1.436±0.187 - 0.291±0.415
MOSES (ours)	-0.315±0.016	-0.298±0.027	-0.063±0.049

Table 7: Experiment on varying observation and forecast horizons. Evaluation metric-mNLL, Lower the better. Best results are in bold.

	36/12	24/24	12/36
NeuralFlows	0.709±0.047	1.065±0.082	1.439±0.198
ProFITi	1.376±1.764	0.705±0.179	2.977±2.978
MOSES (ours)	-0.083±0.025	-0.020 ± 0.060	0.040 ± 0.131

B DATASETS

4 real-world datasets are used in the experiments.

USHCN Menne et al. (2015). This is a climate dataset consisting of 5 climate variables such as daily temperatures, precipitation and snow measured over 150 years at 1218 meteorological stations in the USA. Following De Brouwer et al. (2019); Yalavarthi et al. (2025), we selected 1114 stations and an observation window of 4 years from 1996 until 2000.

PhysioNet2012 Silva et al. (2012). This physiological dataset consists of the medical records of 12,000 patients who are admitted into ICU. 37 vitals are recorded for 48 hrs. Following the protocol of Yalavarthi et al. (2024); Che et al. (2018), dataset consists of hourly observations in each series.

MIMIC-III Johnson et al. (2016). This is also a physiological dataset. It is a collection of readings of the vitals of the patients admitted to ICU at Beth Israeli Hospital. Dataset consists of 18,000 instances and 96 variables are measured for 48 hours. Following De Brouwer et al. (2019); Biloš et al. (2021); Yalavarthi et al. (2025) observations are rounded to 30 minute intervals.

MIMIC-IV Johnson et al. (2021). The successor of the MIMIC-III dataset. Here, 102 variables from patients admitted to ICU at a tertiary academic medical center in Boston are measured for 48 hours. Following De Brouwer et al. (2019); Biloš et al. (2021); Yalavarthi et al. (2025) observations, are rounded to 1 minute intervals.

Blast distribution (toy dataset). Blast distribution is a bivariate distribution which is created as follows:

$$z \sim \mathcal{N}\left(\begin{bmatrix} 0\\0 \end{bmatrix}, \begin{bmatrix} 1 & 1\\1 & 2 \end{bmatrix}\right)$$
$$y = \operatorname{sign}(z) \odot z \odot z$$

Circle (toy dataset). Circle is also a bi-variate distribution.

$$z \sim \mathcal{N}(0, \mathbb{I}_2)$$
$$y = \frac{z}{\|z\|_2} + 0.05 \cdot \mathcal{N}(0, \mathbb{I}_2)$$

C ADDITIONAL EXPERIMENTS

C.1 Comparing for Point Forecasting

While point forecasting is an important task in time series analysis, the goal of probabilistic forecasting is fundamentally different. Probabilistic forecasting aims to capture the full predictive distribution

Table 8: Comparing for Energy Score. Lower the better. Best results are in bold.

	USHCN	PhysioNet'12	MIMIC-III	MIMIC-IV
NeuralFlows ProFITi	0.661 ± 0.059 $\mathbf{0.452 \pm 0.044}$	1.691 ± 0.001 $\mathbf{0.879 \pm 0.303}$	1.381 ± 0.033 1.606 ± 0.168	0.982 ± 0.009 $\mathbf{0.808 \pm 0.003}$
MOSES	0.552 ± 0.044	1.599 ± 0.013	1.353 ± 0.033	0.906 ± 0.029

Table 9: Comparing models w.r.t. CRPS score on marginals. Lower the better. Best results are in bold.

Model	USHCN	PhysioNet'12	MIMIC-III	MIMIC-IV
Neural-flows ProFITi	0.306 ± 0.028 0.182 ± 0.007	0.277 ± 0.003 0.271 ± 0.003	0.308 ± 0.004 0.319 ± 0.003	0.281 ± 0.004 0.279 ± 0.012
MOSES (ours)	0.220 ± 0.019	0.260 ± 0.002	0.296 ± 0.005	0.245 ± 0.010

rather than just a single-point estimate. Nonetheless, one might intuitively expect that the best probabilistic model would also yield the most accurate point estimates. However, this is not always the case in practice, as noted in prior works (Lakshminarayanan et al., 2017; Seitzer et al., 2021; Rasul et al., 2021; Yalavarthi et al., 2025).

We compare probabilistic models in terms of point prediction accuracy using Mean Squared Error (MSE), as reported in Table 5. Our results show that no single model consistently outperforms the others across all datasets. GMM is the overall best ranked model with MOSES and GRU-ODE being the next best. We believe there are two primary reasons for this phenomenon:

- (1.) MSE is related to the Negative Log-Likelihood (NLL) of a Gaussian distribution with a fixed standard deviation. Therefore, models explicitly trained by minimizing Gaussian Negative Log-Likelihood (even if they predict more than just the mean) are naturally optimized for this metric.
- (2.) Probabilistic models are trained to predict the underlying data distribution, not solely the optimal point estimate (e.g., the conditional mean). Their objective is to accurately capture the uncertainty and dependencies in the data, which involves learning the (co)variance structure. This focus on the full distribution can sometimes lead to point estimates that are not strictly optimized for minimizing the squared error, even if the overall probabilistic forecast is superior.

Except for MOSES, GMM and ProFITi, all the other probabilistic models are designed to predict Gaussian distributions. MOSES and GMM have competing accuracy in all the datasets other than USHCN. In USHCN, MOSES performs worse because of outliers in the data. This is the reason for large standard deviations for all the models in USHCN dataset.

C.2 EXPERIMENT ON VARYING OBSERVATION AND FORECAST HORIZONS

We would like to see if MOSES is scalable to long observations and forecast horizons. For this, we performed an experiment on varying length observation and forecasting horizons on Physionet'12 dataset and compared against the published results from (Yalavarthi et al., 2025) in Table 6. The observation and forecasting horizons are: {(36h, 12h), (24h, 36h), (12h, 26h)}.

Tables 6 and 7 present the njNLL and mNLL results for ProFITi and MOSES. The results follow the trends observed in Tables 1 and 2. ProFITi performs best when predicting joint distributions. However, its lack of marginalization consistency leads to a severe performance drop when predicting marginal distributions. In contrast, MOSES maintains stable performance from njNLL to mNLL. While it performs slightly worse than ProFITi on njNLL, it significantly outperforms ProFITi on mNLL.

C.3 COMPARING FOR ENERGY SCORE

The Energy Score between the ground truth y and predicted distribution \hat{p}_Y is computed as:

$$ES(y, \hat{p}_Y) := \mathbb{E}_{y' \sim \hat{p}_Y} \|y - y'\|_2^p - \frac{1}{2} \mathbb{E}_{y', y'' \sim \hat{p}_Y} \|y' - y''\|_2^p, \tag{25}$$

Table 10: Ablation study on PhysioNet'12

Table 11	: Importance	of mixer components	

Model	njNLL (↓)
MOSES	-0.491 ± 0.041
MOSES-f	1.063 ± 0.002
MOSES-cov	-0.308 ± 0.024
MOSES-w	-0.451 ± 0.038
MOSES (1)	-0.493 ± 0.029

Dataset	MOSES(1)	MOSES(D)
USHCN	-3.255±0.204	-3.357±0.176
Physionet'12	-0.493±0.029	-0.491±0.041
MIMIC-III	0.173±0.332	-0.305±0.027
MIMIVC-IV	-1.686±0.072	-1.668±0.097

Table 12: Comparing the number of parameters and run-time per epoch for results in Table 1 for GMM and MOSES, and ProFITi for reference

	USHCN		Physionet'12		MIMIC-III		MIMIC-IV	
	Parameters	Run Time	Parameters	Run Time	Parameters	Run Time	Parameters	Run Time
ProFITi	1,093.0K	3.8s	75.8K	42.14s	59.7K	66.8s	285.9K	70.2s
GMM	416.0K	0.9s	390.9K	5.9s	33.0K	18.5s	101.1K	21.3s
MOSES (ours)	167.6K	2.4s	134.6K	14.1s	112.6K	25.4s	398.6K	33.3s

where $\|\cdot\|_2$ denotes the Euclidean norm and $p \in (0,2)$ is a parameter. In our evaluation, we set p=1. Marcotte et al. (2023) demonstrated that the Energy Score is not a reliable metric for evaluating multivariate distributions. Additionally, it suffers from the curse of dimensionality, as it requires N^K samples, where K is the number of variables and N is the number of samples required to accurately estimate a univariate distribution.

However, since many regularly sampled, fully observed multivariate time series probabilistic forecasting models use the Energy Score as an evaluation metric, we examine how MOSES compares to the best-performing inconsistent multivariate probabilistic model, ProFITi, and the consistent univariate probabilistic model, NeuralFlows in Table 8. Our results show that MOSES outperforms NeuralFlows across all datasets. As shown by the njNLL metric in Table 1, ProFITi is the best-performing model, outperforming MOSES in 3 out of 4 datasets.

C.4 COMPARING FOR MARGINALS IN TERMS OF CRPS

We compare with CRPS score in Table 9, a widely used evaluation metric in time series forecasting. We see that MOSES outperforms all the consistent models. It performs better than ProFITi in 3 out of 4 dataset. For ProFITi and MOSES, we sampled 1000 instances and computed the CRPS.

C.5 ABLATION STUDY.

Using Physionet'12, we show the importance of different model components. As summarized in Table 10, the performance is reduced by removing the flows (MOSES – f) which is same as GMM. It is expected that normalizing flows are more expressive compared to simple mixture of Gaussians. On the other hand, by using only isotropic Gaussian as the base distribution (MOSES – COV) model performance worsened. Similarly, parameterizing the components weights have a slight advantage over fixing them to 1/D with D being the number of components. One interesting observation is even using single component (MOSES(1)) gives similar results compared to mixture of such components. This could be because the dataset we have may not require multiple components. We note that we have D = 1 in our hyperparameter space, and we select the best D based on validation dataset. Note that (MOSES(1)) can be seen as representative of Copula Processes (Wilson & Ghahramani, 2010) and Non-Gaussian Gaussian Process (Sendera et al., 2021).

When we conduct this study on other datasets (Table 11) we realized that mixer model is indeed helpful if there is a requirement.

C.6 COMPARING THE NUMBER OF PARAMETERS AND RUNTIME FOR GMM AND MOSES

Since MOSES is built upon GMM, Table 12 presents the number of parameters and runtime for both MOSES and GMM. For reference, we also include ProFITi.

The results show that GMM has a relatively low number of parameters for MIMIC-III and MIMIC-IV, whereas for USHCN and PhysioNet'12, the number of parameters is significantly higher. The primary

 difference between GMM and MOSES is the inclusion of flows. Given that all other factors remain the same, MOSES is expected to have a slightly higher number of parameters than GMM due to these additional flows. Also, the parameters for the flows are shared among all the variables and the components, their number does not grow with increase in components or variables. However, differences in the chosen hyperparameters for GMM and MOSES lead to some discrepancies from this expectation. Moreover, the inclusion of flows in MOSES results in a slightly higher runtime compared to GMM.

C.7 UNCONDITIONAL DENSITY ESTIMATION: FIGURE 2

Figure 2 shows the advantage of MOSES compared to GMM. GMM requires more number of components compared to MOSES for predicting similar distributions. In this experiment, we use the unconditional version of the model meaning the flow architecture is independent. Hence, mean and covariance of the Gaussian distributions which are same for both MOSES and GMM are trainable parameters. In addition to them, we have the parameters for the seperable flows which are not shared among the variables and components. Due to this, MOSES have significantly higher number of parameters compared to GMM. Although one could condition the models upon constant vectors and reduce the number of parameters, we did not use this approach and stick to the general practice.