
Conformalised Conditional Normalising Flows for Joint Prediction Regions in time series

Eshant English*

Hasso Plattner Institute for Digital Engineering, Germany
University of California, Irvine, USA
eshant.english@hpi.de

Christoph Lippert

Hasso Plattner Institute for Digital Engineering, Germany
Hasso Plattner Institute for Digital Health at the Icahn School of Medicine at Mount Sinai, USA
christoph.lippert@hpi.de

Abstract

Conformal Prediction offers a powerful framework for quantifying uncertainty in machine learning models, enabling the construction of prediction sets with finite-sample validity guarantees. While easily adaptable to non-probabilistic models, applying conformal prediction to probabilistic generative models, such as Normalising Flows is not straightforward. This work proposes a novel method to conformalise conditional normalising flows, specifically addressing the problem of obtaining prediction regions for multi-step time series forecasting. Our approach leverages the flexibility of normalising flows to generate potentially disjoint prediction regions, leading to improved predictive efficiency in the presence of potential multimodal predictive distributions.

1 Introduction

Recent advances in Conformal Prediction have seen its application extended to diverse domains, including Gaussian Processes [Jaber et al., 2024] and Monte-Carlo methods [Bethell et al., 2024], driven by the attractive finite-sample validity guarantees offered via conformal prediction [Vovk et al., 2005, Angelopoulos and Bates, 2022]. However, direct application to generative models, such as Normalising Flows [Kobyzev et al., 2021, Papamakarios et al., 2021] or Diffusion models [Song et al., 2021, Luo, 2022], has been limited due to their primary focus on generation rather than prediction tasks.

Notably, discrete-time normalising flows offer a unique advantage over other generative models: simulation-free exact likelihood computation. This enables direct density evaluation at a given sample, naturally extending to conditional density computation. In this work, we exploit this property of normalising flows to derive prediction regions for multi-step time series forecasting.

A prediction region is defined as a volume in H -dimensional space, where H is the forecast horizon, capturing the true forecast with probability $1 - \epsilon$. Constructing such regions for time series is challenging due to temporal dependencies. Existing methods, such as Stankeviciute et al. [2021], rely on Bonferroni corrections based on conditional independence assumptions on the errors, leading to overly conservative predictions as H increases. Moreover, these approaches typically yield continuous

*primary contact: eshant.english@hpi.de

The work was done during a research stay at UCI

band or rectangular regions, which may not be optimal for multimodal predictive distributions present in irregular time series.

Traditional conformal methods often rely on univariate conformity scores, with residuals being a natural choice in regression settings. However, this approach does not readily generalise to multivariate prediction cases. We propose the use of conditional density as a conformity score, capitalising on its adaptability to the historical context and its ability to quantify the plausibility of a given prediction naturally.

In summary, our key contribution lies in extending conformal methods to conditional normalising flows, thereby enabling adaptive and efficient predictive inference for multi-step time series forecasting (or more generally multivariate output).

2 Related works

Previous research has explored normalising flows for time series forecasting [Feng et al., 2023, Rasul et al., 2021]. However, existing methods for uncertainty quantification in this context either lack coverage guarantees or rely on unrealistic assumptions when dealing with multi-step or multivariate forecasts. In the univariate case, conformalisation using residuals is straightforward. However, for higher-dimensional predictions, mapping residuals to quantiles is not trivial.

Existing approaches attempt to address this issue through various means, such as transforming multivariate residuals into Euclidean distances [Messoudi et al., 2021], using Chebyshev distance after modulation [English et al., 2024, Zhou et al., 2024], or modelling interdependencies between successive time steps using copulas [Sun and Yu, 2024]. Nevertheless, these methods often result in large, continuous prediction regions that may be inefficient for capturing multimodal predictive distributions in time series. Achieving predictive efficiency, where the volume of the prediction set is minimised, becomes a crucial objective in such scenarios.

A recent work Colombo [2024] also combined Normalising Flows with Conformal Prediction. However, they do not use Normalising Flows to obtain density estimates for conformal scores; instead, they train the calibration process with Normalising Flows acting on the joint distribution of the errors. To the best of our knowledge, no existing conformal method for time series forecasting can produce multiple disjoint regions when the predictive distribution is multimodal. Building upon prior work, we propose a novel approach applicable to scenarios with multiple independent time series. Our method aims to overcome the limitations of previous approaches and provide a more efficient and flexible way to quantify uncertainty in time series forecasting.

3 Preliminaries

3.1 Inductive Conformal Predictors

Inductive conformal predictors provide a computationally efficient alternative to Full conformal predictors for constructing prediction intervals by training the model only once. Given a dataset $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$ drawn from a distribution P , we divide it into a training set $\mathcal{D}_{\text{train}}$ of size m and a calibration set \mathcal{D}_{cal} of size l , where $n = m + l$.

A point predictor h is trained on $\mathcal{D}_{\text{train}}$, and a conformity measure A is used to assess the *conformity* of an example relative to the calibration set. We compute conformity scores (denoted by α) for both the calibration set and a test point x^* with a postulated label $y \in \mathcal{Y}$, the label space:

- $\alpha_i = A(h, \mathcal{D}_{\text{train}}, (x_{m+i}, y_{m+i}))$, for $i = 1, \dots, l$
- $\alpha^* = A(h, \mathcal{D}_{\text{train}}, (x^*, y))$

The prediction set (Γ^ϵ) at significance level ϵ is then defined as:

$$\Gamma^\epsilon(x^*) = \left\{ y : \frac{|\{i : \alpha_i \leq \alpha^*\}| + 1}{l + 1} > \epsilon \right\}. \quad (1)$$

This set represents the values of y that are deemed *conforming* compared to the calibration set, with a probability of at least $1 - \epsilon$. Instead of comparing all the conformity scores, one can invert the quantile function of the conformity scores on the calibration set to get a threshold conformity score.

3.2 Conditional Normalising Flows

Conditional normalising flows (CNFs) are normalising flows that incorporate additional conditioning information, c , into the transformation. This allows the learned distribution to be adapted based on the context provided by c . Formally, a CNF is a function $f : \mathbb{R}^D \times \mathbb{R}^M \rightarrow \mathbb{R}^D$, where M is the dimensionality of the conditioning variable.

Given a dataset $\mathcal{D} = \{(x_1, c_1), \dots, (x_n, c_n)\}$, we aim to maximise the conditional log-likelihood:

$$\max_f \sum_{i=1}^n \log p_{X|C}(x_i | c_i),$$

where $p_{X|C}$ is the conditional probability density function of the data given the conditioning variable, c . Using the change of variables formula, we express the conditional density as:

$$p_{X|C}(x|c) = p_{Z|C}(f^{-1}(x, c)|c) |\det J_{f^{-1}}(x, c)|. \quad (2)$$

Here, J refers to the Jacobian of the transformation and \det refers to the determinant. Then, the maximum log-likelihood optimisation objective can be stated as:

$$\max_f \sum_{i=1}^n [\log p_{Z|C}(f^{-1}(x_i, c_i)|c_i) + \log |\det J_{f^{-1}}(x_i, c_i)|].$$

Similar to standard normalising flows, the first term can be computed analytically by assuming a simple form for $p_{Z|C}$ (e.g., a Gaussian distribution with mean and variance dependent on c). The second term, involving the Jacobian determinant, requires careful architectural choices for efficient computation.

CNFs are typically constructed as a composition of simpler invertible transformations, each of which can be conditioned on c . Various techniques, such as affine coupling layers or autoregressive flows, are easily adaptable to the conditional setting. The choice of transformations and their conditioning mechanism influences the expressivity and flexibility of the CNF model.

4 Conformalised Conditional Normalising Flow for Joint Prediction Region in time series (CCN-JPR)

We consider the problem of probabilistic forecasting for a collection of n independent multivariate time series, each denoted by $\mathbf{x}^{(i)} \in \mathbb{R}^{D \times (T+H)}$, where T is the length of the context provided to the model and H is the forecast horizon. We partition the time series' into a training set proper of m series and a calibration set of l series, such that $n = m + l$.

Let $\mathbf{x}_t^{(i)} \in \mathbb{R}^{D \times (T+H)}$ represent the i -th time series at time step t . Given the context window $\mathbf{x}_{1:T}^{(i)}$, the conditional distribution of future values $\mathbf{x}_{T+1:T+H}^{(i)}$ over the entire dataset $\mathbf{x} = \{\mathbf{x}_{1:T+H}^{(i)}\}_{i=1}^m$ is modeled as:

$$p(\mathbf{x}_{T+1:T+H} | \mathbf{x}_{1:T}; \theta) = \prod_{i=1}^m p(\mathbf{x}_{T+1:T+H}^{(i)} | \mathbf{h}_T^{(i)}; \theta),$$

where $\mathbf{h}_T^{(i)}$ is temporal summarisation for the i -th time series at time step T , and θ represents the model parameters.

Similar to Rasul et al. [2021], it is possible to use an RNN-based architecture or a transformer to summarise the temporal history instead of flow transformation to incorporate it. For RNNs, $\mathbf{h}_T^{(i)}$ could be the hidden state of the RNN for the i -th time series at time step T . For any arbitrary time step t , the hidden state is updated recursively:

$$\mathbf{h}_t^{(i)} = \text{RNN}(\mathbf{x}_{t-1}^{(i)}, \mathbf{h}_{t-1}^{(i)}),$$

To capture the potentially complex, high-dimensional distribution of $\mathbf{x}_{T+1:T+H}^{(i)}$, we use a flow to model the emission distribution $p(\mathbf{x}_{T+1:T+H}^{(i)} | \mathbf{h}_T^{(i)}; \theta)$:

$$\mathbf{x}_{T+1:T+H}^{(i)} = g_\theta(\mathbf{z}_{T+1:T+H}^{(i)}; \mathbf{h}_T^{(i)}),$$

where g_θ is an invertible transformation parameterized by θ , and $\mathbf{z}_{T+1:T+H}^{(i)}$ is a latent variable drawn from a simple base distribution (e.g., standard Gaussian).

Our specific implementation utilises a stack of K conditional affine coupling layers for g_θ , each conditioned on the hidden state $\mathbf{h}_T^{(i)}$ (identity transformation on $\mathbf{x}_{1:T}^{(i)}$ for the experiments) to accommodate the temporal dynamics. The model is trained via maximum likelihood estimation using stochastic gradient descent, and forecasts are generated by sampling from the conditional distribution.

Once a model is trained using the training set proper, we can use the calibration set to calculate the conditional density around the true prediction given the historical context by using the change of the variable formula as stated in Equation 2. We consider this as our conformity score and can apply Equation 1 for any postulated label to know if it is a part of the prediction set. Additionally, the regions formed are adaptive to the historical context by using conditional density as a conformity score.

We can also sort all the conformity scores and obtain the desired quantile q_ϵ of the sorted scores. To construct a region that contains the true prediction with the desired significance level ϵ , we can, then, sample points from the multi-variate label space either through a grid or through other methods such as Monte-Carlo sampling, the sampled points with a conditional density larger than q_ϵ form cluster(s) which can be several if there are multiple modes of the predictive distribution. Unlike other baselines, a region formed like this has no geometrical restrictions.

5 Empirical Results

We assess the *coverage* of our model on two synthetic time-series datasets and one real-world dataset (COVID-UK) from Kipf et al. [2018], following the experimental protocol of Sun and Yu [2024].

For both *synthetic* datasets, we forecast 24 steps ahead ($H = 24$) using a history of 35 steps ($T = 35$), Each time step consists of two-dimensional data (\mathbb{R}^2). The datasets differ in the magnitude of the added Gaussian noise: $\sigma = 0.01$ and $\sigma = 0.05$, respectively. For the COVID-UK dataset, which is a difficult time series due to irregularity, we predict daily cases for the upcoming 10 days based on the preceding 100 days.

Table 1: Comparison of the coverages (nominal) on the test set for $\epsilon = 0.1$. The coverage values closer to 0.9 are highlighted in bold on each dataset. Our method achieves values close to 0.9 whereas CopulaCPTS always undercovers, MC-dropout severely undercovers, and CRNN often overcovers (except on the COVID-UK dataset).

Method	Particle1	Particle5	COVID-UK
MC-Dropout	0.79	0.43	0.00
CRNN	0.99	0.93	0.87
CopulaCPTS	0.86	0.86	0.78
CCNF-JPR	0.91	0.89	0.87

We computed the coverages by checking if the true value is accepted by the conditional density threshold used as a conformity score by our method. Table 1 presents results for all the datasets: ($\sigma = 0.01$) *particle1* and ($\sigma = 0.05$) *particle5* dataset. Our proposed method demonstrates coverage closest to the nominal 0.9 level ($\epsilon = 0.1$). CRNN overcovers(except on the COVID-UK dataset), and other methods, notably MC-dropout, exhibit severe undercoverage.

6 Discussion and Future Work

In this work, we introduced a novel approach to conformalise conditional normalising flows for multi-step time series forecasting, enabling the construction of prediction regions without assuming specific geometric shapes like spheres, bands, or ellipses [Cleveland et al., 2024, Xu et al., 2024]. By conditioning the conformity scores on the time series history, our method also yields adaptive prediction regions, a feature not commonly found in existing approaches (English et al. [2023], Zhou et al. [2024] being notable exceptions, but can only have rectangular regions). Although our work focussed on time series forecasting, the idea is trivially applicable to multivariate regression as

multivariate regression can be considered a special case where the time steps/indices are independent. One bottleneck is that exhaustive sampling might be computationally expensive if the label space is high-dimensional. Efficient sampling schemes can be explored for computational feasibility.

While normalising flows are often employed with neural networks, this might not always be desirable when using classical statistical methods such as ARMA models. Future research could explore extending existing kernel-based flow approaches [Meng et al., 2020, English et al., 2023] to conditional normalising flows (conditioned on prediction from the classical methods), potentially mitigating these challenges.

A further limitation is the lack of standard datasets exhibiting multimodal predictive distributions in time series, which are necessary to showcase the full potential of our approach. Future work will focus on addressing these limitations and conducting comprehensive empirical evaluations.

Acknowledgements

We extend our gratitude to Thomas Gaertner for valuable feedback on an earlier draft and to Prof Vladimir Vovk for helpful discussions. This research was funded by the HPI Research School on Data Science and Engineering.

References

- Anastasios N. Angelopoulos and Stephen Bates. A gentle introduction to conformal prediction and distribution-free uncertainty quantification, 2022.
- Daniel Bethell, Simos Gerasimou, and Radu Calinescu. Robust uncertainty quantification using conformalised monte carlo prediction, 2024.
- Matthew Cleaveland, Insup Lee, George J. Pappas, and Lars Lindemann. Conformal prediction regions for time series using linear complementarity programming, 2024.
- Nicolo Colombo. Normalizing flows for conformal regression, 2024. URL <https://arxiv.org/abs/2406.03346>.
- Eshant English, Matthias Kirchler, and Christoph Lippert. Kernelised normalising flows, 2023.
- Eshant English, Eliot Wong-Toi, Matteo Fontana, Stephan Mandt, Padhraic Smyth, and Christoph Lippert. Janet: Joint adaptive prediction-region estimation for time-series, 2024. URL <https://arxiv.org/abs/2407.06390>.
- Shibo Feng, Chunyan Miao, Ke Xu, Jiayang Wu, Pengcheng Wu, Yang Zhang, and Peilin Zhao. Multi-scale attention flow for probabilistic time series forecasting, 2023.
- Edgar Jaber, Vincent Blot, Nicolas Brunel, Vincent Chabridon, Emmanuel Remy, Bertrand Iooss, Didier Lucor, Mathilde Mougeot, and Alessandro Leite. Conformal approach to gaussian process surrogate evaluation with coverage guarantees, 2024.
- Thomas Kipf, Ethan Fetaya, Kuan-Chieh Wang, Max Welling, and Richard Zemel. Neural relational inference for interacting systems, 2018.
- Ivan Kobyzev, Simon J.D. Prince, and Marcus A. Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):3964–3979, nov 2021. doi: 10.1109/tpami.2020.2992934. URL <https://doi.org/10.1109/2Ftpami.2020.2992934>.
- Calvin Luo. Understanding diffusion models: A unified perspective, 2022.
- Chenlin Meng, Yang Song, Jiaming Song, and Stefano Ermon. Gaussianization flows, 2020.
- Soundouss Messoudi, Sébastien Destercke, and Sylvain Rousseau. Copula-based conformal prediction for multi-target regression, 2021.
- George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference, 2021.

- Kashif Rasul, Abdul-Saboor Sheikh, Ingmar Schuster, Urs Bergmann, and Roland Vollgraf. Multivariate probabilistic time series forecasting via conditioned normalizing flows, 2021.
- Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations, 2021. URL <https://arxiv.org/abs/2011.13456>.
- Kamile Stankeviciute, Ahmed M. Alaa, and Mihaela van der Schaar. Conformal Time-series Forecasting. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P. S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 6216–6228. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/312f1ba2a72318edaaa995a67835fad5-Paper.pdf.
- Sophia Huiwen Sun and Rose Yu. Copula conformal prediction for multi-step time series prediction. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=ojIJZDNIBj>.
- Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic learning in a random world*, volume 29. Springer, 2005.
- Chen Xu, Hanyang Jiang, and Yao Xie. Conformal prediction for multi-dimensional time series by ellipsoidal sets, 2024.
- Yanfei Zhou, Lars Lindemann, and Matteo Sesia. Conformalized adaptive forecasting of heterogeneous trajectories, 2024. URL <https://arxiv.org/abs/2402.09623>.