# RESCP: RESERVOIR CONFORMAL PREDICTION FOR TIME SERIES FORECASTING

**Anonymous authors** 

Paper under double-blind review

# **ABSTRACT**

Conformal prediction (CP) offers a powerful framework to build distribution-free prediction intervals for exchangeable data. Existing methods that extend CP to sequential data rely on fitting a relatively complex model to capture temporal dependencies. However, these methods can fail if the sample size is small and often require expensive retraining when the underlying data distribution changes. To overcome these limitations, we propose <u>Reservoir Conformal Prediction</u> (RESCP), a novel training-free CP for time series. Our approach leverages the efficiency and representation learning capabilities of Reservoir Computing to dynamically reweight conformity scores. In particular, we compute similarity scores across reservoir states and use them to adaptively reweight the observed residuals. RESCP enables us to account for local temporal dynamics when modeling the error distribution without compromising computational scalability. We prove that, under reasonable assumptions, RESCP achieves asymptotic conditional coverage, and we empirically demonstrate its effectiveness across diverse forecasting tasks.

# 1 Introduction

Despite deep learning often achieving state-of-the-art results in time series forecasting, widely used methods do not provide a way to quantify the uncertainty of the predictions (Benidis et al., 2022), which is crucial for their adoption in risk-sensitive scenarios, such as healthcare (Makridakis et al., 2019), load forecasting (Hong & Fan, 2016), and weather forecasting (Palmer & Hagedorn, 2006). Moreover, in many cases, prediction intervals (PIs) must not only be reliable but also fast to compute. Existing uncertainty quantification approaches for time series often rely on strong distributional assumptions (Benidis et al., 2022; Salinas et al., 2020), which do not always fit real-world data. Moreover, they often require long and expensive training procedures, as well as modifications to the underlying forecasting model, which limits their applicability to large datasets. A framework that recently gained attention in uncertainty quantification is CP (Vovk et al., 2005; Angelopoulos et al., 2023). CP assumes exchangeability between the data used to build the PIs and the test data points, meaning that the joint distribution of the associated sequence of random variables does not change when indices are permuted. Clearly, this does not hold for time series, as the presence of temporal dependencies and, possibly, non-stationarity violates the assumption. Moreover, temporal dynamics can be heterogeneous and often result in heteroskedastic errors. This requires mechanisms to make PIs locally adaptive (Lei et al., 2018; Guan, 2022).

To extend CP to time series data, a popular approach is to reweight the observed residuals to handle non-exchangeability and/or temporal dependencies (Barber et al., 2023; Tibshirani et al., 2019; Auer et al., 2023). In particular, Auer et al. (2023) proposes an effective reweighting mechanism based on a learned model and soft-attention operators, which, however, suffers from the high computational cost typical of Transformer-like architectures. Other methods build PIs by directly learning the quantile function of the error distribution (Jensen et al., 2022; Xu & Xie, 2023a; Lee et al., 2024; 2025; Cini et al., 2025). This, however, often results in high sample complexity and might require frequent model updates to adapt to changes in the data distribution.

To address these limitations, we propose <u>Reservoir Conformal Prediction</u> (RESCP), a novel CP method that leverages a reservoir, i.e., a randomized recurrent neural network (Jaeger, 2001; Lukoševičius & Jaeger, 2009), driven by the observed prediction residuals to efficiently compute data-dependent weights adaptively at each time step. To capture locally similar dynamics, the

weights are derived from the similarity between current and past states of the reservoir, which is implemented as an echo state network (ESN) (Jaeger, 2001; Gallicchio & Scardapane, 2020). This enables RESCP to model local errors through a weighted empirical distribution of past residuals, where residuals associated with similar temporal dynamics receive higher weights. The ESN does not require any training and is initialized to yield stable dynamics and informative representations of the input time series at each time step. This makes our approach extremely scalable and easy to implement. RESCP can be applied on top of *any* point forecasting model, since it only requires residuals from a disjoint calibration set.

Our main contributions are summarized as follows.

- We provide the first assessment of reservoir computing for CP in time series analysis and show that it is a valid and scalable alternative to existing methods.
- We introduce RESCP, a novel, scalable, and theoretically grounded tool for distribution-free uncertainty quantification in sequential prediction tasks.
- We prove that RESCP can achieve asymptotic conditional coverage under reasonable assumptions on the data-generating process.
- We introduce variants of RESCP to handle distribution shifts and leverage exogenous variables.

We evaluate RESCP on data from real-world applications and show its robustness and effectiveness by comparing it to state-of-the-art baselines.

## 2 BACKGROUND AND RELATED WORK

We consider a forecasting setting where the goal is to predict future values of a time series based on past observations. Let  $\{y_t\}_{t=1}^T$  be the sequence of scalar target observations at each time step t. We denote by  $y_{1:T}$  the entire history of observations up to time T. Moreover, let  $\{u_t\}_{t=1}^T$  be the sequence of exogenous variables at each time step t of dimension  $D_u$ . In this context, we are also given a forecasting model  $\hat{y}_{t+H} = \hat{f}(y_{t-W:t}, u_{t-W:t})$  that produces point forecasts of H-stepsahead, with  $H \geq 1$ , given a window  $W \geq 1$  of past observations. The model  $\hat{f}$  can be any kind of forecasting model, e.g., a recurrent neural network (RNN). The objective is to construct valid PIs that reflect the uncertainty in the forecasts. Ideally, PIs should be conditioned on the current state of the system, i.e., we want to achieve conditional coverage:

$$\mathbb{P}\left(y_{t+H} \in \hat{C}_T^{\alpha}(\hat{y}_{t+H}) \middle| y_{\leq t}, u_{\leq t}\right) \ge 1 - \alpha \tag{1}$$

where  $\hat{C}_T^{\alpha}(\hat{y}_{t+H})$  is the prediction interval for the forecast  $\hat{y}_{t+H}$  at significance level  $\alpha$ .

## 2.1 CONFORMAL PREDICTION IN TIME SERIES FORECASTING

CP (Vovk et al., 2005; Angelopoulos et al., 2023) is a framework for building valid *distribution-free* PIs, possibly with finite samples. CP provides a way to quantify uncertainty in predictions by leveraging the empirical quantiles of *conformity* scores, which in our regression setting we define as the prediction residuals

$$r_t = y_t - \hat{y}_t. \tag{2}$$

These scores measure the discrepancy between the model predictions and the observed values. The CP framework we consider in this work is split CP (SCP) (Vovk et al., 2005), which constructs valid PIs post-hoc by leveraging a disjoint calibration set to estimate the distribution of prediction errors. In its standard formulation, CP treats all calibration points symmetrically when building PIs. In heterogeneous settings, this can lead to overly conservative intervals as conformal scores are often heteroskedastic. To address heteroskedasticity, CP can be made *locally adaptive* by providing PIs whose width can adaptively shrink and inflate at different regions of the input space (Lei et al., 2018). In this context, Guan (2022) proposed localized CP (LCP), which reweights conformity scores according to the similarity between samples in feature space, given a localization function. Hore & Barber (2023) introduced a randomized version of LCP, with approximate conditional coverage guarantees and more robust under covariate shifts. Although this reweighting breaks exchangeability, validity is recovered through a finite-sample correction of the target coverage level.

Other approaches rely on calibrating estimates of the spread of the data (Lei & Wasserman, 2013) or quantile regression (Romano et al., 2019; Jensen et al., 2022; Feldman et al., 2023), but often require ad-hoc base predictors (i.e., they cannot be applied on top of generic point predictors).

**Conformal prediction for time series** Under the assumption of data exchangeability, CP methods provide, in finite samples, valid intervals given a specified level of confidence (Angelopoulos et al., 2023). However, in settings where exchangeability does not hold, standard CP methods can fail in providing valid coverage (Barber et al., 2023; Tibshirani et al., 2019). This is the case of time series data where temporal dependencies break the exchangeability assumption. Barber et al. (2023) show that reweighting the available residuals (in a data-independent fashion) can handle nonexchangeability and distribution shifts, and introduce NexCP, a weighting scheme that exponentially decays the past residuals over time. Barber & Tibshirani (2025) provides a unified view of several weighted CP methods as approaches to condition uncertainty estimates on the available information on the target data point. Other recent approaches, instead, fit a model to account for temporal dependencies and rely on asymptotic guarantees under some assumptions on the data-generating process (Xu & Xie, 2023a; Auer et al., 2023; Lee et al., 2025). For example, SPCI (Xu & Xie, 2023a) uses a quantile random forest trained at each step on the most recent scores. Conversely, HopCPT leverages a Hopfield network to learn a data-dependent reweighting scheme based on soft attention (Auer et al., 2023). Lee et al. (2025) use a reweighted Nadaraya-Watson estimator to perform quantile regression over the past nonconformity scores and derive the weights from the kernel function. Cini et al. (2025) proposes instead to use graph neural networks (Bacciu et al., 2020; Bronstein et al., 2021) to additionally condition uncertainty estimates on correlated time series. As already mentioned, we focus on methods that estimate uncertainty on the prediction of a pre-trained model; we refer to Benidis et al. (2022) for a discussion of probabilistic forecasting architectures.

# 2.2 ECHO STATE NETWORKS

ESNs are among the most popular reservoir computing (RC) methods (Jaeger, 2001; Lukoševičius & Jaeger, 2009). In particular, ESNs are RNNs that are randomly initialized to ensure the stability and meaningfulness of the representations and left untrained. An ESN encodes input sequences into nonlinear, high-dimensional state representations through the recurrent component, called *reservoir*, which extrapolates from the input a rich pool of dynamical features. The state update equation of the reservoir is

$$h_t = (1 - l)h_{t-1} + l \sigma (W_x x_t + W_h h_{t-1} + b),$$
 (3)

where  $x_t$  denotes the input at time t,  $W_x \in \mathbb{R}^{D_h \times D_x}$  and  $W_h \in \mathbb{R}^{D_h \times D_h}$  are fixed random weight matrices,  $b \in \mathbb{R}^{D_h}$  is a random bias vector,  $h_t \in \mathbb{R}^{D_h}$  is the reservoir state,  $l \in (0,1]$  is the leak rate which controls how much of the current state is retained at each update, and  $\sigma$  is a nonlinear activation, typically the hyperbolic tangent. When  $W_x$  and  $W_h$  are properly initialized, the reservoir states  $h_t$  provide expressive embeddings of the past inputs  $x_{\leq t}$ , which can be processed by a trainable readout to perform downstream tasks such as classification, clustering, or time-series forecasting (Bianchi et al., 2020). Since the reservoir embeddings already model nonlinear dynamics, the readout is usually implemented as a simple linear layer that can be trained efficiently.

Echo state property The reservoir dynamics should be neither too contractive nor chaotic. In the first case, the reservoir produces representations that are not sufficiently expressive. In the latter, the reservoir is unstable and responds inconsistently to nearly identical input sequences, resulting in poor generalization. ESNs are usually configured to ensure the Echo State Property (ESP), a necessary condition for global asymptotic stability under which the reservoir state asymptotically forgets its initial conditions and the past inputs. Formally, for any initial states  $h_0, h'_0$  and input sequence  $x_{1:T}$ , by calling  $f_R(h_0, x_{1:T})$  the final state of the reservoir initialized with  $h_0$  and being fed with  $x_{1:T}$ , then the ESP is defined as

$$||f_R(\mathbf{h}_0, \mathbf{x}_{1:T}) - f_R(\mathbf{h}_0', \mathbf{x}_{1:T})|| \to 0 \text{ as } T \to \infty$$
 (4)

The ESP can be achieved by properly initializing  $W_h$  and by setting its spectral radius  $\rho(W_h) < 1$  (Yildiz et al., 2012; Gallicchio & Scardapane, 2020).

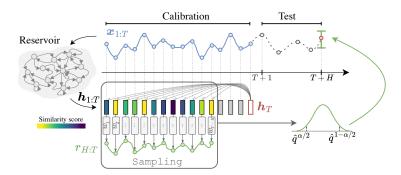


Figure 1: Let  $\{x_t\}_{t=1}^T$  and  $\{r_t\}_{t=1}^T$  be the time series and residuals of the calibration set, respectively. Each sequence  $x_{1:t}$  generates the state of the reservoir  $h_t$ . The last state  $h_T$  is the query state for the prediction of  $\hat{y}_{T+H}$ . We compute similarity scores between  $h_T$  and the calibration states  $\{h_t\}_{t=1}^{T-H}$ , which are used to reweight with  $\{w_t\}_{t=1}^{T-H}$  associated residuals  $\{r_t\}_{t=H}^T$  by resampling them. Quantiles are computed from the sampled residuals and used to build the PI for  $\hat{y}_{T+H}$ .

# 3 RESERVOIR CONFORMAL PREDICTION

We propose a novel approach for computing CP intervals for time series forecasting based on the representations generated by RC to model temporal dependencies. Our method leverages CP to provide accurate uncertainty estimates while maintaining the computational efficiency of RC.

## 3.1 RESERVOIR SAMPLING CONFORMAL PREDICTION

We construct PIs from a calibration set consisting of a time series of residuals of length T. In particular, we aim at reweighting the residuals based on similarities with the local dynamics at the target time step. We leverage the reservoir capability of embedding the observed dynamics  $x_{\leq t}$  into its state  $h_t$ . An overview of the whole procedure is depicted in Figure 1. In general, the input  $x_t$  can include any set of endogenous and exogenous variables available at time step t. For simplicity, we consider  $x_t$  to be the prediction residual ( $x_t = r_t$  obtained by training a generic forecasting model on a disjoint training set); the discussion on how to effectively incorporate exogenous variables is deferred to Section 3.2.

Given an ESN with dynamics defined in Equation 3 and hyperparameters  $\theta$ , the time series  $x_{1:t} = (x_1, \dots, x_t)$  is embedded into a sequence of t reservoir states which encode observed dynamics at each time step as

$$h_t = \mathrm{ESN}_{\boldsymbol{\theta}}(\boldsymbol{x}_{1:t}), \qquad h_t \in \mathbb{R}^{D_h}.$$
 (5)

In particular, to compute the PI for the H-step-ahead prediction  $\hat{y}_{t+H}$ , we take as *query state* the current state  $h_t$  (since we do not have access to states  $h_{>t}$ ) and compute similarity scores between  $h_t$  and the states in the calibration set. Note that in most scenarios, the calibration set can be iteratively updated as new observations become available (i.e., we can include in the calibration set all the observations up to the current time step t). Weights are computed as

$$z_s(\mathbf{h}_t) = \text{SIM}(\mathbf{h}_t, \mathbf{h}_s), \qquad 1 \le s \le T - H \tag{6}$$

$$\{w_1(\boldsymbol{h}_t), \dots, w_{T-H}(\boldsymbol{h}_t)\} = \text{SOFTMAX}\left\{\frac{z_1(\boldsymbol{h}_t)}{\tau}, \dots, \frac{z_{T-H}(\boldsymbol{h}_t)}{\tau}\right\}$$
(7)

where SIM is a similarity score (e.g., dot product) and  $\tau > 0$  is a temperature hyperparameter. Our goal is to approximate the conditional cumulative distribution function (CDF) of the residual at time t+H with

$$\widehat{F}(r \mid \boldsymbol{h}_t) := \sum_{s=1}^{T-H} w_s(\boldsymbol{h}_t) \mathbb{1}(r_{s+H} \le r) \qquad \approx \mathbb{P}(r_{t+H} \le r \mid \boldsymbol{x}_{\le t})$$
(8)

where  $w_s(\mathbf{h}_t)$  is a weight proportional to the similarity of  $\mathbf{h}_s$  to  $\mathbf{h}_t$  and such that  $\sum_s w_s(\mathbf{h}_t) = 1$ . Quantiles of the conditional error distribution can then be estimated as:

$$\widehat{Q}_{\beta}(\mathbf{h}_t) := \inf\{r \in \mathbb{R} : \widehat{F}(r \mid \mathbf{h}_t) \ge \beta\}. \tag{9}$$

In practice, instead of computing the weighted quantile in Eq. 9 directly, we approximate it through Monte Carlo sampling akin to Auer et al. (2023). More specifically, we sample residuals accordingly to the weights in Eq. 7 and compute the standard empirical  $\beta$ -quantile on the sampled residuals:

$$\hat{q}_{t+H}^{\alpha/2} = \widehat{Q}_{\alpha/2}(\mathbf{h}_t), \qquad \hat{q}_{t+H}^{1-\alpha/2} = \widehat{Q}_{1-\alpha/2}(\mathbf{h}_t).$$
 (10)

PIs with the desired confidence level  $\alpha$  are obtained as:

$$\widehat{C}_{T}^{\alpha}(\hat{y}_{t+H}) = \left[ \hat{y}_{t+H} + \hat{q}_{t+H}^{\alpha/2}, \hat{y}_{t+H} + \hat{q}_{t+H}^{1-\alpha/2} \right]. \tag{11}$$

To account for skewed error distributions, we follow the approach of Xu & Xie (2023b) and refine the interval in 11 by selecting the level  $\beta^*$  that minimizes its width:

$$\begin{split} \beta^* &= \underset{\beta \in [0,\alpha]}{\arg\min} \left[ \widehat{Q}_{1-\alpha+\beta}(\boldsymbol{h}_t) - \widehat{Q}_{\beta}(\boldsymbol{h}_t) \right] \\ \widehat{C}_T^{\alpha}(\hat{y}_{T+H}) &= \left[ \hat{y}_{T+H} + \widehat{Q}_{\beta^*}(\boldsymbol{h}_t), \hat{y}_{T+H} + \widehat{Q}_{1-\alpha+\beta^*}(\boldsymbol{h}_t) \right]. \end{split}$$

This can provide narrower PIs, at the cost of searching for the optimal level  $\beta$  (Xu & Xie, 2023b).

**Discussion** Essentially, RESCP is a local CP method for time series where local similarity is gauged by relying on reservoir states. By sampling more residuals associated with similar states, we condition the estimates on specific dynamics of interest in the time series. This allows us to build locally adaptive intervals. While Guan (2022) provides asymptotic guarantees for standard local LCP for i.i.d. data, such results cannot be applied in the context of time series. In the following section, we discuss the conditions that allow RESCP to provide (asymptotically) valid intervals. Notably, this will require assumptions on the nature of the process generating the data and on the reservoir dynamics. Finally, note that with finite samples, even with i.i.d. observation, data-dependent weights might introduce a bias that should be accounted for (Guan, 2022). In practice, we found that tuning RESCP hyperparameters on a validation set was sufficient to provide accurate coverage in most scenarios (see Sec. 4).

#### 3.1.1 THEORETICAL ANALYSIS

We analyze the theoretical properties of RESCP, and show in Corollary 3.7 that it asymptotically guarantees coverage under some regularity assumptions. As already mentioned, guarantees of standard CP methods break down for time series unless one relies on highly unrealistic assumptions (e.g., exchangeable or i.i.d. observations) (Barber et al., 2023). Moreover, even assuming exchangeability, it is impossible to construct finite-length PIs with distribution-free conditional coverage guarantees (Lei & Wasserman, 2013; Vovk, 2012). We start by introducing the assumptions needed to prove the consistency of the weighted empirical CDF in Eq. 8.

**Assumption 3.1** (Time-invariant and mixing process). We assume that the process  $\{Z_t = (\boldsymbol{x}_t, r_{t+H})\}_{t=1}^{\infty}$  is time-invariant and strongly mixing  $(\alpha$ -mixing) with coefficient  $\alpha(k) \to 0$  as  $t \to \infty$ .

Intuitively, Assumption 3.1 implies that the system's dynamics do not change over time and that it "forgets" its initial conditions and past structure, thereby allowing observations that are far apart in time to be treated as if they were independent.

**Assumption 3.2** (ESP and Lipschitz properties in ESN). Let  $\mathbf{x}_{\leq t} = (\dots, \mathbf{x}_{t-1}, \mathbf{x}_t)$  denote the history. Define the fading-memory metric for  $\gamma \in (0,1)$  by

$$d_{\mathrm{fm}}(\boldsymbol{x}_{\leq t}, \boldsymbol{x}'_{\leq t}) \coloneqq \sum_{k=1}^{\infty} \gamma^{k} \| \boldsymbol{x}_{t-k+1} - \boldsymbol{x}'_{t-k+1} \|_{2}.$$

Then the ESN state map  $ESN_{\theta}: \mathbf{x}_{\leq t} \mapsto \mathbf{h}_{t} \in \mathbb{R}^{D_{h}}$  is well-defined and causal with fading memory (i.e., the output only depends on the past and present inputs), and there exist a constant  $L_{X} > 0$  such that for all histories  $\mathbf{x}_{\leq t}, \mathbf{x}'_{\leq t}$ :

$$\|\mathit{ESN}_{m{ heta}}(m{x}_{\leq t}) - \mathit{ESN}_{m{ heta}}(m{x}_{\leq t}')\|_{2} \leq L_{X} d_{\mathrm{fm}}(m{x}_{\leq t}, m{x}_{\leq t}').$$

Assumption 3.2 tells us that the reservoir state evolves in a stable and predictable manner: past inputs influence the state in a controlled, decaying way, and small perturbations are eventually washed out from the reservoir's state. These assumptions are commonly used in RC (Grigoryeva & Ortega, 2019; Gallicchio & Micheli, 2011).

**Assumption 3.3** (Continuity of the conditional law). Let  $\mathbb{P}(r_{t+H} \in \mathbb{R} \mid \mathbf{x}_{\leq t})$  denote the conditional law of  $r_{t+H}$  given the past. Then for every  $r \in \mathbb{R}$  the map

$$\boldsymbol{x}_{\leq t} \; \mapsto \; F(r \mid \boldsymbol{x}_{\leq t})$$

is continuous in  $x_{\leq t}$  with respect to  $d_{\mathrm{fm}}$ , i.e.

$$\bar{d}_{\text{fm}}(\boldsymbol{x}_{\leq t}, \boldsymbol{x}'_{\leq t}) \to 0 \implies \sup_{b} |F(r \mid \boldsymbol{x}_{\leq t}) - F(r \mid \boldsymbol{x}'_{\leq t})| \to 0$$

Intuitively, Assumption 3.3 ensures that if two histories are similar (in the fading-memory sense), then the distributions of the future residuals conditioned on these histories are also close. This is a reasonable assumption to enable learning.

**Definition 3.4** (Effective sample size). Let n be the number of available calibration samples, the effective sample size (Kong et al., 1994; Liu, 2008) is defined in our setting as

$$m_n := \left(\sum_{i=1}^n (w_i^{(n)})^2\right)^{-1}.$$
 (12)

Note that at high temperature, weights are approximately uniform  $(w_i^{(n)} \approx \frac{1}{n})$  and  $\sum_i (w_i^{(n)})^2 \approx \frac{1}{n}$ , hence  $m_n \approx n$ , which means that the method collapses to vanilla SCP as all calibration points are treated symmetrically. If instead the temperature is too low, then all the mass concentrates on only one point, meaning that  $\sum_i (w_i^{(n)})^2 = 1$  and hence  $m_n = 1$ .

**Assumption 3.5** (Softmax weighting scheme). Let n be the number of available calibration samples, and let  $\tau_n = \tau(n)$  be the temperature parameter of the SOFTMAX of equation 7,  $m_n$  the associated effective sample size, and  $\mathbf{w}^{(n)}$  the output weights. Assume:

- (i)  $\tau_n$  is configured to slowly decrease as  $n \to \infty$ , so that  $m_n \to \infty$ ;
- (ii) for every  $\delta > 0$ .

$$\sum_{s: \|\boldsymbol{h}_s - \boldsymbol{h}_t\| \geq \delta} w_s^{(n)}(\boldsymbol{h}_t) \ \xrightarrow[n \to \infty]{\mathbb{P}} \ 0.$$

Condition (ii) states that for sufficiently small temperature  $\tau_n$ , the softmax normalization should concentrate almost all mass on those calibration points whose states lie in a shrinking neighborhood of  $h_t$ . On the other hand, we also need these points within the shrinking neighborhood of  $h_t$  to increase as the size of the calibration set increases. This can be seen as a bias-variance tradeoff:  $\tau_n$  must be set to a small enough value to localize weights at points similar to the query state (reducing bias), but at the same time it should be large enough to guarantee a good effective sample size. Theoretically, this requirement translates into Assumption 3.5 (i) which prescribes that the temperature has to shrink at a reasonably slow rate: one that allows the number of effective neighbors  $m_n$  to diverge as n grows.

**Theorem 3.6** (Consistency of the weighted empirical CDF). Let  $\widehat{F}_n(\cdot \mid \mathbf{h}_t)$  denote the conditional weighted empirical CDF in Eq. 8 with calibration data  $\{(\mathbf{x}_i, r_i)\}_{i=1}^n$  and  $F_n(r \mid \mathbf{h}_t) \coloneqq \mathbb{P}(r_{t+H} \le r \mid \mathbf{h}_t)$ . Under Assumptions 3.1–3.5, we have for any query state  $\mathbf{h}_t \in \mathbb{R}^{D_h}$ ,

$$\sup_{b \in \mathbb{R}} |\widehat{F}_n(b \mid \boldsymbol{h}_t) - F(b \mid \boldsymbol{h}_t)| \xrightarrow[n \to \infty]{\mathbb{P}} 0.$$

The proof can be found in App. A. Theorem 3.6 implies the consistency of the empirical conditional quantile estimator  $\widehat{Q}_{\beta}(h_t)$  as defined in Equation 9 as estimator to the true conditional quantile  $Q_{\beta}(h_t)$ . Finally, it is trivial to show that the asymptotic coverage of RESCP is guaranteed.

**Corollary 3.7** (Asymptotic conditional coverage guarantee). *Under Assumptions 3.1–3.5*, for any  $\alpha \in (0,1)$  and for  $n \to \infty$ ,

$$\mathbb{P}\left(y_{t+h} \in \hat{C}_t^{\alpha}(\hat{y}_{t+h}) \mid \boldsymbol{h}_t\right) \xrightarrow[n \to \infty]{\mathbb{P}} (1 - \alpha).$$

Remark In the theoretical analysis, we model the distribution of future residuals conditioned on the state  $h_t$ , i.e., on the reservoir state at time t. This is clearly weaker than conditioning on the entire history  $x_{\leq t}$ . Recalling Eq. 8, whether or not  $\widehat{F}(r \mid h_t)$  is a good approximation of  $F(r \mid x_{\leq t})$  entirely depends on the ability of the ESN to encode all relevant information from the past in its state. In particular, if the state representation captures all the relevant information, then RESCP provides asymptotic conditional coverage given such a history. Conversely, in the extreme case where every sequence is mapped by the reservoir into the same uninformative state, RESCP would simply provide marginal coverage, as one would expect. We refer the reader to the rich body of literature on the expressiveness of ESNs, particularly regarding their universal approximation properties (Grigoryeva & Ortega, 2018a;b; Li & Yang, 2025) and their effectiveness in extracting meaningful representations from time series data (Bianchi et al., 2020).

#### 3.1.2 TIME-DEPENDENT WEIGHTS

RESCP, by default, compares the query state with *any* time step from the calibration state. Note that state representations do not include any positional encoding that accounts for how far back in time a certain sample is, treating recent and distant residuals alike. This is not problematic for time-invariant and stable processes, which we assumed in 3.1.1 to provide coverage guarantees. To deal, instead, with distribution shifts, we 1) update the calibration set over time, keeping its size N fixed using a first-in-first-out approach, and 2) make the weights time-dependent, following an approach similar to NexCP (Barber et al., 2023). In particular, we condition the weights on the distance between time steps as

$$w_i(\mathbf{h}_t, t) = \gamma(\Delta(t, i))w_i(\mathbf{h}_t), \tag{13}$$

where  $\gamma:\mathbb{N}\to\mathbb{R}$  is a discount function that maps the distance between time steps  $\Delta(t,i)$  to a decay factor. The discount schedule can be chosen in different ways, e.g., exponential like NexCP (Barber et al., 2023) or linear. In our settings, we observed that a linear decay–i.e.,  $\gamma(\Delta(t,i)) \coloneqq 1/\Delta(t,i)$ – allowed to keep PIs up-to-date without reducing too much the effective sample size. This is particularly important if assumption 3.1 is violated and the underlying dynamics change over time. It is worth noting that RESCP is a non-parametric approach without learnable parameters that is much more robust to non-stationarity compared to approaches that train a model (Auer et al., 2023; Cini et al., 2025), which must necessarily be updated when time-invariance is lost.

#### 3.2 HANDLING EXOGENOUS INPUTS

Adding exogenous variables as input to the reservoir can affect the internal dynamics of the network and, thus, its states. Depending on the characteristics of specific exogenous variables, they can harm the effectiveness of the reservoir in localizing prediction w.r.t. the dynamics of the target variable. Since the network is not trained, adjusting to the relevance and characteristics of the covariates is not possible with the completely unsupervised approach discussed in Sec. 3.1. As an alternative to RESCP in these scenarios, we consider a variant of our original approach, called Reservoir Conformal Quantile Regression (RESCQR), which can account for exogenous inputs by relying on quantile regression. In practice, we use a linear readout to map the state at each time step to a set of quantiles of interest. Specifically, we use states  $h_{1:T-H}$  and residuals  $r_{H+1:T}$  in the calibration set to fit a linear model  $\hat{Q}_{\alpha}(h_t)$  as a quantile predictor. To do so, we train the readout to minimize the pinball loss for the target quantiles  $\{\beta_1,\ldots,\beta_M\}$  of the conformity scores:

$$\mathcal{L}_{\beta_i}(\widehat{Q}_{\beta_i}(\boldsymbol{h}_t), r) = \begin{cases} (1 - \beta_i)(\widehat{Q}_{\beta_i}(\boldsymbol{h}_t) - r), & \widehat{Q}_{\beta_i}(\boldsymbol{h}_t) \ge r \\ \beta_i(r - \widehat{Q}_{\beta_i}(\boldsymbol{h}_t)), & \widehat{Q}_{\beta_i}(\boldsymbol{h}_t) < r \end{cases}$$

We can then construct PIs as discussed in Sec. 3.1. As we will show in Sec. 4, RESCQR provides a practical methods approach that can work well in scenarios where the availability of enough calibration data allows for training the readout and informative exogenous variables are available.

# 4 EXPERIMENTS

We compare the performances of RESCP against state-of-the-art conformal prediction baselines on time series data coming from several applications.

Table 1: Performance comparison for  $\alpha = 0.1$ .  $\Delta$ Cov values are color-coded for undercoverage cases: yellow (1-2%), orange (2-4%), red (>4%). The top three Winkler scores for each scenario are highlighted: **bold+underlined** (1st), <u>underlined</u> (2nd), dot-underlined (3rd).

				Non-learning				
	Metric	SPCI	HopCPT	ning CP-QRNN	RESCQR	SCP	NexCP	RESCP
7	$\Delta Cov$	$0.05 \pm 0.17$	$-1.64\pm1.18$	$-0.26\pm0.92$	$-1.10\pm0.91$	0.37	1.46	$0.74 \pm 0.24$
RNN	PI-Width	$70.41 \pm 1.73$	$60.49 \pm 2.10$	$55.74 \pm 0.98$	$59.99 \pm 1.72$	79.15	89.56	$62.25 \pm 0.75$
~	Winkler	151.50±1.18	$112.46 \pm 9.34$	$78.42 \scriptstyle{\pm 0.20}$	$82.76{\scriptstyle\pm0.26}$	171.41	164.96	$104.24 \pm 0.79$
	$\Delta \text{Cov}$	-0.16±0.46	$1.32 \pm 0.62$	1.37±2.09	$-3.51\pm 16.26$	0.35	1.46	$3.09\pm0.35$
Solar Transf	PI-Width	$71.36 \pm 3.21$	$61.49 \pm 1.74$	$55.70{\scriptstyle\pm0.95}$	$59.56{\scriptstyle\pm1.59}$	79.03	89.21	$63.34 \pm 1.11$
$\sim$	Winkler	$152.85 \pm 0.67$	$107.59 \pm 6.07$	$\textbf{77.61} {\scriptstyle\pm0.25}$	$82.16 \pm 0.32$	169.64	163.34	$103.13 \pm 0.58$
[₹	$\Delta \text{Cov}$	$0.51 \pm 0.36$	2.88±2.13	$-0.41 \pm 0.42$	$-2.03\pm0.62$	0.16	1.76	$0.68 \pm 0.95$
ARIMA	PI-Width	91.71±1.18	$143.32 \pm 7.36$	$59.17 \pm 0.41$	$66.19{\scriptstyle\pm0.81}$	124.19	137.53	$77.17{\scriptstyle\pm2.07}$
AF	Winkler	$148.86 \pm 0.34$	$173.49 \pm 5.15$	$\underline{77.34{\scriptstyle\pm0.30}}$	$85.38 \pm 0.45$	215.77	207.58	$110.38 \pm 4.03$
7	$\Delta \text{Cov}$	$-1.73 \pm 0.67$	$-5.18\pm12.67$	$-1.86\pm 2.16$	$-1.21 \pm 1.65$	-0.32	0.00	$-0.70 \pm 0.77$
RNN	PI-Width	$67.93 \pm 1.43$	$68.47 \pm 13.30$	$61.71 \pm 4.91$	$65.53 \pm 4.09$	67.99	69.79	$65.96{\scriptstyle\pm2.50}$
	Winkler	$124.51 \pm 2.25$	$140.50 \pm 43.64$	$104.03 \pm 0.99$	$105.43 \pm 0.85$	126.41	124.10	$106.07 \pm 0.47$
Beijing Transf	$\Delta \text{Cov}$	$-0.98\pm1.04$	$-8.05 \pm 16.41$	$-1.07 \pm 0.69$	$-1.43\pm1.10$	-0.41	0.03	$-0.49\pm0.59$
Beijing Transf	PI-Width	$69.96 \pm 4.53$	$61.76 \pm 14.39$	$62.41 \pm 1.66$	$64.41 \pm 2.72$	67.46	69.64	$64.06 \pm 1.74$
ğ   L		$125.41 \pm 1.43$	$140.30 \pm 36.01$	$\underline{102.81 \scriptstyle{\pm 0.44}}$	$105.97 \pm 1.21$	126.63	124.35	$103.64 \pm 0.21$
₹	$\Delta \text{Cov}$	$-0.23\pm0.40$	$-1.37 \pm 0.26$	$-1.54 \pm 0.77$	$-1.42\pm1.15$	-0.24	-0.16	$0.63 \pm 0.22$
ARIMA	PI-Width	$74.68 \pm 1.21$	$67.78 \pm 0.50$	$61.80 \pm 1.68$	$66.01 \pm 3.01$	75.72	76.45	$70.43 \pm 0.86$
A	Winkler	$130.59 \pm 0.59$	$122.48 \pm 4.36$	$\underline{101.84 {\pm} 0.67}$	$107.20 \pm 1.21$	135.07	132.03	$108.75 \pm 0.31$
7	$\Delta \text{Cov}$	$2.98 \pm 0.65$	$2.75 \pm 0.08$	$-1.07 \pm 2.52$	$3.18\pm 1.25$	2.29	1.64	$1.13 \pm 0.27$
RNN	PI-Width	$0.0241 \pm 0.0007$	$0.0404 {\scriptstyle\pm0.0001}$	$0.0341 {\scriptstyle \pm 0.0018}$	$0.0383 {\pm} 0.0008$	0.0444	0.0405	$0.0210 \pm 0.0001$
	Winkler				$0.0464 \scriptstyle{\pm 0.0005}$	0.0517	0.0492	$\underline{\textbf{0.0264} {\pm 0.0002}}$
Exchange Transf	$\Delta \text{Cov}$	4.44±0.35	$2.98 \pm 0.07$	$-0.57 \pm 1.58$	$0.82 \pm 1.89$	4.57	3.25	1.46±0.18
xchang	PI-Width				$0.0365 {\scriptstyle\pm0.0013}$			$0.0229 \pm 0.0001$
	William				$0.0475 \pm 0.0008$		0.0602	$\underline{\textbf{0.0294} {\pm 0.0001}}$
₹	$\Delta \text{Cov}$	$3.49 \pm 0.41$	$2.07 \pm 0.08$	$-1.22\pm 1.78$	$0.68 \pm 1.58$	3.08	2.13	$0.38 \pm 0.41$
ARIMA	PI-Width				$0.0351 \!\pm\! 0.0009$	0.0387		$0.0207 \pm 0.0001$
					$0.0455 \pm 0.0006$			$0.0268 \pm 0.0001$
z	$\Delta \text{Cov}$	$-0.78 \pm 1.88$	$-2.07 \pm 0.26$	$-11.84 \pm 8.93$	$-17.24 \pm 5.56$	-0.99	-0.33	$1.56 \pm 0.62$
RNN	PI-Width	$8.99{\scriptstyle\pm0.68}$	$18.98{\scriptstyle\pm0.19}$	$15.59 \pm 1.71$	$15.44 \pm 1.41$	19.63	20.15	$9.61{\scriptstyle\pm0.26}$
"	Winkler	$14.27 \pm 0.19$	$27.56 \pm 0.04$	$31.48 \pm 5.20$	$33.25 \pm 3.48$	27.60	26.83	$12.91 \pm 0.23$
F. St	$\Delta \text{Cov}$	-1.41±1.29	$-2.40 \pm 0.27$	-21.92±9.88	$-23.98 \pm 7.16$	-5.52	-0.45	$3.54 \pm 0.32$
ACEA Transf	PI-Width	$9.10 \pm 0.23$	$18.40{\scriptstyle\pm0.24}$	$13.61 \pm 1.49$	$14.25{\scriptstyle\pm1.22}$	16.53	20.20	$10.10 \pm 0.16$
	WIIIKICI	$14.58 \pm 0.36$	$27.52{\scriptstyle\pm0.02}$	$37.59 \pm 5.45$	$37.54 \pm 4.44$	29.24	27.47	$12.90 \pm 0.16$
₹	$\Delta \text{Cov}$	$1.41 \pm 0.90$	$-3.55 \pm 0.06$	$-26.91 \pm 11.91$	$-22.21 \pm 4.57$	-0.75	-0.40	$5.02 \pm 0.40$
ARIMA	PI-Width	$12.46 \pm 0.35$	$34.83{\scriptstyle\pm0.00}$	$18.48{\scriptstyle\pm2.59}$	$19.80{\scriptstyle\pm1.35}$	38.13	36.08	$13.63{\scriptstyle\pm0.55}$
	Winkler	$17.36 \pm 0.13$	$44.70 \pm 0.03$	$53.47 \pm 11.95$	$49.49 \pm 3.00$	45.99	43.70	$16.21 \pm 0.53$

Datasets and baselines We evaluated our method across four datasets. 1) The Solar dataset comes from the US National Solar Radiation Database (Sengupta et al., 2018). We used the dataset containing 50 time series from different locations over a period of 3 years, as done in previous work (Auer et al., 2023). 2) The Beijing dataset contains air quality measurements taken over a period of 4 years from 12 locations in the city of Beijing, China (Zhang et al., 2017). 3) The Exchange dataset consists of a collection of the daily exchange rates of eight countries from 1990 to 2016 (Lai et al., 2017). 4) ACEA contains electricity consumption data coming from the backbone of the energy supply network in the city of Rome (Bianchi et al., 2015). More details are reported in App. B. We compare RESCP and its RESCQR variant against the competitors presented in Sec. 2: 1) vanilla SCP (Vovk et al., 2005), 2) NexCP (Barber et al., 2023), 3) SPCI (Xu & Xie, 2023a), and 4) HopCPT Auer et al. (2023). We also include 5) CP-QRNN, an RNN trained to perform quantile regression calibration set using the pinball loss, which is similar to the model proposed by Cini et al. (2025). For RESCP, we use time-dependent weights (with a linear decay schedule) and the cosine similarity between reservoir states as the similarity score.

Table 2: Runtime (in seconds) using the RNN point forecasting baseline.

	Learning					Non-learning		
Dataset	SPCI	HopCPT	CP-QRNN	RESCQR	SCP	NexCP	RESCP	
Solar	1039.8±2	4574.6±1356.5	172.4±17	$82.1 \pm 6.7$	18	66	52.9±14.7	
Beijing	351.4±2.1	1838.5±202.4	81.6±3	46±2.1	9	29	34.6±1.4	
Exchange	50.6±0.5	318±0	$36.9{\scriptstyle\pm1.2}$	$16 \pm 0.9$	2	2	$6.5{\scriptstyle\pm0.5}$	
ACEA	227.6±5.8	$2262.8 \pm 918.9$	$95.3{\scriptstyle\pm9.5}$	$56.6{\scriptstyle\pm2.9}$	7	57	$70.5{\scriptstyle\pm0.5}$	

Experimental setup and evaluation metrics In our experiments, we adopted a 40%/40%/20% split for training, calibration, and testing sets, respectively. As base models, we consider three different point predictors: a simple RNN with gated recurrent cells (Cho et al., 2014), a decoder-only Transformer (Vaswani et al., 2017), and an ARIMA model (Box & Jenkins, 1970). After training, we evaluated each model on the calibration set and saved the residuals, which we then used in all the baselines to compute the PIs. As evaluation metrics, we considered the  $\Delta$ Cov, i.e., the difference between the specified confidence level  $1-\alpha$  and the achieved coverage on the test set, the width of the PIs, and the Winkler score (Winkler, 1972), which penalizes the PI width whenever the observed value falls outside the computed interval, with the penalty scaled proportionally to the magnitude of the deviation. Model selection for all methods is done by minimizing the Winkler score over a validation set, except for HopCPT, which follows a custom procedure (Auer et al., 2023). More details on model selection for our methods are reported in App. C.

## 4.1 RESULTS

Results across the datasets and base models are summarized in Tab. 1, where CP methods are grouped based on whether they rely on learning or not. RESCP achieves competitive performance across all settings and datasets, both in terms of coverage and Winkler score, while remaining highly scalable. Notably, RESCPs outperforms HopCPT in almost all scenarios. This is particularly remarkable since both methods use weighted empirical distributions to model uncertainty, but HopCPT requires training an attention-based architecture end-to-end, while RESCPs does not perform any form of training. Moreover, similarly to NexCP and differently from most of the other methods, RESCP provides approximately valid coverage across all scenarios, but with large improvements in PI width, which is reduced up to 60%. The CP-QRNN baseline achieves strong performance on the large datasets with informative exogenous variables, outperforming all methods in Solar, but fails in achieving good results in the smaller datasets, such as ACEA and Exchange. Moreover, on those datasets, trainable methods obtain worse performance, likely due to distribution changes that require the trained models to be updated over time. RESCOR achieves competitive performance against the competitors, including more complex models such as CP-QRNN in Solar, while being more scalable. All other baselines obtain good coverage in most settings, but produce prediction intervals that are much more conservative. The runtime of each CP method is reported in Tab. 2 and shows the advantage of RESCP in terms of scalability against methods that require fitting a model. Moreover, as RESCP does not need centralized training on a GPU, it can be easily scaled in a distributed setting to process big data. Finally, App. D reports additional results on the sensitivity of RESCP to different hyperparameter configurations.

# 5 CONCLUSION AND FUTURE WORKS

We introduced <u>Reservoir Conformal Prediction</u> (RESCP), a simple, fast, and effective method for uncertainty quantification in time series forecasting that provides accurate uncertainty estimation across several tasks. We showed that RESCP is principled and theoretically sound. In particular, assuming strong mixing conditions and reasonable regularity conditions, we demonstrated that RESCP provides asymptotically valid conditional coverage guarantees. Results on a diverse suite of benchmarks show that RESCP achieves state-of-the-art performance while being more scalable and computationally efficient than its competitors. There are several directions for future work. In particular, it would be interesting to explore extensions to modeling joint distributions in multistep forecasting (Sun & Yu, 2024), treat multidimensional time series (Feldman et al., 2023; Xu et al., 2024), and spatiotemporal data (Cini et al., 2025).

# REFERENCES

- Anastasios N Angelopoulos, Stephen Bates, et al. Conformal prediction: A gentle introduction. Foundations and Trends® in Machine Learning, 16(4):494–591, 2023.
- Andreas Auer, Martin Gauch, Daniel Klotz, and Sepp Hochreiter. Conformal prediction for time series with modern hopfield networks. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), <u>Advances in Neural Information Processing Systems</u>, volume 36, pp. 56027–56074. Curran Associates, Inc., 2023.
  - Davide Bacciu, Federico Errica, Alessio Micheli, and Marco Podda. A gentle introduction to deep learning for graphs. Neural Networks, 129:203–221, 2020.
  - Rina Foygel Barber and Ryan J Tibshirani. Unifying different theories of conformal prediction. arXiv preprint arXiv:2504.02292, 2025.
  - Rina Foygel Barber, Emmanuel J. Candès, Aaditya Ramdas, and Ryan J. Tibshirani. Conformal prediction beyond exchangeability. <u>The Annals of Statistics</u>, 51(2):816 845, 2023. doi: 10. 1214/23-AOS2276.
  - Konstantinos Benidis, Syama Sundar Rangapuram, Valentin Flunkert, Yuyang Wang, Danielle Maddix, Caner Turkmen, Jan Gasthaus, Michael Bohlke-Schneider, David Salinas, Lorenzo Stella, François-Xavier Aubet, Laurent Callot, and Tim Januschowski. Deep Learning for Time Series Forecasting: Tutorial and Literature Survey. ACM Comput. Surv., 55(6), dec 2022. ISSN 0360-0300. doi: 10.1145/3533382. URL https://doi.org/10.1145/3533382.
  - Filippo Maria Bianchi, Enrico De Santis, Antonello Rizzi, and Alireza Sadeghian. Short-term electric load forecasting using echo state networks and pca decomposition. <u>IEEE Access</u>, 3:1931–1943, 2015. doi: 10.1109/ACCESS.2015.2485943.
  - Filippo Maria Bianchi, Simone Scardapane, Sigurd Løkse, and Robert Jenssen. Reservoir computing approaches for representation and classification of multivariate time series. <u>IEEE Transactions on Neural Networks and Learning Systems</u>, 2020.
  - George EP Box and Gwilym M Jenkins. <u>Time Series Analysis: Forecasting and Control.</u> Holden-Day, 1970.
  - Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. <a href="mailto:arXiv:2104.13478">arXiv:2104.13478</a>, 2021.
  - Andrea Ceni and Claudio Gallicchio. Edge of stability echo state network. <u>IEEE Trans. Neural Netw. Learn. Syst.</u>, 36(4):7555–7564, April 2025.
  - Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. Learning phrase representations using rnn encoder–decoder for statistical machine translation. <a href="mailto:arXiv:1406.1078"><u>arXiv:1406.1078</u></a>, 2014.
  - Andrea Cini, Alexander Jenkins, Danilo Mandic, Cesare Alippi, and Filippo Maria Bianchi. Relational Conformal Prediction for Correlated Time Series. <u>International Conference on Machine</u> Learning, 2025.
- Jonathan Dong, Erik Börve, Mushegh Rafayelyan, and Michael Unser. Asymptotic stability in reservoir computing. In 2022 International Joint Conference on Neural Networks (IJCNN), volume 148.34, pp. 01–08. IEEE, July 2022.
- Paul Doukhan. Mixing: Properties and Examples. Information Systems Development: Reflections, Challenges and New Directions №85. Springer, 1 edition, 1994. ISBN 1461226422; 9780387942148; 0387942149; 9781461226420.
- Shai Feldman, Stephen Bates, and Yaniv Romano. Calibrated multiple-output quantile regression with representation learning. Journal of Machine Learning Research, 24(24):1–48, 2023.
- Claudio Gallicchio and Alessio Micheli. Architectural and markovian factors of echo state networks. Neural Networks, 24(5):440–456, June 2011. ISSN 0893-6080. doi: 10.1016/j.neunet.2011.02. 002. URL http://dx.doi.org/10.1016/j.neunet.2011.02.002.

- Claudio Gallicchio and Simone Scardapane. <u>Deep Randomized Neural Networks</u>, pp. 43–68. Springer International Publishing, Cham, 2020. ISBN 978-3-030-43883-8. doi: 10.1007/978-3-030-43883-8\_3. URL https://doi.org/10.1007/978-3-030-43883-8\_3.
  - Lyudmila Grigoryeva and Juan-Pablo Ortega. Echo state networks are universal. Neural Networks, 108:495–508, 2018a. ISSN 0893-6080. doi: https://doi.org/10.1016/j.neunet. 2018.08.025. URL https://www.sciencedirect.com/science/article/pii/S089360801830251X.
  - Lyudmila Grigoryeva and Juan-Pablo Ortega. Universal discrete-time reservoir computers with stochastic inputs and linear readouts using non-homogeneous state-affine systems. <u>Journal of Machine Learning Research</u>, 19(1):892–931, January 2018b. ISSN 1532-4435.
  - Lyudmila Grigoryeva and Juan-Pablo Ortega. Differentiable reservoir computing. <u>Journal of Machine Learning Research</u>, 20(179):1–62, 2019. URL http://jmlr.org/papers/v20/19-150.html.
  - Leying Guan. Localized conformal prediction: a generalized inference framework for conformal prediction. Biometrika, 110(1):33–50, 07 2022. ISSN 1464-3510. doi: 10.1093/biomet/asac040.
  - Tao Hong and Shu Fan. Probabilistic electric load forecasting: A tutorial review. <a href="International Journal of Forecasting">International Journal of Forecasting</a>, 32(3):914–938, 2016. ISSN 0169-2070. doi: <a href="https://doi.org/10.1016/j.ijforecast.2015.11.011">https://doi.org/10.1016/j.ijforecast.2015.11.011</a>. URL <a href="https://www.sciencedirect.com/science/article/pii/S0169207015001508">https://doi.org/10.1016/j.ijforecast.2015.11.011</a>. URL <a href="https://www.sciencedirect.com/science/article/pii/S0169207015001508">https://doi.org/10.1016/j.ijforecast.2015.11.011</a>. URL <a href="https://www.sciencedirect.com/science/article/pii/S0169207015001508">https://doi.org/10.1016/j.ijforecast.2015.11.011</a>. URL <a href="https://www.sciencedirect.com/science/article/pii/S0169207015001508">https://www.sciencedirect.com/science/article/pii/S0169207015001508</a>.
  - Rohan Hore and Rina Foygel Barber. Conformal prediction with local weights: randomization enables local guarantees. arXiv preprint arXiv:2310.07850, 2023.
  - Herbert Jaeger. The "echo state" approach to analysing and training recurrent neural networks. GMD Report 148, GMD German National Research Institute for Computer Science, 2001. URL http://www.faculty.jacobs-university.de/hjaeger/pubs/EchoStatesTechRep.pdf.
  - Vilde Jensen, Filippo Maria Bianchi, and Stian Normann Anfinsen. Ensemble conformalized quantile regression for probabilistic time series forecasting. <u>IEEE Transactions on Neural Networks</u> and Learning Systems, 35(7):9014–9025, 2022.
  - Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun (eds.), <u>ICLR (Poster)</u>, 2015.
  - Augustine Kong, Jun S. Liu, and Wing Hung Wong. Sequential imputations and bayesian missing data problems. <u>Journal of the American Statistical Association</u>, 89(425):278–288, 1994. ISSN 01621459, 1537274X. URL http://www.jstor.org/stable/2291224.
  - Guokun Lai, Wei-Cheng Chang, Yiming Yang, and Hanxiao Liu. Modeling long- and short-term temporal patterns with deep neural networks. arXiv preprint arXiv:1703.07015, 2017.
  - Jonghyeok Lee, Chen Xu, and Yao Xie. Kernel-based optimally weighted conformal time-series prediction. The Thirteenth International Conference on Learning Representations, 2025. URL https://openreview.net/forum?id=oP7arLOWix.
  - Junghwan Lee, Chen Xu, and Yao Xie. Transformer conformal prediction for time series. <u>arXiv</u> preprint arXiv:2406.05332, 2024.
  - Jing Lei and Larry Wasserman. Distribution-free prediction bands for non-parametric regression.

    Journal of the Royal Statistical Society Series B: Statistical Methodology, 76(1):71–96, 07 2013.

    ISSN 1369-7412. doi: 10.1111/rssb.12021. URL https://doi.org/10.1111/rssb.12021.
    - Jing Lei, Max G'Sell, Alessandro Rinaldo, Ryan J Tibshirani, and Larry Wasserman. Distribution-free predictive inference for regression. <u>Journal of the American Statistical Association</u>, 113 (523):1094–1111, 2018.

- Zhen Li and Yunfei Yang. Universality and approximation bounds for echo state networks with random weights. <u>IEEE Transactions on Neural Networks and Learning Systems</u>, 36(2):2720–2732, 2025. doi: 10.1109/TNNLS.2023.3339512.
- Jun S. Liu. Monte Carlo Strategies in Scientific Computing. Springer Publishing Company, Incorporated, 2008. ISBN 0387763694.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In <u>International Conference on Learning Representations</u>, 2019. URL https://openreview.net/forum?id=Bkg6RiCqY7.
- Mantas Lukoševičius and Herbert Jaeger. Reservoir computing approaches to recurrent neural network training. Computer Science Review, 3(3):127–149, 2009. ISSN 1574-0137. doi: https://doi.org/10.1016/j.cosrev.2009.03.005. URL https://www.sciencedirect.com/science/article/pii/S1574013709000173.
- Spyros Makridakis, Richard Kirkham, Ann Wakefield, Maria Papadaki, Joanne Kirkham, and Lisa Long. Forecasting, uncertainty and risk; perspectives on clinical decision-making in preventive and curative medicine. <a href="International Journal of Forecasting">International Journal of Forecasting</a>, 35(2):659-666, 2019. ISSN 0169-2070. doi: <a href="https://doi.org/10.1016/j.ijforecast.2017.11.003">https://doi.org/10.1016/j.ijforecast.2017.11.003</a>. URL <a href="https://www.sciencedirect.com/science/article/pii/S0169207017301346">https://www.sciencedirect.com/science/article/pii/S0169207017301346</a>.
- Tim Palmer and Renate Hagedorn (eds.). <u>Predictability of Weather and Climate</u>. Cambridge University Press, Cambridge, United Kingdom, 2006.
- Yaniv Romano, Evan Patterson, and Emmanuel Candes. Conformalized quantile regression. Advances in neural information processing systems, 32, 2019.
- David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. <u>International Journal of Forecasting</u>, 36(3): 1181–1191, 2020.
- Manajit Sengupta, Yu Xie, Anthony Lopez, Aron Habte, Galen Maclaurin, and James Shelby. The national solar radiation data base (nsrdb). Renewable and Sustainable Energy Reviews, 89:51–60, 2018. ISSN 1364-0321. doi: https://doi.org/10.1016/j.rser.2018.03.003. URL https://www.sciencedirect.com/science/article/pii/S136403211830087X.
- 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.
- Ryan J Tibshirani, Rina Foygel Barber, Emmanuel Candes, and Aaditya Ramdas. Conformal prediction under covariate shift. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper\_files/paper/2019/file/8fb21ee7a2207526da55a679f0332de2-Paper.pdf.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In <u>Advances in Neural Information</u> Processing Systems (NeurIPS), volume 30, 2017.
- Vladimir Vovk. Conditional validity of inductive conformal predictors. In Steven C. H. Hoi and Wray Buntine (eds.), Proceedings of the Asian Conference on Machine Learning, volume 25 of Proceedings of Machine Learning Research, pp. 475–490, Singapore Management University, Singapore, 04–06 Nov 2012. PMLR. URL https://proceedings.mlr.press/v25/vovk12.html.
- Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. <u>Algorithmic learning in a random world</u>, volume 29. Springer, 2005.
- Robert L. Winkler. A decision-theoretic approach to interval estimation. <u>Journal of the American</u> Statistical Association, 67(338):187–191, 1972.

Chen Xu and Yao Xie. Sequential predictive conformal inference for 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. 38707–38727. PMLR, 23–29 Jul 2023a.

Chen Xu and Yao Xie. Conformal prediction for time series. <u>IEEE Transactions on Pattern Analysis</u> and Machine Intelligence, 45(10):11575–11587, 2023b. doi: 10.1109/TPAMI.2023.3272339.

Chen Xu, Hanyang Jiang, and Yao Xie. Conformal prediction for multi-dimensional time series by ellipsoidal sets. In Forty-first International Conference on Machine Learning, 2024. URL https://openreview.net/forum?id=uN39Tt9P8b.

Izzet B Yildiz, Herbert Jaeger, and Stefan J Kiebel. Re-visiting the echo state property. <u>Neural</u> networks, 35:1–9, 2012.

Shuyi Zhang, Bin Guo, Anlan Dong, Jing He, Ziping Xu, and Song Xi Chen. Cautionary tales on air-quality improvement in beijing. <a href="Proc. Math. Phys. Eng. Sci.">Proc. Math. Phys. Eng. Sci.</a>, 473(2205):20170457, September 2017

# A PROOFS FOR THE ASYMPTOTIC CONDITIONAL COVERAGE

We start the proof with the following lemma.

**Lemma A.1** (Regularity of the conditional CDF). *Under Assumptions 3.2–3.3*, the conditional CDF  $F(r \mid h_t)$  is continuous in  $h_t \in \mathbb{R}^{d_h}$ , uniformly in  $r \in \mathbb{R}$ .

*Proof.* By Assumption 3.2, the map  $x_{\leq t} \mapsto h_t$  is Lipschitz with respect to  $d_{\text{fm}}$ . By Assumption 3.3, the map  $x_{\leq t} \mapsto F(r \mid x_{\leq t})$  is continuous. Therefore the composition  $x_{\leq t} \mapsto h_t \mapsto F(r \mid h_t)$  is continuous in  $h_t$ .

We then continue with the proof of Theorem 3.6, which we start by proving point wise convergence.

*Proof of Theorem 3.6.* Define

$$\bar{F}_n(r \mid \boldsymbol{h}_t) \coloneqq \sum_{i=1}^n w_i^{(n)}(\boldsymbol{h}_t) F(r \mid \boldsymbol{h}_i).$$

Decompose

$$\widehat{F}_n(r \mid \boldsymbol{h}_t) - F(r \mid \boldsymbol{h}_t) = \underbrace{(\widehat{F}_n(r \mid \boldsymbol{h}_t) - \overline{F}_n(r \mid \boldsymbol{h}_t))}_{\text{(I)}} + \underbrace{(\overline{F}_n(r \mid \boldsymbol{h}_t) - F(r \mid \boldsymbol{h}_t))}_{\text{(II)}}.$$

We handle the two terms separately and show both going uniformly to 0 in probability.

# **Term (I):**

$$\widehat{F}_n(r \mid \mathbf{h}_t) - \bar{F}_n(r \mid \mathbf{h}_t) = \sum_{i=1}^n w_i^{(n)}(\mathbf{h}_t) \left( \mathbb{1}\{r_{i+H} \le r\} - F(r \mid \mathbf{h}_i) \right) = \sum_{i=1}^n w_i^{(n)}(\mathbf{h}_t) Y_i(r) = S_n(r)$$

with 
$$Y_i(r) = \mathbb{1}\{r_{i+H} \le r\} - F(r \mid h_i)$$
.

It's easy to see that since  $\mathbb{E}[Y_i(r) \mid h_i] = 0$  as

$$\mathbb{E}[\mathbb{1}\{r_i < r\} \mid \boldsymbol{h}_i] = 1 \cdot \mathbb{P}(r_i < r \mid \boldsymbol{h}_i) + 0 \cdot \mathbb{P}(r_i > r \mid \boldsymbol{h}_i) = \mathbb{P}(r_i < r \mid \boldsymbol{h}_i) =: F(r \mid \boldsymbol{h}_i),$$

then  $\mathbb{E}\left[S_n(r)\right]=0$ . Moreover,  $|Y_i(r)|\leq 1$ . To prove that  $S_n(r)\xrightarrow[n\to\infty]{\mathbb{P}}0$ , we leverage Chebyshev's inequality

$$\mathbb{P}(|S_n(r)| > \epsilon) \le \frac{\operatorname{Var}(S_n(r))}{\epsilon^2}.$$

and then expand the variance as

$$\operatorname{Var}(S_n(r)) = \operatorname{Var}\left(\sum_{i=1}^n w_i^{(n)} Y_i\right) = \sum_{i=1}^n (w_i^{(n)})^2 \operatorname{Var}(Y_i) + \sum_{i \neq j} w_i^{(n)} w_j^{(n)} \operatorname{Cov}(Y_i, Y_j)$$
(14)

$$\leq \sum_{i=1}^{n} (w_i^{(n)})^2 \operatorname{Var}(Y_i) + \sum_{i \neq j} w_i^{(n)} w_j^{(n)} |\operatorname{Cov}(Y_i, Y_j)|.$$
 (15)

Starting by the variances term, since  $|Y_i| \le 1$ , we have  $\mathrm{Var}(Y_i) \le 1$ , thus the first term of Eq. 14 is bounded  $\sum_{i=1}^n (w_i^{(n)})^2 \mathrm{Var}(Y_i) \le \sum_{i=1}^n (w_i^{(n)})^2 \to 0$  thanks to Assumption 3.5 (i). We now move to the covariances. From Lemma 3 of Doukhan (1994), we know that since the process generating  $(x_i, r_{i+H})$  is mixing by Assumption 3.1, for bounded variables  $Y_i, Y_j$  that are function of this process the following holds

$$|Cov(Y_i, Y_j)| \le 4\alpha(|i - j|),$$

where  $\alpha(k)$  is the mixing coefficient. Hence, in our case,

$$\sum_{i \neq j} w_i^{(n)} w_j^{(n)} |\text{Cov}(Y_i, Y_j)| \le 4 \sum_{i \neq j} w_i^{(n)} w_j^{(n)} \alpha(|i - j|).$$
(16)

We can split the sum in two terms

$$4\underbrace{\sum_{1 \le |i-j| \le K} w_i^{(n)} w_j^{(n)} \alpha(|i-j|)}_{(*)} + 4\underbrace{\sum_{|i-j| > K} w_i^{(n)} w_j^{(n)} \alpha(|i-j|)}_{(**)}$$
(17)

such that for all lags k > K the mixing coefficient is arbitrary small, i.e., for any  $\epsilon > 0, \alpha(k) < \epsilon$  (this is always possible).

$$(*) \le \left(\max_{1 \le k \le K} \alpha(k)\right) \sum_{1 \le |i-j| \le K} w_i^{(n)} w_j^{(n)} \tag{18}$$

$$\leq A_{max} \sum_{1 \leq |i-j| \leq K} w_i^{(n)} w_j^{(n)}$$
(19)

$$=2A_{max}\sum_{k=1}^{K}\sum_{i=1}^{n-k}w_i^{(n)}w_{i+k}^{(n)}$$
(20)

Using the inequality  $2ab \le a^2 + b^2$ :

$$2A_{max} \sum_{k=1}^{K} \sum_{i=1}^{n-k} w_i^{(n)} w_{i+k}^{(n)} \le A_{max} \sum_{k=1}^{K} \left( \sum_{i=1}^{n-k} (w_i^{(n)})^2 + \sum_{i=1}^{n-k} (w_{i+k}^{(n)})^2 \right)$$
(21)

$$\leq 2A_{max} \sum_{k=1}^{K} \sum_{i=1}^{n} (w_i^{(n)})^2$$
(22)

$$\leq 2A_{max}K\sum_{i=1}^{n}(w_i^{(n)})^2$$
(23)

$$=\frac{2A_{max}K}{m}\to 0. (24)$$

For the second term

$$(**) = \sum_{|i-j| > K} w_i^{(n)} w_j^{(n)} \alpha(|i-j|)$$
 (25)

$$\leq \epsilon \sum_{|i-j|>K} w_i^{(n)} w_j^{(n)} \tag{26}$$

$$<\epsilon$$
 (27)

where  $\epsilon$  is arbitrarily small.

#### Term (II):

$$\bar{F}_n(r \mid \boldsymbol{h}_t) - F(r \mid \boldsymbol{h}_t) = \sum_{i=1}^n w_i^{(n)}(\boldsymbol{h}_t) \big( F(r \mid \boldsymbol{h}_i) - F(r \mid \boldsymbol{h}_t) \big)$$

Because weights are a convex combination (nonnegative, sum to 1),

$$|\bar{F}_n(r \mid \boldsymbol{h}_t) - F(r \mid \boldsymbol{h}_t)| \leq \sum_{i=1}^n w_i^{(n)}(\boldsymbol{h}_t) |F(r \mid \boldsymbol{h}_i) - F(r \mid \boldsymbol{h}_t)|$$

Now we use Lemma A.1. For any  $\epsilon > 0$  we can choose  $\delta > 0$  so that  $\|\boldsymbol{h}_i - \boldsymbol{h}_t\| < \delta$  implies  $\sup_b \big| F(r \mid \boldsymbol{h}_i) - F(r \mid \boldsymbol{h}_t) \big| < \epsilon$ . We can split the sum into indices with  $\|\boldsymbol{h}_i - \boldsymbol{h}_t\| < \delta$  and  $\|\boldsymbol{h}_i - \boldsymbol{h}_t\| \ge \delta$ :

$$\sum_{i=1}^{n} w_{i}^{(n)}(\boldsymbol{h}_{t}) \big| F(r \mid \boldsymbol{h}_{i}) - F(r \mid \boldsymbol{h}_{t}) \big| = \sum_{i: \|\boldsymbol{h}_{i} - \boldsymbol{h}_{t}\| < \delta} w_{i}^{(n)}(\boldsymbol{h}_{t}) \big| F(r \mid \boldsymbol{h}_{i}) - F(r \mid \boldsymbol{h}_{t}) \big|$$

$$+ \sum_{i: \|\boldsymbol{h}_{i} - \boldsymbol{h}_{t}\| \ge \delta} w_{i}^{(n)}(\boldsymbol{h}_{t}) \big| F(r \mid \boldsymbol{h}_{i}) - F(r \mid \boldsymbol{h}_{t}) \big|$$

The first term is smaller than  $\epsilon$ ,

$$\sum_{i: \|\boldsymbol{h}_i - \boldsymbol{h}_t\| < \delta} w_i^{(n)}(\boldsymbol{h}_t) \big| F(r \mid \boldsymbol{h}_i) - F(r \mid \boldsymbol{h}_t) \big| < \epsilon \sum_{i: \|\boldsymbol{h}_i - \boldsymbol{h}_t\| < \delta} w_i^{(n)}(\boldsymbol{h}_t) < \epsilon$$

which we can make arbitrarily small. The second term goes to 0 in probability under Assumptions 3.5-(i) and 3.5-(ii): as the temperature goes to zero, the weight concentrates on observations  $h_i$  close to  $h_t$ , and the weight of points not in the neighborhood of  $h_t$  vanishes

$$\sum_{i:\|\boldsymbol{h}_i-\boldsymbol{h}_t\|>\delta} w_i^{(n)}(\boldsymbol{h}_t) \big| F(r\mid \boldsymbol{h}_i) - F(r\mid \boldsymbol{h}_t) \big| \leq 2 \sum_{i:\|\boldsymbol{h}_i-\boldsymbol{h}_t\|>\delta} w_i^{(n)}(\boldsymbol{h}_t) \xrightarrow[n\to\infty]{\mathbb{P}} 0.$$

**Uniform convergence:** The proof here is analogous to proofs of the Glivenko–Cantelli theorem. Consider points

$$-\infty = r_0 < r_1 < \dots < r_{K-1} < r_K = \infty$$

such that for all  $\epsilon > 0$ 

$$F(r_k \mid \boldsymbol{h}_t) - F(r_{k-1} \mid \boldsymbol{h}_t) < \epsilon.$$

This implies

$$\sup_{r \in \mathbb{R}} \left| \widehat{F}_n(r \mid \boldsymbol{h}_t) - F(r \mid \boldsymbol{h}_t) \right| \leq \max_{1 \leq k \leq K} \left| \widehat{F}_n(r_k \mid \boldsymbol{h}_t) - F(r_k \mid \boldsymbol{h}_t) \right| + \epsilon$$

Since we already proven pointwise convergence, we have that

$$\max_{1 \le k \le K} \left| \widehat{F}_n(r_k \mid \boldsymbol{h}_t) - F(r_k \mid \boldsymbol{h}_t) \right| \xrightarrow[n \to \infty]{\mathbb{P}} 0$$

and since  $\epsilon$  can be made arbitrary small

$$\sup_{r \in \mathbb{R}} |\widehat{F}_n(r \mid \boldsymbol{h}_t) - F(r \mid \boldsymbol{h}_t)| \xrightarrow[n \to \infty]{\mathbb{P}} 0.$$

*Proof of Corollary* 3.7.

812 813

810

811

Starting from 
$$\mathbb{P}\left(y_{t+H} \in \hat{C}^{\alpha}_t(\hat{y}_{t+H}) \mid \boldsymbol{h}_t\right) = \mathbb{P}\left(r_{t+H} \in \left[\widehat{Q}_{\alpha/2}(\boldsymbol{h}_t), \widehat{Q}_{1-\alpha/2}(\boldsymbol{h}_t)\right] \mid \boldsymbol{h}_t\right)$$

814 815

818 819 820

821 822

823 824

825 826 827

828 829

830

831

832 833

834

835

836

837

838 839

840

841

842

843

844 845

846

847

848 849

850

851

852 853 854

855

856

857

858

859

860

861

862

863

816 so we have 817

$$\mathbb{P}\left(r_{t+H} \in \left[\widehat{Q}_{\alpha/2}(\boldsymbol{h}_t), \widehat{Q}_{1-\alpha/2}(\boldsymbol{h}_t)\right] \mid \boldsymbol{h}_t\right) = F\left(\widehat{Q}_{1-\alpha/2}(\boldsymbol{h}_t) \mid \boldsymbol{h}_t\right) - F\left(\widehat{Q}_{\alpha/2}(\boldsymbol{h}_t) \mid \boldsymbol{h}_t\right)$$

and since Theorem 3.6 gives us the consistency of the quantile estimator to the true conditional quantile, we have that

$$F\left(\widehat{Q}_{1-\alpha/2}(\boldsymbol{h}_t)\big|\ \boldsymbol{h}_t\right) - F\left(\widehat{Q}_{\alpha/2}(\boldsymbol{h}_t)\big|\ \boldsymbol{h}_t\right) \xrightarrow[n \to \infty]{\mathbb{P}} 1 - \alpha.$$

#### DATASETS AND POINT PREDICTORS

# B.1 DATASETS

In the experiments that we performed, we considered four datasets coming from different real-world scenarios and application domains.

Solar data The solar data comes from the US National Solar Radiation Database (Sengupta et al., 2018). We worked with the dataset of 50 time series generated by Auer et al. (2023), each consisting of 52,609 half-hourly measurements. These were resampled to obtain 26,304 hourly observations. Solar radiation was designated as the target variable, and the dataset includes eight additional environmental features.

Air Quality data The dataset was first introduced in Zhang et al. (2017) and comprises 35,064 measurements from 12 locations in Beijing, China. It includes two air-quality measurements (PM10 and PM2.5); for our experiments, we used PM10 as the target and excluded PM2.5 from the feature set. The dataset also contains ten additional environmental variables. Wind direction was encoded following the method in Auer et al. (2023).

**Exchange rate data** The collection of the daily exchange rates of eight foreign countries, including Australia, Britain, Canada, Switzerland, China, Japan, New Zealand, and Singapore. The data span 1990-2016, yielding 7,588 samples.

ACEA data The data was first introduced in Bianchi et al. (2015). The time series represents electricity load provided by ACEA (Azienda Comunale Energia e Ambiente), the company that manages the electric and hydraulic distribution in the city of Rome, Italy. It was sampled every 10 minutes and contains 137,376 observations.

#### B.2 POINT PREDICTORS

For each dataset, we trained three distinct point-forecasting models that span different modelling paradigms. First, a RNN built with gated recurrent unit (GRU) cells and a single layer with a hidden state size of 32. Second, a decoder-only Transformer configured with hidden size 32, a feedforward size of 64, two attention heads, three stacked layers, and a dropout rate of 0.1. Third, an Autoregressive Integrated Moving Average (ARIMA) model of order (3, 1, 3). The first two models were trained and configured similarly to Cini et al. (2025), i.e. by minimizing the mean absolute error (MAE) loss with the Adam optimizer for 200 epochs and a batch size of 32.

For each dataset, every model was trained independently using 40% of that dataset's samples as the training set. The first 25% of the calibration split is used as a validation set for early stopping.

# C HYPERPARAMETERS AND EXPERIMENTAL SETUP

#### C.1 EVALUATION METRICS

Given a prediction interval  $\widehat{C}^{\alpha}(\hat{y}_t) = \left[\hat{y}_t + \hat{q}_t^{\alpha/2}, \hat{y}_t + \hat{q}_t^{1-\alpha/2}\right]$  computed for a desired confidence level  $\alpha$  and true observation  $y_t$ , we evaluated our method against the baselines using three fundamental metrics in uncertainty quantification.

Coverage gap is defined as

$$\Delta \text{Cov} = 100 \left( \mathbb{1}(y_t \in \widehat{C}^{\alpha}(\hat{y}_t)) - (1 - \alpha) \right).$$

and indicates the difference between the target coverage  $1-\alpha$  and the observed one.

**Prediction interval width** This metric quantifies the width of the prediction intervals, i.e., how uncertain the model is about prediction  $\hat{y}_t$ 

PI-Width = 
$$\hat{q}_t^{1-\alpha/2} - \hat{q}_t^{\alpha/2}$$
.

**Winkler score** This is a combined metric that adds to the PI-Width a penalty if the actual value  $y_t$  falls outside the PI. The penalty is proportional by a factor  $\frac{2}{\alpha}$  to the degree of error of the PI, i.e. how far the actual value is from the closest bound of the interval. The Winkler score is defined as:

$$W = \begin{cases} (\hat{q}_t^{1-\alpha/2} - \hat{q}_t^{\alpha/2}) + \frac{2}{\alpha}(\hat{q}_t^{\alpha/2} - y_t) & \text{if } y_t < \hat{q}_t^{\alpha/2}, \\ (\hat{q}_t^{1-\alpha/2} - \hat{q}_t^{\alpha/2}) & \text{if } \hat{q}_t^{\alpha/2} \leq y_t \leq \hat{q}_t^{1-\alpha/2}, \\ (\hat{q}_t^{1-\alpha/2} - \hat{q}_t^{\alpha/2}) + \frac{2}{\alpha}(y_t - \hat{q}_t^{1-\alpha/2}) & \text{if } y_t > \hat{q}_t^{1-\alpha/2}. \end{cases}$$

# C.2 ESN HYPERPARAMETERS

If properly configured, ESNs can provide rich and stable representations of the temporal dependencies in the data. The main hyperparameters that we considered are: the **reservoir size**, which represents the dimension  $D_h$  of the high-dimensional space where reservoir's states evolve (the "capacity" of the reservoir); the **spectral radius**  $\rho(\boldsymbol{W}_h)$  of the state update weight matrix  $W_h$ , which governs the stability of the internal dynamics of the reservoir; the **leak rate**, which controls the proportion of information from the previous reservoir state that is preserved and added to the new state after the non-linearity; and the **input scaling**, which regulates the magnitude of the input signal before it is projected into the reservoir, and thereby control the degree of non-linearity that will be applied to the state update.

Other hyperparameters that define how the ESN behaves are the connectivity of the reservoirs, which denotes the proportion of non-zero connections in the state update weight matrix and thereby determines the sparsity structure of the reservoir, the amount and magnitude of noise injected during state update, and the activation function for the non-linearity. Sometimes, also the topology of the reservoir can lead to meaningful dynamical properties (a ring-like arrangement of the reservoir's internal connection as opposed to random connections).

# C.3 MODEL SELECTION

We performed hyperparameter tuning for all combinations of point predictors and datasets on a validation set, depending on the model.

**RESCP** For our method, we used 10% of the calibration set for validation. To perform model selection, we run a grid search over the spectral radius for values in the interval [0.5, 1.5], for the leak rate [0.5, 1] (1 meaning no leak), for the input scaling [0.1, 3], for the temperature [0.01, 2] and for the size of the calibration window [100, ``all''], where "all' considers all states in the calibration set. The remaining ESN hyperparameters where kept fixed (reservoir size 512, connectivity 0.2, tanh activation function, no injected noise during state update, random connections). Given the high scalability of the method, performing such a grid search was very fast. The same was done for

Table 3: Best hyperparameter values found with the grid search described in the previous section for each dataset and base model combination.

		Spectral Radius		Leak Rate		Input Scaling		Temperature	Sliding Window
	Model	RESCP	RESCQR	RESCP	RESCQR	RESCP	RESCQR	RESCP	RESCP
Solar	RNN	0.9	0.9	0.75	1.0	0.7	0.1	0.1	3900
	Transformer	0.9	0.9	0.8	1.0	0.4	0.1	0.1	7600
	ARIMA	1.0	0.95	0.8	0.9	0.2	0.1	0.1	1500
Beijing	RNN	1.3	0.99	0.95	1.0	0.25	0.1	0.1	3200
	Transformer	1.0	0.9	0.8	1.0	0.25	0.1	0.1	10500
	ARIMA	1.45	0.99	0.65	1.0	0.75	0.1	0.15	3800
Exch.	RNN	0.95	0.9	0.8	0.9	0.5	0.25	0.1	1000
	Transformer	0.99	0.99	0.8	0.9	0.5	0.1	0.1	1000
	ARIMA	0.95	0.99	0.8	0.8	0.5	0.1	0.1	1000
ACEA	RNN	0.95	0.99	0.8	0.8	0.25	0.25	0.1	1000
	Transformer	0.99	0.99	0.8	0.8	0.25	0.25	0.1	3000
	ARIMA	0.95	0.9	0.8	1.0	0.25	0.1	0.1	1000

RESCQR, but only on the ESN hyperparameters (reservoir's size, spectral radius, leak rate and input scaling). In this setting, to train the readout we used the Adam optimizer (Kingma & Ba, 2015) with an initial learning rate of 0.003, which was reduced by 50% every time the loss did not decrease for 10 epochs in a row. Each epoch consisted of mini-batches of 64 samples per batch.

**NexCP** The search was made for the parameter  $\rho$  used to control the decay rate for the weights. The search was over the values  $\{0.999, 0.99, 0.95, 0.9\}$ .

**SPCI** Since SPCI (Xu & Xie, 2023a) required fitting a quantile random forest at each time step, the computational demand to use the official implementation was limiting. Therefore, we used the implementation provided by Auer et al. (2023), and followed the same protocol for the training. SPCI was run using a fixed window of 100 time steps, which corresponds to the longest window considered in Xu & Xie (2023a).

**HopCPT** We followed the same model selection procedure and hyperparameter search of the original paper (Auer et al., 2023). The model was trained for 3,000 epochs in each experiment, validating every 5 epochs on 50% of the calibration data. AdamW (Loshchilov & Hutter, 2019) with standard parameters ( $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\delta = 0.01$ ) was used to optimize the model. A batch size of 4 time series was used. The model with smallest PI-Width and non-negative  $\Delta$ Cov was selected. If no models achieved non-negative  $\Delta$ Cov, then the one with the highest  $\Delta$ Cov was chosen.

**CP-QRNN** Also for this baseline, we followed the same model selection procedure as in the original paper (Cini et al., 2025). The number of neurons and the number of GRU layers was tuned with a small grid search on 10% of the calibration data.

# C.4 RESULTING HYPERPARAMETERS

The best hyperparameters found with the grid search described in the previous section can be found in Tab. 3.

#### D SENSITIVITY ANALYSIS ON ESN HYPERPARAMETERS AND TEMPERATURE

After performing model selection and evaluating our model with the resulting hyperparameters, we conducted a small sensitivity analysis to determine how each part of the reservoir influenced model performance in RESCP, allowing us to assess the robustness of our chosen hyperparameters and identify the key aspects of the reservoir architecture that contribute most to quantification accuracy. We considered the RNN point forecasting baseline.

Fig. 2 shows how the value of the Winkler score and of the coverage varies as the temperature used in the SOFTMAX in equation Eq. 7 increases. As can be seen, there is a well-defined optimal range

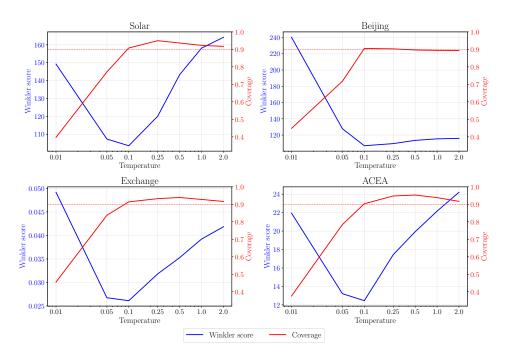


Figure 2: Sensitivity analysis of the temperature.

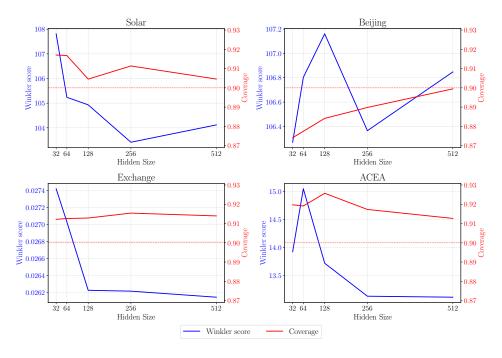


Figure 3: Sensitivity analysis of the reservoir size.

for the temperature in which the Winkler score is minimal and the coverage is close to the targeted one. Notice that for a small temperature, the method is unable to achieve the nominal coverage. This is the *bias-variace tradeoff* that we mentioned in Section 3.1.1, more specifically in Assumption 3.5: the temperature should be small enough to localize only the meaningful residuals close to the query one, yet high enough to avoid collapsing the neighborhood to a single or limited set of similar samples.

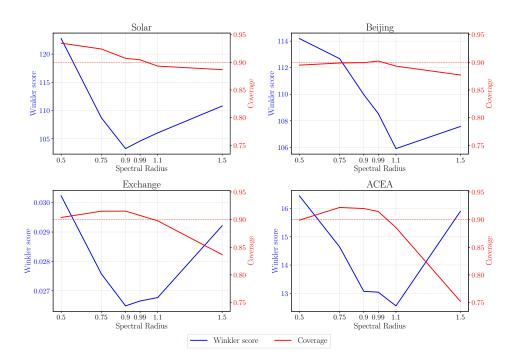


Figure 4: Sensitivity analysis of the spectral radius.

Fig. 3 shows the behaviour of RESCP depending on the capacity of the reservoir. Apart from the Beijing dataset, the method can achieve the nominal coverage also when the dimensionality of the reservoir space is contained, at the price of a higher Winkler score. For the Beijing dataset, a bigger reservoir dimensionality is required to reach the target coverage.

The sensitivity analysis of the spectral radius can be seen in Fig. 4. It is interesting to see that in the Solar and Beijing datasets, the best performance is achieved even with  $\rho(W_h) > 1$ . This is due to the condition being only sufficient, not necessary: while  $\rho(W_h) < 1$  guarantees contractive dynamics in the autonomous linearized reservoir (Gallicchio & Scardapane, 2020), driven reservoirs with input scaling, leaky units, and bounded nonlinearities (Dong et al., 2022; Ceni & Gallicchio, 2025).

In Fig. 5, we observe how varying the leak rate affects the reservoir's performance. In particular, a low value of the leaking rate slows the reservoir's state updates, causing the network to retain past information for longer but reducing its responsiveness to new inputs. We can see how this heavily affects the observed coverage in almost all datasets.

Finally, the effect of different values of input scaling can be found in Fig. 6. High values of input scaling lead the tanh activation function to saturation during state update, yielding loss of sensitivity to input variations. On the other hand, a small input scaling keeps the activation function in the linear part, limiting the richness of nonlinear dynamics necessary to capture complex temporal patterns.

# E USE OF LARGE LANGUAGE MODELS

We acknowledge the use of Large Language Models as a writing tool for minor edits in single sentences.

# F ETHICS STATEMENT

The work presented in this paper is about basic machine learning research and we perform experiments on standard, publicly available datasets. The authors have read and adhere to the ICLR Code of Ethics and do not foresee any direct ethical concerns or potential for misuse.

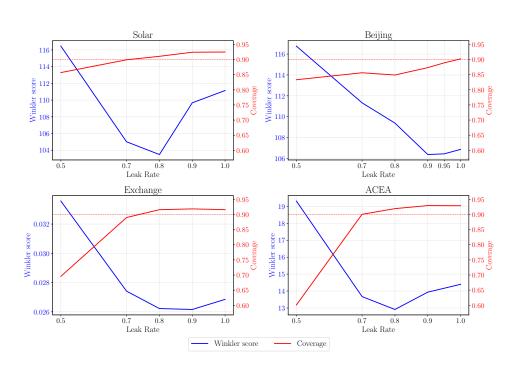


Figure 5: Sensitivity analysis of the leak rate.

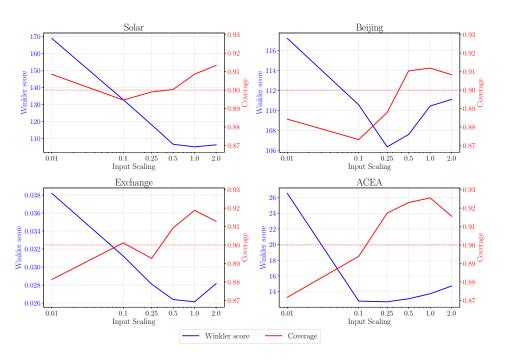


Figure 6: Sensitivity analysis of the input scaling.

# G REPRODUCIBILITY STATEMENT

The code to reproduce the experiments presented in this paper is included in the supplementary material and will be made publicly available upon acceptance.