

Received 4 October 2023, accepted 29 November 2023, date of publication 15 December 2023, date of current version 29 December 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3343467

APPLIED RESEARCH

Probabilistic Load Forecasting With Reservoir Computing

MICHELE GUERRA¹, SIMONE SCARDAPANE², (Member, IEEE),
AND FILIPPO MARIA BIANCHI^{1,3}

¹Department of Mathematics and Statistics, UiT The Arctic University of Norway, 9037 Tromsø, Norway

²Department of Information Engineering, Electronics and Telecommunications, Sapienza Università di Roma, 00185 Rome, Italy

³NORCE Norwegian Research Centre, 5008 Bergen, Norway

Corresponding author: Michele Guerra (michele.guerra@uit.no)

This work was supported by the Publication Fund of the UiT The Arctic University of Norway.

ABSTRACT Some applications of deep learning require not only to provide accurate results but also to quantify the amount of confidence in their prediction. The management of an electric power grid is one of these cases: to avoid risky scenarios, decision-makers need both *precise* and *reliable* forecasts of, for example, power loads. For this reason, point forecasts are not enough hence it is necessary to adopt methods that provide an *uncertainty quantification*. This work focuses on reservoir computing as the core time series forecasting method, due to its computational efficiency and effectiveness in predicting time series. While the reservoir computing literature mostly focused on point forecasting, this work explores the compatibility of some popular uncertainty quantification methods with the reservoir setting. Both Bayesian and deterministic approaches to uncertainty assessment are evaluated and compared in terms of their prediction accuracy, computational resource efficiency and reliability of the estimated uncertainty, based on a set of carefully chosen performance metrics.

INDEX TERMS Time series forecasting, reservoir computing, probabilistic forecasting, Bayesian inference, quantile regression.

I. INTRODUCTION

Models for time-series forecasting are crucial tools in many fields, such as climate modelling, finance, meteorology and energy analytics [39]: an accurate prediction of the energy demand is fundamental to managing a power grid and preventing disruptions in the energy supply, which might cause serious and costly consequences [19]. Especially with a larger penetration of intermittent energy resources [3], [54], load forecasting plays a key role in the efficient planning of electricity supplies [40].

Other than producing a forecast, it is important to quantify the uncertainty in the predictions, since that affects how reliable the forecast is for critical decision-making on power management policies [22]. For example, if the model is fed with an input that lies far from the distribution of training

data, it is desirable that it provides a high uncertainty about its prediction, to alert the user about a possibly unreliable prediction. In this regard, a method that computes point estimates is not enough. On the other hand, probabilistic forecasting methods provide the whole distribution of the forecast or, at least, some confidence intervals [17], [34].

Recurrent Neural Networks (RNNs) are one of the most effective models for processing data with causal dependencies such as time series, but they come with some drawbacks. Those include the difficulty of modelling long-term dependencies when performing gradient descent through time and the long training times due to their sequential structure that prevents parallelisation on modern hardware [7]. Reservoir Computing (RC) overcomes these problems by leaving the recurrent part of the model untrained [42], which allows bypassing the propagation of the gradient through time, greatly reducing the training time. Remarkably, RC models perform as well as deep learning

The associate editor coordinating the review of this manuscript and approving it for publication was Àngel F. García-Fernández.

methods on several tasks [26]. This work uses an Echo State Network (ESN) as a core forecasting model for power demand time series. An ESN is a RC model particularly suited for systems that display complex and chaotic behaviour [49]. As explained in more detail later, the ESN embeds the input in a high dimensional vector state which can then be fitted to the desired output.

So far, most of the existing forecasting approaches based on ESNs have focused on computing point estimates, limiting their use in real-world applications that require estimating confidence intervals for the predictions. To the best of our knowledge, a comprehensive study that thoroughly evaluates uncertainty quantification in Reservoir Computing does not exist. We contribute to fill this gap with the following:

Uncertainty for RC We adapt popular uncertainty quantification methods to the Reservoir Computing framework, filling a gap in the existing literature of ESN.

High-dimensional challenge We address the challenge of handling high-dimensional embeddings generated by the reservoir in a probabilistic programming framework.

Comprehensive evaluation We conduct a comprehensive evaluation by examining the forecasting accuracy, computational efficiency and robustness of uncertainty estimates across a diverse spectrum of metrics. In particular, we carried out our experiments on two real-world datasets, representing the electric load of two power grids.

Practical insights Our findings are positioned to provide practical, nuanced insights for practitioners seeking models that strike an optimal balance between predictive performance, training efficiency and dependable uncertainty assessment.

Empowering decision-making By offering practical solutions and expanding the possibilities of RC applications, our research aims to empower more informed decision-making and facilitate the widespread deployment of robust models.

Our findings suggest that popular Bayesian methods, such as Markov chain Monte Carlo and Variational Inference, suffer considerably from the high dimensionality of the reservoir embeddings, while Quantile Regression proves to be accurate and lightweight.

The paper is organised as follows. Section II presents a summary of the most recent literature on the subject, then a background section introduces Reservoir Computing III-A and the problem of quantifying uncertainties III-B. All the methods we are comparing are presented in Section IV, finally in Sections V and VI we describe the experiments and provide an evaluation of the results.

II. RELATED WORKS

In the last few decades, machine learning methods earned increasing popularity for time series forecasting. In particular, RNNs became one of the most popular options when working with data characterised by time dependencies [12].

Gated RNNs, such as LSTM [30], were introduced to address some limitations of vanilla RNNs [21]. The reader can find comparisons of machine learning methods for time series forecasting in [12] and [39]. RC, introduced in [32], proves to perform better than other machine learning approaches on time series stemming from complex dynamical systems [49].

The deep learning models mentioned so far were originally designed to perform point forecasting, which is the prediction of the most likely value, given a particular choice of the loss function. The simplest approach to model uncertainty is to train an ensemble of RNNs and build confidence intervals from the statistics of the predictions in the ensemble [53]. This approach makes strong assumptions about the underlying data distribution and often generates overconfident intervals if the models in the ensemble are not different enough [36]. A more elegant approach to model the uncertainty is Quantile Regression (QR), which was first introduced in [35] and later further developed in [25]. QR predicts pre-defined quantiles of the output distribution and can be easily implemented in a neural network trained by minimising a particular loss function [55]. Another approach, called conformal prediction, generates prediction intervals from the statistics obtained by a trained model on a calibration set. Originally designed for independent data [2], conformal prediction has been recently extended to handle time series data [51] and paired with RNNs [34], [56].

The interest in probabilistic forecasting with deep learning has recently increased, as evidenced by the presence of open-source libraries such as GluonTS [1] and PyTorch Forecasting¹ that provide several models out of the box. One of the deep learning models for probabilistic forecasting that gained popularity in recent years is DeepAR [47]. DeepAR uses an LSTM that produces an embedding of the input time series and is paired with a probabilistic model chosen by the user, which forecasts future values in the form of Monte Carlo samples. A similar approach is followed by DeepTCN [15], which uses temporal convolutional networks and outputs the parameters of a chosen probability distribution from which one can sample the forecast. In [45], the authors use a Bayesian model trained with Markov chain Monte Carlo (MCMC) and a sparsity-inducing peak-and-slab prior over regression parameters. A recent overview of probabilistic forecasting based on deep learning models can be found in the related works section of [20].

Despite the prolific research on probabilistic forecasting with deep learning, only a few efforts have been devoted to implementing probabilistic forecasting with RC approaches. ESGP [14] combines ESN with Gaussian processes, in [38] Bayesian regression with Gaussian and Laplace priors is used to train the ESN, while [13] employs stochastic Variational Inference (VI). Another work that goes in this direction is [44], where the authors introduced a Bayesian deep ESN [33] and used MCMC to account for the prediction's uncertainty. To try to make the model lighter, they also

¹Source: <https://pytorch-forecasting.readthedocs.io/en/stable/>

make use of stochastic search variable selection priors. They also compare their proposed framework against ensemble forecast methods. These four works consider a single Bayesian approach to model the uncertainty in the prediction of an ESN. However, a systematic analysis and comparison between popular Bayesian and non-Bayesian uncertainty quantification techniques within an RC framework are still missing. The aim of the present work is to fill this gap.

III. BACKGROUND

A. RESERVOIR COMPUTING

The most popular RC model in machine learning is the ESN; the terms RC and ESN are often used interchangeably [32]. An RC model is composed of three elements: the input layer and the recurrent layer called *reservoir*, which are randomly initialised and remain untrained, and a linear output layer, which is the only part that is trained, usually by linear regression. Denoting with $x_t \in \mathbb{R}^K$, $y_t \in \mathbb{R}^L$ and $s_t \in \mathbb{R}^N$ respectively the input, output and state of the reservoir at time t , the ESN dynamics is governed by the following equations:

$$s_{t+1} = f(W_{\text{in}}x_{t+1} + Ws_t) \quad (1)$$

$$y_t = g(Rs_t), \quad (2)$$

where $W_{\text{in}} \in \mathbb{R}^{N \times K}$ and $W \in \mathbb{R}^{N \times N}$ are fixed and $R \in \mathbb{R}^{L \times N}$ is the only trainable part; the two functions f and g introduce non-linearities. In general, an ESN could also include feedback from the output in the equation that updates the reservoir state, but we are not considering it.

In order for the ESN to produce meaningful embeddings, the random matrix W must be initialised ensuring that the echo state property is satisfied [24], [57]. Loosely speaking, this requires that the state of the reservoir s_t asymptotically should not depend on its initial state, but only on the input sequence.

A widely used rule of thumb to initialise the reservoir is to set its spectral radius to less than one. However, being the ESN a non-autonomous system, it has been proved that this criterion is neither necessary nor sufficient to ensure the echo state property and it could also be far from the optimal choice [57]. For this reason, more advanced methods have been proposed to encourage the reservoir to generate rich dynamics of the input [5], [6]. Nevertheless, it has been proven that limiting the spectral radius to be less than one is “sufficient in practice” since a reservoir built in this way satisfies the echo state property with high probability [58]. Therefore, even if potentially sub-optimal, in this work we use the aforementioned criterion for the sake of simplicity and practicality.

Finally, while the vanilla ESN implements $g(s_t; R)$ as a linear readout, more recent works showed that a non-linear readout implemented by a Multi-Layer Perceptron (MLP) usually achieves better performances [8]. In order to train g to accurately map each state s_t into an output $y_t = g(s_t; R)$, where R represents the model’s trainable parameters,

a training dataset must be constructed. The input time-series $\{x_t\}_{t=1}^T$, with $x_t \in \mathbb{R}$, is used by the reservoir f to update its inner state $S = \{s_t\}_{t=1}^T$; hence, the training dataset is defined as $\{(s_t, y_t)\}_{t=1}^T$, where $y_t = x_{t+h}$, with h the forecast horizon. As it is, g is a deterministic function and deterministic models return only one data point for each input, so they can’t possibly produce an estimate for the uncertainty.

B. BAYESIAN APPROACH

In the present work, we focus on the readout $g(s_t; R)$ and how we can change it so as to provide uncertainties. One approach is to modify the above setting so that the output of the model provides y together with information on its uncertainty. This can be done, for example, with QR, which will be introduced in Section IV-D. Alternatively, one can resort to Bayesian Neural Network (BNN), which naturally embody information about uncertainty by generating entire distributions over the output domain.

In a BNN, R is described by a distribution that represents what are the values of R that most likely generate data. In this probabilistic setting, we are interested in searching for the posterior distribution $p(R|S, Y)$, where S and Y are random variables representing respectively the reservoir state and the output, which we could then use to sample an output y^* for a previously unseen state s^* from

$$p(y^*|s^*, S, Y) = \int_{\Omega} p(y^*|s^*, R)p(R|S, Y) dR, \quad (3)$$

where Ω is the domain of R and $p(y|s, R)$ is the likelihood, and it is given by the model itself.²

Bayes’ theorem gives the explicit form of the posterior as

$$p(R|S, Y) = \frac{p(Y|S, R)p(R)}{p(Y|S)}, \quad (4)$$

where at the numerator we have the likelihood and the prior $p(R)$, which are known and selected by us. A standard option for the likelihood that we use in our work is a Gaussian with variance Σ

$$Y \sim p(y|s, R) = \mathcal{N}(g(s; R), \Sigma). \quad (5)$$

As priors for the parameters R ,³ it is common to use a Gaussian or uniform distribution, according to the standard approach to use semi-informative priors [27], [48]. A more specialised choice for $p(R)$ is the horseshoe distribution, discussed in Section IV-B.

The problem about (4) lies in the normalisation constant at the denominator

$$p(Y|S) = \int_{\Omega} p(Y|S, R)p(R) dR, \quad (6)$$

which is often intractable and cannot be computed analytically. This issue can be overcome by approximating the

²Notice that we use the same symbol, i.e. R , to indicate both the random variable and a possible value that it might assume.

³Given our choice of the likelihood in (5), Σ is also a parameter of the model. For ease of notation, we are not specifying it every time, and we only refer to R .




FIGURE 1. Overview of the used framework. On the left, in green, is the input time series used for training. Each point x_t is fed into the reservoir and used to update its internal state s_t . The series of states is then used to train the readout $g(s_t; R)$ using one of the represented methods. On the right, we have the outcome, with ground truth values in orange, and a prediction represented by a solid blue line, with its uncertainty pictured as a shaded band around it.

posterior with VI (see Section IV-A) or by using numerical routines, like MCMC (see Section IV-B). Another way to create a BNN and avoid computing (6) is by transforming a deterministic model into a probabilistic one using dropout (see Section IV-C). This latter approach sacrifices the fine control over the priors, which naturally allows the embedding of previous knowledge and acts as regularisation when dealing with small datasets [27].

IV. METHODS

What follows is a description of the methods to produce an uncertainty estimate that we compare in this work. We start with the three approaches for training a BNN, then we present QR (Section IV-D), and we conclude by discussing the problem of calibration in Section IV-E. Refer to Figure 1 for a schematic view of the framework.

A. VARIATIONAL INFERENCE

Variational Inference (VI) offers a way to tackle the intractability of (4). The idea is to substitute the true posterior $p(R|S, Y)$ with an approximation $q_\theta(R)$, called *variational distribution*, which depends on some parameters θ . In order to identify among all the possible $q_\theta(R)$ the one that best approximates the posterior, one has to minimise the Kullback-Leibler divergence between the true and approximate posterior

$$\begin{aligned} \text{KL}(q_\theta(R) || p(R|S, Y)) &:= \\ &= \int_{\Omega} q_\theta(R) \log \frac{q_\theta(R)}{p(R|S, Y)} dR. \end{aligned} \quad (7)$$

One can prove that

$$\begin{aligned} \text{KL}(q_\theta(R) || p(R|S, Y)) &\leq \\ &\leq \text{KL}(q_\theta(R) || p(R)) - \mathcal{L}_{\text{LL}}, \end{aligned} \quad (8)$$

where $\mathcal{L}_{\text{LL}} := \int q_\theta(R) \log p(Y|S, R) dR$ is the expected log-likelihood. Hence an equivalent approach to minimising the KL on the left is to minimise the right-hand side (which does not depend anymore on $p(R|S, Y)$), i.e. maximising the Evidence Lower Bound (ELBO), defined as

$$\mathcal{L}_{\text{VI}}(\theta) := \mathcal{L}_{\text{LL}} - \text{KL}(q_\theta(R) || p(R)), \quad (9)$$

where, by maximising the first term, we are pushing $q_\theta(R)$ to better approximate the data distribution, while the second term keeps it close to the prior $p(R)$. VI, in other words, translates the task of solving (4) into the optimisation problem (8), which requires computing derivatives rather than integrals.

As variational distribution $q_\theta(R)$, we use a multivariate Gaussian distribution on the whole parameter space, constructed using the Cholesky decomposition, where we don't assume any independence between parameters.

In our case, the parameters R are mainly stemming from the MLP $g(s; R)$ used as a readout of the ESN, so their number N increases with the number of units within each layer and with the number of layers. For example, if we consider a linear layer 500×100 , the dimension of R for that layer alone will be $N = 50\,000$. A multivariate Gaussian as variational distribution would then be defined by a covariance matrix of size N^2 , which reserves a lot of computing memory, leading easily to its saturation.

To avoid this issue we adopted a sparsity-inducing distribution, like the *low-rank* multi-variate Gaussian, as variational distribution $q_\theta(R)$. If we assume that $R \in \mathbb{R}^N$ depends on a smaller number of latent parameters $\phi \in \mathbb{R}^r$, we have $R = C\phi + \epsilon$, where C is a $N \times r$ matrix with $r < N$, ϕ is a r -dimensional random variable from a Gaussian distribution with 0 mean and covariance \mathbb{I}_r , and ϵ is a N -dimensional vector of independently distributed error terms with zero mean and finite variance $\text{var}(\epsilon) = \Psi$. Then one can prove that

$$\text{cov}(R) = CC^T + \Psi, \quad (10)$$

so the covariance matrix can be expressed as the sum of a low-rank matrix CC^T of rank r and a diagonal matrix Ψ . A common choice is $r = \sqrt{N}$. Notice that when $C = 0$ the covariance matrix becomes diagonal, so the mean-field approximation, where all parameters R are assumed independent, is a special case.

B. MARKOV CHAIN MONTE CARLO

Another way to tackle the posterior in (4) is to use Markov chain Monte Carlo (MCMC), which allows sampling from the posterior without the need to compute the integral in (6).

Indeed, under certain assumptions, it is possible to define a stochastic process $\{R_0, R_1, \dots\}$, where R_i is a discrete random variable describing the model's parameters R , that converges to the true posterior distribution $p(R|S, Y)$, by only knowing that $p(R|S, Y) \propto p(Y|S, R)p(R)$. MCMC provides different routines to simulate such a stochastic process and allows the collection of samples from the posterior without knowing explicitly its normalisation. In particular, in this work, we use the No-U-Turn Sampler (NUTS) [31], which is a variant of the Hamiltonian Monte Carlo (HMC) [18].

One of the most important issues in MCMC is understanding whether the chain has converged or not to the posterior distribution [43]. There are a few methods to evaluate convergence, but none of them is perfect, since without knowing the target distribution in advance it's impossible to guarantee that convergence has been reached. Even assuming that convergence is reached, it could take a lot of iterations depending on a number of factors, among which are the dimension of the parameters' space and the complexity of the target distribution. A large number of parameters negatively affects the use of MCMC also from a mere memory perspective, since to infer properties of the target distribution it is necessary to collect a relatively large number of samples. Given the high dimensionality of the reservoir space, a straightforward application of MCMC is difficult.

We explored two possible solutions to overcome this limitation. The first one is to apply Principal Component Analysis (PCA) to reduce the dimension of the reservoir state space before running MCMC. Besides reducing the computational burden in training the readout, this procedure also offers a regularisation that can improve the overall performance in the downstream task [41].

The second solution is to implement Stochastic Search Variable Selection (SSVS) [28], which was introduced to speed up convergence, following the intuition that more promising parts of the parameter space should be more likely and so MCMC would converge there faster. In a Bayesian regression problem, we consider a likelihood distribution of the form $p(Y|S, \beta, \sigma) = \mathcal{N}(S^T \beta, \sigma^2)$,⁴ with Y real-valued, and S and $\beta \in \mathbb{R}^p$ -valued random variables, and we are interested in the posterior distribution of σ and of the regression parameters β that best fit the inputs S to Y . The idea of SSVS is that some of the β_i 's might be negligible. For this reason, the prior over all parameters is $p(\beta, \sigma) = p(\beta|\sigma)p(\sigma)$, where we chose $p(\sigma)$ to be a uniform distribution and $p(\beta|\sigma)$ is a *shrinking* distribution [52], inducing sparsity over β . As the shrinking distribution we chose the *horseshoe* distribution [111], which introduces a local shrinkage parameter λ_i and a global one τ , distributed like follows

$$\beta_i \sim \mathcal{N}(0, \lambda_i^2 \tau^2) \quad (11)$$

$$\lambda_i, \tau \sim \mathcal{C}^+(0, 1), \quad (12)$$

where \mathcal{C}^+ is a half-Cauchy distribution.

⁴In this case we consider a linear readout $g(s; \beta) = s^T \beta$.

C. DROPOUT

The dropout technique [50] can be used to easily morph the readout g , which in our case is an MLP, into a Bayesian model. Consider an MLP model with just one hidden layer and without the bias term, for simplicity. Let us denote Q the dimension of the input, K the number of hidden units and D the output dimension. Dropout is applied by sampling two vectors b_1 and b_2 of dimensions Q and K from a Bernoulli distribution of parameter p and by changing the action of the MLP as follows:

$$y = \sigma(((s \circ b_1)W_1 \circ b_2)W_2), \quad (13)$$

where \circ stands for the Hadamard product and σ is the activation function.

Dropout is commonly used during training as a regularisation technique [50], but, when used during inference, it makes the MLP in (13) a fully Bayesian model [23] that can be applied iteratively to build an ensemble of values $\{y_i\}$, allowing to estimate uncertainties without the need to solve (6).

D. QUANTILE REGRESSION

Another approach, other than using BNNs, is to change the training of g so that for each reservoir state s_t it outputs both the forecast y_t and a measure q of its uncertainty. In QR [35] such q can be one or more quantiles of choice. Given a probability distribution $P(Y|S)$, the quantile of level τ is that value $q_\tau(Y|S)$ such that $\Pr(Y \leq q_\tau) = \tau$. By simply redefining the loss function of a regression task, QR infers directly the quantile we are interested in.

In a standard regression problem, we are able to find the sample mean $\mathbf{E}[Y|S] = g(S; R)$ by minimising the sum of squared residuals, so we solve

$$\arg \min_R \sum_{t=1}^T (y_t - g(s_t; R))^2. \quad (14)$$

In QR the optimisation problem is modified as follows

$$\arg \min_{R(\tau)} \sum_{t=1}^T \rho_\tau(y_t - g(s_t; R(\tau))), \quad (15)$$

where ρ_τ is called check loss (or pinball loss) and is defined as

$$\rho_\tau(r) = \tau \max(r, 0) + (1 - \tau) \max(-r, 0), \quad (16)$$

and the regression parameters $R(\tau)$ will depend on the chosen quantile.

By performing QR for multiple values of τ , we are able to infer a more detailed view of the distribution $P(Y|S)$, than just its mean value.

E. CALIBRATION

In all the above scenarios the model does not yield a simple real value, but it has a more structured output containing

information about uncertainty. For this reason, we can't simply evaluate its performance by some accuracy measure alone. Indeed, the mean value or the median of the output can accurately reproduce y_t and yet the output distribution can still be far off from the values spanned in the dataset.

The idea behind calibration is that if the output says that $I \subset \mathbb{R}$ is a 90% confidence interval for \tilde{y}_t , then y_t should fall in I roughly 90% of the time [29], [37]. See Figure 2 for an example. Ideally, we need to check that the condition above holds for every possible confidence interval. In practice, we pick a set of K quantile levels $\{\tau_i\}_{i=1}^K$, that span the interval $[0, 1]$ and check the calibration on each of them. To do so we first need to compute for each time step the quantiles $q_{\tau,t}$, that is those values that satisfy

$$\begin{aligned} \tau &= \int_{-\infty}^{q_{\tau,t}} p(y|s_t) dy = F_t(q_{\tau,t}) \\ \implies q_{\tau,t} &= F_t^{-1}(\tau), \end{aligned} \quad (17)$$

where F_t is the Cumulative density distribution (CDF). Notice that quantiles need to be computed from the sampled values of \tilde{y}_t for VI, MCMC and dropout, while they are provided directly as output for QR.

Ideally, for each τ_i , y_t should be smaller than $q_{\tau_i,t}$ approximately with a frequency of τ_i , that is

$$\tilde{\tau}_i := \frac{1}{T} \sum_{t=1}^T |\{y_t : y_t \leq q_{\tau_i,t}\}| \simeq \tau_i, \quad (18)$$

for $i = 1, \dots, K$, where $|\cdot|$ represents the cardinality of a set. The left-hand side of (18), that we called $\tilde{\tau}$, is the empirical CDF, while on the right-hand side we have the theoretical one.

We can quantify how uncalibrated our model is by computing the calibration error

$$\text{cal} = \sum_{i=1}^K w_i (\tau_i - \tilde{\tau}_i)^2, \quad (19)$$

where w_i are weights that can be used to give less importance to some quantiles; in our case, we used $w_i = 1$.

In order to calibrate a model, we use the method described in [37]: with a separate portion of the dataset, we create a recalibration dataset $\{(\tau_i, \tilde{\tau}_i)\}_{i=1}^K$ and we use it to train a calibration model $\mu : [0, 1] \rightarrow [0, 1]$, which is then used to correct the uncalibrated quantiles. Consider for example Figure 2b: the horizontal axis represents the predicted CDF τ , while the vertical axis is the empirical CDF $\tilde{\tau}$. The purple points are the couples $\{(\tau_i, \tilde{\tau}_i)\}_{i=1}^K$ obtained from the test set before calibration, which mostly lay far from the dashed line where $\tau = \tilde{\tau}$. After applying the calibration model μ , trained on the calibration set so that $\mu(\tau)$ approximates $\tilde{\tau}$, the results are much more calibrated, as shown in red in Figure 2b.

V. EXPERIMENTAL FRAMEWORK

Our experiments aim at comparing in a time series forecasting task the performance of a probabilistic readout implemented

with different uncertainty quantification techniques. The methods we employed are VI, MCMC (with and without the use of PCA to reduce the dimensionality of the reservoir state), SSVS, dropout and QR. Our software implementation⁵ makes use of PyTorch and the probabilistic programming library Pyro [9]. The different methods are compared according to common criteria quantifying accuracy in the prediction, quality of the uncertainty and computational complexity. For each method, we also discuss the simplicity of usage in terms of expertise required by the practitioner.

All experiments were run on NVIDIA RTX A6000 graphics cards.

A. DATASETS

To carry out the experiments, two datasets of univariate time series have been used.

The first time series represents the electric load measured by Azienda Comunale Energia e Ambiente (ACEA), the company managing the electricity distribution in Rome. Data were collected every 10 minutes for 3 years of activity, from 2009 to 2011 [4] with a precision up to the second decimal digit. From the original dataset, we took a point every 6, so to have an hourly resolution, and removed 11 weeks (being careful not to disrupt the pattern of the time series) due to a strong anomaly that would affect disproportionately the statistics of the training set compared to the calibration and test sets. The final dataset consists of 21048 time steps.

The second dataset (Spain) represents the daily electric demand in the Spanish market between 2014 and 2018, with a precision up to the fourth decimal digit.⁶ The dataset has a length of 1825 time steps.

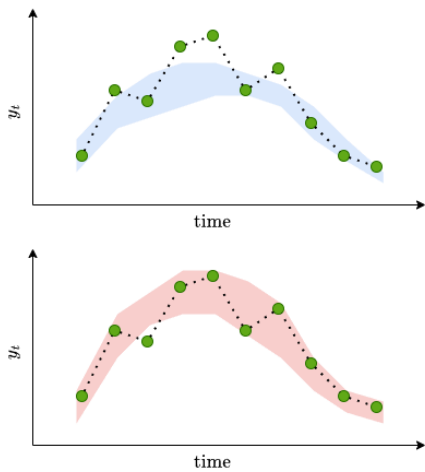
In both cases, the training dataset was formed by the first 70% of the time series, while the test and validation datasets both take half of the remaining 30%. We normalized the datasets by subtracting the mean and by dividing them by the standard deviation of the training set.

When working on time series forecasting, a crucial element to take into consideration is the presence of seasonalities. As expected, both ACEA and Spain datasets have a strong weekly seasonality, i.e. $s = 168$ for ACEA and $s = 7$ for Spain. While seasonalities with longer lags are of course present, they are much less prominent.

The dataset will be of the form $\{(x_t, y_t)\}_{t=1}^T$, with $y_t = x_{t+h}$, where h is the forecast horizon. The input x_t is first processed by the reservoir to produce the states s_t , so the readout g is trained with a dataset $\{(s_t, y_t)\}_{t=1}^T$. The choice of the forecast horizon h is constrained by the particular application at hand, but it is also limited by the removed seasonality s : h must be smaller or equal to s , or, if multiple seasonalities are removed, it must be smaller than the smallest seasonality. Indeed, after removing the seasonality, each data point in the dataset

⁵Code available at <https://github.com/MicheleUIT/Probabilistic-load-forecasting-with-Reservoir-Computing>

⁶Source: <https://www.esios.ree.es/en>



(a) In the upper plot, the 90% confidence interval fails to contain 90% of the data points. Below, in red, is the corresponding calibrated interval.

FIGURE 2. Depiction of the calibration procedure.

will become

$$\begin{aligned} \tilde{x}_t &= x_t - x_{t-s} \\ \tilde{y}_t &= \tilde{x}_{t+h} = x_{t+h} - x_{t-s+h}. \end{aligned} \quad (20)$$

To reconstruct the output is necessary to reintroduce the seasonality by undoing the differencing. In particular, the reconstructed value is $x_{t+h} = \tilde{x}_{t+h} + x_{t-s+h}$. If $h > s$ we can see that $t - s + h > t$, so x_{t-s+h} is part of the forecast. The uncertainty of the forecaster will affect x_{t+h} twice, from both terms. This is an artefact stemming from the seasonal differencing procedure that inflates the uncertainty of the predictions. If, instead, $h \leq s$, then x_{t-s+h} is part of the input, so its value is known with certainty and it will not contribute to the uncertainty of the output. For this reason, we chose $h = 24$ for ACEA and $h = 1$ for Spain, i.e. one day in both cases.

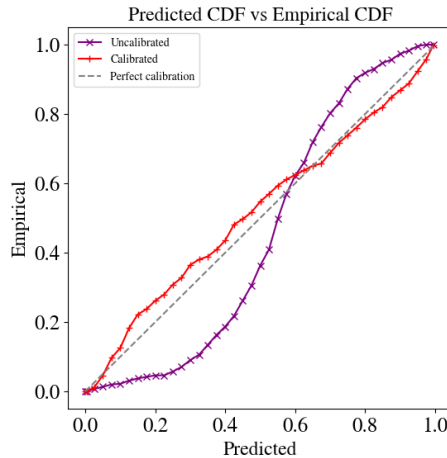
B. METRICS

To perform our analyses, we need metrics that

- measure the capability of each method to produce accurate point forecasts and
- ascertain the quality of the produced uncertainty.

To perform a comparison, we necessitate these metrics to be applicable to each method.

As a measure of forecasting performance, we use the Mean squared error (MSE) between the true value y_t and the median of the forecast $g(s_t; R)$. Evaluating the uncertainty would not be very meaningful if the underlying forecast is not accurate, so the MSE is the main metric considered for hyperparameter selection. Referring to the median, instead of the average, allows the computation of the MSE homogeneously for all methods (in particular, with QR is not possible to compute the average), so it can be used to compare them.



(b) Example of the relationship between predicted τ and empirical CDF $\hat{\tau}$ before calibration (in purple), after calibration (in red) and the ideal case (dashed grey line).

TABLE 1. Schematic view of the possible values and intervals searched for each hyperparameter. The number of layers, units per layer and activation functions refer to the readout's configuration, so they are common to all methods. On the other hand, the prior refers only to VI, MCMC and SSVS; the dropout probability to the dropout method; the learning rate to VI, dropout and QR.

Hyperparameters	Searched values
# layers	1, 2, 3
units per layer	8, 16, 32, 128, 256, 512
activation	tanh, ReLU
prior	$\mathcal{N}(0, 1)$, $\mathcal{N}(0, 10)$, Unif(0, 1), Unif(0, 10), horseshoe
learning rate	[0.0001, 0.1]
p (dropout)	[0.1, 0.9]

Another important gauge of performance is provided by the training time used by the different methods. It must be noted that MCMC strictly speaking does not undergo a training process. The closest notion would be the time of convergence of the chains, but assessing when those converge is a difficult problem and an open research question [16], [46]. For MCMC then we measure the total time elapsed while running the single chain.

To analyse the quality of the uncertainty estimate, we use four different metrics. The first one is the calibration error (cal), see equation (19), which quantifies how much the empirical CDF veers from the theoretical one (ideally, it would be zero). Then, we consider the width and the coverage of the 95% confidence interval, which measure, respectively, how wide that confidence interval is and if it actually covers 95% of the points. Namely, a smaller width reflects a less uncertain forecast, while the coverage should be close to 0.95. We notice that the coverage is redundant since the information is already contained in the calibration error

TABLE 2. This table reports the optimal hyperparameters we used, for each method, to produce the results in Table 3. The optimal values are found by performing a grid search among all the possible combinations reported in Table 1 and by selecting the best-performing one in terms of the accuracy of the forecast (MSE) obtained on the validation set. The colour blue refers to ACEA dataset, while red refers to Spain dataset. Those in black are common to both datasets. A hyperparameter is marked as not applicable (N/A) if it doesn't apply to that particular method.

Parameters	Methods					
	VI	Dropout	MCMC	MCMC (PCA)	SSVS	QR
# layers	1	2	1	3 2	1	3
units per layer	512	512, 256	512	32, 16, 8 32, 16	512	512, 256, 128
activation	ReLU	ReLU	tanh	tanh	tanh ReLU	ReLU tanh
prior $p(R)$	$\mathcal{N}(0, 1)$ Unif(0, 1)	N/A	$\mathcal{N}(0, 1)$ Unif(0, 1)	$\mathcal{N}(0, 1)$ Unif(0, 1)	horseshoe	N/A
prior $p(\Sigma)$	Unif(0, 1)	N/A	Unif(0, 10)	Unif(0, 10)	Unif(0, 10)	N/A
learning rate	~ 0.001 ~ 0.01	~ 0.007 ~ 0.01	N/A	N/A	N/A	~ 0.001 ~ 0.0006
p (dropout)	N/A	~ 0.12 ~ 0.57	N/A	N/A	N/A	N/A

TABLE 3. Results are obtained by averaging 10 runs after selecting the best-performing hyperparameters of each method. The best values are in bold, with the colours encoding the datasets.

Method	Dataset	width (95%)	coverage (95%)	cal	mCRPS	MSE	training time (s)
VI	ACEA	3.11 ± 0.01	0.97 ± 0.00	0.57 ± 0.05	0.32 ± 0.00	0.39 ± 0.01	10.41 ± 0.16
	Spain	3.55 ± 0.09	0.97 ± 0.00	0.53 ± 0.28	0.39 ± 0.02	0.45 ± 0.05	2.78 ± 0.17
Dropout	ACEA	0.32 ± 0.04	0.15 ± 0.03	5.46 ± 0.45	0.20 ± 0.02	0.58 ± 0.11	0.10 ± 0.14
	Spain	0.48 ± 0.03	0.25 ± 0.02	2.43 ± 0.15	0.26 ± 0.02	0.86 ± 0.04	0.08 ± 0.14
MCMC	ACEA	9.64 ± 5.12	0.99 ± 0.00	1.70 ± 0.54	0.67 ± 0.12	0.69 ± 0.40	(3 ± 4) × 10 ³
	Spain	5.95 ± 2.89	0.97 ± 0.02	0.44 ± 0.30	0.58 ± 0.24	1.02 ± 1.07	(3 ± 2) × 10 ³
MCMC (PCA)	ACEA	4.80 ± 1.97	0.98 ± 0.01	0.47 ± 0.06	0.44 ± 0.11	0.62 ± 0.17	(6 ± 6) × 10 ³
	Spain	4.77 ± 1.76	0.91 ± 0.02	0.10 ± 0.10	0.68 ± 0.15	1.34 ± 0.40	(8 ± 5) × 10 ³
SSVS	ACEA	3.08 ± 1.93	0.65 ± 0.21	2.95 ± 0.78	2.03 ± 1.31	1.36 ± 0.43	(17 ± 5) × 10 ³
	Spain	8.13 ± 9.20	0.94 ± 0.03	0.51 ± 0.35	0.75 ± 0.28	0.98 ± 0.24	(13 ± 7) × 10 ³
QR	ACEA	1.11 ± 0.03	0.95 ± 0.00	0.00 ± 0.00	0.16 ± 0.00	0.11 ± 0.00	2.92 ± 0.23
	Spain	1.74 ± 0.05	0.91 ± 0.01	0.04 ± 0.01	0.29 ± 0.01	0.29 ± 0.00	2.69 ± 0.14

on each quantile. Nevertheless, we report also the coverage on the 95% confidence interval as it is a popular performance measure that is adopted in many applications. The last metric is the CRPS and its mean in particular (mCRPS), which measures the average distance between the CDF and the Heaviside function θ ,

$$\text{CRPS}(F, x) = \int (F(y) - \theta(y - x))^2 dy \quad (21)$$

$$\text{mCRPS}(F) = \mathbb{E}_x[\text{CRPS}(F, x)]. \quad (22)$$

The CDF F in (21) is known exactly when working with VI or MCMC, while is only known discretely with QR, through the predicted quantiles. In order to homogeneously compute the CRPS for all methods then, we approximate F with a step function defined using quantiles, which can be easily computed in all cases. The integral can then be approximated

with a sum. Referring to (17), we have

$$F(y) = \tau_i \quad \text{if } y \in (q_{\tau_i}, q_{\tau_{i+1}}), \quad (23)$$

for $i = 1, \dots, K$. CRPS can be interpreted as a combined measure of both the accuracy (MSE) and calibration (cal).

VI. RESULTS

To ensure a fair comparison, the reservoir is generated once, configured with hyperparameters that satisfy the echo state property, and then used for all methods. In addition, for each method, we optimised the hyperparameters using cross-validation based on grid search. The optimised hyperparameters involve both the structure of the MLP used for the readout g and some method-specific features too. Table 1 contains a list of all the possible values for each hyperparameter, while Table 2 reports the optimal

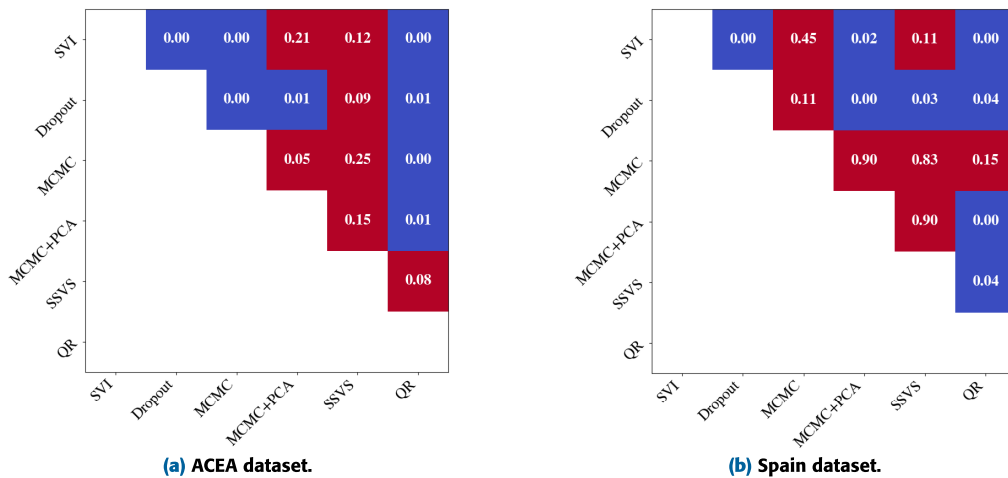


FIGURE 3. Outputs of Games-Howell post hoc test of the mContinuous Rank Probability Score (CRPS) results reported in Table 3: p -values higher than 0.05 are highlighted in red.

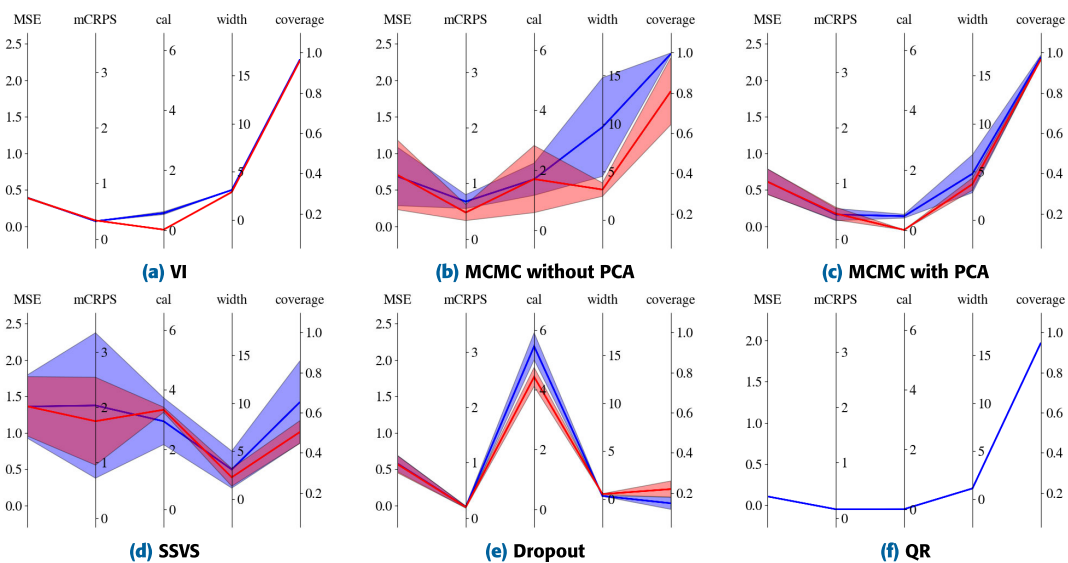


FIGURE 4. Visual representation of the results in Table 3 for ACEA dataset, before (in blue) and after calibration (in red). Notice that calibration is not performed on QR because the predicted CDF is discontinuous.

hyperparameter configuration for each method on the two analysed datasets.

Once the best-performing hyperparameters have been determined, we perform 10 runs for each experiment and collect the resulting means and standard deviations in Table 3: it shows the performance of each method in terms of the computational burden (training time), the accuracy of the forecast (MSE) and the robustness of the uncertainty (cal, width, coverage and mCRPS). In Figures 3 are reported the p -values resulting from Games-Howell post hoc test applied on the mCRPS of Table 3, to show the compatibility between the different methods. Plots in Figures 4 and 5 present, tersely, the same results, with the additional effect of calibration on each metric. Specifically, the course of a blue streak moves through the metrics before calibration, with a solid

line representing the mean and the shaded area the standard deviation, while the red band depicts those same metrics after calibration occurred.

Figures 4 and 5 show that the effect of calibration is different in each method. For example, Figure 4a shows that, for the ACEA dataset, VI takes advantage of the calibration process: the MSE gets a bit smaller even though not in a statistically significant way (from 0.3940.010 to 0.3910.009), but the calibration error drops almost to zero (from 0.5680.049 to 0.0170.004), with better coverage of the 95% confidence interval (from 0.9670.001 to 0.9610.003).

Post-calibration bands for QR are not reported in Figures 4 and 5. Since QR encompasses a discrete number of quantiles ($K = 42$) and the predicted CDF is discontinuous, the calibration model μ could potentially point to quantiles that

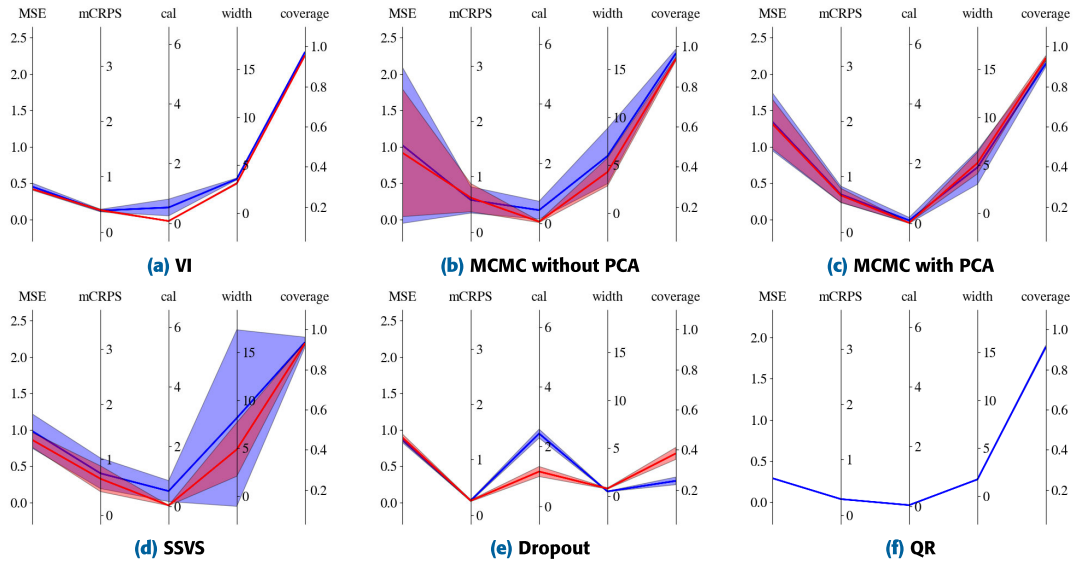


FIGURE 5. Visual representation of the results in Table 3 for Spain dataset, before (in blue) and after calibration (in red). Notice that calibration is not performed on QR because the predicted CDF is discontinuous.

were not computed. Indeed, the 95% confidence interval is defined by $(q_{\tau'}, q_{\tau''})$, where $\tau' = 0.025$ and $\tau'' = 0.975$, but, after calibrating, $\mu(\tau')$ and $\mu(\tau'')$ will, in general, be different, so the corresponding confidence interval will not refer to 95% anymore. All the other methods provide a distribution as output, so it is possible to compute new quantiles $q_{\tilde{\tau}'}$ and $q_{\tilde{\tau}''}$ corresponding to the calibrated confidence interval. QR instead outputs the quantiles directly, so if the calibrated quantiles were not computed beforehand, they can not be obtained afterwards (it would be necessary to re-train the model anew, but it will be uncalibrated again).

Figure 6 shows some examples of forecasts produced by each method for both datasets, after calibration (except for QR).

By looking at Table 3 and the plots in Figures 4, 5 and 6, we notice that it is difficult, at first glance, to identify which is the best-performing method, since there are different methods that achieve the best performance, according to some metric, on each dataset. This underlines the complexity and possible ambiguities in evaluating the performance in probabilistic forecasting. Nevertheless, there are important considerations that can be made, which are reported in the following.

- The most pronounced difference regards training times, where methods based on MCMC show times from 2 to 4 orders of magnitude longer than other methods. The difference is further stressed in Figure 7a. Although it is improper to talk about training time with MCMC, it does not change the fact that the MCMC method is extremely expensive and impractical to use in this scenario. Even by compressing the dimensionality of the reservoir states via PCA or by using SSVS, the computing times do not decrease.

- Overall, QR is the method that achieves better performance, especially on ACEA dataset. Even if post hoc calibration cannot be applied for the reasons discussed above, QR still achieves the lowest calibration error.
- All methods based on MCMC behave generally worse according to all metrics. In particular, the MSE in MCMC is higher than in other methods. The standard deviation of the results is also larger than other methods: this may be due to poor convergence of the chains.
- The fastest technique is dropout, which also achieves an MSE comparable with the other approaches. However, the calibration error, out of the box, is definitely higher due to the extremely narrow confidence intervals that yield poor coverage. Even after calibration, the uncertainty estimation does not improve by much as we can see from Figures 4e and 5e). The poor performance in terms of uncertainty quantification is apparent by looking at Figures 6i and 6j: on both datasets, the uncertainty is generally very low, meaning that the randomness induced by dropout is not strong enough. With a higher value of p the predicted uncertainty improves, but the MSE gets worse.
- Figure 7b depicts the relationship between uncertainty (cal) and forecast accuracy (MSE), with datasets shown with different colours. The better-performing methods are QR and VI: QR achieves the best performance on the Spain dataset, while VI is the best-performing method on ACEA and the second-best on Spain.

Overall, the approach that prevails when jointly considering the low training times, the MSE and the calibration error is QR. Its performance in terms of prediction accuracy and uncertainty estimation is comparable to those of VI, but its computational complexity is significantly lower. Another important advantage of QR is the simplicity of

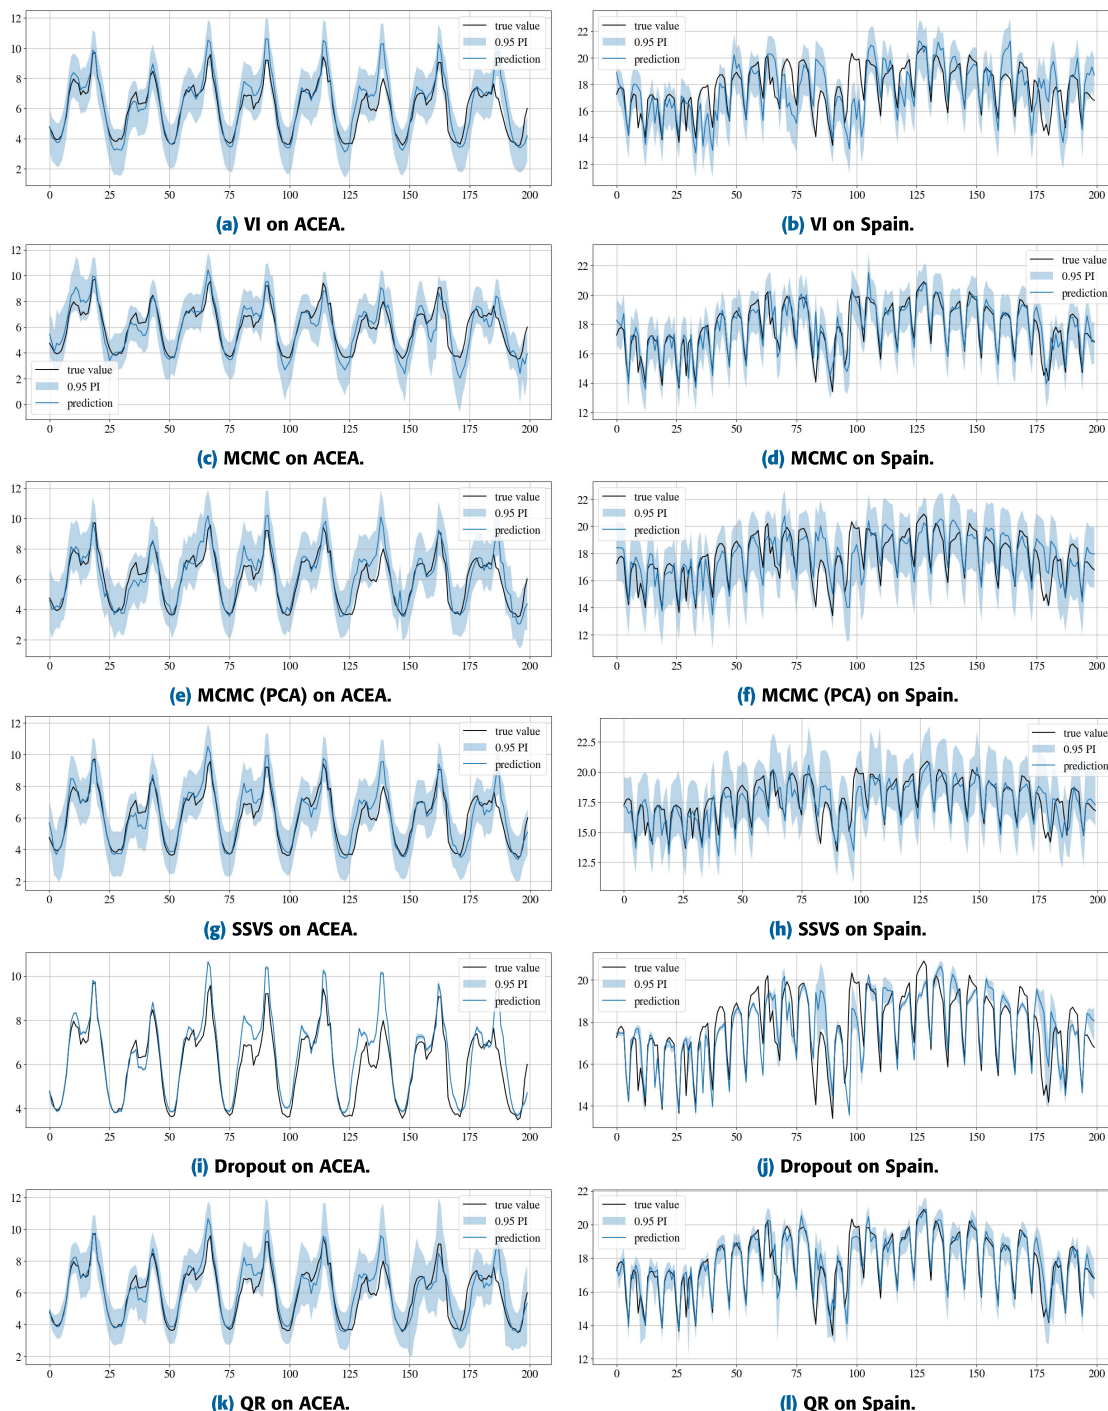


FIGURE 6. Examples of forecast trajectories for both datasets using all the discussed methods, after calibration. The black line represents the true value, the solid blue line is the median of the prediction and the shaded area is the 95% confidence interval.

implementing it: it is sufficient to replace the standard MSE loss used for training a regression model with the pinball loss, making QR a straightforward choice to be used in machine learning applications. On the other hand, Bayesian approaches such as MCMC and VI rely on additional software libraries for probabilistic programming that adds an additional layer of complexity in terms of implementation.

Additionally, while Bayesian inference offers more flexibility thanks to the possibility of specifying the data distribution and introducing prior information in the model, it also poses additional challenges. Indeed, the user is given the additional task of identifying the optimal distributions and the most meaningful priors, which is often not straightforward in many real-world applications, in addition to practical

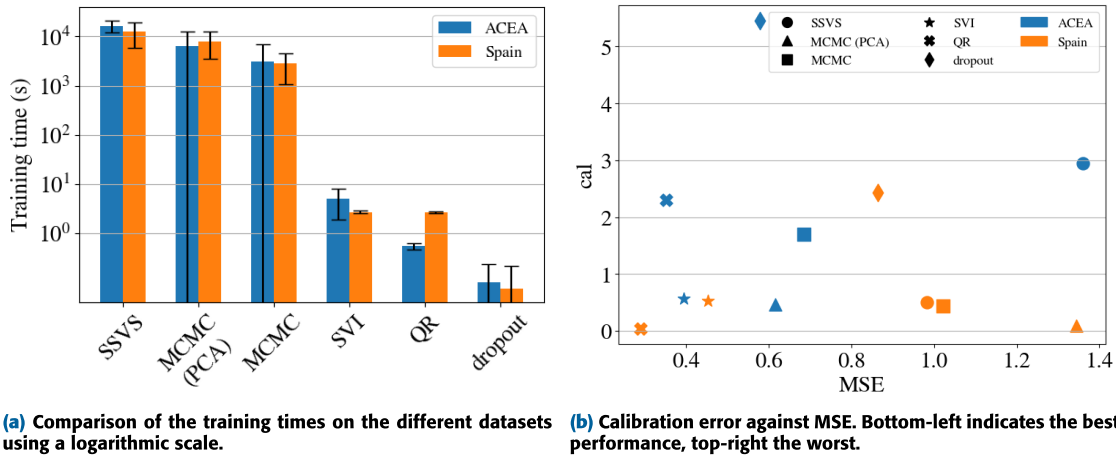


FIGURE 7. Comparison of the most relevant metrics.

TABLE 4. Results for DeepAR are obtained by averaging 10 runs after selecting the best-performing hyperparameters for each dataset. On the other hand, since there is no stochasticity in its fitting procedure, ARIMA is executed only once and we do not report standard deviation in the results.

Method	Dataset	width (95%)	coverage (95%)	cal	mCRPS	MSE	training time (s)
ARIMA	ACEA	3.08	0.94	0.30	0.35	0.77	6.7×10^3
	Spain	2.35	0.84	0.05	0.39	0.95	30.44
DeepAR	ACEA	2.00 ± 0.25	0.93 ± 0.03	0.37 ± 0.53	0.26 ± 0.53	0.21 ± 0.03	$(1.8 \pm 0.5) \times 10^3$
	Spain	2.29 ± 0.56	0.91 ± 0.04	0.22 ± 0.34	0.36 ± 0.08	0.58 ± 0.12	196.90 ± 73.21

challenges such as checking the convergence of the chains in MCMC.

Comparison With Baselines: To understand how the reservoir-based methods perform with respect to other baselines, we compare them against two popular methods for probabilistic time-series forecasting, ARIMA [10] and DeepAR [47]. To determine the ARIMA model we relied on Auto-ARIMA from the pmdarima library,⁷ which identified ARIMA(3, 0, 1) and ARIMA(9, 0, 0) as the best-performing models for ACEA and Spain, respectively. For DeepAR we followed the same hyperparameters’ search procedure as for the reservoir-based models. The best-performing DeepAR model has 2 RNN layers with a hidden size of 32 for ACEA and 8 RNN layers with a hidden size of 16 for Spain. Table 4 presents the same metrics obtained by ARIMA and DeepAR on the two datasets. As expected, the performance of ARIMA according to the different metrics is worse than most reservoir-based methods. The accuracy of DeepAR is comparable with the other methods but with a much longer training time. Finally, we notice how in both ARIMA and DeepAR the calibration error is aligned with those from the other methods, and that training times are much more sensitive than reservoir methods to the size of the dataset.

VII. CONCLUSION

In the present paper, we delved into the task of time series forecasting in an energy analytics setting. As the core model,

we deployed an Echo State Network, which, via its reservoir’s dynamics, is able to unravel the possibly complex inner dependencies of the input time series and embed it in its state space. We argued that a point estimate, even if accurate, is often insufficient for policymakers in high-risk contexts like power market management, so it needs to be accompanied by a measure of how likely the forecast is. For this reason, we combined the Echo State Network with popular methods for uncertainty quantification. Specifically, we considered deterministic methods, like Quantile Regression, and Bayesian approaches, like dropout, Variational Inference and Markov chain Monte Carlo. The goal of our work was first to adapt them to handle the high-dimensional states of the reservoir and then compare them, by identifying their advantages and disadvantages.

Our experiments show that advanced methodologies like VI and MCMC, although they provide fine control over parameters’ priors, do not substantiate a significant advantage in terms of forecasting accuracy or calibration, in fact proving to be demanding in terms of computational resources. Quantile regression, on the other hand, does not seem to be affected by the high dimensionality of the reservoir states, emerging as fast to train, while achieving performances on par with the other methods on other indicators, without even using a post hoc calibration step.

ACKNOWLEDGMENT

The authors gratefully acknowledge NVIDIA Corporation for the donation of two RTX A6000 that were used in this project.

⁷Source: <https://alkaline-ml.com/pmdarima/index.html>

REFERENCES

- [1] A. Alexandrov, K. Benidis, M. Bohlke-Schneider, V. Flunkert, J. Gasthaus, T. Januschowski, D. C. Maddix, S. Rangapuram, D. Salinas, J. Schulz, L. Stella, A. C. Türkmen, and Y. Wang, "GluonTS: Probabilistic and Neural Time Series Modeling in Python," *J. Mach. Learn. Res.*, vol. 21, no. 116, pp. 1–6, 2020.
- [2] A. N. Angelopoulos and S. Bates, "A gentle introduction to conformal prediction and distribution-free uncertainty quantification," 2021, *arXiv:2107.07511*.
- [3] J. Antonanzas, N. Osorio, R. Escobar, R. Urraca, F. J. Martinez-de-Pison, and F. Antonanzas-Torres, "Review of photovoltaic power forecasting," *Sol. Energy*, vol. 136, pp. 78–111, Oct. 2016.
- [4] F. M. Bianchi, E. De Santis, A. Rizzi, and A. Sadeghian, "Short-term electric load forecasting using echo state networks and PCA decomposition," *IEEE Access*, vol. 3, pp. 1931–1943, 2015.
- [5] F. M. Bianchi, L. Livi, and C. Alippi, "Investigating echo-state networks dynamics by means of recurrence analysis," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 29, no. 2, pp. 427–439, Feb. 2018.
- [6] F. M. Bianchi, L. Livi, C. Alippi, and R. Jenssen, "Multiplex visibility graphs to investigate recurrent neural network dynamics," *Sci. Rep.*, vol. 7, no. 1, Mar. 2017, Art. no. 44037.
- [7] F. M. Bianchi, E. Maiorino, M. C. Kampffmeyer, A. Rizzi, and R. Jenssen, *Recurrent Neural Networks for Short-Term Load Forecasting: An Overview and Comparative Analysis*. Cham, Switzerland: Springer, 2017.
- [8] F. M. Bianchi, S. Scardapane, S. Løkse, and R. Jenssen, "Bidirectional deep-readout echo state networks," in *Proc. 26th Eur. Symp. Artif. Neural Netw. (ESANN)*, Bruges, Belgium, Apr. 2018, pp. 425–430.
- [9] E. Bingham, J. P. Chen, M. Jankowiak, F. Obermeyer, N. Pradhan, T. Karaletsos, R. Singh, P. Szerlip, P. Horsfall, and N. D. Goodman, "Pyro: Deep universal probabilistic programming," *J. Mach. Learn. Res.*, vol. 20, pp. 1–6, Jan. 2019.
- [10] G. E. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time Series Analysis: Forecasting and Control*. Hoboken, NJ, USA: Wiley, 2015.
- [11] C. M. Carvalho, N. G. Polson, and J. G. Scott, "Handling sparsity via the horseshoe," in *Proc. Artif. Intell. Statist.*, 2009, pp. 73–80.
- [12] R. Chandra, S. Goyal, and R. Gupta, "Evaluation of deep learning models for multi-step ahead time series prediction," *IEEE Access*, vol. 9, pp. 83105–83123, 2021.
- [13] S. P. Chatzis, "Sparse Bayesian recurrent neural networks," in *Proc. Joint Eur. Conf. Mach. Learn. Knowl. Discovery Databases*, Porto, Portugal. Cham, Switzerland: Springer, Sep. 2015, pp. 359–372.
- [14] S. P. Chatzis and Y. Demiris, "Echo state Gaussian process," *IEEE Trans. Neural Netw.*, vol. 22, no. 9, pp. 1435–1445, Sep. 2011.
- [15] Y. Chen, Y. Kang, Y. Chen, and Z. Wang, "Probabilistic forecasting with temporal convolutional neural network," *Neurocomputing*, vol. 399, pp. 491–501, Jul. 2020.
- [16] M. K. Cowles and B. P. Carlin, "Markov chain Monte Carlo convergence diagnostics: A comparative review," *J. Amer. Stat. Assoc.*, vol. 91, no. 434, pp. 883–904, Jun. 1996.
- [17] H. Du, S. Du, and W. Li, "Probabilistic time series forecasting with deep non-linear state space models," *CAAI Trans. Intell. Technol.*, vol. 8, no. 1, pp. 3–13, Mar. 2023.
- [18] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth, "Hybrid Monte Carlo," *Phys. Lett. B*, vol. 195, pp. 216–222, Sep. 1987.
- [19] O. F. Eikeland, I. S. Holmstrand, S. Bakkejord, M. Chiesa, and F. M. Bianchi, "Detecting and interpreting faults in vulnerable power grids with machine learning," *IEEE Access*, vol. 9, pp. 150686–150699, 2021.
- [20] O. F. Eikeland, F. D. Hovem, T. E. Olsen, M. Chiesa, and F. M. Bianchi, "Probabilistic forecasts of wind power generation in regions with complex topography using deep learning methods: An Arctic case," *Energy Convers. Manag.*, X, vol. 15, Aug. 2022, Art. no. 100239.
- [21] J. L. Elman and D. Zipser, "Learning the hidden structure of speech," *J. Acoust. Soc. Amer.*, vol. 83, no. 4, pp. 1615–1626, Apr. 1988.
- [22] Y. Gal, "Uncertainty in deep learning," Ph.D. thesis, Dept. Eng., Univ. Cambridge, Cambridge, U.K., 2016.
- [23] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Insights and applications," in *Proc. Deep Learn. Workshop (ICML)*, vol. 1, 2015, p. 2.
- [24] C. Gallicchio and A. Micheli, "Architectural and Markovian factors of echo state networks," *Neural Netw.*, vol. 24, no. 5, pp. 440–456, Jun. 2011.
- [25] J. Gasthaus, K. Benidis, Y. Wang, S. S. Rangapuram, D. Salinas, V. Flunkert, and T. Januschowski, "Probabilistic forecasting with spline quantile function RNNs," in *Proc. 22nd Int. Conf. Artif. Intell. Statist.*, 2019, pp. 1901–1910.
- [26] D. J. Gauthier, E. Bollt, A. Griffith, and W. A. S. Barbosa, "Next generation reservoir computing," *Nature Commun.*, vol. 12, no. 1, p. 5564, Sep. 2021.
- [27] A. Gelman, D. Simpson, and M. Betancourt, "The prior can often only be understood in the context of the likelihood," *Entropy*, vol. 19, no. 10, p. 555, Oct. 2017.
- [28] E. I. George and R. E. McCulloch, "Approaches for Bayesian variable selection," *Statistica sinica*, vol. 7, pp. 339–373, Apr. 1997.
- [29] T. Gneiting and M. Katzfuss, "Probabilistic forecasting," *Annu. Rev. Stat. Appl.*, vol. 1, pp. 125–151, Jun. 2014.
- [30] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
- [31] M. D. Homan and A. Gelman, "The no-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1593–1623, Jan. 2014.
- [32] H. Jaeger, "The 'echo state' approach to analysing and training recurrent neural networks—with an erratum note," German Nat. Res. Center Inf. Technol. GMD, Bonn, Germany, Tech. Rep. 148, 2001, p. 13, vol. 148, no. 34.
- [33] H. Jaeger, "Discovering multiscale dynamical features with hierarchical echo state networks," School Eng. Sci., Jacobs Univ., Bremen, Tech. Rep. 10, 2007.
- [34] V. Jensen, F. M. Bianchi, and S. N. Anfinsen, "Ensemble conformalized quantile regression for probabilistic time series forecasting," *IEEE Trans. Neural Netw. Learn. Syst.*, early access, doi: 10.1109/TNNLS.2022.3217694.
- [35] R. Koenker and G. Bassett Jr., "Regression quantiles," *Econometrica*, vol. 46, no. 1, pp. 33–50, Jan. 1978.
- [36] T. Konstantinou and N. Hatzigiorgiouri, "Day-ahead parametric probabilistic forecasting of wind and solar power generation using bounded probability distributions and hybrid neural networks," *IEEE Trans. Sustain. Energy*, vol. 14, no. 4, pp. 2109–2120, Oct. 2023, doi: 10.1109/TSTE.2023.3270968.
- [37] V. Kuleshov, N. Fenner, and S. Ermon, "Accurate uncertainties for deep learning using calibrated regression," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 2796–2804.
- [38] D. Li, M. Han, and J. Wang, "Chaotic time series prediction based on a novel robust echo state network," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 23, no. 5, pp. 787–799, May 2012.
- [39] B. Lim and S. Zohren, "Time-series forecasting with deep learning: A survey," *Phil. Trans. Roy. Soc. A*, vol. 379, no. 2194, 2021, Art. no. 20200209.
- [40] Y. Liu, W. Wang, and N. Ghadimi, "Electricity load forecasting by an improved forecast engine for building level consumers," *Energy*, vol. 139, pp. 18–30, Nov. 2017.
- [41] S. Løkse, F. M. Bianchi, and R. Jenssen, "Training echo state networks with regularization through dimensionality reduction," *Cognit. Comput.*, vol. 9, no. 3, pp. 364–378, Jun. 2017.
- [42] M. Lukoševičius and H. Jaeger, "Reservoir computing approaches to recurrent neural network training," *Comput. Sci. Rev.*, vol. 3, no. 3, pp. 127–149, Aug. 2009.
- [43] S. M. Lynch, *Introduction to Applied Bayesian Statistics and Estimation for Social Scientists*. Cham, Switzerland: Springer, 2007.
- [44] P. L. McDermott and C. K. Wikle, "Deep echo state networks with uncertainty quantification for spatio-temporal forecasting," *Environmetrics*, vol. 30, no. 3, May 2019, Art. no. e2553.
- [45] J. Qiu, S. R. Jammalamadaka, and N. Ning, "Multivariate time series analysis from a Bayesian machine learning perspective," *Ann. Math. Artif. Intell.*, vol. 88, no. 10, pp. 1061–1082, Oct. 2020.
- [46] V. Roy, "Convergence diagnostics (don't short) for Markov chain Monte Carlo," *Annu. Rev. Statist. Appl.*, vol. 7, no. 1, pp. 387–412, Mar. 2020.
- [47] D. Salinas, V. Flunkert, J. Gasthaus, and T. Januschowski, "DeepAR: Probabilistic forecasting with autoregressive recurrent networks," *Int. J. Forecast.*, vol. 36, no. 3, pp. 1181–1191, 2020.
- [48] A. Sarma and M. Kay, "Prior setting in practice: Strategies and rationales used in choosing prior distributions for Bayesian analysis," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, Apr. 2020, pp. p1–12.
- [49] S. Shahi, F. H. Fenton, and E. M. Cherry, "Prediction of chaotic time series using recurrent neural networks and reservoir computing techniques: A comparative study," *Mach. Learn. With Appl.*, vol. 8, Jun. 2022, Art. no. 100300.

- [50] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, pp. 1929–1958, Sep. 2014.
- [51] K. Stankeviciute, A. M. Alaa, and M. van der Schaar, "Conformal time-series forecasting," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, 2021, pp. 6216–6228.
- [52] R. van de Schoot, S. Depaoli, R. King, B. Kramer, K. Märtens, M. G. Tadesse, M. Vannucci, A. Gelman, D. Veen, and J. Willemsen, "Bayesian statistics and modelling," *Nature Rev. Methods Primers*, vol. 1, no. 1, p. 1, 2021.
- [53] J. Wang, S. Wang, B. Zeng, and H. Lu, "A novel ensemble probabilistic forecasting system for uncertainty in wind speed," *Appl. Energy*, vol. 313, May 2022, Art. no. 118796.
- [54] X. Wang, P. Guo, and X. Huang, "A review of wind power forecasting models," *Energy Proc.*, vol. 12, pp. 770–778, Jan. 2011.
- [55] Y. Wang, D. Gan, M. Sun, N. Zhang, Z. Lu, and C. Kang, "Probabilistic individual load forecasting using pinball loss guided LSTM," *Appl. Energy*, vol. 235, pp. 10–20, Feb. 2019.
- [56] C. Xu and Y. Xie, "Conformal prediction interval for dynamic time-series," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 11559–11569.
- [57] I. B. Yildiz, H. Jaeger, and S. J. Kiebel, "Re-visiting the echo state property," *Neural Netw.*, vol. 35, pp. 1–9, Nov. 2012.
- [58] B. Zhang, D. J. Miller, and Y. Wang, "Nonlinear system modeling with random matrices: Echo state networks revisited," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 23, no. 1, pp. 175–182, Jan. 2012.



MICHELE GUERRA is currently pursuing the Ph.D. degree with the Department of Mathematics and Statistics, UiT The Arctic University of Norway, under the supervision of Prof. Filippo Maria Bianchi and Simone Scardapane. His research interests include explainability for graph neural networks, uncertainty estimation, and dynamical systems.



SIMONE SCARDAPANE (Member, IEEE) is currently a tenure-track Assistant Professor with Sapienza Università di Roma. His research interests include graph neural networks, explainability, continual learning, and more recently modular and efficient deep networks. He has published more than 100 papers on these topics in top-tier journals and conferences. He is a member of multiple groups and societies, including the ELLIS Society, the IEEE Task Force on Reservoir Computing, the Machine Learning in Geodesy Joint Study Group of the International Association of Geodesy, and the Statistical Pattern Recognition Techniques TC of the International Association for Pattern Recognition. He is an Associate Editor of IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS (IEEE), *Neural Networks* (Elsevier), *Industrial Artificial Intelligence* (Springer), and *Cognitive Computation* (Springer).



FILIPPO MARIA BIANCHI is currently an Associate Professor with the Department of Mathematics and Statistics, UiT The Arctic University of Norway, and a Senior Researcher with the NORCE Norwegian Research Centre. His research interests include machine learning, dynamical systems, complex networks, energy analytics, and remote sensing. He is the Vice-Chair of the IEEE Task Force on Learning for Structured Data and a member of other groups, such as the IEEE Task Force on Reservoir Computing and the Northernmost Graph Machine Learning Group.

...