

Quantiles for Curve-to-Curve Regression and Probabilistic Forecasting for Daily Electricity Load Curves*

Xiuqin Xu

Institute of Data Science, National University of Singapore, Singapore
xiuqin.xu@u.nus.edu

Ying Chen

Department of Mathematics and Risk Management Institute
National University of Singapore, Singapore, matcheny@nus.edu.sg

Yannig Goude

Électricité de France R&D, Paris, France, yannig.goude@edf.fr

Qiwei Yao

Department of Statistics, London School of Economics and Political Science,
London, United Kingdom, q.yao@lse.ac.uk

August 30, 2020

Abstract

Probabilistic forecasting of electricity load curves is of fundamental importance for effective scheduling and decision making in the increasingly volatile and competitive energy markets. We propose a novel approach to construct probabilistic predictors for curves (PPC), which leads to a natural and new definition of quantiles in the context of curve-to-curve linear regression. There are three types of PPC: a predict set, a predictive band and a predictive quantile, and all of them are defined at a pre-specified nominal probability level. In the simulation study, the PPC achieve promising coverage probabilities under a variety of data generating mechanisms. When applying to one day ahead

*The authors gratefully acknowledge *please remember to list all relevant funding sources in the unblinded version*

forecasting for the French daily electricity load curves, PPC outperform several state-of-the-art predictive methods in terms of forecasting accuracy, coverage rate and average length of the predictive bands. For example, PPC achieve up to 2.8-fold of the coverage rate with much smaller average length of the predictive bands. The predictive quantile curves provide insightful information which is highly relevant to hedging risks in electricity supply management.

Keywords: Dimension reduction via SVD; Electricity load forecasting; Linear curve-to-curve regression; Curve quantile; Probabilistic predictors for curves.

1 Introduction

Electricity load forecasting is an essential element for effective scheduling and decision making in energy markets. The conventional methods of load forecasting are fundamentally deterministic and focus on the mean level of future consumption. However, future is uncertain and costs are more driven by extreme events when, for example, the electricity storage capacities are exhausted. The modern upgrades of power grids with integration of renewable energy and development of efficient electricity management system further introduce uncertainty and fluctuations in the volatile markets. This advocates for probabilistic forecasting methods, aiming at not only accurate point prediction of load, but also insightful predictive intervals and predictive quantiles at pre-determined nominal probability levels.

One-day-ahead high temporal resolution (hourly or half-hourly) load forecasting plays a key role in power system planning and operation, where many operating decisions rely on the load forecasts, such as dispatch scheduling of production transformation, reliability analysis and demand management, see Rolnick et al. (2019). On one hand, the development and integration of renewable production such as photovoltaic panel or wind turbines increases the proportion of electricity production units dependent on meteorological conditions, making the supply of electricity

more volatile and unpredictable, see Gielen et al. (2019). On the other hand, the electricity storage capacities are still expensive and limited, though the smart grid infrastructures (Wang et al., 2018) and smart charging programs (García-Villalobos et al., 2014) allow for more information to transit between end users and producers, increasing the efficiency of demand response (Wang et al., 2015).

While a point prediction is most frequently used in forecasting future electricity loads, the associated risk and the uncertainty are not immediately clear. A predictive interval with a pre-specified coverage probability is more informative, which is arguably the most frequently used probabilistic forecaster. Hong and Fan (2016) surveys the available models/methods for probabilistic forecasting for electricity loads, including evaluation methods and common misunderstanding. See also Hong et al. (2016, 2019). Nevertheless, most available methods and analysis are for forecasting individual loads separately, as simultaneous probabilistic forecasting for multiple values (such as the loads on each 30 minute interval over a day) imposes extra complication and challenges. Simply using individual predictive intervals together, though commonly used in practice, loses the probability interpretation immediately. Adjusting each individual coverage probability requires sophisticated modeling for the intricate dependence among different predictive intervals, which is almost always formidable. The approximation based on the Bonferroni correction is often too crude to be useful. Direct construction of a joint predictive region, with a pre-specified coverage probability, is only possible in some simple cases. Even then, one faces the difficulties in choosing the geometric shape of the region (Polonik and Yao, 2000).

Within the context of forecasting daily loads with high temporal resolution, it is attractive to review the loads over a day as a curve. It embeds nonstationary daily patterns into a stationary framework in a Hilbert space, and, therefore, provides

competitive and reliable pointwise forecasting; see Cho et al. (2013, 2015); Chen and Li (2017); Chen et al. (2020). The goal of this paper is to develop some probabilistic predictors for curves (PPC) with a pre-specified coverage probability. We advocate three types of PPC: a predictive set which consists of a bundle of curves, a predictive band which is a continuous region, and a predictive quantile. The key idea of the new approach is to transfer a curve-to-curve linear regression into several scalar linear regressions (Cho et al., 2013). Then a joint predictive region for the error terms across all scalar linear regression models is constructed via a χ^2 -type statistic, based on which a predictive set for the original targeted curve is induced. The calibration to the nominal coverage probability is achieved either via a residual-based χ^2 -approximation or the empirical distributions of the residuals. For the latter, a bootstrap adjustment should be applied when the sample size is small. We define the envelope of the predictive set as the predictive band, and the most ‘outside’ curves in the predictive set as the predictive quantiles. The ‘outsideness’ is quantified using the concept of the extremal depth for curve data due to Narisetty and Nair (2016).

Our simulation study indicates that PPC achieve accurate coverage rates. When applying to the French electricity loads data with a temporal resolution of 30 minutes, PPC provide the probabilistic forecasts for the daily loads as a whole. The predictive quantiles at different probability levels also deliver insightful information on prospective future scenarios, which is valuable for hedging risks in electricity management. In the context of forecasting each individual loads, it also outperforms several state-of-the-art predictive methods in terms of not only more accurate pointwise forecasts, but also more accurate predictive intervals in the sense of larger coverage probability and shorter length of the intervals.

To our best knowledge, this is the first attempt to construct probabilistic forecasting predictors at a given nominal probability for a curve. The proposed quantile

curves also admit natural and explicit probability interpretation. The literature on forecasting electricity loads is large. We only list a selection of those on predictive interval forecasting for individual values. Taylor and Buizza (2003) constructed interval forecasting based on weather ensemble prediction consisting of 51 weather scenarios. Petiau (2009) proposed interval forecasts based on empirical quantiles of the relative forecasting errors in the past. Kou and Gao (2014) proposed a heteroscedastic Gaussian model for predictive distribution of one day ahead electricity loads. Recently, there are some approaches extending quantile regression based on pin-ball loss optimization of Koenker and Bassett (1978) to produce interval forecasts, such as quantile additive models (Gaillard et al., 2016; Dordonnat et al., 2016; Fasiolo et al., 2020), tree-based ensemble with gradient boosting models (Roach, 2019), and ensemble of experts neural network, quantile random forest and tree-based ensemble (Smyl and Hua, 2019). Apart from constructing predictive intervals, Cabrera and Schulz (2017) derived daily quantile curves based on pointwise quantile estimation and forecasted the future load curves based on functional principal component analysis. Unfortunately, the quantile curves do not inherit the probability interpretation of pointwise quantiles. Therefore, the coverage probabilities of the resulting forecasting bands are unknown. Antoniadis et al. (2016) proposed a nonparametric function-valued model which combines kernel regression and wavelet transformation to produce simultaneous loads predictions at multiple time horizons, where construction of predictive interval for the whole daily curves is considered difficult; see Section 4 of Antoniadis et al. (2016).

The rest of the article is structured as follows. Section 2 introduces the curve linear regression framework to model the dependence structures between curves and proposes the method to construct predictive sets, predictive bands and the predictive quantiles for curves. Several simulation studies are conducted in Section 3 to illus-

trate the finite sample performance of the proposed PPC methodology. In Section 4, PPC are applied to predict day-ahead French electricity load curves in 2019 based on the historical values from 2012 to 2018. Section 5 concludes.

2 Methodology

2.1 Curve regression and dimension reduction

Let $Y_t(u), u \in \mathcal{I}_1$, be the electricity load curve on the t -th day. Associated with each $Y_t(\cdot)$, there is a regressor curve $X_t(v), v \in \mathcal{I}_2$, which may be $Y_{t-1}(\cdot)$, or multiple lagged curves such as $Y_{t-1}(\cdot)$ and $Y_{t-2}(\cdot)$ together, or even contain multiple exogenous variables such as (predicted) temperature curve for the day. We assume that the first two moments of $\{Y_t(\cdot), X_t(\cdot)\}$ are time-invariant. Consider the curve linear regression

$$Y_t(u) = \mu_y(u) + \int_{\mathcal{I}_2} \{X_t(v) - \mu_x(v)\} \beta(u, v) dv + \varepsilon_t(u), \quad u \in \mathcal{I}_1, \quad (1)$$

where $\mu_y(u) = E\{Y_t(u)\}$, $\mu_x(v) = E\{X_t(v)\}$, $\varepsilon_t(\cdot)$ is zero-mean independent noise curve. This is the same setting as Cho et al. (2013, 2015). Perform the singular-value-decomposition (SVD):

$$\Sigma_{yx}(u, v) \equiv \text{Cov}(Y_t(u), X_t(v)) = \sum_{j=1}^{\infty} \sqrt{\lambda_j} \varphi_j(u) \psi_j(v),$$

where $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$ are singular values, then it holds that

$$Y_t(u) - \mu_y(u) = \sum_{j=1}^{\infty} \xi_{tj} \varphi_j(u), \quad X_t(v) - \mu_x(v) = \sum_{j=1}^{\infty} \eta_{tj} \psi_j(v), \quad (2)$$

where

$$\xi_{tj} = \int_{\mathcal{I}_1} \{Y_t(u) - \mu_y(u)\} \varphi_j(u) du, \quad \eta_{tj} = \int_{\mathcal{I}_2} \{X_t(v) - \mu_x(v)\} \psi_j(v) dv.$$

It follows from Theorem 1 of Cho et al. (2013) that the curve regression (1) is effectively equivalent to

$$\xi_{tj} = \sum_{\ell=1}^{\infty} b_{j\ell} \eta_{t\ell} + \varepsilon_{tj}, \quad \varepsilon_{tj} \sim (0, \sigma_j^2), \quad j = 1, 2, \dots, \quad (3)$$

where

$$b_{j\ell} = \int_{\mathcal{I}_1 \times \mathcal{I}_2} \varphi_j(u) \psi_\ell(v) \beta(u, v) du dv, \quad \varepsilon_{tj} = \int_{\mathcal{I}_1} \varepsilon_t(u) \varphi_j(u) du.$$

To simplify the exploration, we assume from now on

$$Y_t(u) - \mu_y(u) = \sum_{j=1}^d \xi_{tj} \varphi_j(u), \quad (4)$$

where $d \geq 1$ is an unknown but finite integer. We will specify how to estimate d in practice below. Furthermore, we assume that (3) admits the finite expression

$$\xi_{tj} = \sum_{\ell \in \pi_j} b_{j\ell} \eta_{t\ell} + \varepsilon_{tj}, \quad \varepsilon_{tj} \sim (0, \sigma_j^2), \quad j = 1, \dots, d, \quad (5)$$

where π_j is a set containing the indices of the finite number of the regressors $\eta_{t\ell}$ for ξ_{tj} . By the virtue of SVD, it holds that $j \in \pi_j$.

With available data $\{(Y_t(\cdot), X_t(\cdot)), 1 \leq t \leq N\}$, put

$$\hat{\mu}_y(u) = \frac{1}{N} \sum_{t=1}^N Y_t(u), \quad \hat{\mu}_x(u) = \frac{1}{N} \sum_{t=1}^N X_t(u), \quad (6)$$

$$\hat{\Sigma}_{yx}(u, v) = \frac{1}{N} \sum_{t=1}^N \{Y_t(u) - \hat{\mu}_y(u)\} \{X_t(v) - \hat{\mu}_x(v)\}.$$

Performing SVD on $\hat{\Sigma}_{yx}(u, v)$, we obtain

$$\hat{\Sigma}_{yx}(u, v) = \sum_{j=1}^{\infty} \hat{\lambda}_j^{\frac{1}{2}} \hat{\varphi}_j(u) \hat{\psi}_j(v), \quad (7)$$

where $\widehat{\lambda}_1 \geq \widehat{\lambda}_2 \geq \dots$ are the singular values of $\widehat{\Sigma}_{yx}(u, v)$. Now replacing $\{\xi_{tj}, \eta_{tj}\}$ in (5) by

$$\widehat{\xi}_{tj} = \int_{\mathcal{I}_1} \{Y_t(u) - \widehat{\mu}_y(u)\} \widehat{\varphi}_j(u) du \quad \text{and} \quad \widehat{\eta}_{tj} = \int_{\mathcal{I}_2} \{X_t(u) - \widehat{\mu}_x(u)\} \widehat{\psi}_j(u) du, \quad (8)$$

we select regressors in model (5), for each fixed j , using stepwise regression controlled by AIC, leading to an estimated index set $\widehat{\pi}_j$. Other methods such as the least squares estimation with ℓ_1 (Lasso) or ℓ_2 (Ridge) penalty can also be used, which produce similar performance in our analysis and are omitted in the manuscript. The fitted model is then of the form

$$\widehat{\xi}_{tj} = \sum_{\ell \in \widehat{\pi}_j} \widehat{b}_{j\ell} \widehat{\eta}_{t\ell} + \widehat{\varepsilon}_{tj}, \quad j = 1, \dots, \widehat{d}, \quad (9)$$

where \widehat{d} is an estimator for d to be specified below, and

$$\widehat{\varepsilon}_{tj} = \widehat{\xi}_{tj} - \sum_{\ell \in \widehat{\pi}_j} \widehat{b}_{j\ell} \widehat{\eta}_{t\ell}, \quad \widehat{\sigma}_j^2 = \frac{1}{N - |\widehat{\pi}_j|} \sum_{t=1}^N (\widehat{\varepsilon}_{tj})^2. \quad (10)$$

In the above expression, $|\widehat{\pi}_j|$ denotes the cardinality of $\widehat{\pi}_j$.

When estimating the predictive curve $E\{Y_t(\cdot)|X_t(\cdot)\}$, Cho et al. (2013) chose d by

$$\widetilde{d}_1 = \arg \min_{1 \leq j \leq d_0} \widehat{\lambda}_{j+1} / \widehat{\lambda}_j,$$

where d_0 is a pre-specified positive integer. Given that $\lambda_1 \geq \dots \geq \lambda_d > 0 = \lambda_{d+1} = \dots$, this selection attempts to ensure that $\xi_{t1}, \dots, \xi_{td}$ catch all the information on $Y_t(\cdot)$ from $X_t(\cdot)$. Hence \widetilde{d}_1 is appropriate for constructing a confidence set for the expectation curve $E\{Y_t(\cdot)|X_t(\cdot)\}$. This is different from our goal of constructing a predictive set for $Y_t(\cdot)$ given $X_t(\cdot)$, for which the noise term $\varepsilon_t(\cdot)$ in (1) also matters. Letting \widetilde{d}_2 be the minimum value of d such that the total variation of the RHS of (4)

over $t = 1, \dots, N$ accounts for more than a certain percentage of the total variation, e.g. 99.9%, of $\{Y_t(\cdot)\}$, we estimate d by

$$\widehat{d} = \max(\widetilde{d}_1, \widetilde{d}_2). \quad (11)$$

2.2 Predictive sets for $Y(\cdot)$ given $X(\cdot)$

Given a new value of $X(\cdot)$, our goal is to predict $Y(\cdot)$ defined by the RHS of (1) with $\{X_t(\cdot), \varepsilon_t(\cdot)\}$ replaced by $\{X(\cdot), \varepsilon(\cdot)\}$, where $\varepsilon(\cdot)$ is unobservable. Then it follows from (2) and (4) that

$$Y(u) - \mu_y(u) = \sum_{j=1}^d \xi_j \varphi_j(u), \quad X(v) - \mu_x(v) = \sum_{j=1}^{\infty} \eta_j \psi_j(v). \quad (12)$$

Denote $\boldsymbol{\varepsilon}_t(d) = (\varepsilon_{t1}, \dots, \varepsilon_{td})'$ and $\boldsymbol{\Sigma} = \text{Var}\{\boldsymbol{\varepsilon}_t(d)\}$. Write $\boldsymbol{\varepsilon}(d) = (\varepsilon_1, \dots, \varepsilon_d)'$. For any $\alpha \in (0, 1)$, define

$$\mathcal{E}_{1-\alpha} = \{\boldsymbol{\varepsilon}(d) \mid \boldsymbol{\varepsilon}(d)' \boldsymbol{\Sigma}^{-1} \boldsymbol{\varepsilon}(d) \leq C_{\alpha,d}\},$$

where $0 < C_{\alpha,d} < \infty$ is a constant determined by

$$P\{\boldsymbol{\varepsilon}_t(d) \in \mathcal{E}_{1-\alpha}\} = 1 - \alpha. \quad (13)$$

Put

$$\mathcal{C}_{1-\alpha}(X) = \left\{ \mu_y(\cdot) + \sum_{j=1}^d \xi_j \varphi_j(\cdot) \mid \left(\xi_1 - \sum_{\ell \in \pi_1} b_{1\ell} \eta_\ell, \dots, \xi_d - \sum_{\ell \in \pi_d} b_{d\ell} \eta_\ell \right)' \in \mathcal{E}_{1-\alpha} \right\}.$$

It follows from (5), (12) and (13) that

$$P\{Y(\cdot) \in \mathcal{C}_{1-\alpha}(X) \mid X(\cdot)\} = 1 - \alpha,$$

i.e. $\mathcal{C}_{1-\alpha}(X)$ is a true predictive set for $Y(\cdot)$ based on $X(\cdot)$ with the nominal coverage probability $1 - \alpha$. In practice, we replace d by \widehat{d} in (11), $\mu_y(\cdot)$ by $\widehat{\mu}_y(\cdot)$ in (6), $\varphi_j(\cdot)$

by $\widehat{\varphi}_j(\cdot)$ in (7), and $(b_{j\ell}, \pi_j)$ by $(\widehat{b}_{j\ell}, \widehat{\pi}_j)$ in (9). A realistic predictive set, i.e. an estimator for $\mathcal{C}_{1-\alpha}(X)$ can be defined as

$$\widehat{\mathcal{C}}_{1-\alpha}(X) = \left\{ \widehat{\mu}_y(\cdot) + \sum_{j=1}^{\widehat{d}} \left(\varepsilon_j + \sum_{\ell \in \widehat{\pi}_j} \widehat{b}_{j\ell} \widehat{\eta}_\ell \right) \widehat{\varphi}_j(\cdot) \mid \boldsymbol{\varepsilon}(\widehat{d})' \widehat{\boldsymbol{\Sigma}}^{-1} \boldsymbol{\varepsilon}(\widehat{d}) \leq C_{\alpha, \widehat{d}} \right\}, \quad (14)$$

where $\{\widehat{\eta}_\ell\}$ are obtained in (8) with $X_t(\cdot)$ replaced by $X(\cdot)$, $\widehat{\boldsymbol{\Sigma}}$ is the sample covariance matrix of $\widehat{\boldsymbol{\varepsilon}}_t \equiv (\widehat{\varepsilon}_{t1}, \dots, \widehat{\varepsilon}_{t\widehat{d}})'$ for $t = 1, \dots, N$, and $\widehat{\varepsilon}_{tj}$ is given in (10). The (i, j) -th element of $\widehat{\boldsymbol{\Sigma}}$ is defined as $\sum_{t=1}^N \varepsilon_{ti} \varepsilon_{tj} / (N - \gamma_{ij})$, where γ_{ij} is the cardinality of $\widehat{\pi}_i \cup \widehat{\pi}_j$. Now, we propose two methods to determine the constant $C_{\alpha, \widehat{d}}$.

- (i) χ^2 -approximation: Assuming that $\varepsilon_{t1}, \dots, \varepsilon_{t\widehat{d}}$ are jointly normal, then $\boldsymbol{\varepsilon}_t(d)' \boldsymbol{\Sigma}^{-1} \boldsymbol{\varepsilon}_t(d) \sim \chi_d^2$. Let $C_{\alpha, \widehat{d}}$ be the $(1 - \alpha)$ -th percentile of the χ^2 distribution with \widehat{d} degrees of freedom. We use in (14)

$$(\varepsilon_1, \dots, \varepsilon_{\widehat{d}})' = \widehat{\boldsymbol{\Sigma}}^{1/2} (z_1, \dots, z_{\widehat{d}})', \quad (15)$$

where $z_1, \dots, z_{\widehat{d}}$ are independent and $\mathcal{N}(0, 1)$, and $\sum_i z_i^2 \leq C_{\alpha, \widehat{d}}$. In principle $\widehat{\mathcal{C}}_{1-\alpha}(X)$ consists of infinite number of curves.

- (ii) Empirical distribution for residuals: An alternative is to let $C_{\alpha, \widehat{d}}$ be the $(1 - \alpha)$ -th percentiles of the empirical distribution of

$$(\widehat{\varepsilon}_{t1}, \dots, \widehat{\varepsilon}_{t\widehat{d}})' \widehat{\boldsymbol{\Sigma}}^{-1} (\widehat{\varepsilon}_{t1}, \dots, \widehat{\varepsilon}_{t\widehat{d}}) \quad t = 1, \dots, N. \quad (16)$$

Then $\widehat{\mathcal{C}}_{1-\alpha}(X)$ consists of the $[N(1-\alpha)]$ curves generated via $\boldsymbol{\varepsilon}_t(\widehat{d}) = (\widehat{\varepsilon}_{t1}, \dots, \widehat{\varepsilon}_{t\widehat{d}})'$ for which the inequality in (14) holds and $1 \leq t \leq N$.

2.3 Predictive bands for $Y(\cdot)$ given $X(\cdot)$

In spite of the clear probability interpretation, the predictive set $\widehat{\mathcal{C}}_{1-\alpha}(X)$ consists of a bundle of curves. In practice, it is more convenient to use a band or a region which

covers the target curve $Y(\cdot)$ with probability $1 - \alpha$. A natural predictive band is the envelope of $\widehat{\mathcal{C}}_{1-\alpha}(X)$:

$$\widehat{\mathcal{B}}_{1-\alpha}(X) = (\widehat{Y}_{\text{low}}(\cdot), \widehat{Y}_{\text{up}}(\cdot)),$$

where

$$\widehat{Y}_{\text{low}}(u) = \min\{Z(u) \mid Z(\cdot) \in \widehat{\mathcal{C}}_{1-\alpha}(X)\}, \quad \widehat{Y}_{\text{up}}(u) = \max\{Z(u) \mid Z(\cdot) \in \widehat{\mathcal{C}}_{1-\alpha}(X)\}. \quad (17)$$

For $\widehat{\mathcal{C}}_{1-\alpha}(X)$ constructed based on the χ^2 distribution, the infinite number of curves in $\widehat{\mathcal{C}}_{1-\alpha}(X)$ fills in every space in $\widehat{\mathcal{B}}_{1-\alpha}(X)$ due to the continuity of the normal distribution; see also (15). It holds that

$$P\{Y(\cdot) \in \widehat{\mathcal{B}}_{1-\alpha}(X) \mid X\} \geq P\{Y(\cdot) \in \widehat{\mathcal{C}}_{1-\alpha}(X) \mid X\} \approx 1 - \alpha. \quad (18)$$

Some adjustment may be applied to make the coverage probability more in line with the nominal level. On the other hand, there is a potential problem in $\widehat{\mathcal{B}}_{1-\alpha}(X)$ when $\widehat{\mathcal{C}}_{1-\alpha}(X)$ is constructed based on the empirical distribution, as then $\widehat{\mathcal{C}}_{1-\alpha}(X)$ consists of merely $[N(1 - \alpha)]$ curves, and the width $\widehat{Y}_{\text{up}}(\cdot) - \widehat{Y}_{\text{low}}(\cdot)$ increases as N increases. For small N , the band $\widehat{\mathcal{B}}_{1-\alpha}(X)$ could be too narrow to cover $Y(\cdot)$ with probability $1 - \alpha$. For forecasting electricity load curves as well as other real world data, the sample sizes are often small in order to retain the ‘stationarity’.

We propose a bootstrap method to adjust the coverage probability of $\widehat{\mathcal{B}}_{1-\alpha}(X)$. For a given positive integer K , let $\varepsilon_{1j}^*, \dots, \varepsilon_{Kj}^*$ be drawn independently (with replacement) from $\widehat{\varepsilon}_{1j}, \dots, \widehat{\varepsilon}_{Nj}$, $j = 1, \dots, \widehat{d}$. Put $\boldsymbol{\varepsilon}_i^* = (\varepsilon_{i1}^*, \dots, \varepsilon_{i\widehat{d}}^*)'$ and

$$\mathcal{C}_{1-\alpha}^*(X) = \left\{ \widehat{\mu}_y(\cdot) + \sum_{j=1}^{\widehat{d}} \left(\varepsilon_{ij}^* + \sum_{\ell \in \widehat{\pi}_j} \widehat{b}_{j\ell} \widehat{\eta}_\ell \right) \widehat{\varphi}_j(\cdot) \mid (\boldsymbol{\varepsilon}_i^*)' \widehat{\boldsymbol{\Sigma}}^{-1} \boldsymbol{\varepsilon}_i^* \leq C_{\alpha, \widehat{d}}, 1 \leq i \leq K \right\}. \quad (19)$$

See also (14). To define $\widehat{\mathcal{B}}_{1-\alpha}(X)$, we replace $\widehat{\mathcal{C}}_{1-\alpha}(X)$ by $\widehat{\mathcal{C}}_{1-\alpha}(X) \cup \mathcal{C}_{1-\alpha}^*(X)$ in (17).

To determine the bootstrap sample size K , we adopt a leave-one-out procedure as follows. For each $1 \leq i \leq N$, we construct a predictive band for $Y_i(\cdot)$ conditionally on $X_i(\cdot)$ in a similar way as above, i.e. we construct $\widehat{\mathcal{C}}_{1-\alpha}(X_i)$ and $\mathcal{C}_{1-\alpha}^*(X_i)$ as in (14) and (19) respectively, replacing $\widehat{\eta}_j$ by $\widehat{\eta}_{ij}$. Leave out the term $(\widehat{\varepsilon}_{i1}, \dots, \widehat{\varepsilon}_{id})$ in (14), and take the bootstrap sample from the other $(N - 1)$ residual vectors. We calculate the relative frequency for the occurrence of the event that the resulting envelop contains $Y_i(\cdot)$ for $i = 1, \dots, N$. Then choose K such that the corresponding relative frequency is closest to the nominal level $1 - \alpha$. In practice, we choose K among a finite set, for example, 0, 200, 400, 600, 800 and 1000. Also note that in the above cross validation approach, we use $\widehat{\mu}_y(\cdot)$, $\widehat{b}_{j\ell}$, $\widehat{\eta}_{ij}$ and $\widehat{\varphi}_j(\cdot)$ estimated from the whole sample of the available observations. The leave-one-out strategy only applies to the residuals.

We can also develop a bootstrap procedure to determine the number of random curves to be included in the χ^2 -based predict set $\widehat{\mathcal{C}}_{1-\alpha}(X)$ such that the resulting $\widehat{\mathcal{B}}_{1-\alpha}(X)$ has the coverage probability $1 - \alpha$; see (18).

2.4 Predictive quantiles for $Y(\cdot)$ given $X(\cdot)$

For forecasting a univariate random variable, it is informative to look at predictive (i.e. conditional) quantiles at different levels to gauge the associated risk and uncertainty. This is effectively equivalent to looking at how a predictive interval varies with respect to its coverage probability. Unfortunately this analogue is no longer available in forecasting a random vector or a random curve, for which the concept of quantiles is not well-defined. Nevertheless it remains attractive to look at some ‘typical’ scenarios among the curves in $\widehat{\mathcal{C}}_{1-\alpha}(X)$ and to observe how they vary with respect to the values of α .

One plausible way forward is to define the most ‘outside’ curve (or curves) in $\widehat{\mathcal{C}}_{1-\alpha}(X)$ as the $(1 - \alpha)$ -th quantile(s). However the concept of ‘outsideness’ needs to be relaxed as it is unlikely that one single curve lies completely on the one side of all the other curves in $\widehat{\mathcal{C}}_{1-\alpha}(X)$, as the curves often cross over with each other. We adopt the so-called extremal depth curves of Narisetty and Nair (2016) to quantify the degree of ‘outsideness’ for each curve. We introduce the concept of the extremal depth first.

Let $\mathcal{S} = \{f_1, \dots, f_p\}$ be a bundle of curves defined on set \mathcal{I} . For any curve g defined on \mathcal{I} and $u \in \mathcal{I}$, a pointwise depth of $g(u)$ with respect to \mathcal{S} is defined as

$$D_g(u, \mathcal{S}) = 1 - \frac{1}{p} \left| \sum_{i=1}^p [I\{f_i(u) < g(u)\} - I\{f_i(u) > g(u)\}] \right|,$$

where $I(\cdot)$ denotes the indicator function. Obviously, D_g takes values $0, 1/p, 2/p, \dots, 1$, and the larger D_g is, the more central $g(u)$ is with respect to $\{f_i(u)\}$. Note that g may or may not be a member of \mathcal{S} . A depth cumulative distribution function (d-CDF) is defined as

$$\Phi_g(r) = \frac{1}{A(\mathcal{I})} \int_{\mathcal{I}} I\{D_g(u, \mathcal{S}) \leq r\} du, \quad r \in [0, 1],$$

where $A(\mathcal{I}) = \int_{\mathcal{I}} du$. Note that if Φ_g has most of its mass close to 0 (or 1), g is away from (or close to) the ‘center’ of \mathcal{S} .

Narisetty and Nair (2016) adopts the left-tail stochastic ordering for d-CDFs to rank the ‘outsideness’ or ‘extremeness’ of two curves, which differs from the other definitions for curve data depth in literature (López-Pintado and Romo, 2009; Fraiman and Muniz, 2001):

Let $0 \leq u_1 \leq \dots \leq u_m \leq 1$ be m pre-specified points. For two curves g and h defined on \mathcal{I} , g is said to be more extremal than h , denoted by

$g \prec h$, if there exists $1 \leq k \leq m$ for which

$$\Phi_g(u_k) > \Phi_h(u_k), \quad \text{and} \quad \Phi_g(u_i) = \Phi_h(u_i) \quad \text{for } 1 \leq i < k.$$

Now the extremal depth (ED) of function g in relation to curve bundle \mathcal{S} is defined as

$$\text{ED}(g, \mathcal{S}) = |\{i : f_i \preceq g, 1 \leq i \leq p\}|/p,$$

where $|\cdot|$ denotes the cardinality of a set, and $f_i \preceq g$ if either $f_i \prec g$ or $f_i(u_j) = g(u_j)$ for all $j = 1, \dots, m$. Note that the smaller $\text{ED}(g, \mathcal{S})$ is, the more outside g is from \mathcal{S} . We refer readers to Narisetty and Nair (2016) for further elaboration of the ED concept.

Now we are ready to define the predictive quantile curves of $Y(\cdot)$ given $X(\cdot)$.

Definition 1. For any $\alpha \in (0, 1)$, a curve $g \in \mathcal{C}_{1-\alpha}(X)$ is called the $(1 - \alpha)$ -th predictive quantile curve of $Y(\cdot)$ given $X(\cdot)$ if

$$\text{ED}\{g, \mathcal{C}_{1-\alpha}(X)\} = \min_{h \in \mathcal{C}_{1-\alpha}(X)} \text{ED}\{h, \mathcal{C}_{1-\alpha}(X)\}. \quad (20)$$

Intuitively, the $(1 - \alpha)$ -th conditional quantile is the most outside curve, on the most possible points in \mathcal{I}_1 , among all the curves in $\mathcal{C}_{1-\alpha}(X)$. An estimator for the quantile curve can be obtained from replacing $\mathcal{C}_{1-\alpha}(X)$ by $\widehat{\mathcal{C}}_{1-\alpha}(X)$ in the above definition. In addition, we recommend also take the 2nd or even the 3rd minimizers of (20) as the $(1 - \alpha)$ -th quantiles, as the most outside can be either above, or below, the other curves in $\mathcal{C}_{1-\alpha}(X)$. Furthermore, when $C_{\alpha, \widehat{d}}$ is estimated by the empirical distribution of (16), we use the union $\widehat{\mathcal{C}}_{1-\alpha}(X) \cup \mathcal{C}_{1-\alpha}^*(X)$ in place of $\widehat{\mathcal{C}}_{1-\alpha}(X)$; see (19).

Remark 1. The definitions for $D_g(u, \mathcal{S})$ and $\text{ED}(g, \mathcal{S})$ presented above are for \mathcal{S} with a finite cardinality. It serves the purpose since in practice we always use $\widehat{\mathcal{C}}_{1-\alpha}(X)$

with finite members. The extension to the cases with $|\mathcal{S}| = \infty$ can be formulated in terms of distribution on \mathcal{S} such as

$$D_g(u, \mathcal{S}) = 1 - |E[I\{f(u) < g(u)\} - I\{f(u) > g(u)\}]|, \quad \text{ED}(g, \mathcal{S}) = P(f \preceq g),$$

where both the expectation and the probability are taken with respect to the distribution of f on \mathcal{S} .

3 Simulation

We illustrate the proposed PPC by simulation. In line with electricity load forecasting, we set $X_t(\cdot) = Y_{t-1}(\cdot)$, and $Y_t(\cdot)$ follows the functional autoregressive (FAR) model

$$Y_t(u) = \sum_{j=1}^d \xi_{tj} \varphi_j(u), \quad \xi_{tj} = \sum_{\ell \in \pi_j} b_{j\ell} \xi_{t-1, \ell} + \varepsilon_{tj}, \quad \varepsilon_{tj} \sim \text{IID}(0, \sigma^2), \quad (21)$$

where $\varphi_1(u) = 1/\sqrt{2}$ and $\varphi_j(u) = \cos((j-1)\pi u)$ for $j \geq 2$, $u \in [-1, 1]$. We sample d characteristic roots for an AR model from the interval $[2, 5] \cup [-5, -2]$ uniformly, based on which the parameters $b_{j\ell}$, $\ell = 1, \dots, d$ are determined for each $j = 1, \dots, d$.

Experiment 1: Let $d = 4$, $\pi_j = \{1, \dots, 4\}$, ε_{tj} be normal distributed, and $\sigma = 0.25$ or 0.50 in (21). We draw $Y_1(\cdot), \dots, Y_{N+N_s}(\cdot)$ from the model, and use the first N curves for estimation and the last N_s curve to check the forecasting accuracy. We set the nominal coverage probability at $1 - \alpha = 0.9$. For $N = 100, 200, 400, 800$ or $1, 600$, and $N_s = 200$, we calculate the predictive set $\widehat{\mathcal{C}}_{0.9}(\cdot)$ and the predictive band $\widehat{\mathcal{B}}_{0.9}(\cdot)$ based on three methods: χ^2 -approximation, empirical distribution for residuals (ECDF), and ECDF with bootstrap (ECDF-B). With the χ^2 -approximation, we include 1,500 curves in $\widehat{\mathcal{C}}_{0.9}(\cdot)$. For ECDF-B, we choose the bootstrap sample size as a multiple of 200 between 0 and 1,000. We check the coverage rates of $\widehat{\mathcal{B}}_{0.9}(\cdot)$ for the

200 post-sample curves and draw illustrative predictive quantile curves at the 40% and 90% confidence levels. For more comprehensive evaluation, we also calculate the pointwise mean absolute errors (MAE) on a grid $u_i = -1 + 0.04(i-1)$ ($i = 1, \dots, 51$) over the 200 post-sample curves $\{Y_k(\cdot)\}$:

$$\mathbf{MAE} = \frac{1}{51 \times 200} \sum_{k=1}^{200} \sum_{i=1}^{51} |\widehat{Y}_k(u_i) - Y_k(u_i)|. \quad (22)$$

For each setting, we replicate the above exercise 400 times.

Figure 1 displays the box plots of the coverage rates of $\widehat{\mathcal{B}}_{0.9}(\cdot)$ based on, χ^2 -approximation, ECDF and ECDF-B respectively, for the 200 post-sample curves over 400 replications. Since ε_{tj} are normal, χ^2 -based $\widehat{\mathcal{B}}_{0.9}(\cdot)$ performs the best. The sample size has minimal impact for predictive region with χ^2 , where the coverage rates are all close to 0.9 except for $N = 100$. A simple trial shows that, when using 10,000 instead of 1,500 simulated curves, the coverage rate can be corrected to be around the nominal rate for $N = 100$. It is also clear that the bootstrap adjustment for the ECDF-based method is necessary, as the coverage rates of $\widehat{\mathcal{B}}_{0.9}(\cdot)$ based on ECDF-B are clearly closer to the nominal coverage probability 0.9 than those based on ECDF, especially for $N \leq 400$. Also clearly noticeable is the improvement of performance as the sample size N increases. On the other hand, the different noise levels $\sigma = 0.25$ and 0.5 has no impact on the performance, as the signal-to-noise ratio of a stationary AR process is invariant with respect to the noise level.

Table 1 lists the means and standard errors (in parentheses), over 400 replications, of MAE in (22), the average length (AvL) of $\widehat{\mathcal{B}}_{0.9}(\cdot)$ over the 51 grid points, and the proportion of the overlapping area:

$$\text{POA} = \text{Area}\{\widehat{\mathcal{B}}_{0.9}(\cdot) \cap \mathcal{B}_{0.9}(\cdot)\} / \text{Area}\{\mathcal{B}_{0.9}(\cdot)\}, \quad (23)$$

where $\text{Area}\{\mathcal{B}_{0.9}(\cdot)\}$ is evaluated by simulation from the true model. The AvL for

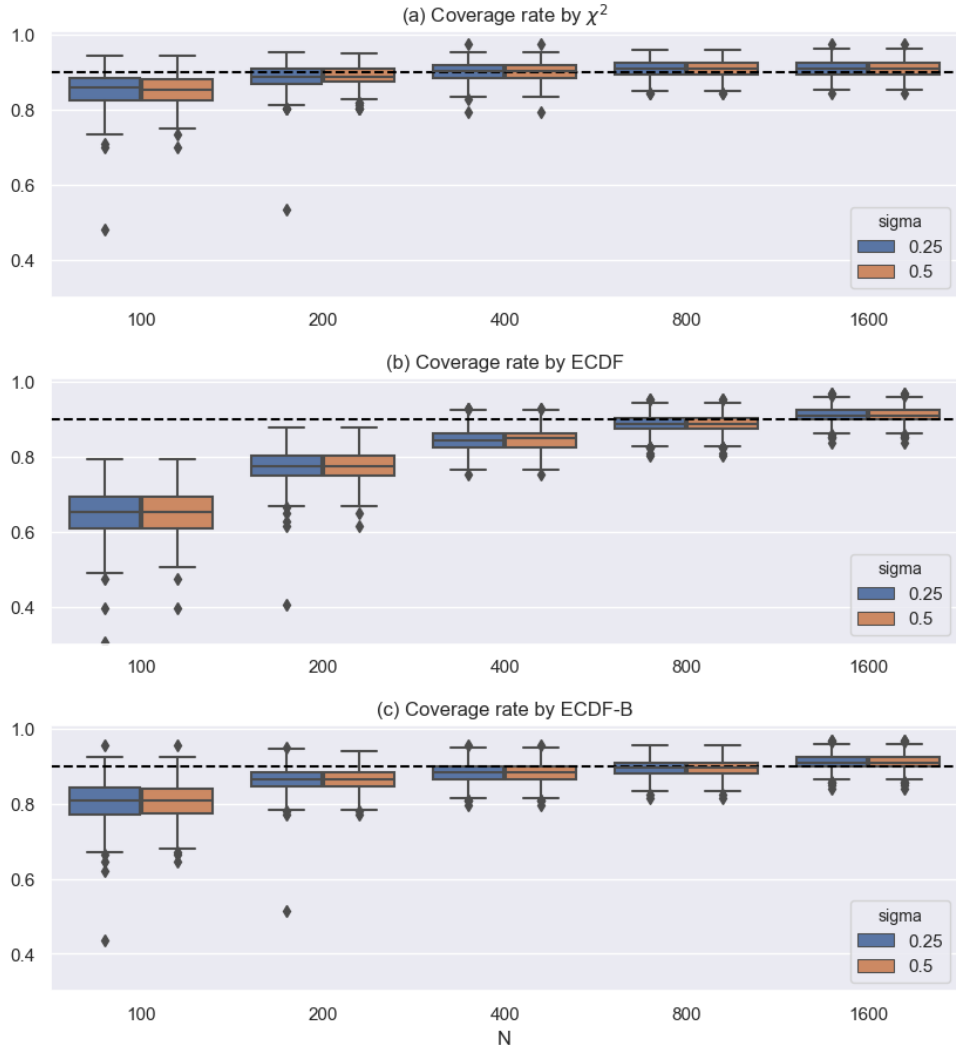


Figure 1: Experiment 1 – Box plots of post sample coverage rates of $\hat{\mathcal{B}}_{1-\alpha}(X)$ based on (a) χ^2 , (b) ECDF, and (c) ECDF-B for $N = 100, 200, 400, 800, 1600$ and $\sigma = 0.25, 0.5$. The horizontal dash lines mark the positions of the nominal coverage probability 0.9.

$\hat{\mathcal{B}}_{0.9}(\cdot)$ based on χ^2 or ECDF-B tends to be larger than that based on ECDF. This is due to the fact that $\hat{\mathcal{B}}_{0.9}(\cdot)$ based on ECDF tends to be smaller, reflected by lower coverage rates (see Figure 1) and smaller overlapping areas. Since the overlapping

areas are always at least, respectively, 94.8%, 82.9%, 91.9% based on χ^2 , ECDF, ECDF-B, $\widehat{\mathcal{B}}_{0.9}(\cdot)$ provides an accurate estimator for $\mathcal{B}_{0.9}(\cdot)$. The average length of $\widehat{\mathcal{B}}_{0.9}(\cdot)$ with $\sigma = 0.5$ is significant larger than that with $\sigma = 0.25$, reflecting more uncertainty in forecasting due to large noise. Note that the true AvL is 1.66 with $\sigma = 0.25$, and 3.32 with $\sigma = 0.50$. The improvement due to the increase of N is evident. For $\widehat{\mathcal{B}}_{0.9}(\cdot)$ based on ECDF-B, we also report the mean and standard errors for B , where $200B$ is the selected bootstrap sample size. Also included in the table are the average estimated values \widehat{d} . In fact, \widehat{d} is always equal to the true value 4 in the 400 replications.

As illustration, Figure 2 depicts 2 randomly selected post-sample curves together with their 90% predictive regions and quantile curves at 40% and 90% confidence levels based on χ^2 with $N = 100$ and $\sigma = 0.25$. Even for the small sample size case, the 90% predictive regions of the 2 curves are overlapping with the true regions to a large extent. In addition, the 90% predictive quantile curves derived from ED coincide with some parts of the bounds of 90 % predictive regions, and indeed represent the most ‘outside’ curve in the associated predictive regions. The predictive quantile curves at a higher confidence level (90%) are more ‘outside’ than those at a lower level (40%) for most parts of the domain of the curves.

Experiment 2: We set in (21) $d = 6$, $\pi_j = \{1, \dots, 6\}$, and consider three distributions for ε_{tj} : $N(0, \sigma^2)$, the centered and rescaled standard exponential distribution, and the rescaled t -distribution with 5 degrees of freedom. The rescaling makes the standard deviation $\sigma = 0.25$ for all the three distributions. Since PPC based on ECDF-B are clearly superior than those based on ECDF in Experiment 1, we drop the results based on ECDF.

We adopt the similar setting as in Experiment 1. The results are reported in

Table 1: Experiment 1 – The means and standard errors (in parentheses) of MAE in (22), the average length (AvL) of $\widehat{\mathcal{B}}_{0.9}(\cdot)$, the proportion of the overlapping area (POA) in (23), and \widehat{d} over 400 replications. For $\widehat{\mathcal{B}}_{0.9}(\cdot)$ based on ECDF-B, the selected bootstrap sample size is $200B$, and the means and standard errors of B are also included.

σ	N	\widehat{d}	MAE	χ^2		ECDF		ECDF-B		
				AvL	POA	AvL	POA	AvL	B	POA
0.25	100	4	.297(.035)	1.726(.063)	.948(.028)	1.420(.065)	.829(.038)	1.626(.075)	4.01(.95)	.919(.034)
	200	4	.288(.016)	1.753(.045)	.970(.015)	1.549(.054)	.900(.026)	1.690(.056)	3.89(.96)	.954(.020)
	400	4	.284(.008)	1.765(.036)	.980(.009)	1.648(.044)	.947(.017)	1.722(.045)	3.20(1.1)	.971(.013)
	800	4	.282(.008)	1.771(.029)	.985(.007)	1.723(.034)	.975(.010)	1.736(.033)	.87(.92)	.978(.010)
	1600	4	.282(.008)	1.777(.024)	.988(.006)	1.779(.024)	.989(.006)	1.779(.024)	.00(.00)	.989(.006)
0.5	100	4	.595(.098)	3.447(.122)	.949(.030)	2.835(.128)	.829(.039)	3.250(.149)	4.07(.92)	.920(.036)
	200	4	.575(.016)	3.506(.087)	.970(.012)	3.097(.103)	.901(.024)	3.379(.109)	3.92(.97)	.954(.018)
	400	4	.567(.016)	3.530(.071)	.980(.008)	3.296(.087)	.947(.016)	3.443(.090)	3.19(1.11)	.971(.012)
	800	4	.564(.016)	3.542(.057)	.985(.007)	3.444(.068)	.975(.010)	3.471(.066)	.87(.93)	.978(.010)
	1600	4	.563(.015)	3.553(.048)	.988(.006)	3.558(.049)	.989(.006)	3.558(.049)	.00(.00)	.989(.006)

Tables 2. The performance of both the pointwise and probabilistic forecasting are satisfactory, and are about the same with the three different distributions for ε_{ti} . The coverage rates (CR) of $\widehat{\mathcal{B}}_{0.9}(\cdot)$ are almost as good as in Experiment 1, especially for $N \geq 400$. In terms of CR, AvL and POA, there are hardly any substantial differences between $\widehat{\mathcal{B}}_{0.9}(\cdot)$ based on χ^2 or on ECDF-B. It is somehow surprising that the χ^2 based method works fine with the heavy-tailed t_5 distribution and the highly skewed exponential distribution. Note that the comparison of AvL can only be made for the two cases with about the same CR. When $N = 1,600$ with exponential distributed ε_{ti} , the CR for $\widehat{\mathcal{B}}_{0.9}(\cdot)$ based on the two methods are the same, the AvL based on ECDF-B is 2.318 which is smaller than 2.359, the AvL based on χ^2 . This may be due to the fact that the residuals from the fitted model capture the skewed exponential distribution better than the χ^2 approximation, though one may

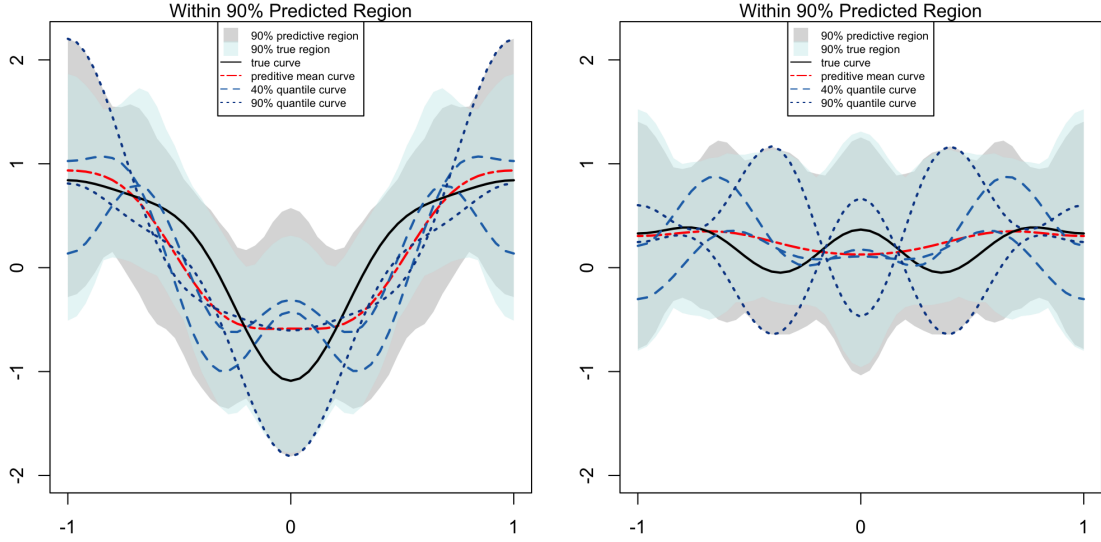


Figure 2: Experiment 1– Randomly selected 2 post-sample curves (solid lines), their true 90% regions (light blue shadow areas), predictive mean curves (dash-dotted lines), 90% predictive regions (grey shadow areas) and quantile curves at 90% and 40% confidence levels (dotted and dashed lines respectively) based on χ^2 , for $N = 100$ and $\sigma = 0.25$.

argue if such a difference is really substantial. The true AvL is 2.17, 2.13 and 2.10, respectively, with normal, t_5 and exponential distributed innovations.

Experiment 3: We investigate the performance of PPC for higher order curve regressions. Consider an FAR(3) process $Y_t(\cdot)$ defined as in (21) in which the second equation is replaced by one of the three equations below.

$$\text{Non sparse: } \xi_{tj} = \sum_{l=1}^d b_{jl} \xi_{t-1,l} + \sum_{l=1}^d b_{j,l+d} \xi_{t-2,l} + \sum_{l=1}^d b_{j,l+2d} \xi_{t-3,l} + \varepsilon_{tj},$$

$$\text{Lag sparse: } \xi_{tj} = \sum_{\ell \in \pi_j} b_{j\ell} \xi_{t-3,\ell} + \varepsilon_{tj}, \quad \pi_j = \{1, \dots, 6\},$$

$$\text{Diagonal sparse: } \xi_{tj} = b_{j1} \xi_{t-1,j} + b_{j,2} \xi_{t-2,j} + b_{j,3} \xi_{t-3,j} + \varepsilon_{tj}.$$

Table 2: Experiment 2 – The means and standard errors (in parentheses) of MAE in (22), the coverage rate (CR) and the average length (AvL) of $\widehat{\mathcal{B}}_{0.9}(\cdot)$, the proportion of the overlapping area (POA) in (23), and \widehat{d} over 400 replications. For $\widehat{\mathcal{B}}_{0.9}(\cdot)$ based on ECDF-B, the selected bootstrap sample size is $200B$, and the means and standard errors of B are also included.

N	dist	\widehat{d}	MAE	χ^2			ECDF-B			
				CR	AvL	POA	CR	AvL	B	POA
100	norm	6	.364(.010)	.849(.040)	2.298(.069)	.952(.016)	.787(.053)	2.138(.081)	4.20(.86)	.920(.024)
	t_5	6	.355(.014)	.851(.045)	2.288(.133)	.952(.022)	.806(.055)	2.146(.128)	4.09(.94)	.926(.032)
	exp	6	.349(.015)	.850(.040)	2.284(.128)	.951(.025)	.820(.052)	2.171(.151)	3.99(.94)	.931(.033)
200	norm	6	.353(.009)	.890(.029)	2.332(.054)	.969(.011)	.859(.039)	2.237(.073)	4.05(.85)	.953(.017)
	t_5	6	.344(.011)	.886(.031)	2.325(.092)	.972(.014)	.866(.034)	2.251(.093)	4.05(.93)	.960(.019)
	exp	6	.339(.012)	.879(.031)	2.324(.106)	.969(.018)	.870(.035)	2.289(.130)	3.99(.96)	.963(.022)
400	norm	6	.349(.008)	.907(.023)	2.347(.042)	.978(.009)	.886(.025)	2.278(.052)	3.23(1.01)	.968(.012)
	t_5	6	.340(.010)	.902(.024)	2.346(.070)	.980(.012)	.888(.028)	2.301(.075)	3.36(1.01)	.972(.014)
	exp	6	.333(.011)	.892(.025)	2.344(.079)	.977(.014)	.891(.025)	2.337(.094)	3.28(.95)	.976(.016)
800	norm	6	.348(.007)	.911(.022)	2.352(.036)	.983(.007)	.890(.027)	2.291(.040)	.82(.68)	.974(.010)
	t_5	6	.338(.009)	.905(.023)	2.355(.053)	.983(.011)	.893(.024)	2.313(.055)	1.74(.72)	.976(.014)
	exp	6	.332(.013)	.895(.025)	2.356(.057)	.981(.012)	.893(.028)	2.332(.071)	1.96(.85)	.980(.013)
1600	norm	6	.346(.008)	.916(.022)	2.359(.032)	.985(.007)	.918(.022)	2.365(.034)	.00(.00)	.986(.007)
	t_5	6	.338(.011)	.907(.022)	2.360(.045)	.986(.009)	.902(.022)	2.345(.044)	.03(.19)	.983(.010)
	exp	6	.331(.012)	.898(.023)	2.359(.042)	.983(.011)	.898(.023)	2.318(.046)	.18(.39)	.982(.011)

Similar to Experiment 1, we set $d = 4$ and let $\varepsilon_{tj} \sim N(0, \sigma^2)$ with $\sigma = 0.25$. But now $X_t = [Y_{t-1}, Y_{t-2}, Y_{t-3}]'$

Table 3 shows the detailed predictive results based on either χ^2 or ECDF-B, for the non-sparse FAR(3) and the two sparse FAR(3) models above. The first 10 singular value components of the regressor curve are included for the AIC selection. The true AvL is 1.66 for all the three FAR(3) models. In general, PPC work well for the higher order curve regressions, with CRs close to the nominal coverage probability

Table 3: Experiment 3 – The means and standard errors (in parentheses) of MAE in (22), the coverage rate (CR) and the average length (AvL) of $\widehat{\mathcal{B}}_{0.9}(\cdot)$, the proportion of the overlapping area (POA) in (23), and \widehat{d} over 400 replications. For $\widehat{\mathcal{B}}_{0.9}(\cdot)$ based on ECDF-B, the selected bootstrap sample size is $200B$, and the means and standard errors of B are also included.

sparsity	N	\widehat{d}	MAE	χ^2			ECDF-B			
				CR	AvL	POA	CR	AvL	B	POA
non sparse	100	4	.310(.017)	.834(.047)	1.768(.090)	.939(.020)	.773(.061)	1.645(.093)	4.08(1.00)	.905(.029)
	200	4	.296(.016)	.882(.031)	1.787(.096)	.962(.013)	.849(.037)	1.712(.095)	3.90(.91)	.944(.019)
	400	4	.289(.013)	.900(.026)	1.790(.062)	.976(.010)	.881(.032)	1.740(.061)	2.98(1.14)	.964(.015)
	800	4	.286(.010)	.905(.024)	1.785(.043)	.980(.008)	.889(.025)	1.744(.042)	.86(.91)	.972(.011)
	1600	4	.285(.012)	.910(.023)	1.792(.059)	.986(.007)	.911(.022)	1.795(.058)	.00(.00)	.986(.007)
lag sparse	100	4	.302(.014)	.832(.046)	1.714(.071)	.936(.022)	.780(.058)	1.613(.082)	3.97(.93)	.906(.031)
	200	4	.292(.011)	.880(.030)	1.756(.054)	.963(.013)	.851(.037)	1.689(.062)	3.93(.99)	.945(.018)
	400	4	.286(.009)	.898(.025)	1.772(.043)	.976(.008)	.881(.029)	1.730(.052)	3.12(1.13)	.967(.013)
	800	4	.284(.009)	.905(.023)	1.779(.037)	.982(.007)	.892(.025)	1.743(.038)	.97(.92)	.975(.009)
	1600	4	.284(.008)	.910(.023)	1.785(.028)	.987(.006)	.909(.023)	1.786(.030)	.00(.00)	.987(.007)
diag sparse	100	4	.332(.175)	.803(.104)	1.705(.070)	.920(.062)	.750(.107)	1.606(.079)	3.71(1.09)	.891(.065)
	200	4	.297(.033)	.864(.053)	1.746(.050)	.955(.026)	.838(.060)	1.685(.060)	3.92(.98)	.940(.030)
	400	4	.290(.033)	.893(.029)	1.777(.146)	.973(.011)	.876(.035)	1.734(.134)	3.09(1.06)	.964(.014)
	800	4	.289(.046)	.904(.023)	1.801(.286)	.981(.008)	.893(.026)	1.763(.246)	.86(.85)	.973(.010)
	1600	4	.282(.010)	.911(.020)	1.781(.041)	.986(.007)	.913(.021)	1.784(.031)	.00(.00)	.987(.007)

0.9, POA above 90% and AvL close to the true AvL. The patterns of CRs and POAs are very similar to those in Experiment 1, though with slightly worse performance for small sample size ($N \leq 400$). Performance for the lag sparse models is better than the non-sparse one, producing narrower AvLs for similar coverage rates. When comparing the diag sparse model with non-sparse model, there is no clear pattern. Note that our estimation method makes no use of the information of the particular sparse structures.

4 Probabilistic forecasting for daily electricity loads

In this section, we apply the proposed PPC to a real data set consisting of French daily electricity load curves from January 1, 2012 to December 31, 2019.

4.1 Data

The French electricity consumption data are collected from the website of the system operator RTE (Réseau et Transport d'Électricité): <https://opendata.rte-france.com>) at a temporal resolution of every half-hour (i.e. 48 points on each day). We remove the data on public holidays, the day before and the day after the holidays, and also in the Christmas periods. Empirical experience indicates that the electricity demand on those days behaves differently, and requires different treatment.

Temperature is a key exogenous factor for electricity consumption in France due to electrical heating and cooling. We obtained data from 96 meteorostations in France from the website of the French weather forecaster MétéoFrance (<https://donneespubliques.meteofrance.fr/>). Temperature are provided at a three hours resolution and interpolated with natural cubic splines at a half-hour resolution. Figure 3 displays the dynamic evolution of the daily load curves from 2012 to 2019, the corresponding daily temperature curves and a scatter plot showing the strong dependence between the load and the temperature.

Electricity loads, like other energy data, highly depend on economic and meteorological factors. The changes of temperature introduce seasonality in the demand which is higher in winter and lower in summer. The shift of working routines causes varying diurnal patterns between weekdays and weekends. Therefore, the profiles of daily load curves differ in months and days. As an illustration, Figure 4 depicts the daily curves on Tuesday-Wednesday in June, Saturday-Sunday in June and

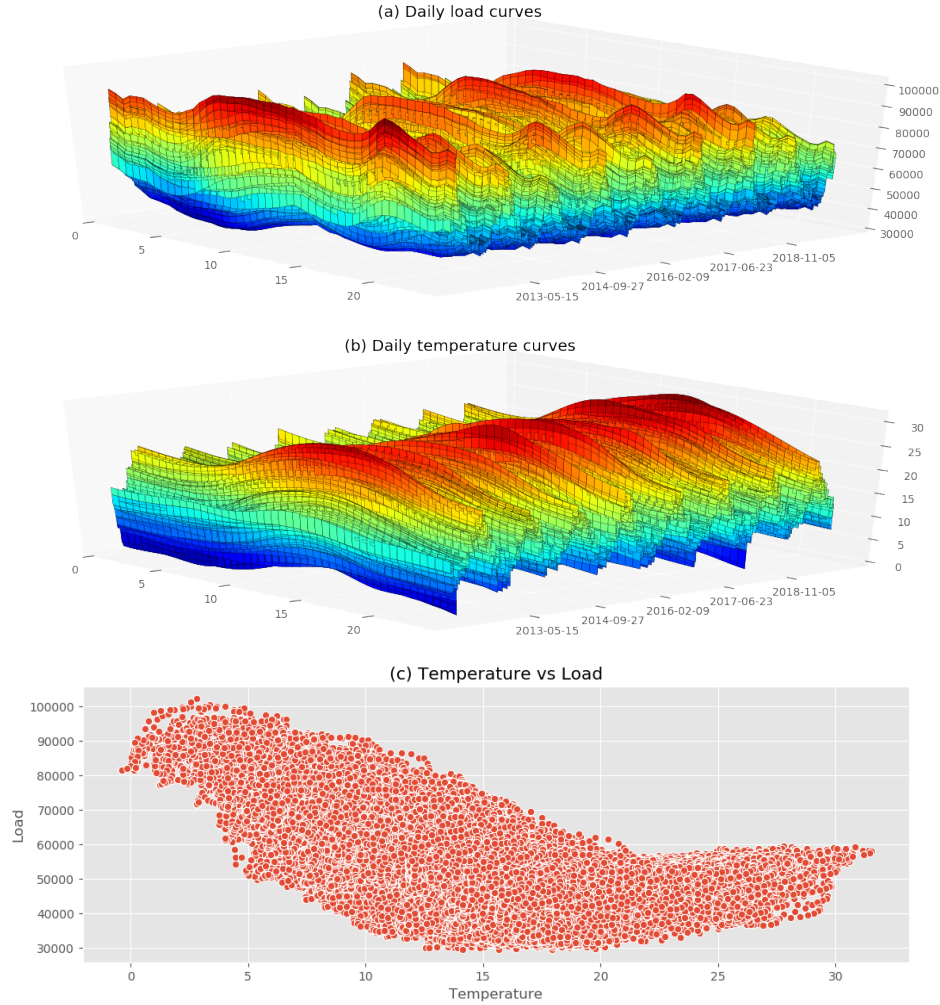


Figure 3: (a) Daily loads curves (b) Daily temperature curves (c) Scatter plot of temperature vs loads at half-hour frequency. Data are from 2012 to 2019.

Tuesday-Wednesday in November between 2012 to 2018. In June, the curves on Tuesday-Wednesday are similar, but differ from those on Saturday-Sunday. Furthermore, the demand is higher on weekdays than that in weekends. In addition, the diurnal pattern varies over different months. For example, peaks occur at noon in June (due to cooling consumption) while in the evening in November (due to heat-




Figure 4: Electricity load curves on Tuesday-Wednesday in June (black dashed lines), Saturday-Sunday in June (red dotted lines) and Tuesday-Wednesday in November (blue solid lines) between 2012 to 2018.

Table 4: Classification rule for electricity load data

Day Class	1	2	3	4	5		
Day type	Mon	Tue, Wed, Thu	Fri	Sat	Sun		
Month Class	1	2	3	4	5	6	7
Month	Dec, Jan, Feb	Mar	Apr, May	Jun, Jul, Sep	Aug	Oct	Nov

ing consumption). The inhomogeneous phenomena dictate the need to segment days into different homogeneous groups, as the proposed PPC are developed under a stationary framework. Table 4 lists the segmentation rule adopted by EDF: each week is divided into 5 groups and the 12 months are classified into 7 groups, leading to in total 35 groups. We will fit a separate model for each of the 35 groups.

4.2 Probabilistic forecasting

For each of the 35 groups, we pair the daily load curve $Y_t(\cdot)$ on day t and $X_t(\cdot)$ consisting of three curves: the load curve on the previous day $Y_{t-1}(\cdot)$, the load curve one week earlier $Y_{t-7}(\cdot)$, and the daily temperature curve $T_t(\cdot)$. The temperature forecasts can be obtained from e.g. Météo-France, which maintain a stable high accuracy. We thus directly use the actual temperature in our analysis, which is commonly adopted in the literature of electricity load forecasting. To make the three curves on the same scale, we standardize each of them (within each of the 35 groups) first before combining them into $X_t(\cdot)$. The standardization uses the training data only, see below.

We use the data in 2019 for testing the post-sample forecasting performance. More precisely, for each day in 2019 we use all the data in the same group before that day to fit the model. The day-ahead forecasts are produced at noon of each day, in alignment with the forecasting operation in EDF. The sample size N varies among the 35 groups from 24 to 284. For determining $\hat{\pi}_j$ in (9) by AIC, we include the first $\min(\lceil N/2 \rceil, 48)$ $\hat{\eta}_{tj}$ as the candidate regressors to avoid over-fitting problem for small N . We set the nominal coverage probability at $1 - \alpha = 0.9$. The predict set $\hat{\mathcal{C}}_{0.9}(\cdot)$ based on χ^2 -approximation contains K randomly generated curves, where $K = 5,000$ or $20,000$.

For the comparison purpose, we also include three state-of-art models popular in electricity load forecasting. Those models are designed for forecasting the load at time point h on day t , denoted by $y_{t,h}$. Fitting these models separately for each hour of the day is a common practice in electricity load forecasting as hour of the day plays a prominent role.

1. The generalized additive model (GAM):

$$y_{t,h} = \psi_{D_{t,h}} + f_{1,h}(t) + f_{2,h}(S_t) + f_{3,h}(y_{t-1,h}, D_t) + f_{4,h}(y_{t-7,h}) + f_{5,h}(t, T_{t,h}) \\ + f_{6,h}(T_{t,h}^s(0.95)) + f_{7,h}(T_{t,h}^s(0.99)) + f_{8,h}(T_{t,h}^{\max}, T_{t,h}^{\min}) + \epsilon_{t,h},$$

where $T_{t,h}$ denotes the temperature at time h on day t , $T_{t,h}^{\max}, T_{t,h}^{\min}$ are the daily maximum, minimum temperatures, $T_{t,h}^s(v) = vT_{t-1}^s(v) + (1-v)T_t$ is an exponentially smoothed version of $T_{t,h}$, S_t represents an annual cycling effect, D_t denotes the day effect and each $f_{j,h}$ is a smooth function of the covariates with the thin plate regression splines basis functions.

2. The seasonal autoregressive (SAR) model:

$$y_{t,h} - \mu_{y,h} = \omega_{1,h}(y_{t-1,h} - \mu_{y,h}) + \omega_{2,h}(y_{t-2,h} - \mu_{y,h}) + \omega_{3,h}(y_{t-7,h} - \mu_{y,h}) + \epsilon_{t,h}.$$

3. The SAR with exogenous variable (SARX) model:

$$y_{t,h} - \mu_{y,h} = \omega_{1,h}(y_{t-1,h} - \mu_{y,h}) + \omega_{2,h}(y_{t-2,h} - \mu_{y,h}) + \omega_{3,h}(y_{t-7,h} - \mu_{y,h}) \\ + \omega_{T,h}(T_{t,h} - \mu_{T,h}) + \epsilon_{t,h}.$$

The GAM was proposed by Pierrot and Goude (2011) and engineered by EDF. It serves here as an industry benchmark. The SAR and SARX are implemented periodically for series attached to each half-hour and fitted separately for each of the 35 groups. For fair comparison to our curve regression approach, a residual correction method is adopted for GAM, SAR and SARX given the intraday dependence among the half-hourly series, in which the residuals from 48 half-hours are jointly modeled. Specifically, an ARMA model is fitted with the 10-fold cross validation residuals on the in-sample data. The forecasted distributions of the residuals for the next 48 half-hours are then used to obtain the mean, 5% quantile and 95% quantile of loads.

Table 5: The mean absolute percentage error (MAPE), coverage rate (CR), pointwise coverage rate (PCR), and the average length (AvL) of the predictive intervals at the 48 half-hour points of the different forecasting methods. The nominal coverage probability is 0.9.

Method	\hat{d}	MAPE	CR	PCR	AvL
$\hat{\mathcal{B}}_{0.9}(\cdot)$ based on χ^2 ($K = 5,000$)			<u>0.563</u>	<u>0.943</u>	<u>3128</u>
$\hat{\mathcal{B}}_{0.9}(\cdot)$ based on χ^2 ($K = 20,000$)	15.8	1.10%	0.669	0.955	3340
$\hat{\mathcal{B}}_{0.9}(\cdot)$ based on ECDF-B			0.484	0.898	2611
GAM	-	1.36%	0.342	0.907	3376
SAR	-	2.03%	0.238	0.823	4065
SARX	-	1.65%	0.238	0.812	3176

For each testing day in 2019, we compute the mean absolute percentage error:

$$\text{MAPE} = \frac{1}{48} \sum_{h=1}^{48} \frac{|\hat{y}_{t,h} - y_{t,h}|}{y_{t,h}}.$$

Table 5 summarizes the results. In terms of forecast accuracy, PPC deliver the best performance with MAPE = 1.10%. Compared to the GAM, SAR and SARX, this corresponds to the reduction in MAPE of 19.1%, 45.8% and 33.3% respectively. It is noticeable that the coverage rate (CR) of the predictive bands are smaller than the nominal level 0.9. One possible reason is the small sample sizes of some groups, for which the variances of noise are likely to be underestimated, as the possible variation in the future is unlikely to be reflected by the small number of the available observations. This is particularly pronounced with the method based on ECDF-B. Nevertheless PPC performs significantly better than the other method, as all the three $\hat{\mathcal{B}}_{0.9}(\cdot)$ listed in Table 5 provide much high coverage rates (CV) for the whole




Figure 5: Bar charts of MAPE, coverage rate, pointwise coverage rate and AvL of χ^2 -based $\hat{\mathcal{B}}_{0.9}$ ($K = 20,000$) over different months.

curve than GAM, SAR and SARX.

One may argue that the comparison above is unfair as the predictive intervals based on GAM, SAR and SARX are constructed in the pointwise manner rather than for a whole curve. Table 5 also lists the pointwise coverage rates (PCR) of the different methods for the 48 points (corresponding the 48 30-minute intervals) on a daily curve. It is interesting to observe that the two $\hat{\mathcal{B}}_{0.9}(\cdot)$ based on χ^2 approximation provide significantly higher PCR than those of GAM, SAR and SARX, and $\hat{\mathcal{B}}_{0.9}(\cdot)$ based ECDF-B offers comparable PCR to GAM with much smaller AvL. It is worth to point out that $\hat{\mathcal{B}}_{0.9}(\cdot)$ based on χ^2 with $K = 5,000$ outperforms GAM, SAR and SARX in terms of all the three measures: delivering higher CR (increases by 0.221 - 0.325), higher PCR (improves by 0.036 - 0.131) and narrower AvL (decreases by 48 - 937).




Figure 6: Bar charts of MAPE, coverage rate, pointwise coverage rate and AvL of χ^2 -based $\hat{\mathcal{B}}_{0.9}$ ($K = 20,000$) over different weekdays.

To appreciate the variation in forecasting performance over different weekdays and different months, Figures 5 and 6 display the bar-charts of MAPE, CR, PCR and AvL of $\hat{\mathcal{B}}_{0.9}(\cdot)$ based on χ^2 -approximation with $K = 20,000$. It is clear that the forecasting in summer is more accurate than that in winter in terms of both MAPE (Figure 5(a)) and AvL (Figure 5(d)), and the forecasting in Monday – Friday is more accurate than that in weekends in terms of MAPE (Figure 6(a)). There is no clear pattern in terms of the two coverage rates.

Moreover, Figure 7 elaborates 2 randomly selected forecasts, showing the predictive mean curves, 90% predictive regions $\hat{\mathcal{B}}_{0.9}(\cdot)$ and the quantile curves at 99% and 40% confidence levels in 2019 based on χ^2 -approximation with $K = 20,000$. As expected, the predictive quantile curves at a higher confidence level is more ‘outside’ than those at a lower level. The 99% predictive quantile curves help to visualize the




Figure 7: Plots of the true load curves (solid lines), predictive mean curves (dash-dotted lines), 90% predictive regions (grey shadow areas) and predictive quantile curves at 99% and 40% confidence level (dotted and dashed lines) for 2 randomly selected curves, one from summer time and one from winter time.

possible extreme cases, which is useful for the optimal controlling of the electricity operation system.

5 Conclusion

In this paper, we propose a novel method to construct three types of probabilistic predictors for curves (PPC): the predictive set, the predictive band and the predictive quantile with probability interpretation for daily electricity load curves in a curve-to-curve linear regression framework. The PPC achieve excellent performance with coverage rates very close to the nominal probabilities for different simulated

data generating processes. When applied to the French load curves, the proposed method provides much smaller forecast errors, with almost half of that of the alternative seasonal autoregressive models. Compared to the powerful generalized additive model, it produces higher coverage rate with narrower average length of the predictive regions. The significant improvement is likely due to the curve regression setting which embeds the nonstationary daily patterns into a stationary framework. The constructed predictive intervals and the predictive quantile curves are meaningful and can be used in the future for risk hedging in the electricity management system.

References

- Antoniadis, A., Brossat, X., Cugliari, J., and Poggi, J.-M. (2016). A prediction interval for a function-valued forecast model: Application to load forecasting. *International Journal of Forecasting*, 32(3):939–947.
- Cabrera, B. L. and Schulz, F. (2017). Forecasting Generalized Quantiles of Electricity Demand: A Functional Data Approach. *Journal of the American Statistical Association*, 112(517):127–136.
- Chen, Y. and Li, B. (2017). An adaptive functional autoregressive forecast model to predict electricity price curves. *Journal of Business & Economic Statistics*, 35(3):371–388.
- Chen, Y., Xu, X., and Koch, T. (2020). Day-ahead high-resolution forecasting of natural gas demand and supply in germany with a hybrid model. *Applied Energy*, 262:114486.

- Cho, H., Goude, Y., Brossat, X., and Yao, Q. (2013). Modeling and forecasting daily electricity load curves: A hybrid approach. *Journal of the American Statistical Association*, 108(501):7–21.
- Cho, H., Goude, Y., Brossat, X., and Yao, Q. (2015). Modelling and forecasting daily electricity load via curve linear regression. In *Modeling and Stochastic Learning for Forecasting in High Dimension* edited by A. Antoniadis and X. Brossat, pages 35–54. Springer.
- Dordonnat, V., Pichavant, A., and Pierrot, A. (2016). GEFCom2014 probabilistic electric load forecasting using time series and semi-parametric regression models. *International Journal of Forecasting*, 32(3):1005–1011.
- Fasiolo, M., Wood, S. N., Zaffran, M., Nedellec, R., and Goude, Y. (2020). Fast calibrated additive quantile regression. *Journal of the American Statistical Association*, 115:1–11.
- Fraiman, R. and Muniz, G. (2001). Trimmed means for functional data. *Test*, 10(2):419–440.
- Gaillard, P., Goude, Y., and Nedellec, R. (2016). Additive models and robust aggregation for GEFCom2014 probabilistic electric load and electricity price forecasting. *International Journal of forecasting*, 32(3):1038–1050.
- García-Villalobos, J., Zamora, I., San Martín, J. I., Asensio, F. J., and Aperribay, V. (2014). Plug-in electric vehicles in electric distribution networks: A review of smart charging approaches. *Renewable and Sustainable Energy Reviews*, 38:717–731.
- Gielen, D., Boshell, F., Saygin, D., Bazilian, M. D., Wagner, N., and Gorini, R.

- (2019). The role of renewable energy in the global energy transformation. *Energy Strategy Reviews*, 24:38 – 50.
- Hong, T. and Fan, S. (2016). Probabilistic electric load forecasting: A tutorial review. *International Journal of Forecasting*, 32(3):914 – 938.
- Hong, T., Pinson, P., Fan, S., Zareipour, H., Troccoli, A., and Hyndman, R. J. (2016). Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond. *International Journal of Forecasting*, 32:896–913.
- Hong, T., Xie, J., and Black, J. (2019). Global energy forecasting competition 2017: Hierarchical probabilistic load forecasting. *International Journal of Forecasting*, 35(4):1389 – 1399.
- Koenker, R. and Bassett, G. (1978). Regression quantiles. *Econometrica*, 46(1):33–50.
- Kou, P. and Gao, F. (2014). A sparse heteroscedastic model for the probabilistic load forecasting in energy-intensive enterprises. *International Journal of Electrical Power and Energy Systems*, 55:144–154.
- López-Pintado, S. and Romo, J. (2009). On the concept of depth for functional data. *Journal of the American Statistical Association*, 104(486):718–734.
- Narisetty, N. N. and Nair, V.N. (2016). Extremal depth for functional data and applications. *Journal of the American Statistical Association*, 111(516):1705–1714.
- Petiau, B. (2009). Confidence interval estimation for short-term load forecasting. *2009 IEEE Bucharest PowerTech: Innovative Ideas Toward the Electrical Grid of the Future*, pages 1–6.

- Pierrot, A. and Goude, Y. (2011). Short-term electricity load forecasting with generalized additive models. *Proceedings of ISAP power*, 2011.
- Polonik, W. and Yao, Q. (2000). Conditional minimum volume predictive regions for stochastic processes. *Journal of the American Statistical Association*, 95(450):509–519.
- Roach, C. (2019). Reconciled boosted models for gefcom2017 hierarchical probabilistic load forecasting. *International Journal of Forecasting*, 35(4):1439–1450.
- Rolnick, D., Donti, P. L., Kaack, L. H., Kochanski, K., Lacoste, A., Sankaran, K., Ross, A. S., Milojevic-Dupont, N., Jaques, N., Waldman-Brown, A., et al. (2019). Tackling climate change with machine learning. *arXiv preprint arXiv:1906.05433*.
- Smyl, S. and Hua, N. G. (2019). Machine learning methods for gefcom2017 probabilistic load forecasting. *International Journal of Forecasting*, 35(4):1424–1431.
- Taylor, J. W. and Buizza, R. (2003). Using weather ensemble predictions in electricity demand forecasting. *International Journal of Forecasting*, 19(1):57–70.
- Wang, Y., Chen, Q., Hong, T., and Kang, C. (2018). Review of smart meter data analytics: Applications, methodologies, and challenges. *IEEE Transactions on Smart Grid*, 10(3):3125–3148.
- Wang, Y., Chen, Q., Kang, C., Zhang, M., Wang, K., and Zhao, Y. (2015). Load profiling and its application to demand response: A review. *Tsinghua Science and Technology*, 20(2):117–129.

SUPPLEMENTARY MATERIAL

1 Additional figures for simulation studies

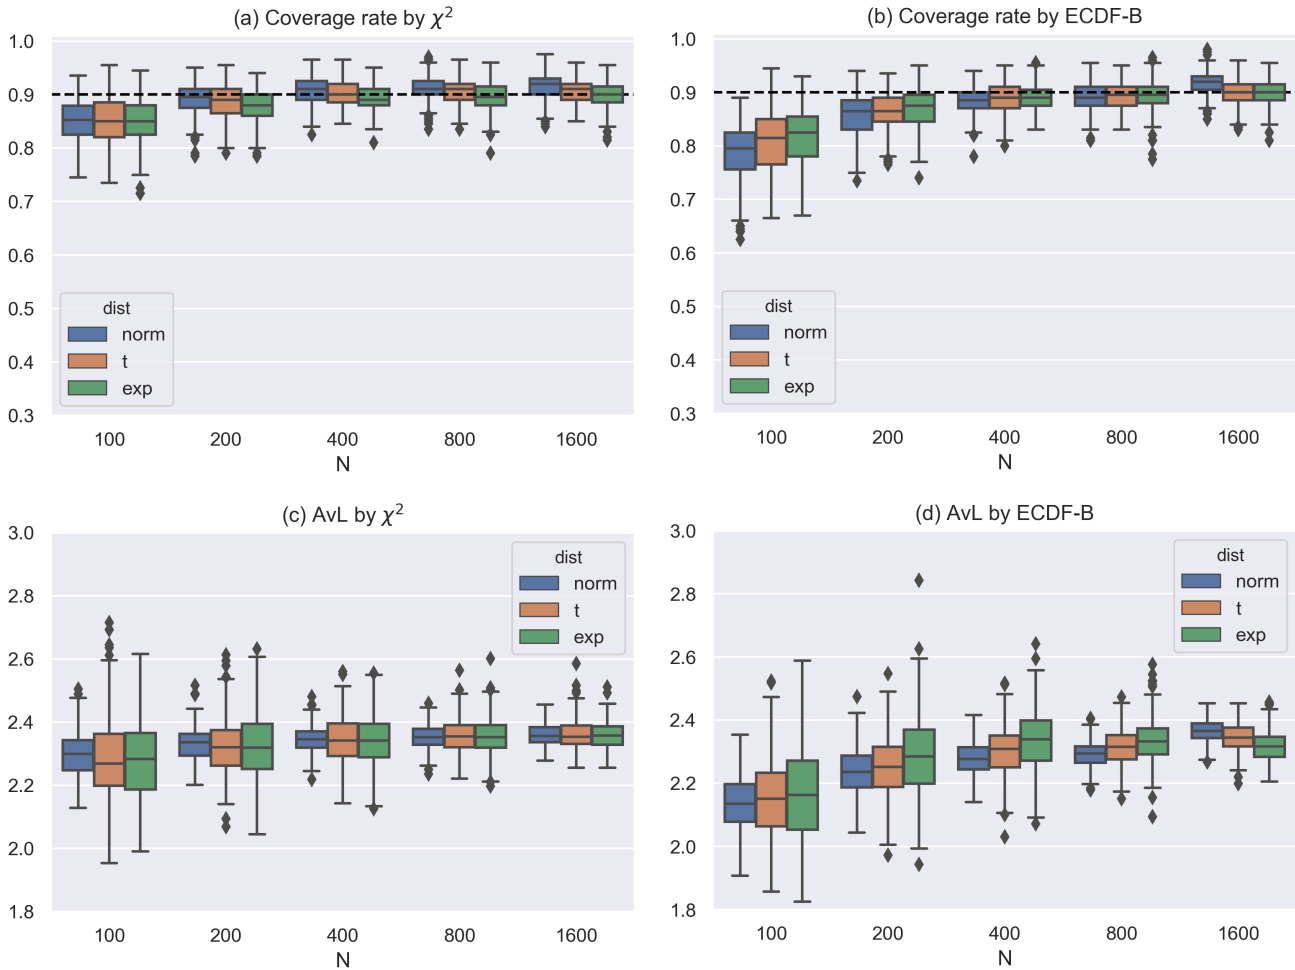


Figure 8: Experiment 2 – Box plots of post-sample coverage rates and average lengths of the predictive region by the envelop of $\widehat{\mathcal{C}}_{1-\alpha}(X)$ based on χ^2 distribution (left column), and those of $\widehat{\mathcal{C}}_{1-\alpha}(X) \cup \mathcal{C}_{1-\alpha}^*(X)$ based on ECDF-B (right column) for the case of $\sigma = 0.25$. The horizontal dash lines mark the positions of the nominal coverage probability 0.9.

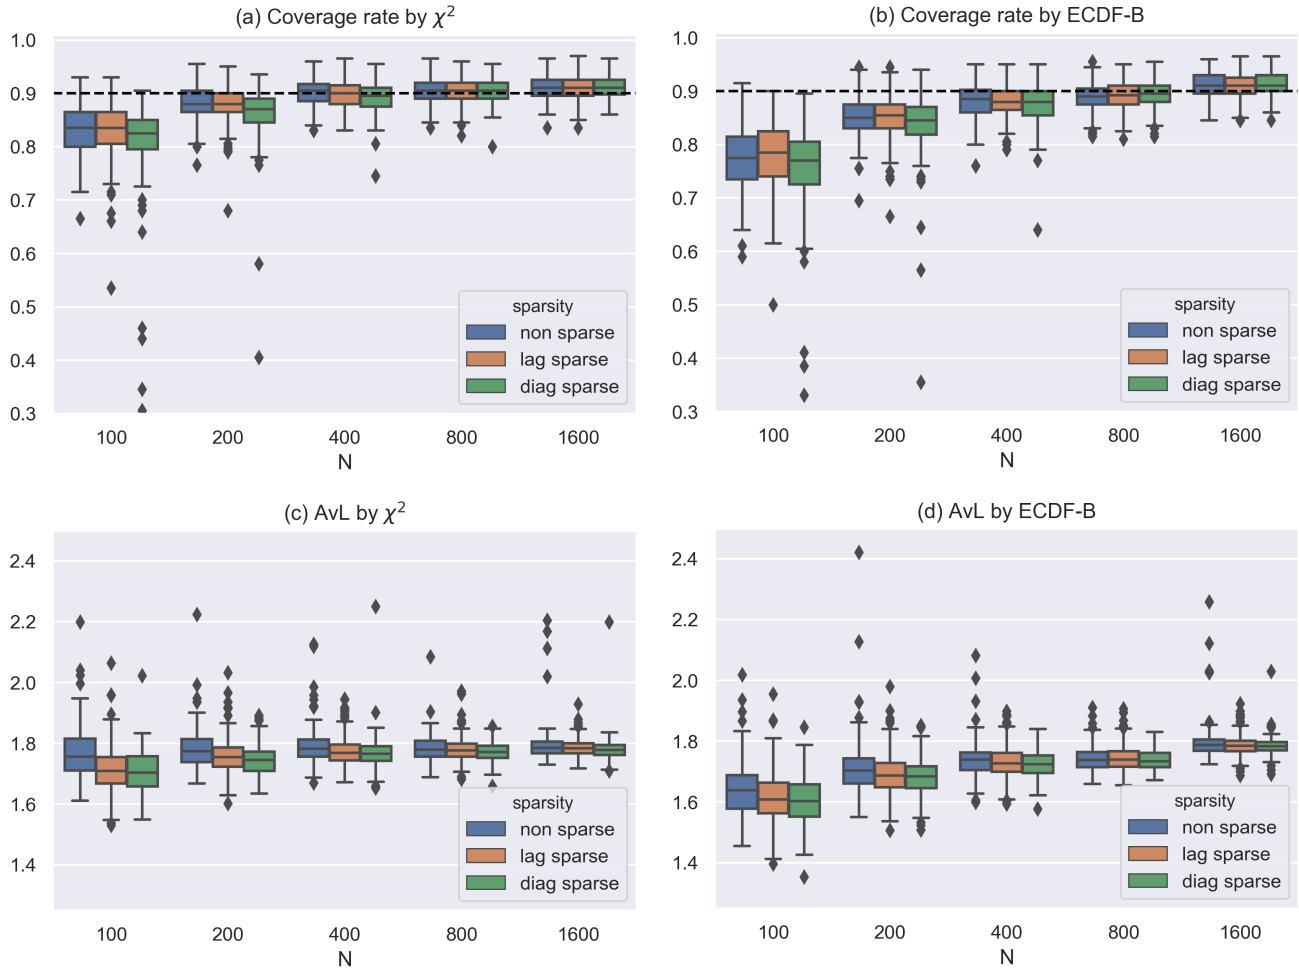


Figure 9: Experiment 3 – Box plots of post-sample coverage rates and average lengths of the predictive regions by the envelop of $\hat{\mathcal{C}}_{1-\alpha}(X)$ based on χ^2 (left column), and those of $\hat{\mathcal{C}}_{1-\alpha}(X) \cup \mathcal{C}_{1-\alpha}^*(X)$ based on ECDF-B (right column) for FAR(3) and sparse FAR(3) data generating processes with $\sigma = 0.25$. The horizontal dash lines mark the positions of the nominal coverage probability 0.9.