

de Vilmarest, J., Browell, J. , Fasiolo, M., Goude, Y. and Wintenberger, O. (2024) Adaptive probabilistic forecasting of electricity (net-)load. *IEEE Transactions on Power Systems*, 39(2), pp. 4154-4163. (doi: [10.1109/TPWRS.2023.3310280](https://doi.org/10.1109/TPWRS.2023.3310280))

Copyright © 2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

This is the author version of the work. There may be differences between this version and the published version. You are advised to consult the published version if you wish to cite from it:
<https://doi.org/10.1109/TPWRS.2023.3310280>

<https://eprints.gla.ac.uk/320295/>

Deposited on 14 March 2024

Adaptive Probabilistic Forecasting of Electricity (Net-)Load

Joseph de Vilmarest, Jethro Browell, Matteo Fasiolo, Yannig Goude and Olivier Wintenberger

Abstract—Electricity load forecasting is a necessary capability for power system operators and electricity market participants. Both demand and supply characteristics evolve over time. On the demand side, unexpected events as well as longer-term changes in consumption habits affect demand patterns. On the production side, the increasing penetration of intermittent power generation significantly changes the forecasting needs. We address this challenge in two ways. First, our setting is *adaptive*; our models take into account the most recent observations available to automatically respond to changes in the underlying process. Second, we consider *probabilistic* rather than point forecasting; indeed, uncertainty quantification is required to operate electricity systems efficiently and reliably. Our methodology relies on the Kalman filter, previously used successfully for adaptive point load forecasting. The probabilistic forecasts are obtained by quantile regressions on the residuals of the point forecasting model. We achieve adaptive quantile regressions using the online gradient descent; we avoid the choice of the gradient step size considering multiple learning rates and aggregation of experts. We apply the method to two data sets: the regional net-load in Great Britain and the demand of seven large cities in the United States. Adaptive procedures improve forecast performance substantially in both use cases for both point and probabilistic forecasting.

Index Terms—Adaptive forecasting, net-load, probabilistic, time series.

I. INTRODUCTION

FORECASTING electricity demand is fundamental in the process of maintaining supply-demand balance. This permanent equilibrium is necessary to maintain a reliable supply of electricity and to avoid damaging infrastructure. As electricity cannot be stored on a large scale, forecasts are crucial to informing production planning. This necessity explains why energy forecasting has gathered so much attention from the time series and forecasting community [1]. The recent increase in electricity prices in Europe further emphasizes the importance of demand forecast quality.

The literature historically focused on point forecasting. However, the expected value of the load is not sufficient for risk management. Forecasting models cannot be perfect

and therefore the production plans cannot match exactly the demand; therefore, grid operators need to call reserves to maintain the equilibrium and traders need to hedge their market positions. It is thus essential to have some information on the distribution around the mean to schedule the reserve needed. In Great Britain, for instance, reserves are tailored to be sufficient in all but 4 hours per year; this corresponds to the estimation of 99.95% confidence intervals. The importance of probabilistic forecasts was highlighted in the last two Global Energy Forecasting Competitions [2], [3]. Indeed, the objective in both competitions was to forecast certain quantiles, and the submissions were evaluated through the sum of quantile losses. Also, a recent challenge called *Net Load Forecasting Prize* was launched by the U.S. Department of Energy and EPRI, focusing on probabilistic net load forecasts [4].

Demand and supply characteristics are both evolving with time. On the demand side, unexpected events such as the coronavirus crisis or the recent increase in electricity prices in Europe can affect it significantly and abruptly [5], [6]. From a longer-term perspective, the evolution of consumption habits such as those induced by the expected growth in electric vehicle penetration will change the demand patterns. On the production side, the increasing penetration of intermittent electricity production strongly changes the forecasting task. Indeed, controllable production units must be employed to meet the difference between the demand and the intermittent electricity production (mainly wind and solar energy) [7], denoted by *net-load*. These changes in the behavior of the net-load can not be captured by classical offline probabilistic methods and motivate the need for adaptive methods, taking into account observations in a streaming fashion to improve the forecasting model.

The recent point forecasting literature has highlighted the improvements yielded by adaptive methods able to learn regime changes. Adaptation is especially crucial to forecast wind and solar power, and it has been successfully applied to online variable selection [8] and online forecast reconciliation [9]. State-space representations have well captured the recent changes of pattern due to the coronavirus crisis [10], [11]. While the majority of recent works have focused on point forecasting, [12] consider adaptive probabilistic load forecasting based on hidden Markov models and is found to perform well for a range of loads; however, this approach relies on Gaussian predictive distributions which leads to relatively poor calibration (quantile bias) on the data we consider here. In [13], [14], quantile GAMs are proposed, as well as variants of GAMLSS (Location Scale and Shape) in [15], for probabilistic demand forecasting. These methods relax the Gaussian assumption

J. de Vilmarest is with Viking Conseil; he was with Électricité de France R&D and Sorbonne Université during part of the work (e-mail: joseph.de-vilmarest@vikingconseil.fr).

J. Browell (Senior Member, IEEE) is with the University of Glasgow (e-mail: Jethro.Browell@glasgow.ac.uk).

M. Fasiolo is with the University of Bristol (e-mail: matteo.fasiolo@bristol.ac.uk).

Y. Goude is with Électricité de France R&D and with the Laboratoire de Mathématique d'Orsay, Université Paris-Saclay (e-mail: yannig.goude@edf.fr).

O. Wintenberger is with Sorbonne Université and with Wolfgang Pauli Institut, c/o Fakultät für Mathematik, Universität Wien (e-mail: olivier.wintenberger@sorbonne-universite.fr).

but are not adaptive. [16] propose a mixture model for short term probabilistic forecasting of individual load demand at different aggregation level, but the distribution of the data is supposed to be invariant with time. [17] propose, similarly to [10], a deep neural network probabilistic forecasting method where the features extracted with a multi-layer RNN-LSTM are mapped to a state-space model for online adaptation. Their work is limited to time-varying Gaussian distributions as they use a Gaussian state-space model to generate density forecasts.

We propose in this article to apply the adaptive framework to quantile regression for probabilistic forecasting. There has been scarce work in that direction. We mention the exciting use of aggregation of experts to obtain an online forecaster [13], [18]. The principle is to combine various predictions (experts) with a weighted average where the weights evolve over time. We exploit the idea of aggregation, but we introduce a more structural adaptation in the sense the experts themselves are adaptive.

The methodology employed and our contributions are summarized in Figure 1. We adopt a two-step procedure. First, we employ a widely used model class, namely that of Generalized Additive Models (GAMs), to forecast the expected value of the variable of interest. This is a non-adaptive model fitted to several years of training data. We consider a state-space representation to obtain an adaptive variant, relying on the Kalman filter for the inference similarly to [10], [11]. Second, we focus on forecasting the quantiles of the net-demand. While the Kalman filter already models it as a Gaussian, we observe that this forecast is not always well-calibrated. Therefore, we use a set of quantile regressions on the residuals of the mean estimates. We apply the online gradient descent (OGD) to obtain adaptive quantile regressions. We select the best gradient step by exploiting an expert aggregation algorithm, Bernstein Online Aggregation [19]; it means we compute the OGD with several step sizes (multiple arrows in Figure 1) and our final forecast is a combination of these forecasts; the quantile loss is optimized online as proposed by [20].

In Section II we present adaptive methods for probabilistic forecasting, and we discuss their evaluation in Section III. Then we apply our approach on two data sets. In Section IV, we consider the regional net-load in Great Britain, and we extend the work of [21] to the adaptive setting, also extending the data set with more recent data (including the coronavirus crisis). Furthermore, we show in Section IV-D that adaptive models have fewer needs for good explanatory variables; indeed, removing two (difficult to obtain) variables from the model has a lower impact on the adaptive variant. In Section V, we apply the methods to the load of seven major cities in the United States [22]. In the latter application, we predict daily, rather than half-hourly, data to determine whether our approach works well at this time granularity.

Our main contributions are :

- a new method based on time-varying quantile regression with OGD for adaptive probabilistic forecasting of time series.
- application of these methods to regional net demand in the UK and urban consumption in the US. Comparison

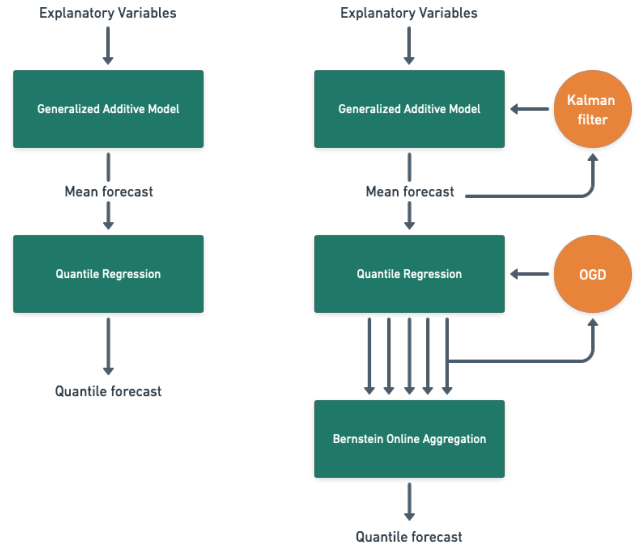


Fig. 1. Our two-step procedure starts from explanatory variables, computes an intermediate mean forecast, and then estimate the quantile forecast. On the left we describe the offline model (Section II-A and [21]). Our contributions are represented by the right diagram, introducing our adaptive model (Sections II-B to II-C4), in which both steps are adapted.

with state of the art online forecasting methods demonstrates the interest of this approach.

II. THEORETICAL FRAMEWORK

We aim to forecast a variable of interest $y_t \in \mathbb{R}$ given some explanatory variables $x_t \in \mathbb{R}^d$, $d > 0$. We formally introduce the impact of time dependence on our approach. Let $\mathcal{F}_t = \sigma(x_1, y_1, \dots, x_t, y_t)$ be the natural filtration; it models the information contained in the observations up to time t . We discriminate between two settings represented with two diagrams in Figure 1:

- *Offline* or *Batch* (Section II-A). The model is learned on a training period, for instance on data up to time n_{train} . We estimate $\mathcal{L}(y_t | x_t, \mathcal{F}_{n_{\text{train}}})$, the conditional distribution of y_t given $x_t, \mathcal{F}_{n_{\text{train}}}$.
- *Online* or *Adaptive* (Sections II-B and II-C). The model is learned sequentially. We estimate $\mathcal{L}(y_t | x_t, \mathcal{F}_{t-1})$ at each time t .

A. Offline Model

Motivated by [21], we decompose our model into two steps. We forecast the conditional mean with a generalized additive model (GAM), and then quantile regressions on the GAM residuals yield quantile forecasts.

- 1) We model y_t as a Gaussian random variable whose mean is a sum of the effects of the covariates:

$$y_t = \sum_{j=1}^d f_j(x_{t,j}) + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma^2), \quad (1)$$

where the effects f_1, \dots, f_d are either linear or nonlinear. In the latter case, the effects are built using linear combinations of spline bases [23]. The optimization of

these effects on the training set yields a mean forecast $\hat{y}_t = \mathbb{E}[y_t | x_t, \mathcal{F}_{n_{\text{train}}}]$.

- 2) The Gaussian assumption with fixed variance is violated in practice. Therefore, we fit a set of quantile regressions [24] on the residuals to predict the distribution around the mean. We use an intermediate vector of covariates $z_t \in \mathbb{R}^{d_0}$, $d_0 > 0$ derived from the original vector x_t ; depending on the application, z_t may contain the GAM prediction, the GAM effects $f_j(x_{t,j}) \dots$. For any quantile level q , we define a vector $\beta_q \in \mathbb{R}^{d_0}$ via the following minimization problem:

$$\beta_q \in \arg \min_{\beta \in \mathbb{R}^{d_0}} \sum_{t=1}^{n_{\text{train}}} \rho_q(y_t - \hat{y}_t, \beta^\top z_t), \quad (2)$$

$$\rho_q(y, \hat{y}_q) = (\mathbb{1}_{y < \hat{y}_q} - q) (\hat{y}_q - y). \quad (3)$$

Then, we predict the quantile for probability level q with $\hat{y}_t + \beta_q^\top z_t$. This is justified by the following well-known lemma [25].

Lemma 1. *Let Y be an integrable real-valued random variable. For any $0 < q < 1$ the q -quantile of Y denoted by Y_q satisfies $Y_q \in \arg \min \mathbb{E}[\rho_q(Y, Y_q)]$.*

This two-step procedure is motivated by computational time and by the fact that the data is sparse in the tails (sparser than around the mean) as in [13]. Indeed, we could also use quantile GAM [26], but it is more time-consuming than simple quantile regressions. Hence the need to use simple linear adjustments to correct the (more richly-parametrized) mean GAM. Also, we explain in the following paragraph that the mean GAM may be adapted using the Kalman filter, and that is not the case for quantile models, for which we adopt a different strategy, see II-C.

The first natural adaptive forecaster consists in re-training the offline model at each time step. We call this procedure *incremental offline*.

B. Adaptation of Generalized Additive Model

We adapt the GAM as done by [10]. More precisely, we freeze the nonlinear effects learned on the training set (Section II-A); with these effects, we define a new covariate vector $f(x_t) = (\bar{f}_1(x_{t,1}), \dots, \bar{f}_d(x_{t,d}), 1)^\top$, where $\bar{f}_j(x_{t,j})$ is the standardized version of $f_j(x_{t,j})$. Then we consider the linear Gaussian state-space model:

$$\theta_t - \theta_{t-1} \sim \mathcal{N}(0, Q), \quad (4)$$

$$y_t - \theta_t^\top f(x_t) \sim \mathcal{N}(0, \sigma^2), \quad (5)$$

where Q is the state noise covariance matrix, and σ^2 the space noise variance. In the preceding set of equations, we assume implicitly that the noises are independent. Estimation in a linear Gaussian state-space model with known variances has been optimally solved by [27]:

Theorem 1 (Kalman Filter). *Provided that the data-generating process is the state-space model with variances Q and σ^2 , and if the prior distribution of θ_1 is $\mathcal{N}(\hat{\theta}_1, P_1)$, then the posterior distribution of the state is a Gaussian whose*

mean and covariance matrix have analytical forms. Precisely, we have $\theta_t | \mathcal{F}_{t-1} \sim \mathcal{N}(\hat{\theta}_t, P_t)$ with:

$$P_{t|t} = P_t - \frac{P_t f(x_t) f(x_t)^\top P_t}{f(x_t)^\top P_t f(x_t) + \sigma^2}, \quad (6)$$

$$\hat{\theta}_{t+1} = \hat{\theta}_t - \frac{P_{t|t}}{\sigma^2} \left(f(x_t) (\hat{\theta}_t^\top f(x_t) - y_t) \right), \quad (7)$$

$$P_{t+1} = P_{t|t} + Q. \quad (8)$$

The unsolved issue of linear Gaussian state-space models concerns the estimation of its variances Q and σ^2 . In our setting, as well as in many applications, the variances are unknown. The choice of the variances can be seen as a parametrization of a gradient algorithm. Indeed, $f(x_t) (\hat{\theta}_t^\top f(x_t) - y_t)$ is the gradient of the quadratic loss with respect to θ ; therefore, the update (7) is a gradient step with a *pre-conditioning* matrix $P_{t|t}/\sigma^2$ which depends crucially on the choice of σ^2 and Q . We propose different settings:

- *Static.* We set $Q = 0$ (constant state vector), and $\sigma^2 = 1$. This yields a detailed link with the gradient community [28]; the constant state is equivalent to the i.i.d. assumption of gradient analyses in the setting $P_{t|t} \rightarrow 0$ (annealing step size).
- *Dynamic.* The natural approach aims at maximizing the likelihood; in that vein, we apply an iterative greedy procedure implemented in the R package `viking` [29], that was applied on electricity load forecasting by [10], [11]. It yields a sparse matrix Q assumed diagonal. In that setting, the gradient step does not converge to 0. Note that we apply this procedure to maximize the likelihood on a training set; therefore the variances σ^2 and Q are learned in a *batch* fashion, but the underlying model assumes a time-varying state θ_t leading to an *adaptive* $\hat{\theta}_t$.

C. Adaptation of the Quantile Forecasts

The novelty of the present approach lies in the adaptation of probabilistic forecasts. We compare different methods.

1) *Gaussian Posterior from Kalman:* We first recall that the Kalman filter does not only yield a mean forecast but already a probabilistic forecast. Indeed, at each time step t , Theorem 1 yields $\mathcal{L}(\theta_t | \mathcal{F}_{t-1}) = \mathcal{N}(\hat{\theta}_t, P_t)$. From the state posterior distribution and the observation distribution (5), we deduce

$$\mathcal{L}(y_t | x_t, \mathcal{F}_{t-1}) = \mathcal{N}(\hat{\theta}_t^\top f(x_t), f(x_t)^\top P_t f(x_t) + \sigma^2). \quad (9)$$

It yields a readily computable probabilistic forecast; for any quantile level q , denoting by N_q the quantile of the standard normal distribution, we predict the q -quantile of y_t as

$$\hat{y}_{t,q} = \hat{\theta}_t^\top f(x_t) + N_q \sqrt{f(x_t)^\top P_t f(x_t) + \sigma^2}. \quad (10)$$

However, this first adaptive forecaster would be an adaptive variant of the offline Gaussian distribution on the offline GAM residuals, where only the GAM is adapted but the variance of the Gaussian distribution is still fixed. The quantile regressions were specifically introduced because the Gaussian assumption with fixed variance is violated in practice. We see in the experiments that this Gaussian forecaster is not always well-calibrated.

2) *Quantile Regressions on Kalman Residuals*: We then remark that the offline quantile regressions should work well in the case where the conditional distribution of the residual $y_t - \hat{y}_t$ given the covariates z_t is fixed. The need for adaptive quantile forecasts is motivated only by changes in the residual distribution. Yet a critical property of the Kalman filter is that the residuals are stationary, provided that the state-space model is well-specified. Consequently, the dependence of the Kalman residuals on the quantile covariates should be more stable than that of the offline GAM. Therefore, we combine state-space adaptation of the GAM and offline quantile regressions.

3) *Online Gradient Descent on the Pinball Loss*: We consider an online quantile regression applying online gradient descent (OGD) on the pinball loss. Precisely, we defined in Section II-A the offline quantile regression for the q -quantile with a vector of covariates $z_t \in \mathbb{R}^{d_0}$, $d_0 > 0$:

$$\beta_q \in \arg \min_{\beta \in \mathbb{R}^{d_0}} \sum_{t=1}^{n_{\text{train}}} \rho_q(y_t - \hat{y}_t, \beta^\top z_t). \quad (11)$$

Motivated by the gradient interpretation of the Kalman filter (Section II-B), we apply the OGD to estimate recursively a vector $\beta_{t,q}$. We start from any $\beta_{1,q} \in \mathbb{R}^{d_0}$ and at each step we update it with a step in the direction opposite to the gradient of the loss:

$$\beta_{t+1,q} = \beta_{t,q} - \alpha \left. \frac{\partial \rho_q(y_t - \hat{y}_t, \beta^\top z_t)}{\partial \beta} \right|_{\beta_{t,q}}. \quad (12)$$

The gradient step size α is the important parameter, analogous to the variances that are determining the gradient step in the gradient interpretation of the Kalman filter. We use a constant α in the OGD, and we standardize the covariates z_t . Note that the gradient is not well defined for the singular point $y_t = \hat{y}_t + \beta_{t,q}^\top z_t$; we choose to set it to 0 in the remote case that this happens in practice.

4) *Bernstein Online Aggregation to Choose the Step Size*:

In order to choose the step size in the previous method, we rely on aggregation of experts as proposed by [20] and similarly by [30]. Our procedure is run separately for each quantile level q using the R package `opera` [31]. First, we run the OGD with different possible values $(\alpha_k)_{1 \leq k \leq K}$; at each time t , it yields a forecast $\hat{y}_{t,q}^{(k)}$ for each step size α_k . Second, we combine these forecasts using Bernstein Online Aggregation [19]; the principle of aggregation is to forecast $\hat{y}_{t,q} = \sum_{k=1}^K p_{t,q}^{(k)} \hat{y}_{t,q}^{(k)}$, where the weights $p_{t,q}^{(k)}$ are obtained sequentially. The properties obtained by the online learning literature guarantee that the total pinball loss of the aggregation has a small regret compared with the total pinball loss of the best expert. A similar approach was applied by [18], where the aggregation weights are estimated jointly for all quantiles; the authors assume the weights are smoothed functions of the quantile level and they optimize directly the CRPS instead of point quantile functions.

D. Computational Complexity of the Methodology

The computations of the Kalman filter are separated in two. The estimation of the variances in the dynamic setting is costly (more costly than a single GAM estimation). However, the

updates detailed in Theorem 1 are very efficient and their computation on the whole data set is lighter than a single GAM estimation. On the other hand, the cost of the *incremental offline* GAM is much bigger because each day, we have a GAM estimation. Therefore, the cost of the mean adaptive mean forecast is lower than that of the *incremental offline* GAM, and is concentrated at the calibration of the variances.

Considering our probabilistic forecasts, note that the OGD and the BOA algorithm are also obtained with a single pass on the data, therefore they are much lighter than re-estimating quantile regressions every day.

Numerical values of computational times are given in the supplementary material [32].

III. EVALUATION

For both mean and probabilistic forecasting tasks, we evaluate the forecasts qualitatively as well as quantitatively.

A. Mean Forecast

We evaluate through the root-mean-square-error (RMSE) and the mean absolute error (MAE), defined on a test set \mathcal{T} by

$$\text{RMSE} = \sqrt{\frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} (y_t - \hat{y}_t)^2}, \quad (13)$$

$$\text{MAE} = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} |y_t - \hat{y}_t|. \quad (14)$$

Forecasting the mean of the variable of interest is a way to minimize the quadratic loss, thus the RMSE is natural. The MAE is known to be more robust to outliers.

We don't use the mean absolute percentage error (MAPE) because in the case of the net-load the variable may be close to 0 or even negative, see Section IV-A. However, our data sets are composed of different time series to predict and to obtain global evaluation we use relative metrics. We choose to divide by the same metric applied to the mean.

Formally, we have N time series $(y_{t,i})_{t \in \mathcal{T}, 1 \leq i \leq N}$ and corresponding forecasts $(\hat{y}_{t,i})_{t \in \mathcal{T}, 1 \leq i \leq N}$. We estimate the means $\bar{y}_i = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} y_{t,i}$. Our aggregate metrics are defined by

$$\text{nRMSE} = \sqrt{\frac{1}{N} \sum_{1 \leq i \leq N} \frac{\sum_{t \in \mathcal{T}} (y_{t,i} - \hat{y}_{t,i})^2}{\sum_{t \in \mathcal{T}} (y_{t,i} - \bar{y}_i)^2}}, \quad (15)$$

$$\text{nMAE} = \frac{1}{N} \sum_{1 \leq i \leq N} \frac{\sum_{t \in \mathcal{T}} |y_{t,i} - \hat{y}_{t,i}|}{\sum_{t \in \mathcal{T}} |y_{t,i} - \bar{y}_i|}. \quad (16)$$

Our normalized metrics may be interpreted as unexplained variations, in the opposite way of the R-squared.

B. Probabilistic Forecast

There are different ways to evaluate quantile forecasts. However, a necessary condition for the forecast to be meaningful is *reliability*, also known as *calibration*. A quantile forecast is reliable if the observed frequency of exceedance coincides

with the quantile level. The forecast of a q -quantile is expected to be empirically higher than the quantity of interest for a fraction q of the data set and smaller for a fraction $1 - q$. Therefore, we compare the observed frequencies with the quantiles in reliability diagrams.

Reliable forecasts may then be ranked by sharpness. Numerical evaluation is obtained by the pinball loss ρ_q defined by Equation (3). A way to combine the pinball losses at different quantile levels is to use the continuous ranked probability score (CRPS) [33], defined equivalently by the two following expressions:

$$\begin{aligned} \text{CRPS}(F, y) &= \int_{-\infty}^{+\infty} (F(x) - \mathbb{1}_{y \leq x})^2 dx \\ &= 2 \int_0^1 \rho_q(y, F^{-1}(q)) dq, \end{aligned} \quad (17)$$

where y is the observation and F the predicted cumulative distribution function. Remark that for any finite y we have $\rho_0(y, F^{-1}(0)) = \rho_1(y, F^{-1}(1)) = 0$, therefore the CRPS is not a good performance indicator for the tail estimation.

We use the discrete approximation of the integral in Equation (17). We have a set of forecasted quantiles $\hat{y}_{q_1}, \dots, \hat{y}_{q_l}$, and we define the RPS as the integral of the piecewise linear function interpolating $0, \rho_{q_1}(y, \hat{y}_{q_1}), \dots, \rho_{q_l}(y, \hat{y}_{q_l}), 0$ at the points $0, q_1, \dots, q_l, 1$. This yields

$$\text{RPS}((\hat{y}_{q_1}, \dots, \hat{y}_{q_l}), y) = \sum_{i=1}^l \rho_{q_i}(y, \hat{y}_{q_i})(q_{i+1} - q_{i-1}), \quad (18)$$

where we define $q_0 = 0, q_{l+1} = 1$. Then we simply use the RPS averaged over time on a test set.

As for mean evaluation, we then aggregate this metric over the multiple time series (regions in IV, cities in V). We divide by the mean absolute error of the mean, which is also the CRPS of a Dirac distribution centered at the mean:

$$\text{nRPS} = \frac{1}{N} \sum_{1 \leq i \leq N} \frac{\sum_{t \in \mathcal{T}} \text{RPS}((\hat{y}_{t,i,q_1}, \dots, \hat{y}_{t,i,q_l}), y_{t,i})}{\sum_{t \in \mathcal{T}} |y_{t,i} - \bar{y}_i|}. \quad (19)$$

IV. REGIONAL NET-LOAD IN GREAT BRITAIN

We first study the data set created by [21]. They studied the data from 2014 to 2018, and we augment the period to range from 2014 to 2021 [32], allowing us to integrate the unstable covid period. We refer to [21] and [34] for more details on the data set. Due to Brexit, the day-ahead electricity price definition changed; we decided to remove this explanatory variable from the data set.

A. Data Presentation and Offline Model

We are interested in forecasting the electricity net-load, defined as the difference between electricity consumption and embedded generation (mostly wind and solar production). We consider the regional data from Great Britain. Great Britain is divided into 14 regions called Grid Supply Point Groups.

The time granularity is half an hour and we assume a one-day delay for data availability: each day at midnight we can update the model with data up to 24 hours before. We display in Figure 2 the evolution of the net-load in each region, as well as the daily profiles. We observe different behaviors. In particular, region P (North Scotland) often has a negative net-load, meaning the consumption is smaller than the embedded production.

For each region independently, the model designed by [21] to predict quantiles of the net-load is decomposed into three steps. First, a GAM is fitted to forecast the mean. Second, a set of quantile regressions are fitted to the GAM residuals (between 2.5% and 97.5%). Third, extreme quantiles are modeled via a generalized Pareto distribution. The first two match our framework introduced in Section II-A; we apply the same additive model for the mean, then the same quantile regression.

1) We use the formula of the *GAM Point* of [21], where electricity prices are removed. The normalized electricity net-load is modeled by (1) with the following covariates: a trend, the time of year, the time of day, the day of the week, school holidays, a moving average of the net-load with one-day delay, the temperature at maximum population density, a moving average of that temperature, solar radiation combined with the embedded solar generation capacity, the wind speed combined with the embedded wind generation capacity, and the precipitation.

Weather variables are weather forecasts. All nonlinear effects are built using cubic regression splines.

2) For the quantile regressions, we also use the model of [21]: the GAM residuals are modeled via linear functions of the GAM prediction, the squared prediction, the product of solar radiation and embedded solar capacity, the wind speed, the temperature, as well as of the categorical versions of the time of day and day of the week. For these two latter variables, we have an additive constant defined for each value of the categorical variable.

Note that the GAM and the quantile regressions are not trained per half an hour; they directly incorporate the time of day as a covariate.

B. Performances of Mean Forecast

We display the evolution of the error in Figure 3. We improve the mean forecast performance in almost all regions and all years using state-space adaptation, considering the nRMSE, see Figure 4. We display the aggregate nRMSE and nMAE in Table I for different methods and different years. When we compare the dynamic Kalman filter to the incremental offline GAM (model re-trained each day), we reduce the nRMSE by approximately 4% in 2019, 7% in 2020, and 8% in 2021; furthermore, we have a much lower computational cost per day. Also, we compare to a benchmark named adaptive probabilistic load forecasting (APLF) [12] in Section 2 of our supplementary material [32].




Fig. 2. On the left: evolution of the net-load of the 14 regions. On the right: daily profiles. We observe that in region P (North Scotland) the embedded generation often exceeds the consumption and the daily profile is close to 0. Also, this high level of renewables implies higher weather-driven volatility, and the net-load does not have a clear yearly profile as in the other regions.




Fig. 3. Evolution of the standardized error for the offline model trained on data up to December 31, 2018 (left), the incremental offline model trained each day (middle), and the GAM adapted by the Kalman filter (right).




Fig. 4. nRMSE for each region in Great Britain. We compare the offline GAM to the incremental offline GAM (trained each day on incremental training data), the Kalman filter in the static setting (degenerate covariance matrix of the state noise $Q = 0$), and the Kalman filter in the dynamic setting. We divide the test set in three (2019, 2020 and 2021). For the last two years, we compare also with the GAM re-trained each year: GAM y is the GAM trained with the data up to year $y - 1$ included.

TABLE I
AGGREGATE METRICS ON THE 14 REGIONS FOR THE DIFFERENT MODELS. WE COMPARE TO PERSISTENCE BENCHMARKS (LAGS OF THE VARIABLE OF INTEREST).

Forecast	2019		2020		2021	
	nRMSE	nMAE	nRMSE	nMAE	nRMSE	nMAE
Persistence (7 days)	0.691	0.589	0.710	0.599	0.737	0.639
Persistence (2 days)	0.767	0.686	0.755	0.668	0.736	0.668
Offline GAM	0.356	0.327	0.485	0.453	0.635	0.601
Incremental offline GAM (yearly)	-	-	0.407	0.376	0.387	0.378
Incremental offline GAM (daily)	0.338	0.307	0.370	0.344	0.377	0.365
Kalman GAM (Static)	0.337	0.307	0.374	0.347	0.380	0.368
Kalman GAM (Dynamic)	0.324	0.292	0.328	0.301	0.332	0.307

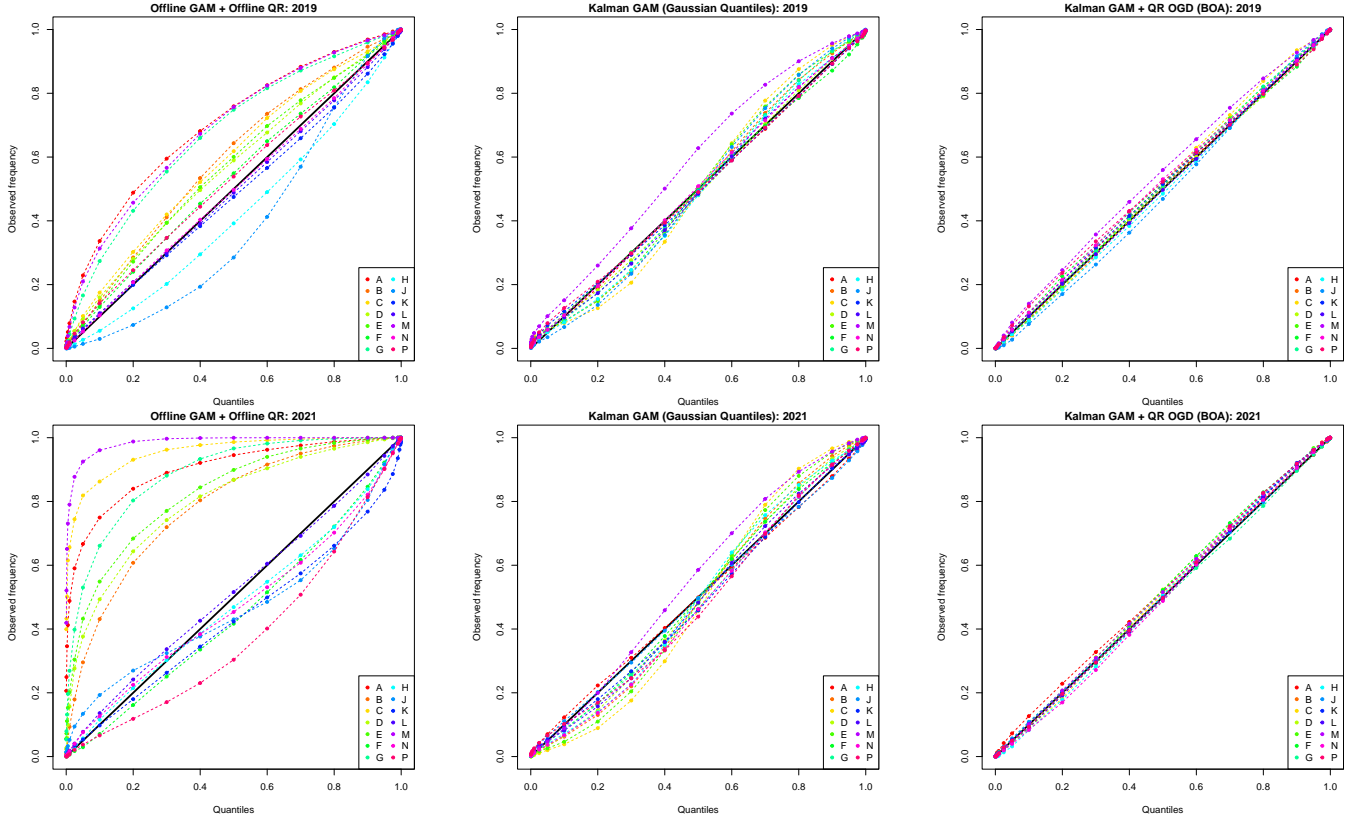


Fig. 5. Reliability diagrams for different methods in 2019 and in 2021.

C. Performances of Quantile Forecast

We display reliability diagrams in Figure 5. We observe that the offline model is not reliable, probably because of the biases in the mean model. The adaptation of the mean model by the Kalman filter yields Gaussian quantiles that are much more reliable. Reliability is then improved further by the adaptation of the quantile regressions. Similarly to the evolution of the bias of the model (Figure 3), we display in Figure 6 the evolution of the observed frequency for a specific quantile. Our adaptive model has a much better evolution of calibration in time than the offline version.

We provide in Table II the RPS averaged on the 14 regions for the different models. We obtain an important gain in nRPS by adapting the GAM using the Kalman filter and keeping an offline quantile regression, even for the stable period (the year 2019). Adapting the quantile regression with an OGD has a very tenuous difference. The step size of this OGD may

TABLE II
AGGREGATE NRPS ON THE 14 REGIONS FOR THE DIFFERENT MODELS.

	2019	2020	2021
Offline Method	0.231	0.338	0.454
GAM Kalman (Gaussian Quantiles)	0.212	0.217	0.222
GAM Kalman + Offline QR	0.206	0.214	0.217
Offline GAM + QR OGD (10^{-3})	0.218	0.270	0.293
Offline GAM + QR OGD (10^{-2})	0.207	0.221	0.218
Offline GAM + QR OGD (10^{-1})	0.250	0.248	0.293
Offline GAM + QR OGD (BOA)	0.204	0.211	0.216
GAM Kalman + QR OGD (10^{-2})	0.205	0.204	0.212
GAM Kalman + QR OGD (BOA)	0.202	0.201	0.209

be selected by Bernstein Online Aggregation, combining step sizes in $\{10^i, -8 \leq i \leq 0\}$. This choice of step size does not only perform as well as the best expert as suggested in Section II-C4, but it outperforms this oracle. Combining both levels of adaptation yields the best performances.




Fig. 6. Evolution of the calibration. We select a specific time of day (midday) and a quantile level 0.8; we display the observed frequency that the net-load exceeds the quantile forecast. The 95% confidence interval is based on the assumption that the residuals are iid. We see that for the offline model (GAM then quantile regression) the confidence interval is not satisfied.




Fig. 7. Weights optimized by BOA for the forecast of region A in Great Britain at quantile level 0.1.

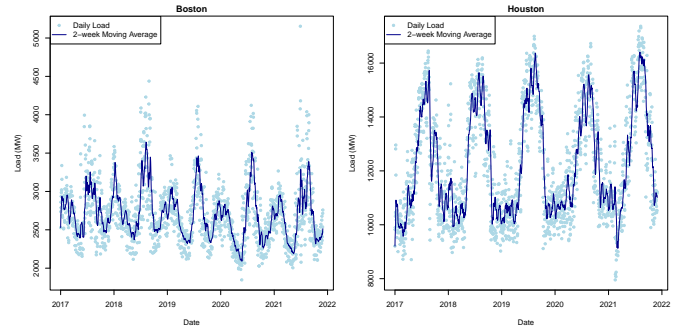


Fig. 8. Daily electricity load in Boston (left) and Houston (right). The consumption has two annual peaks in summer and winter in Boston, whereas there is only a summer peak in Houston.

D. Learning the Embedded Capacities

A remarkable advantage of model adaptation is that it reduces the need for good explanatory variables, which may be difficult to obtain. In particular, we show that our Kalman GAM can *learn* the embedded generation capacities. That is removing these variables does not significantly change the predictions, contrary to the offline method.

More precisely, in the GAM and in the quantile regressions presented in Section IV-A, we remove the solar and wind generation capacities. We consider solar radiation only instead of its product with solar capacity. In the nonlinear GAM effect of the wind speed, we remove the dependence on the embedded wind capacity. We thus evaluate removing the capacities in the offline model, as well as in the Kalman adaptation of the GAM. The evolution of the capacities may be captured by the adaptation of the state coefficient.

We evaluate during 2019 to not include the coronavirus crisis. For mean prediction, removing the capacities increases the nRMSE of the offline method by more than 10%, while it reduces the nRMSE by 0.4% for the Kalman adaptation of the GAM. For probabilistic prediction, removing the capacities increases the nRPS by more than 10% for the offline model and reduces it by 0.02% for the offline quantile regression on the residuals of GAM Kalman.

V. CITY-WIDE LOAD IN THE UNITED STATES

We test our framework also on US data during the coronavirus crisis [22]. To demonstrate the generality of the approach, we consider a very different setting. We forecast the standard load at a daily granularity and with no delay in data availability. Each day, we can adapt on the data up to the current day in order to forecast the next.

A. Data Presentation and Offline Model

The data set covers seven U.S. electricity markets, and the major city for each one: Boston, Chicago, Houston, Kansas City, Los Angeles, New York City, Philadelphia. Consider the city-wide loads shown in Figure 8 and note that load patterns are different across cities.

We follow the forecasting structure presented in Section II.

1) The generalized additive model we consider has the following form:

$$\begin{aligned} Load_t = & \sum_{i=1}^7 \alpha_i \mathbb{1}_{WeekDay_t=i} + \beta_1 B H_t + \beta_2 W B_t \\ & + \beta_3 LoadD_t + f_1(LoadW_t) + f_2(t) \quad (20) \\ & + f_3(Temp_t) + f_4(Hum_t) + f_5(Toy_t) + \varepsilon_t, \end{aligned}$$

TABLE III
AGGREGATE METRICS IN THE SEVEN US CITIES.

Forecast	2020		2021	
	nRMSE	nMAE	nRMSE	nMAE
Persistence (7 days)	0.777	0.688	0.852	0.745
Persistence (1 day)	0.455	0.417	0.464	0.414
GAM	0.407	0.400	0.666	0.616
Incremental GAM (daily)	0.223	0.213	0.214	0.185
GAM KF Static	0.206	0.195	0.204	0.178
GAM KF Dynamic	0.194	0.168	0.198	0.166

TABLE IV
AGGREGATE NRPS IN THE SEVEN US CITIES.

Forecast	2020	2021
Offline method	0.319	0.539
GAM Kalman (Gaussian Quantiles)	0.122	0.123
GAM Kalman + Offline QR	0.230	0.391
Offline GAM + QR OGD (10^{-4})	0.232	0.157
Offline GAM + QR OGD (10^{-3})	0.145	0.109
Offline GAM + QR OGD (10^{-2})	0.109	0.196
Offline GAM + QR OGD (10^{-1})	0.612	1.75
Offline GAM + QR OGD (BOA)	0.099	0.097
GAM Kalman + QR OGD (BOA)	0.094	0.094

where ε_t is a Gaussian i.i.d. noise and for each day t ,

- $WeekDay_t$ is the day of the week,
- BH_t (respectively WB_t) is a Boolean denoting if the day is a bank holiday (respectively in the winter break),
- $LoadD_t$ and $LoadW_t$ are lags of the load with a one-day and one-week delays,
- t is the day (variable growing linearly with time),
- $Temp_t$ and Hum_t are the temperature and humidity,
- Toy_t is the time of year (variable growing linearly from 0 on January 1st to 1 on December 31st).

The nonlinear effects f_1, f_2, f_3, f_4, f_5 are built using thin-plate splines for the first four, and cubic cyclic splines for f_5 .

- 2) The quantile regression we fit on the residuals uses as covariates the effects of the GAM, the mean prediction, and the square of the mean prediction as in Section IV-A:

$$z_t = \left(\hat{y}_t, \hat{y}_t^2, \sum_{i=1}^7 \alpha_i \mathbb{1}_{WeekDay_t=i}, \beta_1 BH_t, \beta_2 WB_t, \beta_3 LoadD_t, f_1(LoadW_t), f_2(t), f_3(Temp_t), f_4(Hum_t), f_5(Toy_t), 1 \right)^\top. \quad (21)$$

B. Evaluation

We display the nRMSE and nMAE aggregated on the seven cities in Table III. The nRPS is displayed in Table IV.

The best models are still the dynamic setting of the Kalman filter for mean forecasting and the combination of Kalman adaptation of GAM with the adaptive quantile regression for the probabilistic task. However, the performances are significantly different from the results of Section IV. In particular, adapting the mean model with Kalman and then applying offline quantile regression models behaves poorly. It is crucial to adapt the quantile models.

VI. CONCLUSION

In this article, we applied an adaptive procedure for probabilistic net-load forecasting. The proposed methodology relies on several steps: an adaptive mean forecast is obtained by the Kalman filter, and adaptive quantile regressions are derived from Online Gradient Descent. To solve the choice of the step size in the OGD, the algorithm is estimated with multiple learning rates and the estimates are combined using Bernstein Online Aggregation.

Evaluation has been performed based on the electricity net-load in the fourteen regions of Great Britain, and on electricity demand in seven major cities in the United States. Online adaptation of the mean model is always important, yielding reductions in RMSE of around 10% in all case studies compared to updating the model periodically. Additional adaptation of quantile regressions added marginal benefits for regional load forecasting in GB, but substantial benefits for US cities, reducing the RPS by over 20%. This may be due to the fact that the coronavirus crisis has had more impact on daily lives in big cities. Therefore, in our applications, adaptive methods have more interest for big cities. Indeed, we also observe a more substantial gain for region C (London) in the GB data set.

A natural extension would be the adaptation of multivariate probabilistic forecasts. Previous works have considered adaptive dependency structures, but how models for marginal distributions and dependency structures could (and should) be adapted simultaneously has not been studied. Also, we focused on non-extreme quantiles and adaptive estimation of extreme values is an exciting and very challenging topic. A special treatment for extremes may be useful to refine reliability as well as numerical performance.

REFERENCES

- [1] T. Hong, P. Pinson, Y. Wang, R. Weron, D. Yang, and H. Zareipour, "Energy forecasting: A review and outlook," *IEEE Open Access Journal of Power and Energy*, vol. 7, pp. 376–388, 2020.
- [2] T. Hong, P. Pinson, S. Fan, H. Zareipour, A. Troccoli, and R. J. Hyndman, "Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond," *International Journal of Forecasting*, vol. 32, no. 3, pp. 896–913, 2016.
- [3] T. Hong, J. Xie, and J. Black, "Global energy forecasting competition 2017: Hierarchical probabilistic load forecasting," *International Journal of Forecasting*, vol. 35, no. 4, pp. 1389–1399, 2019.
- [4] U.S. Department of Energy. Net load forecasting prize. [Online]. Available: <https://www.energy.gov/eere/solar/american-made-net-load-forecasting-prize>
- [5] P. IEA, "Covid-19 impact on electricity," Technical report, Tech. Rep., 2021.
- [6] —, "World energy outlook 2022," Technical report, Tech. Rep., 2022.
- [7] I. Calero, C. A. Cañizares, K. Bhattacharya, and R. Baldick, "Duck-curve mitigation in power grids with high penetration of pv generation," *IEEE Transactions on Smart Grid*, vol. 13, no. 1, pp. 314–329, 2022.
- [8] J. W. Messner and P. Pinson, "Online adaptive lasso estimation in vector autoregressive models for high dimensional wind power forecasting," *International Journal of Forecasting*, vol. 35, no. 4, pp. 1485–1498, 2019.
- [9] C. Di Modica, P. Pinson, and S. B. Taieb, "Online forecast reconciliation in wind power prediction," *Electric Power Systems Research*, vol. 190, p. 106637, 2021.
- [10] P. Obst, J. de Vilmarest, and Y. Goude, "Adaptive methods for short-term electricity load forecasting during covid-19 lockdown in france," *IEEE Transactions on Power Systems*, vol. 36, no. 5, pp. 4754–4763, 2021.

- [11] J. de Vilmaest and Y. Goude, "State-space models for online post-covid electricity load forecasting competition," *IEEE Open Access Journal of Power and Energy*, 2022.
- [12] V. Álvarez, S. Mazuelas, and J. A. Lozano, "Probabilistic load forecasting based on adaptive online learning," *IEEE Transactions on Power Systems*, vol. 36, no. 4, pp. 3668–3680, 2021.
- [13] P. Gaillard, Y. Goude, and R. Nedellec, "Additive models and robust aggregation for gecom2014 probabilistic electric load and electricity price forecasting," *International Journal of forecasting*, vol. 32, no. 3, pp. 1038–1050, 2016.
- [14] M. Fasiolo, S. N. Wood, M. Zaffran, R. Nedellec, and Y. Goude, "Fast calibrated additive quantile regression," *Journal of the American Statistical Association*, vol. 116, no. 535, pp. 1402–1412, mar 2020.
- [15] C. Gilbert, J. Browell, and B. Stephen, "Probabilistic load forecasting for the low voltage network: Forecast fusion and daily peaks," *Sustainable Energy, Grids and Networks*, vol. 34, p. 100998, 2023.
- [16] B. Wang, M. Mazhari, and C. Y. Chung, "A novel hybrid method for short-term probabilistic load forecasting in distribution networks," *IEEE Transactions on Smart Grid*, vol. 13, no. 5, pp. 3650–3661, 2022.
- [17] P. Arora, A. Khosravi, B. K. Panigrahi, and P. N. Suganthan, "Remodelling state-space prediction with deep neural networks for probabilistic load forecasting," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 6, pp. 628–637, 2022.
- [18] J. Berrisch and F. Ziel, "Crps learning," *Journal of Econometrics*, 2021.
- [19] O. Wintenberger, "Optimal learning with bernstein online aggregation," *Machine Learning*, vol. 106, no. 1, pp. 119–141, 2017.
- [20] M. Zaffran, O. Féron, Y. Goude, J. Josse, and A. Dieuleveut, "Adaptive conformal predictions for time series," in *International Conference on Machine Learning*. PMLR, 2022, pp. 25 834–25 866.
- [21] J. Browell and M. Fasiolo, "Probabilistic forecasting of regional net-load with conditional extremes and gridded nwp," *IEEE Transactions on Smart Grid*, vol. 12, no. 6, pp. 5011–5019, 2021.
- [22] G. Ruan, D. Wu, X. Zheng, H. Zhong, C. Kang, M. A. Dahleh, S. Sivaranjani, and L. Xie, "A cross-domain approach to analyzing the short-run impact of covid-19 on the us electricity sector," *Joule*, vol. 4, no. 11, pp. 2322–2337, 2020.
- [23] S. N. Wood, *Generalized additive models: an introduction with R*. CRC press, 2017.
- [24] R. Koenker and G. Bassett Jr, "Regression quantiles," *Econometrica: journal of the Econometric Society*, pp. 33–50, 1978.
- [25] R. Koenker and K. F. Hallock, "Quantile regression," *Journal of economic perspectives*, vol. 15, no. 4, pp. 143–156, 2001.
- [26] M. Fasiolo, S. N. Wood, M. Zaffran, R. Nedellec, and Y. Goude, "Fast calibrated additive quantile regression," *Journal of the American Statistical Association*, vol. 116, no. 535, pp. 1402–1412, 2021.
- [27] R. E. Kalman and R. S. Bucy, "New results in linear filtering and prediction theory," *Journal of Basic Engineering*, vol. 83, no. 1, pp. 95–108, 1961.
- [28] J. de Vilmaest and O. Wintenberger, "Stochastic online optimization using kalman recursion," *J. Mach. Learn. Res.*, vol. 22, pp. 223–1, 2021.
- [29] J. de Vilmaest, *viking: State-Space Models Inference by Kalman or Viking*, 2022, r package version 1.0.0. [Online]. Available: <https://CRAN.R-project.org/package=viking>
- [30] T. van Erven, W. M. Koolen, and D. van der Hoeven, "Metagrad: Adaptation using multiple learning rates in online learning," *Journal of Machine Learning Research*, vol. 22, no. 161, pp. 1–61, 2021. [Online]. Available: <http://jmlr.org/papers/v22/20-1444.html>
- [31] P. Gaillard, Y. Goude, L. Plagne, T. Dubois, and B. Thieurmel, *opera: Online Prediction by Expert Aggregation*, 2021, r package version 1.2.0. [Online]. Available: <https://CRAN.R-project.org/package=opera>
- [32] J. de Vilmaest, J. Browell, M. Fasiolo, Y. Goude, and O. Wintenberger, "Supplementary materials for "Adaptive Probabilistic Forecasting of Electricity (Net-)Load"," Dec. 2022. [Online]. Available: <https://doi.org/10.5281/zenodo.7849665>
- [33] T. Gneiting and A. E. Raftery, "Strictly proper scoring rules, prediction, and estimation," *Journal of the American statistical Association*, vol. 102, no. 477, pp. 359–378, 2007.
- [34] J. Browell, "Supplementary Material for "Probabilistic Forecasting of Regional Net-load with Conditional Extremes and Gridded NWP"," 2021. [Online]. Available: <https://doi.org/10.5281/zenodo.5031704>



Joseph de Vilmaest graduated from École Normale Supérieure de Paris in 2019, then received his Ph.D. degree in Statistics from Sorbonne Université in 2022. He founded *Viking Conseil* to develop adaptive methods for time series forecasting, and to apply them in companies that depend highly on forecasting, especially in the energy sector.



Jethro Browell (Senior Member, IEEE) received a Ph.D. in wind energy systems from the University of Strathclyde in 2015 and is presently a Senior Lecturer in the University of Glasgow's School of Mathematics and Statistics. His research interests span all aspects of data analytics and forecasting with a focus on applications in the energy sector. Dr. Browell is also Chair of the IEEE Power & Energy Society Working Group on Energy Forecasting and Analytics, an editor for *Sustainable Energy, Grids and Networks* and associate editor for *IEEE Access*.



Matteo Fasiolo received a Ph.D. degree in Statistics from the University of Bath, U.K., in 2016. He is a Senior Lecturer in the University of Bristol's Institute for Statistical Science, where he works on developing non-parametric regression modelling methodology, with a particular focus on energy-related applications.



Yannig Goude is senior researcher in the field of Data Science at EDF R&D and associate professor at Université Paris-Sud Orsay, France. He obtained his Ph.D. in statistics and probability in 2008 at Université Paris-Sud Orsay. His research interests are time series forecasting for electricity markets, time series analysis, non-parametric models and aggregation of experts.



Olivier Wintenberger is a professor at Sorbonne University and WPI fellow at the University of Vienna. He obtained his Ph.D. in 2007 at Panthéon-Sorbonne University. His research interests are extreme value analysis, time series forecasting, and online learning. He is an associate editor for the journals *Bernoulli*, *Extremes*, and *Stochastic Models*.