

Hu, Junjie; López Cabrera, Brenda; Melzer, Awdesch

Working Paper

Advanced statistical learning on short term load process forecasting

IRTG 1792 Discussion Paper, No. 2021-020

Provided in Cooperation with:

Humboldt University Berlin, International Research Training Group 1792 "High Dimensional Nonstationary Time Series"

Suggested Citation: Hu, Junjie; López Cabrera, Brenda; Melzer, Awdesch (2021) : Advanced statistical learning on short term load process forecasting, IRTG 1792 Discussion Paper, No. 2021-020, Humboldt-Universität zu Berlin, International Research Training Group 1792 "High Dimensional Nonstationary Time Series", Berlin

This Version is available at:

<https://hdl.handle.net/10419/246491>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Advanced Statistical Learning on Short Term Load Process Forecasting

Junjie Hu *

Brenda López Cabrera *

Awdesch Melzer *



* Humboldt-Universität zu Berlin, Germany

International Research Training Group 1792

This research was supported by the Deutsche
Forschungsgesellschaft through the
International Research Training Group 1792
"High Dimensional Nonstationary Time Series".

Advanced Statistical Learning on Short Term Load Process Forecasting

Junjie Hu, Brenda López Cabrera, Awdesch Melzer
Humboldt-Universität zu Berlin

03.10.2021

Abstract

Short Term Load Forecast (STLF) is necessary for effective scheduling, operation optimization trading, and decision-making for electricity consumers. Modern and efficient machine learning methods are recalled nowadays to manage complicated structural big datasets, which are characterized by having a nonlinear temporal dependence structure. We propose different statistical nonlinear models to manage these challenges of hard type datasets and forecast 15-min frequency electricity load up to 2-days ahead. We show that the Long-short Term Memory (LSTM) and the Gated Recurrent Unit (GRU) models applied to the production line of a chemical production facility outperform several other predictive models in terms of out-of-sample forecasting accuracy by the Diebold-Mariano (DM) test with several metrics. The predictive information is fundamental for the risk and production management of electricity consumers.

JEL classification: C51, C52, C53, Q31, Q41.

Keywords: Short Term Load Forecast, Deep Neural Network, Hard Structure Load Process

1. Introduction

Accurate short term load forecast (STLF) is the key component in energy management, and more and more energy-intensive companies have been demanding new methodologies to reduce their daily costs. This paper is a collaboration between academic research and a real-world case from the industry. So far, much of the literature studies the short term load forecast dealing only with the cases of the aggregated level of electricity load. To our knowledge, there is no study on forecasting the machine-scale electricity consumption, and we propose linear/nonlinear models on a load process generated from a chemical facility. We model this complicated load process by adjusting the restriction on parameters using three different methods, including the classic parametric AR-type model, a semi-parametric model, and state-of-the-art machine learning models.

Short term load forecast on small-scale energy consumers, such as a machine or a household, is not commonly seen in previous literature but getting crucial nowadays. Unlike the load processes in aggregated levels which are often characterized by certain seasonality, the electricity load from a small-scale consumer always evolves with abrupt jumps and incorporates noise. Therefore, the pattern of such a load process is almost unrecognizable. Nevertheless, forecast on the small-scale consumer is demanded in the industry, notwithstanding the difficulties of such modeling. Companies can better optimize their production schedules with the knowledge about the energy consumption of each one of their machines, and retail consumers can also implement the forecasting models as the energy trading market is now expanding to households in many countries thanks to the rapid development of demand-side management (Weber et al. (2017)). This study fits the gap and discusses the predictability of such load processes.

The load process modeled in this paper is characterized by having a complicated structure. More specifically, the dataset we used is filled with jumps or structure breaks and various types of noises. The jumps are generally modeled as a càdlàg process, i.e., the jumps are not predictable without extra information. Such a structure is a typical situation of modeling small-scale energy consumers where exogenous shocks become dominant. Papers discuss the forecasts of the household electricity consumption as the development of smart grids in recent years extensively, for example, Amara et al. (2017) propose an adaptive model to forecast the conditional density of household electricity demand.

Load forecast is a major research topic in energy fields. In terms of time-scale, the literature can be divided into the long term forecast, e.g., monthly load forecast (Chang et al. (2011)), and the short term forecast, which is our focus. In terms of individual-scale, the existing literature mainly studies load process aggregated in a large-scale, such as the

state-level (Bianco et al. (2009); Darbellay and Slama (2000); Do et al. (2016); Oğcu et al. (2012); Castelli et al. (2015); Pardo et al. (2002)). All the load processes in those articles tend to have clear patterns of seasonalities. We obtained a unique dataset consisted of an intraday electric load process with the corresponding production schedules from a chemical facility. On each day, we forecast a 15-minute frequency load process in the following two days.

In our case, one question is how much the recent advanced models can outperform the classic time-series models. One can find literature applying neural network models on STL, such as An et al. (2013); Bianchi et al. (2017); Guan et al. (2012) concluding that modeling nonlinearity does improve the forecasting accuracy. However, the question remains in our case that does not appear with clear cyclic patterns. Hence a classic mode decomposition or noise filtering method is unlikely useful. Many other different approaches have been applied to the STL. A strand of literature is motivated to describe the seasonality with the single/multiple equation(s) models or functional data analysis (Taylor (2003); Cottet and Smith (2003); Bianchi et al. (2017); Clements et al. (2016); López Cabrera and Schulz (2017)). Most of those models are carefully designed and calibrated. However, they do not fit in our case as they are restricted to certain fixed seasonal patterns. Another angle to improve the forecasting accuracy is by accounting for the nonlinear dependence of the load processes using semi-parametric or non-parametric models. We show that the Long-short Term Memory (LSTM) and the Gated Recurrent Unit (GRU) architectures outperform a semi-parametric functional data method named FActorisable Sparse Tail Event Curves (FASTEC) (Chao et al. (2015)) and the ARX model. The forecast does show a consistent result with the literature. However, LSTM or GRU models should be carefully used since it losses forecasting power on tails or jumps, which we believe is the key conclusion that should be emphasized for many practitioners, i.e., the machine learning is *NOT* omnipotent.

The remainder of this paper is organized as follows. Section 2 gives an overview of the Data of a chemical production facility and conducts the empirical analysis of the high frequency data. In section 3, we discussed different forecast models, statistical and machine learning models, and evaluate their performance. Section 4 concludes and make suggestions for further research. All computations for this study were carried out in Python.

2. Data and stylised facts

The dataset is provided by the consulting company Ökotec Energiemanagement GmbH and contains the total electricity consumption and 72 production schedules of a product line from one chemical production facility of an anonymous chemical company in Germany. The

time series ranges from 1st January 2017 to 31st December 2017 and contains 96 observations daily, i.e., a 15-minute resolution, and the 72 production schedules can be divided into a set of 3 major groups of production schedules. Moreover, the production schedules are lagging for 2 hours, meaning that the information about the production schedules is saved 2 hours after the actual production process has started. Hence, the electricity consumption process is matched at the 8th lag with all the production schedules under the assumption that production schedules are known contemporaneously.

The load process is scaled into $[0,1]$ range, shown in Figure 1 in which we can directly see that the frequent jumps on the load process are caused by switching between different products or fully-loaded/minimum-loaded. Hence, then we can see that machinery is working at different states at different periods.

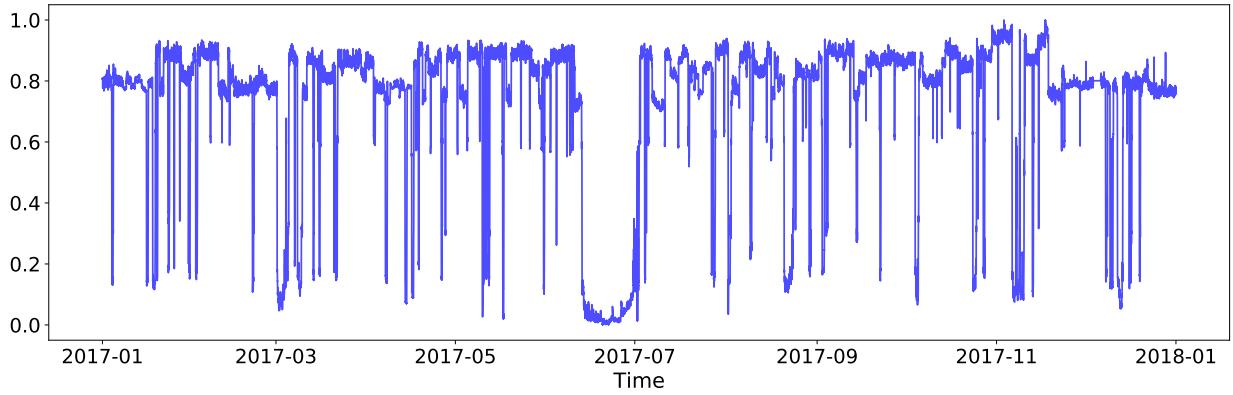


Figure 1
Electricity Load Process

The figure shows the 15-minute frequency electricity consumption process from 1.Jan.2017 to 31.Dec.2017. The load process is scaled down into $[0,1]$.

Another issue of the data is the dimensionality. Having detailed information about the production schedule of all 72 product lines is informative for the producer but leads to inefficiency and inconsistent estimators in a time series regression model due to the large coefficient space. There exist several methods that deal with the problem. Regularized regressions techniques, such as elastic net, scad, ridge regression, and LASSO, avoid this curse of dimensionality and shrink the parameter space by using data aggregation. However, they can be slow and not performing well compared to ARMA or linear regression models. Moreover, given rigid server-time constraints by the consulting company, our production schedules had to be burned down to an aggregate of three groups. Large-scale energy consumers have to give their bids and orders at the energy market every day until 2 pm. By 10 am the current data is fed to the consulting company.

(Weak) Stationarity is one of the prerequisites for econometric models calibration, anal-

ysis, and inference. We found the load process interesting as two of the most commonly used stationarity tests, the Augmented Dickey-Fuller (ADF, Dickey and Fuller (1979, 1981)) test and Kwiatkowski–Phillips–Schmidt–Shin (KPSS, Kwiatkowski et al. (1992)) test, tell differently.

Table 1
Stationarity Test

	test	p-value	lags
<i>ADF</i>	−10.79	0.00	51
<i>KPSS</i>	0.83	0.01	52

The table reports the results of two stationarity tests, ADF and KPSS, on the electricity load process. Both tests suggest rejecting their null hypotheses, implying a conflicted conclusion from the stationarity test.

The null hypothesis of ADF is unit root exists in the process against the alternative that such process is (trend) stationary. On the other hand, the null hypothesis of KPSS is that the observed process is (trend) stationary against the unit root alternative. The results from Table 1 reject both null hypothesis under the 1% significant level. We may conclude that the load process is neither stationary nor unit root process. Lack of stationarity could cause most of the estimation for statistical inference to be biased or inconsistent. The reason for such contradictory might be a significant amount of structural breaks in the process, in which case the assumptions for tests do not hold. On the other hand, the process is by default stationary, as the maximum electricity consumption level of that single machinery cannot increase over time.

In this article, we do not model the structural break explicitly. Modeling structural breaks depend heavily on detection, which is most likely to be ex-post methods. Literature shows that structural break doesn't affect forecast result significantly per se (Stock and Watson (1996)). Also, the benefits of modeling structural break could be canceled out by the estimation errors of break time points and parameters (Elliott and Müller (2007, 2014)). The commonly used forecast metric, Mean Squared Error (MSE), is always facing a bias-variance trade-off which means ignoring structural break could bring even more accurate forecast (Pesaran and Timmermann (2005)).

From Figure 2, we can see that the load is mixture-Gaussian distributed, implying that the load process consists of a certain number of states and stationary process within each state. The load process is characterized by a slow decay of ACF (Long-memory property), which enables us to forecast with lagged load values, and therefore the classic AR-type models can be applied.

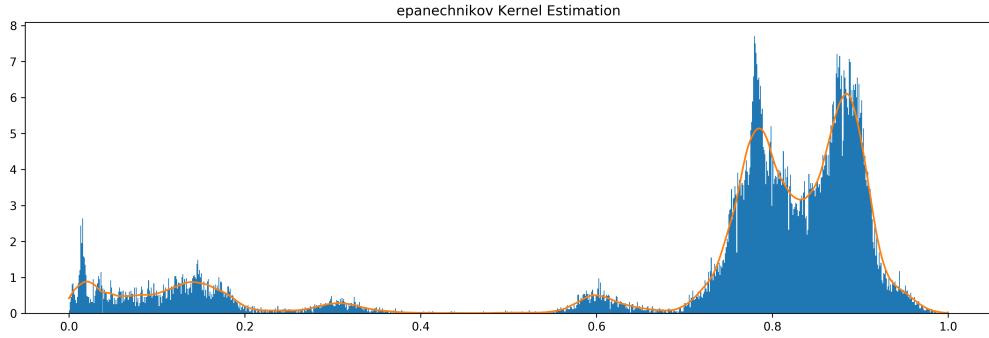


Figure 2

KDE of the Electricity Load Process

Kernel density estimation of full sample electricity consumption values. Epanechnikov kernel with bandwidth 1.6%.

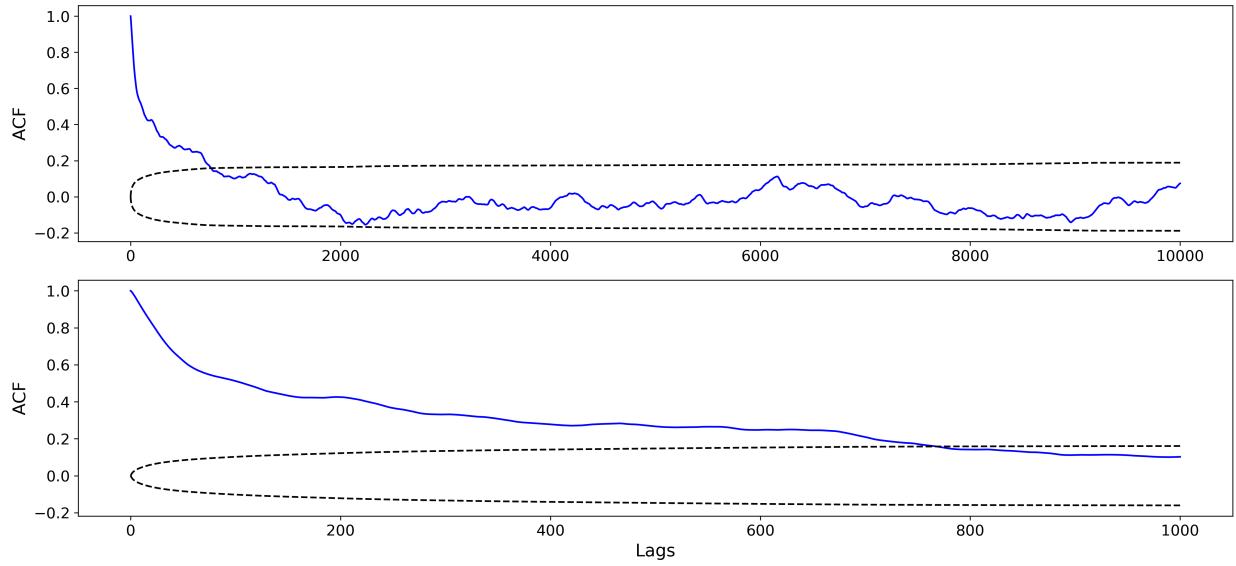


Figure 3

ACF of the Electricity Load Process

The figure shows the Autocorrelation function of the load process. The upper panel shows 10000 maximum lags; the bottom panel has 1000 maximum lags. The solid line is the autocoefficient on different lags, and the dashed line forms the confidence interval at a 5% significant level.

Apart from the load process, this dataset also includes 72 sub-product schedules. There are three major products manufactured in this production line, and each schedule corresponds to one of those three products. Therefore, we divide those schedules into three groups as shown in Figure 4. For further investigation, we aggregate those sub-product schedules into three groups as exogenous variables that can be considered as a relative volume for three independent products.

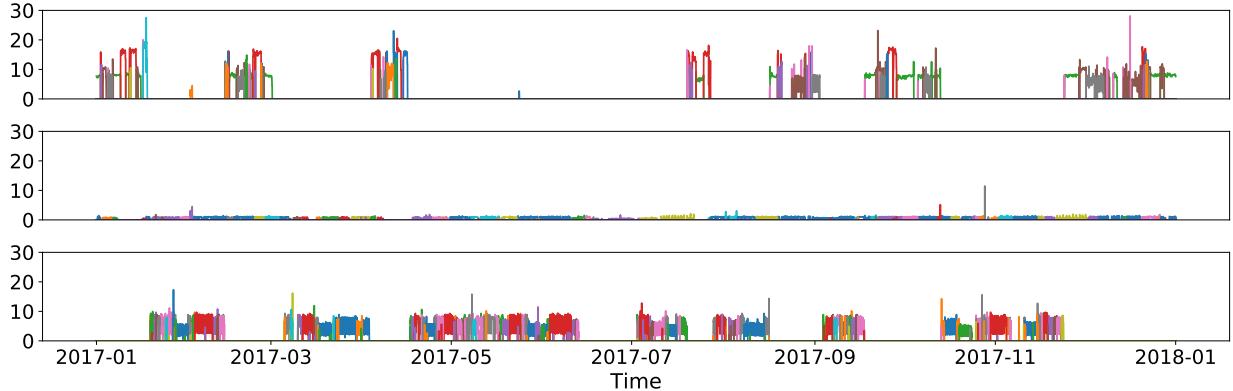


Figure 4

The 3-Group of Production Schedules

The figures show the production schedules grouped into three groups (from top to bottom) based on the products being manufactured. The production schedules are recorded 2-hour late after the actual production. We shift the production schedules 2-hour earlier, assuming that the schedules are known as soon as they start.

3. Forecasting Models

In this section, we first construct and justify the advantages of each forecasting model, followed by analyzing the forecast results with different metrics. Three approaches are taken in this article, including linear parametric model, semi-parametric model, and non-parametric model. The non-parametric approach, neural network models LSTM and GRU, are of our focus to capture the nonlinear temporal dependency of the process. In addition, a functional data model "FASTEC" combined with the regime-switching method detecting different states at each time t is also constructed. Most of the load forecasting purposes are related to trading on the day-ahead energy market, which requires forecasts at least 36-hours ahead. Hence, the forecasting horizon of both models is 192-steps (2-day, 15-minute per step) ahead. In each day t , $t = 1, 2, \dots, T$, The forecasting task can be generally written as,

$$Y_{t+2} = g(\mathbf{X}_t; \Theta), \quad (1)$$

where the forecasting load with 192 steps, $Y_{t+2} \stackrel{\text{def}}{=} \{y_{s,t+1}, y_{s,t+2}\}_{s=1}^S$, $y_{s,t+h} \in \mathbb{R}$, $h = 1, 2$; and the predictor variables $\mathbf{X}_t \stackrel{\text{def}}{=} \{X_{s,t}\}_{s=1}^S$ with feature vector $X_{s,t} \in \mathbb{R}^k$, and $k = 4$. S is fixed as 96 for the 15-minute sampling frequency each day.

For different forecasting models, function g is chosen within a different class of functions, and the parameter Θ is calibrated accordingly. We demonstrate how to prepare the forecasting samples with the following algorithm 1 representation:

Algorithm 1: Forecasting Sample Construction

Input : Production schedules variables: $\{Z_{s,t}\}_{s=1}^{96} \in \mathbb{R}^{72}$; Load process: $\{l_{s,t}\}_{s=1}^{96}$;
Time step s is 15-minute frequency, $t = 1, 2, \dots, T$
Output: Training Set: $\mathcal{S}^{train} = \{(X_{s,t}, y_{s,t+h})\}_{s=1}^{96}, h = 1, 2$, First 7-month sample t
Test Set: $\mathcal{S}^{test} = \{(X_{s,t}, y_{s,t+h})\}_{s=1}^{96}, h = 1, 2$, Last 5-month sample t
1 Aggregate $\{Z_{s,t}\}_{s=1}^{96} \in \mathbb{R}^{72}$ into 3 groups, $\{x_{s,t}^1, x_{s,t}^2, x_{s,t}^3\}_{s=1}^{96}, t = 1, 2, \dots, T$;
2 Predictor Variables: $\{X_{s,t}\}_{s=1}^{96} = \{x_{s,t}^1, x_{s,t}^2, x_{s,t}^3, l_{s,t}\}_{s=1}^{96}, t = 1, 2, \dots, T$;
3 Forecast Load: $\{y_{p,t}\}_{p=1}^{192}, t = 1, 2, \dots, T$;
4 Sample Pairs: $\{(X_{s,t}, y_{s,t+h})\}_{s=1}^{96}, h = 1, 2, t = 1, 2, \dots, T$;
5 Split sample pairs into a training set (first 7-month) and test set (last 5-month)

3.1. Neural Network Models

Nonlinearity has been the ignored truth in statistics for a long time due to the interpretability of many nonlinear models. However, on the other hand, interpretation of the linear parametric models can sometimes be misleading when models are deviating too far away from the true model, likely caused by the poor flexibility of parametric models. To capture the nonlinear dependency of the load process, we exploit the power of machine learning techniques, especially Recurrent Neural Networks (RNN).

RNN is essentially the modified Neural Network (NN) architecture with a hidden state depending on the previous states and present information designed for the sequential problem. In general, the backpropagation algorithm (Rumelhart et al. (1985, 1988)) is used to solve RNN problems. However, the RNN suffers from gradient-vanishing or gradient-exploding as the network. In turn, the network fails to capture long-memory effects(Bengio et al. (1994)). Hochreiter and Schmidhuber (1997) proposed a new design of recurrent neural unit called Long-short term memory unit (LSTM) that has "cell state" and "forget gate" which memorize and forget certain information, respectively. Gated Recurrent Unit (GRU) provided by Cho et al. (2014)) is one of the many variations of LSTM that aims to capture nonlinear temporal dependency but without the memory gate. GRU has performed well, even better than LSTM in some cases. We will describe details of LSTM and empirical forecast results.

3.1.1. Long-Short Term Memory Unit

LSTM models can be constructed in various ways depending on the number of layers and the number of neurons(units) of each layer of that network. On the other hand, there are only very limited versions of the LSTM unit. Here we employ only one layer of the most common LSTM unit finished with a dense layer. Each LSTM unit is shown in Figure 5 and describe it as follows.

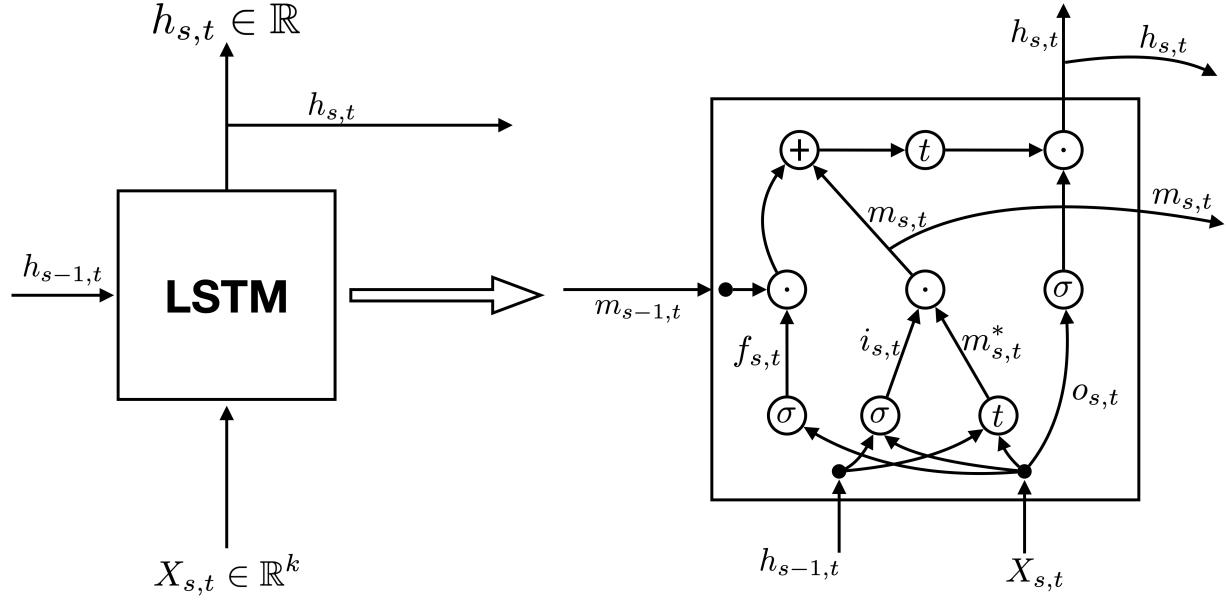


Figure 5
Architecture of LSTM

The figure shows the architecture of a unit in the LSTM model. The left part shows the outside structure, where the input vector $X_{s,t} \in \mathbb{R}^k$ on the bottom and the output vector $h_{s,t} \in \mathbb{R}$ on the top, and the hidden state vector $h_{s-1,t} \in \mathbb{R}$ as a feedback input. The right part is the detailed structure inside the LSTM unit.

One LSTM unit has input of $X_{s,t} \in \mathbb{R}^k$ and output of $h_{s,t} \in \mathbb{R}$, for each time step s , and $t = 1, 2, \dots, T$. The LSTM unit consists of 4 components called “gates”. The first gate is the “input gate” that has its output $i_{s,t} \in \mathbb{R}$ defined in Equation (2) as a function of the the input vector $X_{s,t}$, the output from last step (memory), $h_{s-1,t} \in \mathbb{R}$, and the intercepts $\alpha_i \in \mathbb{R}$. Coefficients $\gamma^{s,i} \in \mathbb{R}^k$ and $\theta^{(s,i)} \in \mathbb{R}$ are estimated in each time step s . Note that k represents the dimensionality of each sample at each time step. In our case, each sample is constructed with the lagged periods load value, 3 groups of lagged production schedules values, hence $k = 4$.

$$i_{s,t} = N^\sigma (X_{s,t}^\top \gamma^{(s,i)} + h_{s-1,t} \theta^{(s,i)} + \alpha^{(s,i)}) \quad (2)$$

The neural network activation function sigmoid function N^σ and tanh function N^{tanh} are element-wise operators defined as Equation (3), respectively

$$\begin{aligned} N^\sigma (x) &= \frac{1}{1 + e^{-x}} \\ N^{tanh} (x) &= \frac{e^x - e^{-x}}{e^x + e^{-x}} \end{aligned} \quad (3)$$

Then the forget gate is defined in Equation (4) which has the same structure as the input gate. This gate controls the amount of information that should be forgotten from the input $X_{s,t}$, previous output $h_{s-1,t}$.

$$f_{s,t} = N^\sigma (X_{s,t}^\top \gamma^{(s,f)} + h_{s-1,t} \theta^{(s,f)} + \alpha^{(s,f)}) \quad (4)$$

The memory gate is defined as Equation (5), where two parts of information are added up. The first part $f_{s,t} \odot m_{s-1,t}$ decides how much information should be kept from memory of last step $m_{s-1,t}$, while the second part is the new memory that is determined by the input $X_{s,t}$ and the output $h_{s-1,t}$ from previous step and the input $i_{s,t}$. The \odot represents the element-wise product.

$$m_{s,t} = f_{s,t} \cdot m_{s-1,t} + i_{s,t} \cdot N^{\tanh} (X_{s,t}^\top \gamma^{(s,m)} + h_{s-1,t} \theta^{(s,m)} + \alpha^{(s,m)}) \quad (5)$$

The $o_{s,t}$ is N^σ function of linear combination of input $X_{s,t}$, output from last period $h_{s-1,t}$. The output vector $h_{s,t}$ is determined by $o_{s,t}$ and $m_{s,t}$.

$$\begin{aligned} o_{s,t} &= N^\sigma (X_{s,t}^\top \gamma^{(s,o)} + h_{s-1,t} \theta^{(s,o)} + \alpha^{(s,o)}) \\ h_{s,t} &= o_{s,t} \cdot N^{\tanh} (m_{s,t}) \end{aligned} \quad (6)$$

For each sample, the LSTM unit will produce output vector $H_t = (h_{1,t}, h_{2,t}, \dots, h_{S,t})^\top$, $h_{s,t} \in \mathbb{R}$ for $s = 1, 2, \dots, S$, and $S = 96$. Finally, a dense layer is employed to transform the H_t to y_{t+2} ,

$$y_{t+2} = D \cdot H_t, \quad (7)$$

where D is a matrix, $D \in \mathbb{R}^{P \times S}$.

Each gate has three parameter vectors to estimated. Take the input gate for example,

$$\begin{aligned} \gamma^i &= (\gamma^{1,i}, \gamma^{2,i}, \dots, \gamma^{S,i})^\top \\ \theta^i &= (\theta^{1,i}, \theta^{2,i}, \dots, \theta^{S,i})^\top \\ \alpha^i &= (\alpha^{1,i}, \alpha^{2,i}, \dots, \alpha^{S,i})^\top \end{aligned} \quad (8)$$

Similarly, the forget gate has $\gamma^f, \theta^f, \alpha^f$; the memory gate has $\gamma^m, \theta^m, \alpha^m$; and the output gate has $\gamma^o, \theta^o, \alpha^o$. Therefore, the parameter matrix can be written as,

$$\begin{aligned}
\Gamma &= (\gamma^i, \gamma^f, \gamma^m, \gamma^o) \in \mathbb{R}^{S \times k*4} \\
\Theta &= (\theta^i, \theta^f, \theta^m, \theta^o) \in \mathbb{R}^{S \times 4} \\
\Lambda &= (\alpha^i, \alpha^f, \alpha^m, \alpha^o) \in \mathbb{R}^{S \times 4}
\end{aligned} \tag{9}$$

The number of parameters estimated in the LSTM model heavily depends on the number of time steps, S . We employ a 1-day lagged sample of a 15-min sampling frequency, i.e., $S = 96$.

3.1.2. Gated Recurrent Unit

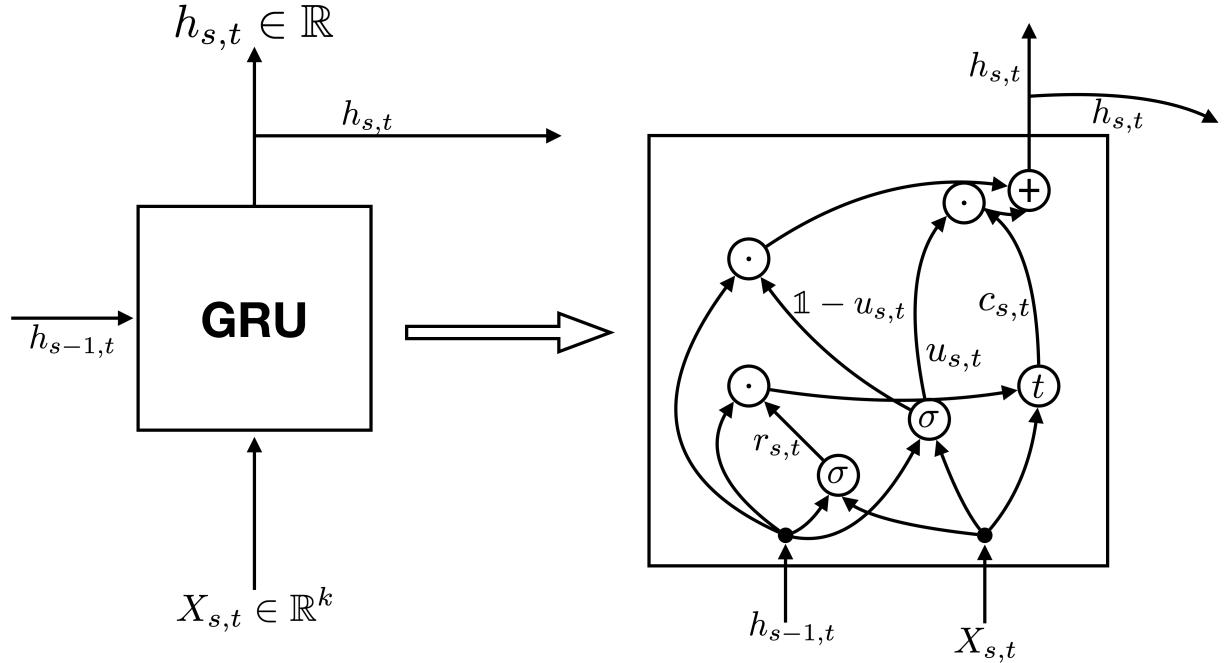


Figure 6
Architecture of GRU

The figure shows the architecture of a unit in the GRU model. The left part shows the outside structure, where the input vector $X_{s,t} \in \mathbb{R}^k$ on the bottom and the output vector $h_{s,t} \in \mathbb{R}$ on the top, and the hidden state vector $h_{s-1,t} \in \mathbb{R}$ as a feedback input. The right part is the detailed structure inside the GRU unit.

We can see the graphic representation of GRU architecture clearly from Figure 6. Instead of having memory gate and forget gate, GRU has a candidate activation $c_{s,t}$. And the output of GRU $h_{s,t} \in \mathbb{R}^d$ is linear interpolation of $c_{s,t}$ and $h_{s-1,t}$, as shown in Equation (13), the update weight is determined by $u_{s,t}$ in Equation (12).

Similar to LSTM, The input gate for GRU defined by Equation (10) is also linear combination of input variable $X_{s,t} \in \mathbb{R}^k$ and previous output $h_{s-1,t}$,

$$i_{s,t} = N^\sigma (X_{s,t}^\top \gamma^{(s,i)} + h_{s-1,t} \theta^{(s,i)} + \alpha^{(s,i)}) \quad (10)$$

Differently, the candidate activation $c_{s,t}$ defined in Equation (11) is a function of $X_{s,t}$ and $i_{s,t} \cdot h_{s-1,t}$. The candidate activation $c_{s,t}$ determines how much information will be kept or reduced from the input and previous result via the N^{tanh} neural network layer.

$$c_{s,t} = N^{tanh} (X_{s,t}^\top \gamma^{(s,c)} + (i_{s,t} \theta^{(s,c)} \cdot h_{s-1,t}) + \alpha^{(s,c)}) \quad (11)$$

The interpolation weighting vector $u_{s,t} \in \mathbb{R}$ defined in Equation (12) controls the degree of updating from candidate activation $c_{s,t}$.

$$u_{s,t} = N^\sigma (X_{s,t}^\top \gamma^{(s,u)} + h_{s-1,t} \theta^{(s,u)} + \alpha^{(s,u)}) \quad (12)$$

The output gate keeps part of the information from the previous output and part of the information from candidate activation $c_{s,t}$.

$$h_{s,t} = (1 - u_{s,t}) \cdot h_{s-1,t} + u_{s,t} \cdot c_{s,t} \quad (13)$$

For each sample, the GRU unit output vector is $H_t = (h_{1,t}, h_{2,t}, \dots, h_{S,t})^\top$, $h_{s,t} \in \mathbb{R}$ for $s = 1, 2, \dots, S$, and $S = 96$. Same to LSTM model, a dense layer with linear transform $D \in \mathbb{R}^{P \times S}$ is employed to transform the H_t to y_{t+2} finally,

$$y_{t+2} = D \cdot H_t, \quad (14)$$

As a GRU unit has one less gate compared to a LSTM unit, the number of parameters to be estimated in the GRU unit is less than in the GRU unit,

$$\begin{aligned} \Gamma &= (\gamma^i, \gamma^c, \gamma^u) \in \mathbb{R}^{S \times k*3} \\ \Theta &= (\theta^i, \theta^c, \theta^u) \in \mathbb{R}^{S \times 3} \\ \Lambda &= (\alpha^i, \alpha^c, \alpha^u) \in \mathbb{R}^{S \times 3} \end{aligned} \quad (15)$$

3.1.3. Parameters Estimation

Due to the flexibility and complexity of LSTM/GRU, a closed-form solution for estimating parameters is not possible. In general, all parameters are estimated by the backpropagation (BPP) algorithm described as follows.

Let $\hat{\Psi}$ be the estimated coefficients matrix in a complete constructed LSTM/GRU model which can be formulated in a general form stated as Equation (16). The forecast y_t at step t is determined by function $g(\cdot)$ which is the nonlinear neural network with forecastors $X_{s,t}$.

$$y_t = g(\Psi; X_{s,t}, h_{s-1,t}) + \varepsilon_t \quad (16)$$

In BPP, all parameters are estimated by iteration shown as Equation (17) with random initial values. δ is a prespecified constant that determines the learning rate of such a neural network, and ∇ is the gradient operation.

$$\hat{\Psi}^{(j+1)} = \hat{\Psi}^{(j)} + \delta \cdot \nabla g\left(\hat{\Psi}^{(j)}; X_{s,t}, h_{s-1,t}\right) \left[h_{s,t} - g\left(\hat{\Psi}^{(j)}; X_{s,t}, h_{s-1,t}\right) \right], i = 1, 2, \dots \quad (17)$$

3.2. FASTEC on Regime Switching

3.2.1. FActorisable Sparse Tail Event Curves

In the following, we present a functional approach to model the load process of chemical production. This method models the short-term electricity consumption Y_t around a local mean, trend, or seasonality. The proposed method is based on a sequence of steps: first, we estimate a functional factor model at each production level using expectiles. Expectiles can be referred to as conditional moments at a specific level of the distribution. The functional approach employs B -splines as regressors and, hence, cannot be used directly for forecasting purposes. By employing an asymmetric loss function, it is possible to estimate the model coefficients conditional on being in a specific region of the distribution, such that the approximation of the entire probability distribution of the stochastic variation is achieved without imposing the type of distribution. Within the estimation procedure, iteratively singular value decomposition is procured in the coefficient estimation step. Left singular-vectors and the singular values give us time-invariant factors that capture the within-day dynamic of our functional data, and time-varying loadings, the right-hand singular vectors, can be deployed in a time series regression model. Then, a vector autoregressive time series model is estimated on the time-varying loadings of the first r most relevant factors to capture any short term variation that might occur in the data.

Expectile regressions have become popular functional data analysis (FDA) techniques (Chao et al. (2015), Schnabel and Eilers (2009), López Cabrera and Schulz (2017), Sobotka et al. (2013), Tran et al. (2018), Ramsay and Silverman (2007), Burdejová and Härdle (2019)). The multivariate expectile regression is done using the FActorisable Sparse Tail Event Curves (FASTEC) methodology, a recent technique involving generalized quantiles

in a multivariate regression setup. It incorporates quantile curves (Chao et al., 2015) or expectile curves (Chao et al., 2018) in a linear model that represents functional variables as linear combinations of r B -splines. This approach is in spirit analog to the Karhunen-Loëve expansion or the Fourier series that also employ linear combinations of orthogonal functions. An approximation of each load curve is then given by the sum of the first r of these functions.

In a general setup, suppose $\{\mathbf{X}_s\}_{s=1}^S \in \mathbb{R}^k$, $\{\mathbf{Y}_s\}_{s=1}^S \in \mathbb{R}^q$ i.i.d.; The multivariate linear model for a τ -expectile curve of the daily curve Y_t , consisting of S intra-day observations, with $t = 1, \dots, T$, and the level $\tau \in (0, 1)$, is given by

$$Y_t = e_t(\tau | \mathbf{X}_s) + U_{st,\tau} = \mathbf{X}_s^\top \Gamma_{*t}(\tau) + U_{st,\tau} \quad (18)$$

where $\Gamma_{*t}(\tau) \in \mathbb{R}^k$ is the coefficient matrix for day t at τ -level and $e_t(\tau | \mathbf{X}_s)$ is the conditional expectile of Y_t given \mathbf{X}_s and $U_{st,\tau}$ is τ -level white noise. k is the number linear combinations of B-splines required to represent the Y_t and is chosen as $\lfloor \sqrt{n} \rfloor$.

The multivariate expectile regression formulation with penalised loss is given by

$$\Gamma_\lambda(\tau) = \arg \min_{\Gamma \in \mathbb{R}^{k \times q}} \left\{ (TS)^{-1} \sum_{s=1}^S \sum_{t=1}^T \rho_\tau(Y_{st} - \mathbf{X}_s^\top \Gamma_{*t}) + \lambda \|\Gamma\|_* \right\}, \quad (19)$$

where $\|\Gamma\|_* = \sum_{t=1}^{\min(k,q)} \sigma_t(\Gamma)$ is the nuclear norm of Γ , \mathbf{X}_s is a $(S \times k)$ -matrix, Γ is the coefficient matrix and λ the penalisation parameter. Please note, that $\Gamma_\lambda(\tau) = \Gamma$ depends on the tuning parameter λ and the τ -level. To ensure efficiency and convergence of the algorithm optimal selection of the tuning parameter λ is of the essence. The n -fold cross-validation method gives a consistent estimate of the tuning parameter (Härdle et al., 2012). In our case, we are required to forecast the load process 2-day ahead, the FASTEC method is applied twice each time to forecast a 2-day result.

The decision to which expectile level the current data belongs and, hence, where the forecast of the load will be, is entirely data-driven and is evaluated by a change-point test. The change-point detection is based on a t -test that captures the mean distance of curve Y_t to all mean levels estimated from the Gaussian Mixture Model described later in Sec.3.2.2, where the load process is estimated in a 15-minute frequency. Thus, the expected distance to one of the levels is expected to be zero. If this is the case, the test-statistic is small, and the null hypothesis of having the same level cannot be rejected.

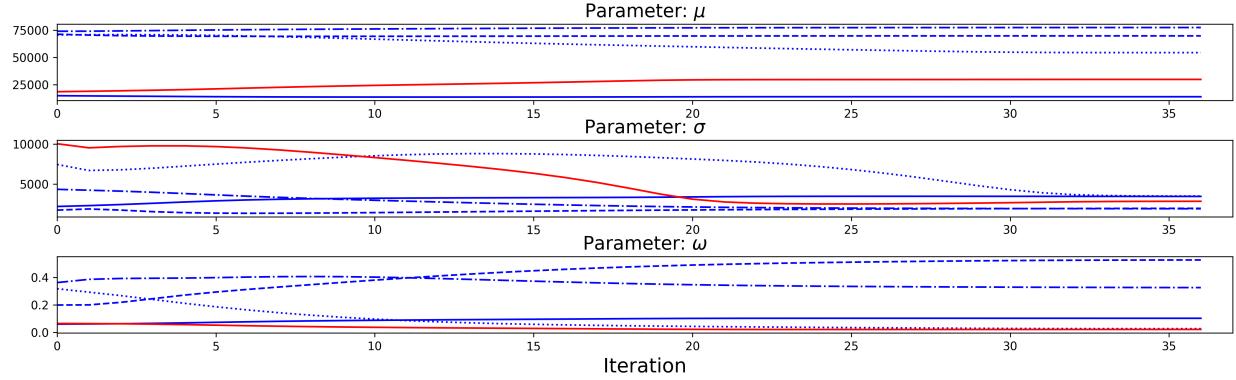


Figure 7
Convergence of EM algorithm on load process

The figure shows how the parameters converge with the EM algorithm.

3.2.2. Gaussian Mixture Model

Motivated by Figure 2, the distribution of load process points can be modeled as Gaussian Mixture Model (GMM) (McLachlan and Peel (2004), Eirola and Lendasse (2013)). Here, we model the GMM using the first 10,000 points, i.e $D = 10,000$. GMM assumes that the probability density function of the load process $\{y_d\}_{d=1}^D$ in a 15-minute frequency can be written as a one-dimensional linear combination of k Gaussian density functions as follows.

$$P(y|\alpha, \mu, \sigma^2) = \sum_{k=1}^K \alpha_k \varphi(y|\mu_k, \sigma_k^2) \quad (20)$$

Where $\Theta_k = (\alpha_k, \mu_k, \sigma_k^2)^\top$ are the weight, mean, and variance of the k_{th} normal density function, respectively. Under the restriction that $\alpha_k \in (0, 1)$, $k = 1, \dots, K$, and $\sum_{k=1}^K \alpha_k = 1$.

In this article, the goal of GMM modeling is to detect the mean values for different states in the load process. Thus we ignore the temporal dependency of the load process. The estimation procedure is carried by the classic Expectation-Maximization (EM) algorithm provide by Dempster et al. (1977) and then extended by Ruud (1991).

To estimate the parameters in Equation (20), we assume that there are $K = 5$ machinery work states, the latent variable $S_t^{(k)}$ represents that machine works on the k_{th} status at time t . $S_t^{(k)}$ is an indicator function that equals 1 with probability of α_k . Conditioning on the working status, the the probability density function is Gaussian, i.e $\varphi(y_t|S_t^{(k)}, \mu_k, \sigma_k^2)$. Then, the log-likelihood function under the assumptions can be written as:

$$\begin{aligned}
\log P(y, S^{(k)} | \Theta) &= \log (P(y | S^{(k)}, \Theta) \cdot P(S^{(k)} | \Theta)) \\
&= \log \left(\prod_{k=1}^K \alpha_k^{n_k} \cdot \prod_{t=1}^T \varphi(y_t | \mu_k, \sigma_k)^{S_t^{(k)}} \right) \\
&= \sum_{k=1}^K \left\{ n_k \log \alpha_k + \sum_{t=1}^T S_t^{(k)} \left[\log \frac{1}{\sqrt{2\pi}} - \log \sigma_k - \frac{(y_t - \mu_k)^2}{2\sigma_k} \right] \right\}
\end{aligned} \tag{21}$$

Where $n_k = \sum_{t=1}^T S_t^{(k)}$ is the number of points that belong to the k_{th} status.

The E-step of the EM algorithm is taking the expectation of the likelihood function in Equation (21) named \mathbb{Q} function.

$$\begin{aligned}
\mathbb{Q}(\Theta, \Theta^{(i-1)}) &= \mathbb{E} \left(\log P(y, S^{(k)} | \Theta) \mid y, \Theta^{(i-1)} \right) \\
&= \sum_{k=1}^K \sum_{t=1}^T \mathbb{E} \left(S_t^{(k)} | y, \Theta^{(i-1)} \right) \log \alpha_k \\
&\quad + \sum_{k=1}^K \sum_{t=1}^T \mathbb{E} \left(S_t^{(k)} | y, \Theta^{(i-1)} \right) \left[\log \frac{1}{\sqrt{2\pi}} - \log \sigma_k - \frac{(y_t - \mu_k)^2}{2\sigma_k} \right]
\end{aligned} \tag{22}$$

The conditional expectation of working status at time t is estimated as follows.

$$\mathbb{E} \left(S_t^{(k)} | y, \Theta^{(i-1)} \right) = \frac{\alpha_k \varphi(y_t | \Theta_k^{(i-1)})}{\sum_{k=1}^K \alpha_k \varphi(y_t | \Theta_k^{(i-1)})} \tag{23}$$

Then the M-step of the EM algorithm is the maximization of \mathbb{Q} function with iteration,

$$\Theta^i = \arg \max_{\Theta} \mathbb{Q}(\Theta, \Theta^{(i-1)}) \tag{24}$$

As shown in Figure 7, the EM algorithm converges quickly. With the estimated states mean μ , the change-point test of FASTEC can then be performed.

We summarize the FASTEC method with regime switching as the following algorithm representation in Algorithm 2.

3.3. Performance Evaluation

To have a more robust comparison, we employ five metrics described as follows to measure our forecast results.

The Mean Absolute Error is defined as $MAE = \frac{1}{T} \sum_{t=1}^T |y_t - \hat{y}_t|$, where T is the out-of-

Algorithm 2: Algorithm of FASTEC Method with Regime Switching

Input : Training Sample $\mathcal{S}^{train} = \{(X_{s,t}, y_{s,t+h})\}_{s=1}^{s=96}, h = 1, 2$ from Algo. 1

Output: Estimated coefficient matrix: $\Gamma_{*t}(\tau)$

- 1 Estimate conditional moments: certain expectiles as cluster centers;
- 2 Perform the Gaussian Mixture Model with load process in training sample $\{y_{s,t}\}_{s=1}^{96}$ to obtain μ and σ at each level;
- 3 Determine affiliation based on t -statistic with estimated μ and σ ;
- 4 Perform FASTEC (training) on $h = 1, 2$ separately:
 - 5 Select periods at conditional moments with 96 observations $\{y_{s,t+h}^\tau\}_{s=1}^{96}$;
 - 6 Bootstrap to get larger sample $y_{s,t+h}^{\tau*}$;
 - 7 Perform FASTEC on each conditional moment sample;
 - 8 Predict conditional moments $\hat{y}_{s,t+h}^{\tau*}$;
- 9 Estimate the weights vector $\hat{\gamma}$ for half of each sample $\{\hat{y}_{s,t+h}^{\tau*}, x_{s,t}^1, x_{s,t}^2, x_{s,t}^3\}_{s=1+96*h}^{96*h}$ s.t. MAE is minimised;

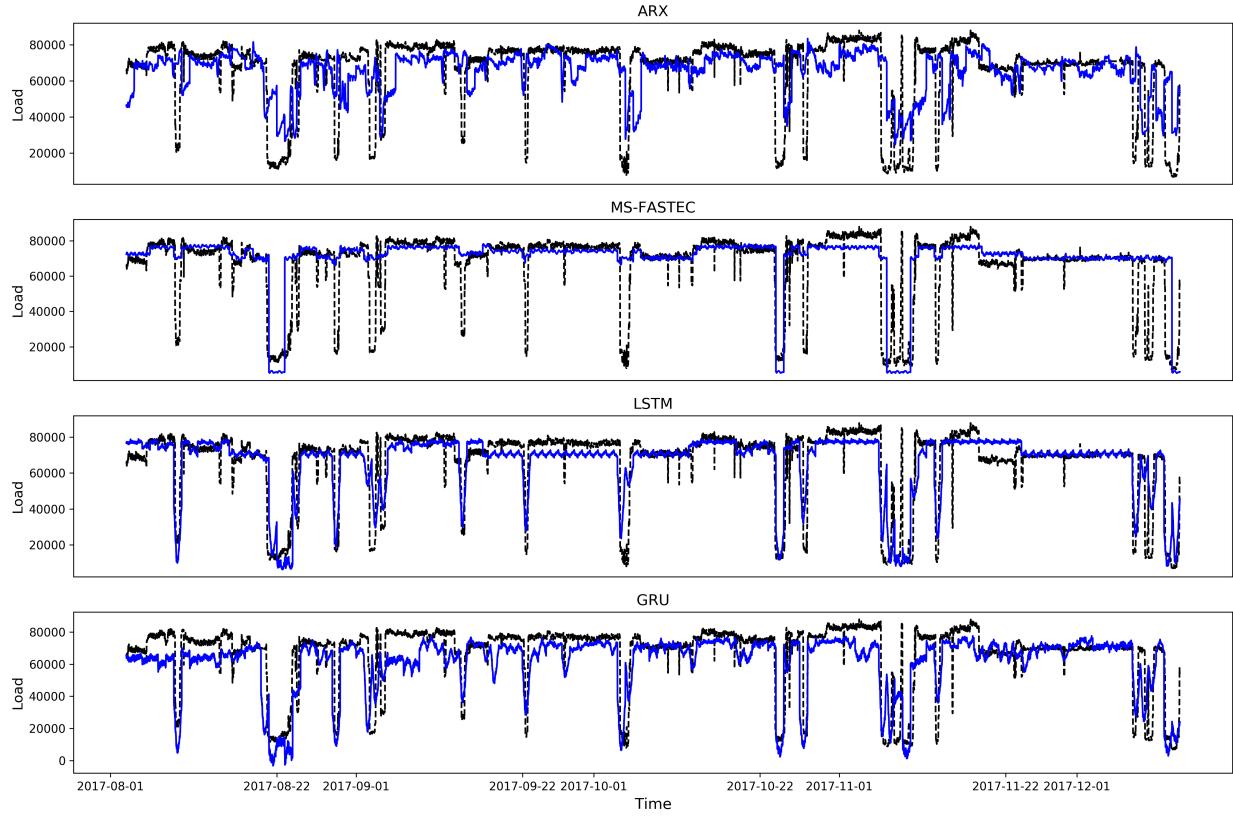


Figure 8

Out-of-Sample Forecasting Comparison

The figure shows the out-of-sample forecasting results (blue line) from all the models compare against the true load process (black line).

sample points. MAE measures the average absolute distance between the true values, y_t ,

and forecasts, \hat{y}_t . Unfortunately, the naive approach (taking value from previous period) would take advantage of MAE metrics if the process is evolving mildly. Using the Mean Absolute Scaled Error, $MASE = \frac{1}{T} \sum_{t=1}^T l^{-1} |y_t - \hat{y}_t|$, overcomes such bias by adjusting MAE with $l = \frac{1}{T-1} \sum_{s=2}^T |y_s - y_{s-1}|$ that accounts the average changing intensity between adjacent periods. To relieve the influence from scale, we also measure with the Mean Absolute Percentage Error $MAPE = \frac{1}{T} \sum_{t=1}^T |(y_t - \hat{y}_t)/y_t|$.

The Mean Squared Error $MSE = T^{-\frac{1}{2}} \sqrt{\sum_{t=1}^T (\hat{y}_t - y_t)^2}$ measures the average Euclidean distance between true values and forecasts. We employ 3 metrics that are essentially MSE adjusted by 3 different factors, Q , namely the Root Mean Squared Error $RMSE = T^{-\frac{1}{2}} Q^{-1} \sqrt{\sum_{t=1}^T (\hat{y}_t - y_t)^2}$. The first one $nrRMSE$ is adjusted by difference between the largest and smallest values, $Q = Q_{nr} = \{max(y_t) - min(y_t)\}$, the second one $nmRMSE$ is adjusted by the mean value $Q = Q_{nm} = T^{-1} \sum_{t=1}^T y_t$, and the third one $nigrRMSE$ is adjusted by empirical quantile difference, $Q = Q_{nigr} = \{q(y_t, 0.75) - q(y_t, 0.25)\}$.

Then Diebold-Mariano test (DM) provided by Diebold and Mariano (2002) is employed to compare forecasting results from different models. D-M test compares the forecasting errors from two models and gives the statistics of whether one error process is significantly smaller or larger than the other one. Here we compare LSTM against GRU, MS-FASTE, and ARX models using quadratic loss, i.e the loss differential $d_t^{LSTM,m} = (y_t^{LSTM} - y_t)^2 - (y_t^m - y_t)^2$ where $m \in \{GRU, MS-FASTE, ARX\}$, and y_t is the true value at time t . Then the sample mean loss differential is $\bar{d}_t^{LSTM,m} = T^{-1} \sum_{t=1}^T d_t^{LSTM,m}$. Under some assumption, the D-M test is formulated as follows:

$$DM_t^{LSTM,m} = \frac{\bar{d}_t^{LSTM,m}}{\hat{\sigma}_{\bar{d}_t^{LSTM,m}}} \quad (25)$$

Where $\hat{\sigma}_{\bar{d}_t^{LSTM,m}}$ is the consistent estimate of standard deviation of $\bar{d}_t^{LSTM,m}$.

Table 2 shows the metrics of different models, and the LSTM model outperforms parametric models and even GRU consistently. This result strongly suggests that using a suitable neural network model can perform much better when the nonlinearity property is captured.

Even though the improvement from the LSTM model is significant, Figure 9 shows that none of the models could generate a forecasting error process that is normally distributed. One explanation of such forecasting results is that in our case, the load process is not surrendered to any clear pattern and is filled with a huge amount of jumps.

Table 2
Out-of-sample 2-days ahead forecast performance evaluation

Measure	LSTM	GRU	MS-FASTEC	ARX
<i>MAPE</i>	0.211 [†]	0.242	0.332	0.415
<i>MASE</i>	0.597 [†]	0.848	0.672	1.040
<i>nRMSE</i>	0.142 [†]	0.175	0.197	0.220
<i>niqRMSE</i>	1.323 [†]	1.628	1.831	2.052
<i>nmRMSE</i>	0.173 [†]	0.212	0.239	0.268
<i>D-M</i> [‡]		-19.138***	-21.556***	-31.297***

[†]: Model with the highest accuracy compared with other models under each metric.

[‡]: DM test with quadratic loss. GRU, MS-FASTEC, and ARX models against the LSTM model. A significant negative testing result implies that the LSTM model outperforms the other models.

***: Significant at 1% level.

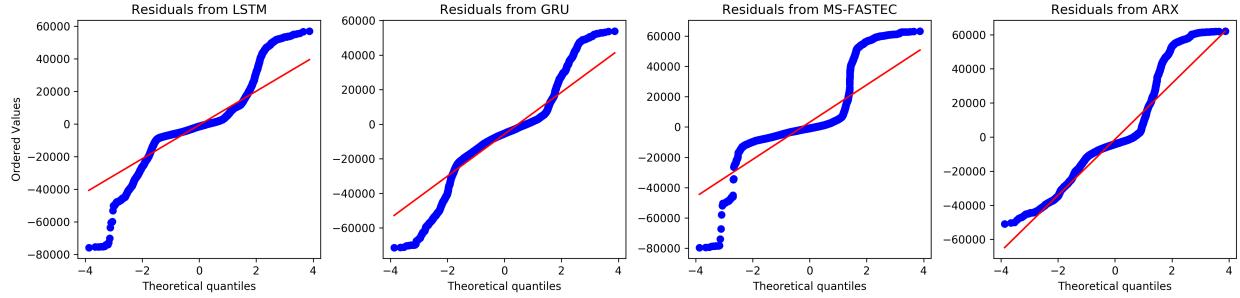


Figure 9
Distribution Residuals from Each Forecasting model

The Q-Q plot shows the distribution of residuals from each model against the standard normal distribution. Models from left to right plots are LSTM, GRU, MS-FASTEC, ARX.

4. Conclusion

Short term load forecast is crucial for energy consumers, especially the big manufacturers, for the forecast is one of the key parts of their decision-making. Previously, most of the studies focus on modeling the load process from large-scale consumers, for example, a countrywide electricity consumption process that evolves with a seasonal or periodic pattern. Forecast on the small-scale consumer, such as a household or a set of machines, is a pragmatic problem for the industries while rarely investigating. Requesting by a chemical production company located in Germany, we are motivated to investigate the electric load process of a chemical production line that involves hard structures.

Despite the fact that the load process has a strong temporal dependency, which can be captured by autoregression models or long-memory models, one can see a huge amount of jumps in the process. Some of the jumps exist only for a short time, and some of them actually represent the load process transfer to another working status. The nonlinear dependence

motivates us to implement state-of-the-art neural network models.

We first calibrate two of the popular deep neural network models, LSTM and GRU. The parameters are estimated by the standard backpropagation algorithm. To capture the tail events of the process, we further construct a functional data analysis based semi-parametric model, FASTEC, combined with the EM algorithm. For comparison, we use the ARX model as a benchmark.

Utilizing the one-year intraday data, all models forecast the electricity consumption in the next two days in 15-min resolution. We train all the models with the first 7-month data and conduct out-of-sample forecasting in the last 5-month. The out-of-sample forecast suggests that the LSTM model outperforms all other models by all metrics significantly evaluated by the D-M test. However, even the LSTM model can not harness the tail risks caused by jumps, which is one crucially problem for further researches.

Appendix A. An Appendix

References

Amara, F., Agbossou, K., Dubé, Y., Kelouwani, S., Cardenas, A., Bouchard, J., 2017. Household electricity demand forecasting using adaptive conditional density estimation. *Energy And Buildings* 156, 271–280.

An, N., Zhao, W., Wang, J., Shang, D., Zhao, E., 2013. Using multi-output feedforward neural network with empirical mode decomposition based signal filtering for electricity demand forecasting. *Energy* 49, 279–288.

Bengio, Y., Simard, P., Frasconi, P., 1994. Learning long-term dependencies with gradient descent is difficult. *Ieee Transactions On Neural Networks* 5, 157–166.

Bianchi, F. M., Maiorino, E., Kampffmeyer, M. C., Rizzi, A., Jenssen, R., 2017. An overview and comparative analysis of recurrent neural networks for short term load forecasting. *Arxiv Preprint Arxiv:1705.04378* .

Bianco, V., Manca, O., Nardini, S., 2009. Electricity consumption forecasting in italy using linear regression models. *Energy* 34, 1413–1421.

Burdejová, P., Härdle, W. K., 2019. Dynamic semi-parametric factor model for functional expectiles. *Computational Statistics* 34, 489–502.

Castelli, M., Vanneschi, L., Felice, M. D., 2015. Forecasting short-term electricity consumption using a semantics-based genetic programming framework: The south italy case. *Energy Economics* 47, 37–41.

Chang, P.-c., Fan, C.-y., Lin, J.-j., 2011. Monthly electricity demand forecasting based on a weighted evolving fuzzy neural network approach. *International Journal Of Electrical Power & Energy Systems* 33, 17–27.

Chao, S.-k., Härdle, W. K., Huang, C., 2018. Multivariate factorizable expectile regression with application to fMRI data. *Computational Statistics & Data Analysis* 121, 1–19.

Chao, S.-k., Härdle, W. k., Yuan, M., 2015. Factorisable Sparse Tail Event Curves. *Arxiv E-prints* .

Cho, K., Van Merriënboer, B., Bahdanau, D., Bengio, Y., 2014. On the properties of neural machine translation: Encoder-decoder approaches. *Arxiv Preprint Arxiv:1409.1259* .

Clements, A. E., Hurn, A., Li, Z., 2016. Forecasting day-ahead electricity load using a multiple equation time series approach. *European Journal Of Operational Research* 251, 522–530.

Cottet, R., Smith, M., 2003. Bayesian modeling and forecasting of intraday electricity load. *Journal Of The American Statistical Association* 98, 839–849.

Darbey, G. A., Slama, M., 2000. Forecasting the short-term demand for electricity: Do neural networks stand a better chance? *International Journal Of Forecasting* 16, 71–83.

Dempster, A. P., Laird, N. M., Rubin, D. B., 1977. Maximum likelihood from incomplete data via the em algorithm. *Journal Of The Royal Statistical Society: Series B (methodological)* 39, 1–22.

Dickey, D. A., Fuller, W. A., 1979. Distribution of the estimators for autoregressive time series with a unit root. *Journal Of The American Statistical Association* 74, 427–431.

Dickey, D. A., Fuller, W. A., 1981. Likelihood ratio statistics for autoregressive time series with a unit root. *Econometrica: Journal Of The Econometric Society* pp. 1057–1072.

Diebold, F. X., Mariano, R. S., 2002. Comparing predictive accuracy. *Journal Of Business & Economic Statistics* 20, 134–144.

Do, L. P. C., Lin, K.-h., Molnár, P., 2016. Electricity consumption modelling: A case of germany. *Economic Modelling* 55, 92–101.

Eirola, E., Lendasse, A., 2013. Gaussian mixture models for time series modelling, forecasting, and interpolation. In: Tucker, A., Höppner, F., Siebes, A., Swift, S. (eds.), *Advances In Intelligent Data Analysis Xii*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 162–173.

Elliott, G., Müller, U. K., 2007. Confidence sets for the date of a single break in linear time series regressions. *Journal Of Econometrics* 141, 1196–1218.

Elliott, G., Müller, U. K., 2014. Pre and post break parameter inference. *Journal Of Econometrics* 180, 141–157.

Guan, C., Luh, P. B., Michel, L. D., Wang, Y., Friedland, P. B., 2012. Very short-term load forecasting: Wavelet neural networks with data pre-filtering. *Ieee Transactions On Power Systems* 28, 30–41.

Härdle, W. K., Müller, M., Sperlich, S., Werwatz, A., 2012. Nonparametric And Semiparametric Models. Springer Science & Business Media.

Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural Computation* 9, 1735–1780.

Kwiatkowski, D., Phillips, P. C., Schmidt, P., Shin, Y., 1992. Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal Of Econometrics* 54, 159–178.

López Cabrera, B., Schulz, F., 2017. Forecasting generalized quantiles of electricity demand: A functional data approach. *Journal Of The American Statistical Association* 112, 127–136.

McLachlan, G., Peel, D., 2004. Finite Mixture Models. John Wiley & Sons.

Öğcu, G., Demirel, O. F., Zaim, S., 2012. Forecasting electricity consumption with neural networks and support vector regression. *Procedia-social And Behavioral Sciences* 58, 1576–1585.

Pardo, A., Meneu, V., Valor, E., 2002. Temperature and seasonality influences on spanish electricity load. *Energy Economics* 24, 55–70.

Pesaran, M. H., Timmermann, A., 2005. Small sample properties of forecasts from autoregressive models under structural breaks. *Journal Of Econometrics* 129, 183–217.

Ramsay, J. O., Silverman, B. W., 2007. Applied Functional Data Analysis: Methods And Case Studies. Springer.

Rumelhart, D. E., Hinton, G. E., Williams, R. J., 1985. Learning internal representations by error propagation. Tech. rep.

Rumelhart, D. E., Hinton, G. E., Williams, R. J., Others, 1988. Learning representations by back-propagating errors. *Cognitive Modeling* 5, 1.

Ruud, P. A., 1991. Extensions of estimation methods using the em algorithm. *Journal Of Econometrics* 49, 305–341.

Schnabel, S. K., Eilers, P. H., 2009. Optimal expectile smoothing. *Computational Statistics & Data Analysis* 53, 4168–4177.

Sobotka, F., Kauermann, G., Waltrup, L. S., Kneib, T., 2013. On confidence intervals for semiparametric expectile regression. *Statistics And Computing* 23, 135–148.

Stock, J. H., Watson, M. W., 1996. Evidence on structural instability in macroeconomic time series relations. *Journal Of Business & Economic Statistics* 14, 11–30.

Taylor, J. W., 2003. Short-term electricity demand forecasting using double seasonal exponential smoothing. *Journal Of The Operational Research Society* 54, 799–805.

Tran, N. M., Burdejová, P., Ospienko, M., Härdle, W. K., 2018. Principal component analysis in an asymmetric norm. *Journal Of Multivariate Analysis* .

Weber, S., Puddu, S., Pacheco, D., 2017. Move it! how an electric contest motivates households to shift their load profile. *Energy Economics* 68, 255–270.

IRTG 1792 Discussion Paper Series 2021



For a complete list of Discussion Papers published, please visit
<http://irtg1792.hu-berlin.de>.

- 001 "Surrogate Models for Optimization of Dynamical Systems" by Kainat Khowaja, Mykhaylo Shcherbatyy, Wolfgang Karl Härdle, January 2021.
- 002 "FRM Financial Risk Meter for Emerging Markets" by Souhir Ben Amor, Michael Althof, Wolfgang Karl Härdle, February 2021.
- 003 "K-expectiles clustering" by Bingling Wang, Yingxing Li, Wolfgang Karl Härdle, March 2021.
- 004 "Understanding Smart Contracts: Hype or Hope?" by Elizaveta Zinovyev, Raphael C. G. Reule, Wolfgang Karl Härdle, March 2021.
- 005 "CATE Meets ML: Conditional Average Treatment Effect and Machine Learning" by Daniel Jacob, March 2021.
- 006 "Coins with benefits: on existence, pricing kernel and risk premium of cryptocurrencies" by Cathy Yi-Hsuan Chen, Dmitri Vinogradov, April 2021.
- 007 "Rodeo or Ascot: which hat to wear at the crypto race?" by Konstantin Häusler, Wolfgang Karl Härdle, April 2021.
- 008 "Financial Risk Meter based on Expectiles" by Rui Ren, Meng-Jou Lu, Yingxing Li, Wolfgang Karl Härdle, April 2021.
- 009 "Von den Mühen der Ebenen und der Berge in den Wissenschaften" by Annette Vogt, April 2021.
- 010 "A Data-driven Explainable Case-based Reasoning Approach for Financial Risk Detection" by Wei Li, Florentina Paraschiv, Georgios Sermpinis, July 2021.
- 011 "Valuing cryptocurrencies: Three easy pieces" by Michael C. Burda, July 2021.
- 012 "Correlation scenarios and correlation stress testing" by Natalie Packham, Fabian Woebbeking, July 2021.
- 013 "Penalized Weighted Competing Risks Models Based on Quantile Regression" by Erqian Li, Wolfgang Karl Härdle, Xiaowen Dai, Maozai Tian, July 2021.
- 014 "Indices on Cryptocurrencies: an Evaluation" by Konstantin Häusler, Hongyu Xia, August 2021.
- 015 "High-dimensional Statistical Learning Techniques for Time-varying Limit Order Book Networks" by Shi Chen, Wolfgang Karl Härdle, Melanie Schienle, August 2021.
- 016 "A Time-Varying Network for Cryptocurrencies" by Li Guo, Wolfgang Karl Härdle, Yubo Tao, August 2021.
- 017 "Green financial development improving energy efficiency and economic growth: a study of CPEC area in COVID-19 era" by Linyun Zhang, Feiming Huang, Lu Lu, Xinwen Ni, September 2021.
- 018 "Robustifying Markowitz" by Wolfgang Karl Härdle, Yegor Klochkov, Alla Petukhina, Nikita Zhivotovskiy, September 2021.
- 019 "Understanding jumps in high frequency digital asset markets" by Danial Saef, Odett Nagy, Sergej Sizov, Wolfgang Karl Härdle, October 2021.
- 020 "Advanced Statistical Learning on Short Term Load Process Forecasting" by Junjie Hu, Brenda López Cabrera, Awdesch Melzer, October 2021.

IRTG 1792, Spandauer Strasse 1, D-10178 Berlin
<http://irtg1792.hu-berlin.de>

This research was supported by the Deutsche Forschungsgemeinschaft through the IRTG 1792.