



**WORMS/21/07**

**Neural basis expansion analysis  
with exogenous variables:  
Forecasting electricity prices  
with NBEATSx**

Kin G. Olivares<sup>1</sup>  
Cristian Challu<sup>1</sup>  
Grzegorz Marcjasz<sup>2</sup>  
Rafał Weron<sup>2</sup>  
Artur Dubrawski<sup>1</sup>

<sup>1</sup> Auton Lab, School of Computer Science, Carnegie Mellon University,  
Pittsburgh, Pennsylvania, USA

<sup>2</sup> Department of Operations Research and Business Intelligence,  
Wrocław University of Science and Technology, Poland

WORMS is a joint initiative of the Management Science departments  
of the Wrocław University of Science and Technology,  
Wyb. Wyspiańskiego 27, 50-370 Wrocław, Poland

# Neural basis expansion analysis with exogenous variables: Forecasting electricity prices with NBEATSx

Kin G. Olivares<sup>a</sup>, Cristian Challu<sup>a</sup>, Grzegorz Marcjasz<sup>b</sup>, Rafał Weron<sup>b</sup>, Artur Dubrawski<sup>a</sup>

<sup>a</sup>*Auton Lab, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA*

<sup>b</sup>*Department of Operations Research and Business Intelligence, Wrocław University of Science and Technology, Wrocław, Poland*

---

## Abstract

We extend the *neural basis expansion analysis* (NBEATS) to incorporate exogenous factors. The resulting method, called NBEATSx, improves on a well performing deep learning model, extending its capabilities by including exogenous variables and allowing it to integrate multiple sources of useful information. To showcase the utility of the NBEATSx model, we conduct a comprehensive study of its application to electricity price forecasting (EPF) tasks across a broad range of years and markets. We observe state-of-the-art performance, significantly improving the forecast accuracy by nearly 20% over the original NBEATS model, and by up to 5% over other well established statistical and machine learning methods specialized for these tasks. Additionally, the proposed neural network has an interpretable configuration that can structurally decompose time series, visualizing the relative impact of trend and seasonal components and revealing the modeled processes' interactions with exogenous factors.

**Keywords:** Deep learning, NBEATS and NBEATSx models, Interpretable neural network, Time series decomposition, Fourier series, Electricity price forecasting

---

## 1. Introduction

In the last decade significant progress has been made in the application of deep learning to forecasting tasks, with models such as the *exponential smoothing recurrent neural network* (ESRNN; Smyl 2019) and the *neural basis expansion analysis* (NBEATS; Oreshkin et al. 2020), outperforming classical statistical approaches in the M3 (Makridakis & Hibon, 2000), Tourism (Athanasopoulos et al., 2011), or M4 competition (Makridakis et al., 2020); redefining the state of the art. Despite the success of these novel methods we still identify two possible improvements, namely the integration of time-dependent exogenous variables as their inputs and the interpretability of the neural network outputs.

---

\*Corresponding author

Email address: [kdgutier@cs.cmu.edu](mailto:kdgutier@cs.cmu.edu) (Kin G. Olivares)

Neural networks have proven powerful and flexible, yet there are several situations where our understanding of the model’s predictions can be as crucial as their accuracy, which constitutes a barrier for their wider adoption. The interpretability of the algorithm’s outputs is critical because it creates trust in its predictions, improves our knowledge of the modeled processes, and provides insights that can improve the method itself.

Additionally, the absence of time-dependent covariates makes these powerful models unsuitable for many applications. For instance, Electricity Price Forecasting (EPF) is a task where covariate features are fundamental to obtain accurate predictions. For this reason, we chose this challenging application as a test ground for our proposed forecasting methods.

In this work, we address the two mentioned limitations by first extending the neural basis expansion analysis, allowing it to incorporate temporal and static exogenous variables. And second, by further exploring the interpretable configuration of NBEATS and showing its use as a time-series variance decomposition tool. We refer to the new method as NBEATSx. The main contributions of this paper include:

- (i) **Incorporation of Exogenous Variables:** We propose improvements to the NBEATS model to incorporate time dependent as well as static exogenous variables. For this purpose, we have designed a special substructure built with convolutions, to clean and encode useful information from these covariates, while respecting time dependencies present in the data.
- (ii) **Interpretable Time Series Decomposition:** We demonstrate the versatility of the NBEATSx approach and its ability to structurally decompose time series. The method combines the flexibility and power of the non-linear transformations of neural networks, with the ability to model multiple seasons. These seasons can have fractional periods and account for interaction events such as holidays and other fundamental covariates, while keeping the resulting model as interpretable as many classic time series decomposition tools.
- (iii) **Time Series Forecasting Comparison:** We showcase the use of NBEATSx model on five EPF tasks achieving state-of-the-art performance on all of the considered datasets. We obtain accuracy improvements of almost 20% in comparison to the original NBEATS and ESRNN architectures, and up to 5% over other well-established machine learning, EPF-tailored methods (Lago et al., 2021).

The remainder of the paper is structured as follows. Section 2 reviews relevant literature on the developments and applications of deep learning to sequence modeling and current approaches to EPF. Section 3 introduces mathematical notation and describes the NBEATSx model. Section 4 explores our model’s application to time series decomposition and forecasting over a broad range of electricity markets and time periods. Finally, Sections 5 and 6 respectively discuss possible directions for future research, and wrap up the results and conclude the paper.

## 2. Literature Review

### 2.1. Deep Learning and Sequence Modeling

The Deep Learning methodology (DL) has demonstrated a significant utility in solving sequence modeling problems with applications to natural language processing, audio signal processing, and computer vision. This subsection summarizes the critical DL developments in sequence modeling, that are building blocks of the NBEATS and ESRNN architectures.

For a long time sequence modeling with neural networks and *Recurrent Neural Networks* (RNNs; Elman 1990; Werbos 1990) were treated as synonym. The hidden internal activations of the RNNs propagated through time provided these models with the ability to encode the observed past of the sequence. This explains their great popularity in building different variants of the *Sequence-to-Sequence* models (Seq2Seq) applied to natural language processing (Graves, 2013; Hermans & Schrauwen, 2013), and machine translation (Sutskever et al., 2014; Bahdanau et al., 2016). Most progress on RNNs was made possible by architectural innovations and novel training techniques that made their optimization easier, and involved popular designs such as the *Long Short Term Memory* (LSTM; Gers et al. 2000) and the *Gated Recurrent Units* (GRU; Chung et al. 2014).

The adoption of convolutions and skip-connections within the recurrent structures were important precursors for new advancements in sequence modeling, as using deeper representations endowed longer effective memory for the models. Examples of such precursors could be found in *WaveNet* for audio generation and machine translation (van den Oord et al. 2016, Dauphin et al. 2017, Kalchbrenner et al. 2016), as well as the *Dilated RNN* (DRNN; Chang et al. 2017) and the *Temporal Convolutional Network* (TCN; Bai et al. 2018).

Nowadays, Seq2Seq models and their derivatives can learn complex nonlinear temporal dependencies efficiently; its use in the time series analysis domain has been a great success. Seq2Seq models have recently showed better forecasting performance than classical statistical methods, while greatly simplifying the forecasting systems into single-box models, such as the *Multi Quantile Convolutional Neural Network* (MQCNN; Wen et al. 2017), the *Exponential Smoothing Recurrent Neural Network* (ESRNN; Smyl 2019), or the *Neural Basis Expansion Analysis* (NBEATS; Oreshkin et al. 2020).

After these breakthroughs, forecasting with neural networks has become the state of the art in many practical applications. Notable examples include large online retailers (Salinas et al. 2020, Wen et al. 2017), weather forecasts (Nascimento et al. 2021), and energy markets (Lago et al., 2018a; Dudek et al., 2021). For quite a while, the academia resisted to broadly adopt these new methods (Makridakis et al., 2018), although their evident success in challenges such as the M4 competition has motivated their wider adoption by the forecasting research community (Långkvist et al., 2014; Benidis et al., 2020).

### 2.2. Electricity Price Forecasting

The Electricity Price Forecasting (EPF) task aims at predicting the spot (balancing, intraday, day-ahead) and forward prices in wholesale markets. Since the workhorse of short-term power trading is the day-ahead market with its conducted once-per-day uniform-price auction (Mayer & Trück, 2018), the vast majority of research has focused on predicting

electricity prices for the 24 hours of the next day, either in a point (Amjady, 2012; Weron, 2014; Lago et al., 2021) or a probabilistic setting (Nowotarski & Weron, 2018). There are, however, also studies on EPF for very short-term (Narajewski & Ziel, 2020), as well as mid- and long-term horizons (Ziel & Steinert, 2018). The recent expansion of renewable energy generation and large-scale battery storage has induced complex dynamics to the already volatile electricity spot prices, turning the field into a prolific subject on which to test novel forecasting ideas and trading strategies (Chitsaz et al., 2018; Gianfreda et al., 2020; Uniejewski & Weron, 2021).

Out of the numerous approaches to EPF developed over the last two decades (Weron & Ziel, 2020), two classes of models are of particular importance when predicting day-ahead prices – statistical (also called econometric or technical analysis) and computational intelligence (also referred to as artificial intelligence, non-linear or machine learning); to provide a comprehensive test ground we will use two well-performing benchmarks, one from each class. Moreover, many – if not most – of the proposed methods are hybrid solutions, that typically comprise data decomposition, feature selection, clustering, forecast averaging and/or heuristic optimization to estimate the (hyper)parameters (Lago et al., 2021).

Inspired by the success of deep learning (DL) in other fields and the outcomes of recent forecasting competitions (Hong et al., 2016; Smyl, 2019; Makridakis et al., 2020), the EPF community has focused on deep neural networks. The study of Lago et al. (2021) provides a comprehensive overview of the deep learning EPF papers published between 2016 and 2019. The consideration is broad and includes various neural network designs ranging from *Multi Layer Perceptrons* (MLPs; Luo & Weng 2019), *Recurrent Neural Networks* (RNNs; Ugurlu et al. 2018; Zhou et al. 2019), to *Convolutional Neural Networks* (CNNs; Afrasiabi et al. 2019), and methods that combine RNNs and CNNs (Kuo & Huang, 2018) or RNNs used to simultaneously forecast price and electric load (Mujeeb et al., 2019).

The method proposed in this paper belongs to the hybrid category which recently has gained substantial attention. These are examples of hybrid methods applied to EPF, Jahangir et al. (2020) use Gray correlation analysis and a stacked denoising autoencoder to preprocess data and feed it into a MLP point forecaster. Zhang et al. (2020) use an isolation forest and LASSO selection preprocessing combined with CNN point forecast that combines with final quantile regression. Zhang et al. (2020) use a categorical boosting feature selection and a bidirectional LSTM network for point forecasts. Rantonen & Korpiahkola (2020) propose a weighted averaging scheme of statistical and ML forecasts with an additional meta learner neural network governing the weights. Among recent extensions to well-established methods, He et al. (2020) combined a CNN with a label-distribution-learning-forest decoder to produce probabilistic predictions. Yang & Schell (2020) proposed a multi-branch three-layered GRU that improves a shallow GRU architecture for short-term probabilistic EPF. Finally, comprehensive comparative studies have recently become available to help the community appreciate attainable benefits and relative limitations of the hybrid approaches to EPF. Hong & Pula (2020) present a case study applying six different methods to the data from the PJM market; the methods include ARIMA and single and multi-stage RNNs. de Simón-Martín et al. (2020) performed a comparison between NAR, NARX, and LSTM in the Spanish MIBEL electricity market.

Unfortunately, the majority of the neural network EPF related research seems to suffer from some of the following issues: too short and limited to one market test periods, lack of well performing and established benchmark methods, and incomplete descriptions of the pipeline and training methodology resulting in a poor reproducibility of the results. From the abundant literature we have found the empirical studies of Mujeeb et al. (2019), Zhou et al. (2019) and Lago et al. (2021) to be perhaps the most rigorous and complete. Hence, our models will be compared against the highly competitive benchmarks recommended in the latter study: the *Lasso Estimated Auto-Regressive* (LEAR) model and a parsimonious *Deep Neural Network* (DNN).

### 3. NBEATSx Model

As a general overview, the NBEATSx framework decomposes the objective signal by performing separate local nonlinear projections of the target data onto basis functions across its different blocks. Figure 1 depicts the general architecture of the model. Each block consists of a *Fully Connected Neural Network* (FCNN) which learns expansion coefficients for the backcast and forecast elements. The backcast model is used to clean the inputs of subsequent blocks, while the forecasts are summed to compose the final prediction. The blocks are grouped in stacks. Each of the potentially multiple stacks specializes in a different variant of basis functions.

To continue the description of the NBEATSx, we introduce the following notation: the objective signal is represented by the vector  $\mathbf{y}$ , the inputs for the model are the backcast window vector  $\mathbf{y}^{back}$  of length  $L$ , and the forecast window vector  $\mathbf{y}^{for}$  of length  $H$ ; where  $L$  denotes the length of the lags available as classic autoregressive features, and  $H$  is the forecast horizon treated as the objective. While the original NBEATS only admits as regressor the backcast period of the target variable  $\mathbf{y}^{back}$ , the NBEATSx incorporates covariates in its analysis denoted with the matrix  $\mathbf{X}$ . Figure 1 shows an example where the target variable is the hourly electricity price, the backcast vector has a length  $L$  of 96 hours, and the forecast horizon  $H$  is 72 hours, in the example, the covariate matrix  $\mathbf{X}$  is composed of wind power production and electricity load. For the EPF comparative analysis of Section 4.3.6 the horizon considered is  $H = 24$  that corresponds to day-ahead predictions, while backcast inputs  $L = 168$  correspond to a week of lagged values.

#### 3.1. Stacks and Blocks

The NBEATSx is composed by  $S$  stacks of  $B$  blocks each, the input  $\mathbf{y}^{back}$  of the first block consists of  $L$  lags of the target time series  $\mathbf{y}$  and the exogenous matrix  $\mathbf{X}$ , while the inputs of the subsequent blocks include residual connections with the backcast output of the previous block. Within each  $l$ -th block, the first component transformation of the data  $(\mathbf{y}_{l-1}^{back}, \mathbf{X}_{l-1})$ , consists of a FCNN that learns hidden units  $\mathbf{h}_l$  that are linearly adapted into the forecast  $\boldsymbol{\theta}_l^{for}$  and backcast  $\boldsymbol{\theta}_l^{back}$  expansion coefficients. These block transformations described in Equation (1) correspond to the blue rectangle of Figure 1.

$$\begin{aligned} \mathbf{h}_l &= \text{FCNN}_l(\mathbf{y}_{l-1}^{back}, \mathbf{X}_{l-1}) \\ \boldsymbol{\theta}_l^{for} &= \text{LINEAR}^{for}(\mathbf{h}_l) \quad \boldsymbol{\theta}_l^{back} = \text{LINEAR}^{back}(\mathbf{h}_l) \end{aligned} \tag{1}$$




Figure 1: Building blocks of the NBEATSx are structured as a system of multilayer fully connected networks with ReLU based nonlinearities. Blocks overlap using the doubly residual stacking principle for the backcast  $\hat{\mathbf{y}}_l^{back}$  and forecast  $\hat{\mathbf{y}}_l^{for}$  outputs of the  $l$ -th block, and the blocks within a stack may share weights.

The second transformation within each  $l$ -th block consists of a basis expansion operation between the coefficients learnt by the FCNN and the block's basis vectors  $\mathbf{V}_l^{back}$  and  $\mathbf{V}_l^{for}$ , this transformation results in the backcast  $\hat{\mathbf{y}}_l^{back}$  and forecast  $\hat{\mathbf{y}}_l^{for}$  components, defined by:

$$\hat{\mathbf{y}}_l^{back} = \sum_{i=1}^{|\boldsymbol{\theta}_l^{back}|} \theta_{l,i}^{back} \mathbf{v}_{l,i}^{back} \equiv \boldsymbol{\theta}_l^{back} \mathbf{V}_l^{back} \quad \text{and} \quad \hat{\mathbf{y}}_l^{for} = \sum_{i=1}^{|\boldsymbol{\theta}_l^{for}|} \theta_{l,i}^{for} \mathbf{v}_{l,i}^{for} \equiv \boldsymbol{\theta}_l^{for} \mathbf{V}_l^{for} \quad (2)$$

### 3.2. Residual Connections:

The connections between blocks are given by the doubly residual stack, from Equation (3) are depicted in the brown and yellow rectangles of Figure 1. The architecture creates a separate residual connection for the backcast  $\hat{\mathbf{y}}^{back}$  and forecast  $\hat{\mathbf{y}}^{for}$ :

$$\mathbf{y}_l^{back} = \mathbf{y}_{l-1}^{back} - \hat{\mathbf{y}}_{l-1}^{back} \quad \text{and} \quad \hat{\mathbf{y}}^{for} = \sum_{l=1}^{S \times B} \hat{\mathbf{y}}_l^{for} \quad (3)$$

The doubly residual stack allows us to express the final prediction  $\hat{\mathbf{y}}^{for}$  as a sum of the partial predictions of each block, which translates into a very intuitive decomposition of the prediction when the bases within the blocks are interpretable. On the other hand, the

residual backcast  $\mathbf{y}_l^{back}$  allows the model to subtract the component associated to the basis of the  $l$ -th block  $\mathbf{V}_l^{back}$  from  $\mathbf{y}^{back}$ , which can be also thought of as a sequential decomposition of the modeled signal. In turn, this methodology helps with the optimization procedure as it allows for an improved flow of the gradients during optimization.

### 3.3. NBEATSx Configurations

The original neural basis expansion analysis method proposed two configurations based on the assumptions encoded in the learning algorithm by selecting the basis vectors  $\mathbf{V}_l^{for}$  used in the blocks from Equation (2). A mindful selection of restrictions to the basis allows the model to output an interpretable decomposition of the forecasts, while setting  $\mathbf{V}_l^{for} = \mathbf{I}_{H \times H}$  with the horizon being the coefficient's cardinality  $|\boldsymbol{\theta}_l^{for}| = H$ , can produce more flexible forecasts by effectively removing any constraints on the form of the basis functions. In this subsection, we present both interpretable and generic configurations in detail, explaining in particular how we propose to include the covariates in each case.

#### 3.3.1. Interpretable Configuration

The choice of basis vectors relies on time series decomposition techniques that are often used to understand the structure of a given time series and patterns of its variation. Work in this area ranges from classical smoothing methods (Macaulay, 1931) and their extensions such as X-11-ARIMA (Shishkin et al., 1967; Dagum, 1980), X-12-ARIMA (Findley et al., 1998), and X-13-ARIMA-SEATS (U.S. Census Bureau, 2013), to modern approaches such as TBATS (Livera et al., 2011), and STR (Dokumentov & Hyndman, 2015). To encourage interpretability, the blocks within each stack may use harmonic functions, polynomial trends, and exogenous variables directly to perform their projections upon. The partial forecasts of the interpretable configuration are described through Equations (4)-(6). Let the time vector  $\mathbf{t} = [0, 1, 2, \dots, H-2, H-1]/H$  is defined on a discrete grid. When the basis  $\mathbf{V}_l^{for}$  is given by  $\mathbf{T} = [\mathbf{1}, \mathbf{t}, \dots, \mathbf{t}^p]$ , where  $p$  is the maximum polynomial degree, the coefficients are those of a polynomial model for the trend. When the bases  $\mathbf{V}_l^{for}$  are harmonic  $\mathbf{S} = [\mathbf{1}, \cos(2\pi\mathbf{t}), \dots, \cos(2\pi\lfloor H/2 - 1 \rfloor \mathbf{t}), \sin(2\pi\mathbf{t}), \dots, \sin(2\pi\lfloor H/2 - 1 \rfloor \mathbf{t})]$ , the coefficients vector  $\boldsymbol{\theta}_l^{for}$  can be interpreted as Fourier transform coefficients. The exogenous basis expansion can be thought as a time-varying local regression when the basis is given by matrix  $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_{N_x}]$ , where  $N_x$  is the number of exogenous variables. The resulting models can flexibly reflect common structural assumptions, in particular using the interpretable bases, as well as their combinations.

$$\hat{\mathbf{y}}_l^{trend} = \sum_{i=0}^p \theta_{l,i}^{trend} \mathbf{t}^i \equiv \boldsymbol{\theta}_l^{trend} \mathbf{T} \quad (4)$$

$$\hat{\mathbf{y}}_l^{seas} = \sum_{i=0}^{\lfloor H/2-1 \rfloor} \theta_{l,i}^{seas} \cos(2\pi i \mathbf{t}) + \theta_{l,i+\lfloor H/2 \rfloor}^{seas} \sin(2\pi i \mathbf{t}) \equiv \boldsymbol{\theta}_l^{seas} \mathbf{S} \quad (5)$$

$$\hat{\mathbf{y}}_l^{exog} = \sum_{i=0}^{N_x} \theta_{l,i}^{exog} \mathbf{x}_i \equiv \boldsymbol{\theta}_l^{exog} \mathbf{X} \quad (6)$$



In this paper, we propose including one more type of stack to specifically represent exogenous variable basis as described in Equation (6) and depicted in Figure 1. In the original NBEATS framework (Oreshkin et al. (2020)), the interpretable configuration usually consists of a trend stack followed by a seasonality stack, each containing three blocks. Our NBEATSx extension of this configuration consists of three stacks, one of each type of factors (trend, seasonal, exogenous). We refer to this extended interpretable configuration as the NBEATSx-I model in the remainder of the paper.

### 3.3.2. Generic Configuration

For the generic configuration, the basis of the non linear projection in Equation (2) corresponds to canonical vectors, as specified in Equation (7). This basis enables NBEATSx to effectively behave like a classic *Fully Connected Neural Network* (FCNN). The output layer of the FCNN inside each block has  $H$  neurons, that correspond to the forecast horizon, each producing the forecast for one particular time point of the forecast period. This can be seen as the basis vectors being learned during optimization, allowing the waveform of the basis of each stack to be freely determined in a data-driven fashion. Compared to the interpretable counterpart described in Section 3.3.1, the constraints on the form of the basis functions are removed. This affords the generic variant more flexibility and power at representing complex data, but it can also lead to less interpretable outcomes and potentially escalated risk of overfitting.

$$\hat{\mathbf{y}}_l^{gen} = \boldsymbol{\theta}_l^{for} \mathbf{V}_l^{for} = \boldsymbol{\theta}_l^{for} \quad (7)$$

For the NBEATSx model with the generic configuration, we propose a new type of exogenous block that learns a context vector  $\mathbf{C}_l$  from the time-dependent covariates with an *encoder* convolutional sub-structure:

$$\hat{\mathbf{y}}_l^{exog} = \sum_{i=1}^{N_c} \theta_{l,i}^f C_{l,i} \equiv \boldsymbol{\theta}_l^f \mathbf{C}_l \quad \text{with} \quad \mathbf{C}_l = \text{TCN}(\mathbf{X}_l) \quad (8)$$

In the previous equation the *Temporal Convolutional Network* (TCN; Bai et al. 2018) is employed as an *encoder*, still any neural network with a sequential structure is compatible with the backcast and forecast branches of the model, and can be used as an *encoder*. For example the *WaveNet* (van den Oord et al., 2016) can be effective an alternative to RNNs as it also able to capture long term dependencies and interactions of covariates by stacking multiple layers, while dilations help keeping the models computationally tractable. In addition, convolutions have a very convenient interpretation as a weighted moving average signal filters. The final linear projection and the additive composition of the predictions can be interpreted as a *decoder*.

The original NBEATS configuration includes only one generic stack with dozens of blocks, while our proposed model includes both the generic and exogenous stacks, with the order determined via data-driven hyperparameter tuning. We refer to this configuration as the NBEATSx-G model.

### 3.3.3. Exogenous Variables

This subsection briefly discusses different types of exogenous variables and how they are incorporated by our proposed NBEATSx method. We distinguish the exogenous variables by whether they reflect static or time-dependent aspects of the modeled data.

The static exogenous variables carry time-invariant information. When the model is built with common parameters to forecast multiple time series, these variables allow sharing information within groups of time series with similar static variable levels. Examples of static variables include designators such as identifiers of regions, groups of products, etc., that can be used to mark agglomerates of series that demonstrate similar behaviors. Since the static variables are used as inputs of the FCNN in each block, the NBEATSx allows for interactions between the static and temporal variables.

As for the time-dependent exogenous covariates, we discern two subtypes. First, we consider seasonal covariates that are linked to the natural frequencies in the data. In the EPF example, these variables include calendar variables to identify hours, days, months, or holidays, among others. They can be implemented using one-hot-encoded vectors of indicators or linear kernels centered around holidays, or days of the week. Linear kernels are inspired by the *Generalized Additive Model* (GAM) heuristics from the time series literature (Evgeniou et al., 2000; Gaillard et al., 2016). These variables are useful for NBEATSx to identify seasonal patterns and special events outside the window lookback periods. They can also help the models reflect sharp contrasts induced by discrete events or apparent change-points, which are hard to represent well using smooth bases such as polynomial or harmonic functions.

Finally, we identify domain-specific temporal covariates which can be unique to each problem. These variables are usually known to affect the target variable fundamentally and need to be incorporated by models to produce accurate forecasts. In the EPF setting, they typically include day-ahead forecasts of electricity load and production levels from renewable energy sources. Time-dependent covariates are included as inputs of the FCNN in all blocks and basis vectors for the blocks in the exogenous stack, as explained in Section 3.

## 4. Empirical Evaluation

### 4.1. Electricity Price Forecasting Datasets

To evaluate our method’s forecasting capabilities, we consider short-term electricity price forecasting tasks, where the objective is to predict day-ahead prices. Five major power markets<sup>1</sup> are used in the empirical evaluation, all comprised of hourly observations of the prices and two influential temporal exogenous variables that extend for 2,184 days (312 weeks, six years). From the six years of available data for each market, we hold two years out, to test the forecasting performance of the algorithms. The length and diversity of the test sets allow us to obtain accurate and highly comprehensive measurements of the robustness and the generalization capabilities of the models.

---

<sup>1</sup>For the sake of reproducibility we only consider datasets that are openly accessible in the EPFtoolbox library <https://github.com/jeslago/epftoolbox> (Lago et al., 2021).

Table 1: Datasets used in our empirical study. For the five day-ahead electricity market considered, we report the test period dates and two influential covariate variables.

Market	Exogenous Variable 1	Exogenous Variable 2	Test Period
NP	day-ahead load	day-ahead wind generation	27-12-2016 to 24-12-2018
PJM	2 day-ahead system load	2 day-ahead COMED load	27-12-2016 to 24-12-2018
EPEX-FR	day-ahead load	day-ahead total France generation	04-01-2015 to 31-12-2016
EPEX-BE	day-ahead load	day-ahead total France generation	04-01-2015 to 31-12-2016
EPEX-DE	day-ahead zonal load	day-ahead wind generation	04-01-2016 to 31-12-2017

Table 1 summarizes the key characteristics of each market. The Nord Pool electricity market (NP), which corresponds to the Nordic countries exchange, contains the hourly prices and day-ahead forecasts of load and wind generation. The second dataset is the Pennsylvania-New Jersey-Maryland market in the United States (PJM), which contains hourly zonal prices in the Commonwealth Edison (COMED) and two day-ahead forecasts of load at the system and COMED zonal levels. The remaining three markets are obtained from the integrated European Power Exchange (EPEX). Belgium (EPEX-BE) and France (EPEX-FR) markets share the day-ahead forecast generation in France as covariates since it is known to be one of the best predictors for Belgian prices (Lago et al., 2018b). Finally, the German market (EPEX-DE) contains the hourly prices, day-ahead load forecasts, and the country level wind and solar generation day-ahead forecast.

Figure 2 displays the NP electricity price time series and its corresponding covariate variables to illustrate the datasets. The NP market is the least volatile among the considered markets, since most of its power comes from hydroelectric generation, renewable source volatility is negligible, and zero spikes are rare. The PJM market is transitioning from coal generation to natural gas and some renewable sources, zero spikes are rare, but the system exhibits higher volatility than NP. In EPEX-BE and EPEX-FR markets, negative prices and spikes are more frequent, and as time passes, these markets begin to show increasing signs of integration. Finally, the EPEX-DE market shows few price spikes, but the most frequent negative and zero price events, due in great part to the impact of renewable sources.

The exogenous covariates are normalized following best practices drawn from the EPF literature (Uniejewski et al., 2018), preprocessing the inputs of neural networks is essential to accelerate and stabilize the optimization (LeCun et al., 1998).

#### 4.2. Interpretable Time Series Decomposition

In this subsection, we demonstrate the versatility of the proposed method and show how a careful selection of the inductive bias, constituted by the assumptions used to learn the modeled signal, endows NBEATSx with an outstanding ability to model complex dynamics while enabling human understanding of its outputs, turning it into a unique and exciting tool for time series analysis. Our method combines the power of non-linear transformations provided by neural networks with the flexibility to model multiple seasons that can be fractional and simultaneously account for interaction events such as holidays and other covariates.




Figure 2: The top panel shows the *day-ahead* electricity price time series for the NordPool market, the vertical dashed lines mark the beginning of the 728-day long held-out sample period. The second and third panels show the day-ahead forecast for the system load and wind generation.

As described earlier, the interpretable configuration of the NBEATSx architecture computes time-varying coefficients for slowly changing polynomial functions to model the trend, harmonic functions to model the cyclical behavior of the signal, and exogenous covariates. Here, we show how this configuration can decompose a time series into the classic set of level, trend, and seasonality components while identifying the effects of exogenous covariates. The time series decomposition is similar to what can be obtained using STL (Cleveland et al., 1990) or X-13-ARIMA-SEATS (U.S. Census Bureau, 2013).

By simply varying the dimensionality of the model’s output  $H$ , we can apply NBEATSx to forecast over short or long horizons, or as an interpolation or extrapolation tool, according to the user’s needs. In Section 4.3.4, we provide the details of the cross-validation based hyperparameter selection approach used in our experiments.

#### 4.2.1. Week-long Hourly Electricity Price Decomposition

Figure 3 shows the NP electricity market’s hourly price, in Euros per megawatt, from May 28, 2013 to June 03, 2013, and the interpretable NBEATSx-I based decomposition from top to bottom: the first panel shows the original time series for the chosen week and the level, that corresponds to the last available value before the prediction; the second panel shows the inferred trend component, the third panel shows the seasonal component, the fourth panel exhibits the effects of exogenous covariates, and the last panel tracks the forecast residuals that can provide useful insights to detect unexpected variations of the signal. The data shows a minor effect of the trend, an evident seasonal effect within the days, and the system load effects with its weekend interactions identified by the model through the exogenous component. This example shows how the seasonal covariates can help NBEATSx identify longer seasonal effects than those included in the backcast period.

#### 4.2.2. Day-ahead Hourly Electricity Price Decomposition

In this time series decomposition application, we show how the NBEATSx model can benefit from explicitly accounting for information carried by exogenous covariates, unlike its original counterpart. We compare the interpretable variants of NBEATS-I and NBEATSx-I, to show the advantageous impact of incorporating exogenous variables on the accuracy of the forecasts and their interpretability.

Figure 4 shows the NP electricity market’s hourly price, in Euros per megawatt, for December 18, 2017 which is a day with high prices due to high load. Other days have a less pronounced difference between the results obtained with the original NBEATS and the NBEATSx. We purposely selected a day with a higher than normal electricity load for exposition purposes, to demonstrate qualitative differences in the forecasts. Section 4.3.6 includes quantitative exploration of the relative benefits of the NBEATSx model.

Here, we apply the interpretable configurations of both NBEATS and NBEATSx for the day-ahead forecast task. We can clearly see a substantial difference in forecast residual magnitudes in the bottom row of graphs in Figure 4. The original model shows a strong negative bias. On the other hand, NBEATSx is able to capture the evidently substantial explanatory value of the exogenous features, resulting in a much more accurate forecast.




Figure 3: NP electricity price week-ahead forecast decomposed using the interpretable NBEATSx. The top graph shows the original signal and the level, the latter is defined as the last available observation before the forecast, the second row corresponds to the polynomial trend, the third and fourth graphs display complex seasonality modeled by non linear Fourier projections and the exogenous effects of the electricity load on the price, and the final panel depicts the unexplained variation of the signal.




Figure 4: NP electricity price day-ahead forecasts decomposed using interpretable variants of NBEATS and NBEATSx. The top row of graphs shows the original signal and the level, the latter is defined as the last available observation before the forecast. The second row shows the polynomial trend components, the third and fourth rows display the complex seasonality modeled by nonlinear Fourier projections and the exogenous effects of the electricity load on the price, respectively. The bottom row graphs show the unexplained variation of the signal. The use of electricity load and production forecasts turns out to be fundamental for accurate price forecasting.

### 4.3. Forecasting Comparative Analysis

#### 4.3.1. Evaluation Metrics

To ensure the comparability of our results with the existing literature, we opted to follow the widely accepted practice of evaluating the accuracy of point forecasts with the following metrics: *mean absolute error* (MAE), *relative mean absolute error* (rMAE)<sup>2</sup>, *symmetric mean absolute percentage error* (sMAPE) and *root mean squared error* (RMSE), defined as follows:

$$MAE = \frac{1}{24N_d} \sum_{d=1}^{N_d} \sum_{h=1}^{24} |y_{d,h} - \hat{y}_{d,h}| \quad rMAE = \frac{\sum_{d=1}^{N_d} \sum_{h=1}^{24} |y_{d,h} - \hat{y}_{d,h}|}{\sum_{d=1}^{N_d} \sum_{h=1}^{24} |y_{d,h} - \hat{y}_{d,h}^{naive}|}$$

$$sMAPE = \frac{200}{24N_d} \sum_{d=1}^{N_d} \sum_{h=1}^{24} \frac{|y_{d,h} - \hat{y}_{d,h}|}{|y_{d,h}| + |\hat{y}_{d,h}|} \quad RMSE = \sqrt{\frac{1}{24N_d} \sum_{d=1}^{N_d} \sum_{h=1}^{24} (y_{d,h} - \hat{y}_{d,h})^2}$$

where  $y_{d,h}$  and  $\hat{y}_{d,h}$  are the actual value and the forecast of the time series at day  $d$  and hour  $h$ . The MAE and RMSE measure the errors in absolute terms, and as such they are not easily comparable across time series. For this reason, we also include the sMAPE and rMAE metrics that are relative. The sMAPE is included as a robust alternative to MAPE which in the presence of values close to zero may degenerate (Hyndman & Koehler, 2006). The difference between the *relative mean absolute error* (rMAE) and *mean absolute scaled error* (MASE) is that the *relative mean absolute error* includes the out-of-sample predictions which helps to compare the relative performance of the model with the actual accuracy of the seasonal naive benchmark (Lago et al., 2021).

#### 4.3.2. Statistical Tests

To assess which forecasting model provides better predictions we rely on the *Giacomini-White test* (GW; Giacomini & White 2006) of the multi-step conditional predictive ability, which can be interpreted as a generalization of the *Diebold-Mariano test* (DM; Diebold & Mariano 2002), widely used in the forecasting literature. Compared with the DM or other unconditional tests, the GW test is valid under general assumptions such as heterogeneity rather than stationarity of data. The GW test examines the null hypothesis of equal accuracy specified in Equation (9), measured by the  $L1$  norm of the daily errors of a pair of models  $A$  and  $B$ , conditioned on the available information to that moment<sup>3</sup> in time  $\mathcal{F}_{d-1}$ .

$$H_0 : \mathbb{E} [||\mathbf{y}_d - \hat{\mathbf{y}}_d^A||_1 - ||\mathbf{y}_d - \hat{\mathbf{y}}_d^B||_1 \mid \mathcal{F}_{d-1}] \equiv \mathbb{E} [\Delta_d^{A,B} \mid \mathcal{F}_{d-1}] = 0 \quad (9)$$

<sup>2</sup>The naïve forecast method in EPF corresponds to a similar day rule, where the forecast for a Monday, Saturday and Sunday equals the value of the series observed on the same weekday of the previous week, while the forecast for Tuesday, Wednesday, Thursday, and Friday is the value observed on the previous day.

<sup>3</sup>In practice, the available information  $\mathcal{F}_{d-1}$  is replaced with a constant and lags of the error difference  $\Delta_d^{A,B}$  and the test is performed using a linear regression with a Wald like test. When the conditional information considered is only the constant variable, one recovers the original DB test.






Figure 5: Partitioning of data for training and evaluation of models. Training data covers the first four years of each data set. The early stopping set is formed by 42 weeks randomly chosen from the training sample (here, marked with blue rectangles). The validation set, used to empirically optimize hyperparameters, spans the last 52 weeks of the training sample, and the test set is the last two years of the data.

#### 4.3.3. Training Methodology

The cornerstone of the training methodology for NBEATSx and the benchmark models included in this work is the definition and use of the training, validation, early stopping, and test datasets summarized in Figure 5. The training set for each of the five markets comprises the first three years of data, the test set includes the last two years of data. The validation set is defined as the last 52 weeks previous to the test set. During the recalibration and prediction phase, the early stopping set is composed of 42 weeks randomly sampled from training data.

During the hyperparameter optimization phase, model performance measured on the validation set is used to guide the exploration of the hyperparameter space defined in Table 2. During the inference over the test set, the optimally selected model is re-calibrated for each day to include newly available information, an early stopping set provides a regularization signal for the retraining optimization.

To train the neural network, we minimize the *mean absolute error* (MAE) using stochastic gradient descent with *Adaptive Moments* (ADAM; Kingma & Ba 2014). As a form of regularization, we use an early stopping strategy that halts the training procedure if a specified number of consecutive iterations occur without improvements of the loss measured on the early stopping set (Yao et al., 2007). The learning rate, batch size, and learning rate decay schedule are determined by the hyperparameter optimization.

The NBEATSx model is implemented and trained in PyTorch (Paszke et al., 2019) and can be run with both CPU and GPU resources. The code is available publicly at the NBEATSx GitHub repository (<https://github.com/nixtla>) to promote reproducibility of the presented results and to support related research.

Table 2: Hyperparameters of NBEATSx networks. They are common to all presented datasets. We list the typical values we considered in our experiments. The configuration that performed best on the validation set was selected automatically.

Hyperparameter	Description	Considered Values
Architecture Parameters		
input_size_multiplier	Input size is a multiple of the output horizon.	{7}
output_size	Output size is the forecast horizon for day ahead forecasting.	{24}
stack_types	Type and number of stacks in the architecture.	Specified in code.
activation	Type of activations used accross the network.	{softplus,selu,prelu,sigmoid}
n_blocks	Number of blocks per stack, residual connections between blocks.	Specified in code.
n_layers	Number of layers within each block.	{2}
n_hidden	Number of neurons on each layer of a block.	{50, ..., 500}
n_x_channels	Number of channels for encoded temporal exogenous variables.	{2, ..., 10}
n_s_channels	Number of channels for encoded static exogenous variables.	{1, ..., 10}
theta_with_x	Whether NBEATSx coefficients receive the covariates as input.	{True, False}
n_polynomials	Only interpretable, degree of trend polynomials.	{2}
n_harmonics	Only interpretable, number of Fourier basis (seasonality smoothness).	{6}
Optimization and Regularization parameters		
initialization	Initialization strategy for network weights.	{orthogonal, he_norm, glorot_norm}
learning_rate	Initial learning rate for regression problem.	[5e-4, 1e-2]
batch_size	The number of samples for each gradient step.	{256, 512}
lr_decay	The decay constant allows large initial lr to escape local minima.	{0.5}
n_lr_decay_steps	Number of times the learning rate is halved during train.	{3}
n_iterations	Maximum number of iterations of gradient descent.	{30000}
early_stopping	Consecutive iterations without validation loss improvement before stop.	{10}
early_stopping_steps	Frequency of validation loss measurements.	{100}
batch_normalization	Whether batch normalization is applied after each activation.	{True, False}
dropout_prob_theta	The probability for dropout of neurons for all in the projection layers.	[0, 1]
dropout_prob_x	The probability for dropout of neurons for the exogenous encoder.	[0, 1]
l1_theta	Constant to control the Lasso penalty used on the coefficients.	[0, 0.1]
weight_decay	Constant that controls the influence of L2 regularization of weights.	[1e-5, 1e-0]
loss	The objective loss function with which the network is trained.	{MAE}
val_weeks	Random weeks from full dataset used to validate.	{42}
hyperopt_iters	Number of iterations of hyperparameter search.	{1500}
random_seed	Random seed that controls initialization of weights.	{1, ..., 1000}
Data Parameters		
sample_freq	Rolling window sample frequency, for data augmentation.	{1, 24}
window_sampling_limit	Number of time windows included in the full dataset.	4 years
n_val_weeks	Number of validation weeks used for early stopping strategy.	{40, 52}
normalizer_x/y	Normalization strategy of model inputs.	{none, median, invariant }

#### 4.3.4. Hyperparameter Optimization

We follow the practice of Lago et al. (2018a) to select the hyperparameters that define the model, input features, and optimization settings. During this phase, the validation dataset is used to guide the search for well performing configurations. To compare the benchmarks and the NBEATSx, we rely on the same automated selection process: a Bayesian optimization technique that efficiently explores the hyperparameter space using tree-structured Parzen estimators (HYPEROPT; Bergstra et al. 2011). The architecture, optimization, and regularization hyperparameters are summarized in Table 2.

#### 4.3.5. Ensembling

In many recent forecasting competitions, and particularly in the M4, competition most of the top-performing models were ensembles (Atiya, 2020), as it has been often shown that in practice, combining a diverse group of models can be a powerful form of regularization to reduce the variance of predictions (Breiman, 1996; Bergmeir et al., 2016; Nowotarski et al., 2014; Hubicka et al., 2018).

The techniques used by the forecasting community to induce diversity in the models are plentiful. The original NBEATS model obtained its diversity from three sources, training with different loss functions, varying the size of the input windows, and bagging models with different random initializations (Oreshkin et al., 2020), and using the median as the aggregation function for 180 different models; interestingly the original model did not rely on

regularization like L2 or dropout as Oreshkin et al. (2020) found it to be good for individual models but detrimental to the ensemble. In our case, we ensemble the NBEATSx model using two sources of diversity, the first being a data augmentation<sup>4</sup> used during the training procedure, and the second being whether a random set was used as validation or the last available sample before the test set; we ensembled four models using the mean as aggregation function. This technique is also used by the DNN benchmark (Lago et al., 2018a, 2021).

#### 4.3.6. Forecasting Results

We conducted an empirical study involving two types of *Autoregressive Models* (AR1 and ARx1; Weron 2014), the *Lasso Estimated Auto-Regressive* (LEARx; Uniejewski et al. 2016), a parsimonious *Deep Neural Network* (DNN; Lago et al. 2018a, 2021), the original *Neural Basis Expansion Analysis* without exogenous covariates (NBEATS; Oreshkin et al. 2020), and the *Exponential Smoothing Recurrent Neural Network* (ESRNN; Smyl 2019). This experiment examines the effects of including the covariate inputs and comparing NBEATSx with state-of-the-art methods for the electricity price day-ahead forecasting task.

Table 3 summarizes the performance of the ensembled models where NBEATSx ensemble shows prevailing performance. It improves 18.77% on average for all metrics and markets when compared with the original NBEATS and 20.6% when compared to ESRNN without time-dependent covariates. For the ensembled models, NBEATSx RMSE improved on average 4.68%, MAE improved 2.53%, rMAE improved 1.97%, and sMAPE improved 1.25%. When comparing NBEATSx ensemble against DNN ensemble on individual markets, NBEATSx improved by 5.38% on the NordPool market, by 2.48% on French market and 2.81% on German market. There was a non-significant difference of NBEATSx performance on PJM and BE markets of 0.24% and 1.1%, respectively.

Figure 6 provides a graphical representation of the statistical significance from the *Giacomini-White* test (GW) for the six ensembled models, across the five markets for the MAE evaluation metric. A similar significance analysis is conducted for the single models. The models included in the significance tests are the same as in Table 3: LEAR, DNN, ESRNN, NBEATS, and our proposed methods, NBEATSx-G and NBEATSx-I. The  $p$ -value of each comparison shows if the performance improvement of the model’s predictions corresponding to the column index of a cell in the grids shown in Figure 6 over the model’s predictions corresponding to the row of this cell of the grid is statistically significant. NBEATSx-G model outperformed DNN model in NP and DE, while NBEATSx-I outperformed it in NP, FR, and DE. Moreover, no benchmark model significantly outperformed NBEATSx-I and NBEATSx-G in any market.

In the Appendix A we observe similar results for the single best models chosen from the four possible configurations of the ensemble components described in Section 4.3.5. Table A1 summarizes the accuracy of the predictions measured with the MAE and Figure A.1 displays the significance of the GW test. Ensembling improves the accuracy of NBEATSx by 3% on average across all markets, when compared to the single best models.

---

<sup>4</sup>Data augmentation is controlled by the sample frequency of the windows used during training.

Table 3: Forecast accuracy measures for day-ahead electricity price predictions of *ensembled models*. The ESRNN and NBEATS do not include time dependent covariates. The reported metrics are *mean absolute error* (MAE), *relative mean absolute error* (rMAE), *symmetric mean absolute percentage error* (sMAPE) and *root mean squared error* (RMSE). The smallest errors in each row are highlighted in bold.

		AR1	ESRNN	NBEATS	ARx1	LEARx	DNN	NBEATSx-G	NBEATSx-I
NP	MAE	2.26	2.09	2.08	2.01	1.74	1.68	<b>1.58</b>	1.62
	rMAE	0.71	0.66	0.66	0.63	0.55	0.53	<b>0.50</b>	0.51
	sMAPE	6.47	6.04	5.96	5.84	5.01	4.88	<b>4.63</b>	4.70
	RMSE	4.08	3.89	3.94	3.71	3.36	3.32	<b>3.16</b>	3.27
PJM	MAE	3.83	3.59	3.49	3.53	3.01	<b>2.86</b>	2.91	2.90
	rMAE	0.79	0.74	0.72	0.73	0.62	<b>0.59</b>	0.60	0.60
	sMAPE	14.5	14.12	13.57	13.64	11.98	<b>11.33</b>	11.54	11.61
	RMSE	6.24	5.83	5.64	5.74	5.13	5.04	5.02	<b>4.84</b>
EPEX-BE	MAE	7.2	6.96	6.84	7.19	6.14	<b>5.87</b>	5.95	6.11
	rMAE	0.88	0.85	0.83	0.88	0.75	<b>0.72</b>	0.73	0.75
	sMAPE	16.26	15.84	15.80	16.11	14.55	<b>13.45</b>	13.86	14.02
	RMSE	18.62	16.84	17.13	18.07	15.97	15.97	<b>15.76</b>	15.80
EPEX-FR	MAE	4.65	4.65	4.74	4.56	3.98	3.87	3.81	<b>3.79</b>
	rMAE	0.78	0.78	0.80	0.76	0.67	0.65	<b>0.64</b>	<b>0.64</b>
	sMAPE	13.03	13.22	13.30	12.7	11.57	10.81	<b>10.59</b>	<b>10.69</b>
	RMSE	13.89	11.83	12.01	12.94	10.68	11.87	11.50	<b>11.25</b>
EPEX-DE	MAE	5.74	5.60	5.31	4.36	3.96	3.41	3.31	<b>3.29</b>
	rMAE	0.71	0.70	0.66	0.54	0.49	0.42	<b>0.41</b>	<b>0.41</b>
	sMAPE	21.37	20.97	19.61	17.73	15.75	14.08	<b>13.99</b>	<b>13.99</b>
	RMSE	9.63	9.09	8.99	7.38	7.08	5.93	5.72	<b>5.65</b>

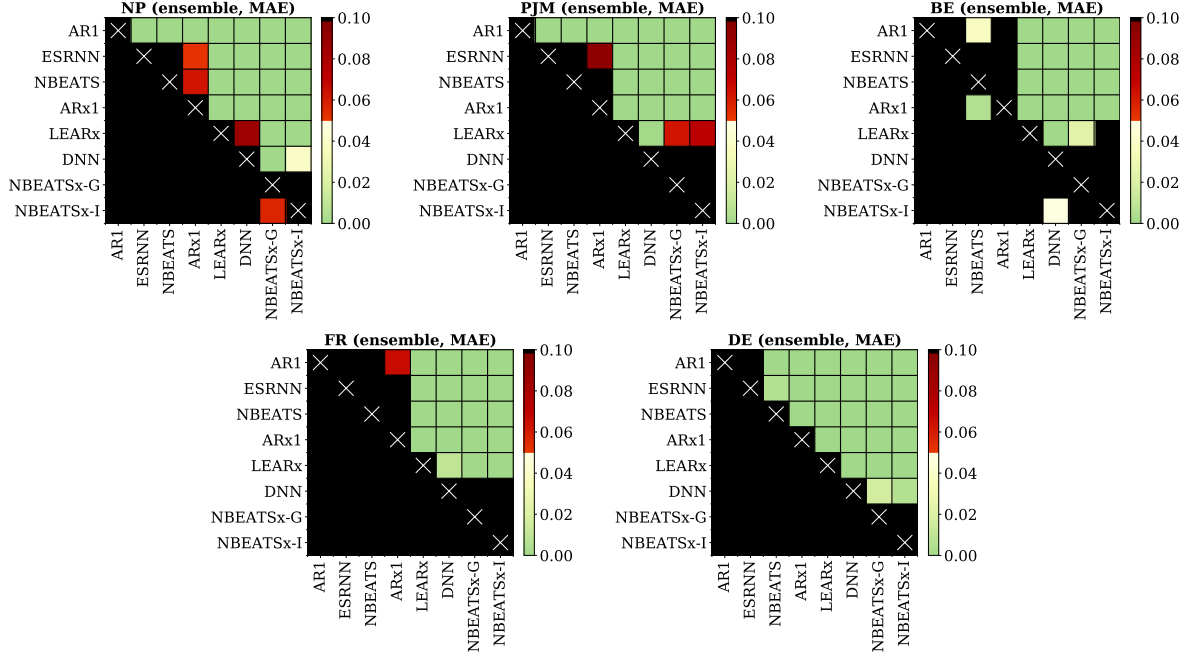


Figure 6: Results of the Giacomini-White test for the day-ahead predictions with *mean absolute error* (MAE) applied to pairs of the ensembled models on the five electricity markets datasets. Each grid represents one market. Each colored cell in a grid is plotted black, unless the predictions of the model corresponding to its column of the grid outperforms the predictions of the model corresponding to its row of the grid. The color scale reflects significance of the difference in MAE, with solid green representing the lowest  $p$ -values.

## 5. Possible Extensions of NBEATSx

Neural basis expansion analysis is a flexible and powerful framework for advanced forecast modeling, yet we identified some limitations and other interesting questions that arose from our empirical evaluations. First, the observed improvements pertain to point forecasts only, while other models, such as the ESRNN, already offer probabilistic forecasts. Second, the generic NBEATSx-G and interpretable NBEATSx-I configurations did not show significant differences in their performance on the selected benchmark tests. This raises an interesting question on the possibility of the simpler interpretable model outperforming the generic counterpart on scarce data settings by being more robust to overfitting and by using specialized alternative basis expansions tailored to the problem at hand. Further research on the use of neural networks for time series forecasting in scarce data settings can overcome their limitations and further their adoption in practice.

Below, we elaborate on a few ideas that we believe could help to further expand the applications of NBEATSx. These extension ideas address first the outputs and functionalities of the model, and second the challenge of data scarcity that is exacerbated by the large number of adaptable parameters in the model.

### 5.1. Probabilistic Forecasts

NBEATSx can be extended towards probabilistic forecasting using a similar approach to that of quantile regression. Without selecting a particular distribution, one can precisely estimate quantiles of the sample distribution by simply changing the training objective function from the mean absolute error (MAE) to the quantile loss (Koenker & Bassett, 1978; Uniejewski & Weron, 2021). This approach has the additional benefit of providing robustness to outliers and noise in the data. In some cases, it may even be used to reduce systemic biases in the forecasts (Smyl, 2019).

### 5.2. Alternative Basis Expansions

As an alternative to the use of one dimensional convolutions in the exogenous covariate encoders, the interpretable configuration of NBEATSx could use other transformations to filter and extract useful features from data. Instead of using the convolutional weighted average procedure from Equation 8, one can use smoothing alternatives such as splines (Hastie & Tibshirani, 1986; Friedman, 1991). As an alternative to Fourier projections from Equation (6) one could also explore the wavelet transforms (Donoho & Johnstone, 1995), that have the benefit of being adaptive to the smoothness or sharpness of the signal. In settings with limited training data, these ideas can replace the automatic featurization, implicit in the representations learned by deeper architectures, with featurization partially driven by time series domain knowledge.

### 5.3. Data Scarcity and Transfer Learning

In applications with limited amounts of training data, all highly parametrized models are vulnerable to overfitting and generalize poorly. One known way to alleviate the problem in the univariate time series setting is by using longer periods of training data, although this may not always be feasible. An interesting alternative idea to explore with NBEATSx model is to optimize it in stages: first on a larger reference dataset, to then fine-tune it for the, perhaps scarce, target data set (Bzinovski & Fulgosi, 1976; Pan & Yang, 2010). The use of a model pre-trained on a large dataset in addition to being able to mitigate overfitting, can also help to predict for time series whose length can be short even for classical statistical methods. This is possible due to the ability of machine learning models to capture common patterns shared between correlated series.

Attempts to implement transfer learning for neural networks in the time-series domain have been limited (Fawaz et al., 2018; Nikolay Laptev & Rajagopal, 2018), but literature has shown that the results can improve or deteriorate depending on the datasets involved in the experiments. Yet, results attainable in computer vision domain (Bengio, 2012; Yosinski et al., 2014) or natural language processing (Ramachandran et al., 2017; Devlin et al., 2019) appear more systematically beneficial. This apparent discrepancy requires more investigation, and we hypothesize that different sets of data in time series domain may show less obvious similarities than databases in those more common domains of application of deep transfer learning. Nonetheless, this avenue of investigation remains promising.

## 6. Conclusion

We have presented NBEATSx: the new method for univariate time series forecasting with exogenous variables. It expands the well-performing neural basis expansion analysis. The resulting neural network based method has several valuable properties that make it suitable for a wide range of forecasting problems. It is fast to optimize as it is mainly composed of fully-connected layers. It can produce interpretable results, and achieves state-of-the-art performance on forecasting tasks where consideration of exogenous variables is fundamental.

We demonstrated the utility of the proposed method using a set of benchmark datasets from electricity price forecasting domain, but it can be straightforwardly applied to forecasting problems in other domains. Qualitative evaluation shows that the interpretable configuration of NBEATSx can provide valuable insights to the analyst, as it explains the variation of the time-series by separating it into trend, seasonality, and exogenous components, in a fashion analogous to classic time series decomposition. Regarding the quantitative forecasting performance comparison, we observed no significant differences between ESRNN and NBEATS without exogenous variables. At the same time, NBEATSx improves over NBEATS by nearly 20% and up to 5% over LEAR and DNN models specialized for the Electricity Price Forecasting tasks. Finally, we found no significant trade-offs between the accuracy and interpretability of NBEATSx-G and NBEATSx-I predictions. We intend, and encourage others, to use the presented method and further tailor it as needed to a wide range of applications in healthcare, seismology, predictive maintenance of equipment, and other areas wherever consideration of exogenous variables could impact forecasts.

## Acknowledgements

This work was partially supported by the Defense Advanced Research Projects Agency (award FA8750-17-2-0130), the National Science Foundation (grant 2038612), the Space Technology Research Institutes grant from NASA’s Space Technology Research Grants Program, the U.S. Department of Homeland Security (award 18DN-ARI-00031), the Ministry of Science and Higher Education (MNiSW, Poland) through grant No. 0219/DIA/2019/48, the National Science Center (NCN, Poland) through grant No. 2018/30/A/HS4/00444, and Abraxas Group. Kin G. Olivares and Cristian Challu want to thank Stefania La Vattiata, Max Mergenthaler and Federico Garza for their support.

## References

- Afrasiabi, M., Mohammadi, M., Rastegar, M., & Kargarian, A. (2019). Multi-agent microgrid energy management based on deep learning forecaster. *Energy*, 186, 115873. URL: <https://www.sciencedirect.com/science/article/pii/S0360544219315452>. doi:<https://doi.org/10.1016/j.energy.2019.115873>.
- Amjady, N. (2012). Short-term electricity price forecasting. In J. Catalão (Ed.), *Electric Power Systems: Advanced Forecasting Techniques and Optimal Generation Scheduling* (pp. 4–1–4–58). CRC Press. doi:10.1201/b11649-4.
- Athanasopoulos, G., Hyndman, R. J., Song, H., & Wu, D. C. (2011). The tourism forecasting competition. *International Journal of Forecasting*, 27, 822–844. URL: <https://www.sciencedirect.com/science/>

- article/pii/S016920701000107X. doi:<https://doi.org/10.1016/j.ijforecast.2010.04.009>. Special Section 1: Forecasting with Artificial Neural Networks and Computational Intelligence. Special Section 2: Tourism Forecasting.
- Atiya, A. F. (2020). Why does forecast combination work so well? *International Journal of Forecasting*, 36, 197–200. URL: <https://www.sciencedirect.com/science/article/pii/S0169207019300779>. doi:<https://doi.org/10.1016/j.ijforecast.2019.03.010>. M4 Competition.
- Bahdanau, D., Cho, K., & Bengio, Y. (2016). Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015*. URL: <https://arxiv.org/abs/1409.0473>. arXiv:1409.0473.
- Bai, S., Kolter, J. Z., & Koltun, V. (2018). An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *Computing Research Repository*, abs/1803.01271. URL: <http://arxiv.org/abs/1803.01271>. arXiv:1803.01271.
- Bengio, Y. (2012). Deep learning of representations for unsupervised and transfer learning. In I. Guyon, G. Dror, V. Lemaire, G. Taylor, & D. Silver (Eds.), *Proceedings of ICML, Workshop on Unsupervised and Transfer Learning* (pp. 17–36). Bellevue, Washington, USA: JMLR Workshop and Conference Proceedings volume 27 of *Proceedings of Machine Learning Research*. URL: <http://proceedings.mlr.press/v27/bengio12a.html>.
- Benidis, K., Rangapuram, S. S., Flunkert, V., Wang, B., Maddix, D., Turkmen, C., Gasthaus, J., Bohlke-Schneider, M., Salinas, D., Stella, L., Callot, L., & Januschowski, T. (2020). Neural forecasting: Introduction and literature overview. *Computing Research Repository*, . arXiv:2004.10240.
- Bergmeir, C., Hyndman, R. J., & Benítez, J. M. (2016). Bagging exponential smoothing methods using stl decomposition and box-cox transformation. *International Journal of Forecasting*, 32, 303–312. URL: <https://www.sciencedirect.com/science/article/pii/S0169207015001120>. doi:<https://doi.org/10.1016/j.ijforecast.2015.07.002>.
- Bergstra, J., Bardenet, R., Bengio, Y., & Kégl, B. (2011). Algorithms for hyper-parameter optimization. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems* (pp. 2546–2554). Curran Associates, Inc. volume 24. URL: <https://proceedings.neurips.cc/paper/2011/file/86e8f7ab32cfd12577bc2619bc635690-Paper.pdf>.
- Breiman, L. (1996). Bagging predictors. *Machine Learning*, 24, 123–140. URL: <https://doi.org/10.1023/A:1018054314350>. doi:10.1023/A:1018054314350.
- Bzinovski, S., & Fulgosi, A. (1976). The influence of pattern similarity and transfer learning upon training a base perceptron B2. Original in Croatian: Utjecaj slicnosti likova i transfera učenja na obucavanje baznog perceptrona B2. *Proceedings of Symposium Informatica*, (pp. 3–121–5). URL: <http://www.informatica.si/index.php/informatica/article/view/2828>.
- Chang, S., Zhang, Y., Han, W., Yu, M., Guo, X., Tan, W., Cui, X., Witbrock, M., Hasegawa-Johnson, M. A., & Huang, T. S. (2017). Dilated recurrent neural networks. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems*. Curran Associates, Inc. volume 30. URL: <https://proceedings.neurips.cc/paper/2017/file/32bb90e8976aab5298d5da10fe66f21d-Paper.pdf>.
- Chitsaz, H., Zamani-Dehkordi, P., Zareipour, H., & Parikh, P. (2018). Electricity price forecasting for operational scheduling of behind-the-meter storage systems. *IEEE Transactions on Smart Grid*, 9, 6612–6622. doi:10.1109/TSG.2017.2717282.
- Chung, J., Gülçehre, Ç., Cho, K., & Bengio, Y. (2014). Empirical evaluation of gated recurrent neural networks on sequence modeling. *NIPS 2014, Workshop on Deep Learning*. URL: <http://arxiv.org/abs/1412.3555>. arXiv:1412.3555.
- Cleveland, R. B., Cleveland, W. S., McRae, J. E., & Terpenning, I. (1990). Stl: A seasonal-trend decomposition procedure based on loess (with discussion). *Journal of Official Statistics*, 6, 3–73.
- Dagum, E. B. (1980). The X-II-ARIMA seasonal adjustment method. *Statistics Canada. Statistical Sciences*, 679. URL: <https://www.census.gov/library/working-papers/1980/adrm/dagum-01.html>.
- Dauphin, Y. N., Fan, A., Auli, M., & Grangier, D. (2017). Language modeling with gated convolutional networks. In D. Precup, & Y. W. Teh (Eds.), *Proceedings of the 34th International Conference on*



- Machine Learning* (pp. 933–941). PMLR volume 70 of *Proceedings of Machine Learning Research*. URL: <http://proceedings.mlr.press/v70/dauphin17a.html>.
- de Simón-Martín, M., Bracco, S., Rosales-Asensio, E., Piazza, G., Delfino, F., & Giribone, P. G. (2020). Electricity spot prices forecasting for mibel by using deep learning: a comparison between nar, narx and lstm networks. In *2020 IEEE International Conference on Environment and Electrical Engineering and 2020 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I CPS Europe)* (pp. 1–6). doi:10.1109/EEEIC/ICPSEurope49358.2020.9160587.
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)* (pp. 4171–4186). Minneapolis, Minnesota: Association for Computational Linguistics. URL: <https://www.aclweb.org/anthology/N19-1423>. doi:10.18653/v1/N19-1423.
- Diebold, F., & Mariano, R. (2002). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 20, 134–44. URL: <https://www.sas.upenn.edu/~fdiebold/papers/paper68/pa.dm.pdf>. doi:10.1080/07350015.1995.10524599.
- Dokumentov, A., & Hyndman, R. J. (2015). *STR: A Seasonal-Trend Decomposition Procedure Based on Regression*. Monash Econometrics and Business Statistics Working Papers 13/15 Monash University, Department of Econometrics and Business Statistics. URL: <https://ideas.repec.org/p/msh/ebswps/2015-13.html>.
- Donoho, D. L., & Johnstone, I. M. (1995). Adapting to unknown smoothness via wavelet shrinkage. *Journal of the American Statistical Association*, 90, 1200–1224. URL: <http://www.jstor.org/stable/2291512>.
- Dudek, G., Pełka, P., & Smył, S. (2021). A hybrid residual dilated lstm and exponential smoothing model for midterm electric load forecasting. *IEEE Transactions on Neural Networks and Learning Systems*, (pp. 1–13). doi:10.1109/TNNLS.2020.3046629.
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14, 179–211. URL: [https://onlinelibrary.wiley.com/doi/abs/10.1207/s15516709cog1402\\_1](https://onlinelibrary.wiley.com/doi/abs/10.1207/s15516709cog1402_1).
- Evgeniou, T., Pontil, M., & Poggio, T. (2000). Regularization networks and support vector machines. *Advances in Computational Mathematics*, 13. URL: <https://doi.org/10.1023/A:1018946025316>.
- Fawaz, H. I., Forestier, G., Weber, J., Idoumghar, L., & Muller, P. (2018). Transfer learning for time series classification. *CoRR*, abs/1811.01533. URL: <http://arxiv.org/abs/1811.01533>. arXiv:1811.01533.
- Findley, D. F., Monsell, B. C., Bell, W. R., Otto, M. C., & Chen, B.-C. (1998). New capabilities and methods of the X-12-ARIMA seasonal-adjustment program. *Journal of Business & Economic Statistics*, 16, 127–152. URL: <http://www.jstor.org/stable/1392565>.
- Friedman, J. H. (1991). Multivariate Adaptive Regression Splines. *The Annals of Statistics*, 19, 1 – 67. URL: <https://doi.org/10.1214/aos/1176347963>. doi:10.1214/aos/1176347963.
- Gaillard, P., Goude, Y., & Nedellec, R. (2016). Additive models and robust aggregation for GEFCom2014 probabilistic electric load and electricity price forecasting. *International Journal of Forecasting*, 32, 1038–1050. URL: <https://www.sciencedirect.com/science/article/pii/S0169207015001545>. doi:https://doi.org/10.1016/j.ijforecast.2015.12.001.
- Gers, F. A., Cummins, F., & Schmidhuber, J. (2000). Learning to forget: continual prediction with LSTM. *Neural Computation*, 12, 2451–2471. URL: [https://digital-library.theiet.org/content/conferences/10.1049/cp\\_19991218](https://digital-library.theiet.org/content/conferences/10.1049/cp_19991218).
- Giacomini, R., & White, H. (2006). Tests of conditional predictive ability. *Econometrica*, 74, 1545–1578. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-0262.2006.00718.x>. doi:https://doi.org/10.1111/j.1468-0262.2006.00718.x.
- Gianfreda, A., Ravazzolo, F., & Rossini, L. (2020). Comparing the forecasting performances of linear models for electricity prices with high RES penetration. *International Journal of Forecasting*, 36, 974–986. URL: <https://www.sciencedirect.com/science/article/pii/S0169207019302596>. doi:https://doi.org/10.1016/j.ijforecast.2019.11.002.
- Graves, A. (2013). Generating sequences with recurrent neural networks. *Computing Research Repository*, abs/1308.0850. URL: <http://arxiv.org/abs/1308.0850>. arXiv:1308.0850.

- Hastie, T., & Tibshirani, R. (1986). Generalized Additive Models. *Statistical Science*, 1, 297–310. URL: <https://doi.org/10.1214/ss/1177013604>. doi:10.1214/ss/1177013604.
- He, H., Lu, N., Jiang, Y., Chen, B., & Jiao, R. (2020). End-to-end probabilistic forecasting of electricity price via convolutional neural network and label distribution learning. *Energy Reports*, 6, 1176–1183. doi:<https://doi.org/10.1016/j.egy.2020.11.057>. 2020 The 7th International Conference on Power and Energy Systems Engineering.
- Hermans, M., & Schrauwen, B. (2013). Training and analysing deep recurrent neural networks. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems 26* (pp. 190–198). Curran Associates, Inc. URL: <http://papers.nips.cc/paper/5166-training-and-analysing-deep-recurrent-neural-networks.pdf>.
- Hong, T., Pinson, P., Fan, S., Zareipour, H., Troccoli, A., & Hyndman, R. J. (2016). Probabilistic energy forecasting: Global Energy Forecasting Competition 2014 and beyond. *International Journal of Forecasting*, 32, 896–913. URL: <https://www.sciencedirect.com/science/article/pii/S0169207016000133>. doi:<https://doi.org/10.1016/j.ijforecast.2016.02.001>.
- Hong, Y. Y., & Pula, R. (2020). Comparative studies of different methods for short-term locational marginal price forecasting. In *2020 5th International Conference on Green Technology and Sustainable Development (GTSD)* (pp. 527–532). doi:10.1109/GTSD50082.2020.9303121.
- Hubicka, K., Marcjasz, G., & Weron, R. (2018). *A note on averaging day-ahead electricity price forecasts across calibration windows*. HSC Research Reports HSC/18/03 Hugo Steinhaus Center, Wrocław University of Technology. URL: <https://ideas.repec.org/p/wuu/wpaper/hsc1803.html>.
- Hyndman, R. J., & Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting*, 22, 679 – 688. URL: <http://www.sciencedirect.com/science/article/pii/S0169207006000239>. doi:<https://doi.org/10.1016/j.ijforecast.2006.03.001>.
- Jahangir, H., Tayarani, H., Baghali, S., Ahmadian, A., Elkamel, A., Golkar, M. A., & Castilla, M. (2020). A novel electricity price forecasting approach based on dimension reduction strategy and rough artificial neural networks. *IEEE Transactions on Industrial Informatics*, 16, 2369–2381. doi:10.1109/TII.2019.2933009.
- Kalchbrenner, N., Espeholt, L., Simonyan, K., van den Oord, A., Graves, A., & Kavukcuoglu, K. (2016). Neural machine translation in linear time. *Computing Research Repository*, abs/1610.10099. URL: <http://arxiv.org/abs/1610.10099>. arXiv:1610.10099.
- Kingma, D. P., & Ba, J. (2014). ADAM: A method for stochastic optimization. URL: <http://arxiv.org/abs/1412.6980> cite arxiv:1412.6980Comment: Published as a conference paper at the 3rd International Conference for Learning Representations (ICLR), San Diego, 2015.
- Koenker, R., & Bassett, G. (1978). Regression quantiles. *Econometrica*, 46, 33–50. URL: <http://www.jstor.org/stable/1913643>.
- Kuo, P.-H., & Huang, C.-J. (2018). An electricity price forecasting model by hybrid structured deep neural networks. *Sustainability*, 10. URL: [www.mdpi.com/2071-1050/10/4/1280](http://www.mdpi.com/2071-1050/10/4/1280). doi:10.3390/su10041280.
- Lago, J., De Ridder, F., & De Schutter, B. (2018a). Forecasting spot electricity prices: Deep learning approaches and empirical comparison of traditional algorithms. *Applied Energy*, 221, 386 – 405. URL: <http://www.sciencedirect.com/science/article/pii/S030626191830196X>. doi:<https://doi.org/10.1016/j.apenergy.2018.02.069>.
- Lago, J., De Ridder, F., Vranckx, P., & De Schutter, B. (2018b). Forecasting day-ahead electricity prices in europe: The importance of considering market integration. *Applied Energy*, 211, 890–903. URL: <https://www.sciencedirect.com/science/article/pii/S0306261917316999>. doi:<https://doi.org/10.1016/j.apenergy.2017.11.098>.
- Lago, J., Marcjasz, G., De Schutter, B., & Weron, R. (2021). Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark. *Applied Energy*, submitted, Working Paper version available from arXiv:2008.08004.
- LeCun, Y., Bottou, L., Orr, G. B., & Müller, K. R. (1998). Efficient backprop. In *Neural Networks: Tricks of the Trade* (pp. 9–50). Berlin, Heidelberg: Springer Berlin Heidelberg. URL: [https://doi.org/10.1007/3-540-49430-8\\_2](https://doi.org/10.1007/3-540-49430-8_2). doi:10.1007/3-540-49430-8\_2.

- Livera, A. M. D., Hyndman, R. J., & Snyder, R. D. (2011). Forecasting time series with complex seasonal patterns using exponential smoothing. *Journal of the American Statistical Association*, 106, 1513–1527. URL: <https://doi.org/10.1198/jasa.2011.tm09771>. doi:10.1198/jasa.2011.tm09771. arXiv:<https://doi.org/10.1198/jasa.2011.tm09771>.
- Luo, S., & Weng, Y. (2019). A two-stage supervised learning approach for electricity price forecasting by leveraging different data sources. *Applied Energy*, 242, 1497 – 1512. URL: 10.1016/j.apenergy.2019.03.129.
- Långkvist, M., Karlsson, L., & Loutfi, A. (2014). A review of unsupervised feature learning and deep learning for time-series modeling. *Pattern Recognition Letters*, 42, 11–24. URL: <https://www.sciencedirect.com/science/article/pii/S0167865514000221>. doi:<https://doi.org/10.1016/j.patrec.2014.01.008>.
- Macaulay, F. (1931). The smoothing of time series. *Journal of the Institute of Actuaries*, 62, 181–182. URL: <https://www.nber.org/books-and-chapters/smoothing-time-series>. doi:10.1017/S0020268100010003.
- Makridakis, S., & Hibon, M. (2000). The M3-competition: results, conclusions and implications. *International Journal of Forecasting*, 16, 451–476. URL: <https://www.sciencedirect.com/science/article/pii/S0169207000000571>. doi:[https://doi.org/10.1016/S0169-2070\(00\)00057-1](https://doi.org/10.1016/S0169-2070(00)00057-1). The M3- Competition.
- Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018). Statistical and machine learning forecasting methods: Concerns and ways forward. *PLoS One*, 13(3), e0194889. URL: <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0194889>.
- Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2020). The M4 competition: 100,000 time series and 61 forecasting methods. *International Journal of Forecasting*, 36, 54–74. URL: <https://www.sciencedirect.com/science/article/pii/S0169207019301128>. doi:<https://doi.org/10.1016/j.ijforecast.2019.04.014>. M4 Competition.
- Mayer, K., & Trück, S. (2018). Electricity markets around the world. *Journal of Commodity Markets*, 9, 77–100. URL: <https://doi.org/10.1016/j.jcomm.2018.02.001>.
- Mujeeb, S., Javaid, N., Ilahi, M., Wadud, Z., Ishmanov, F., & Afzal, M. (2019). Deep long short-term memory: A new price and load forecasting scheme for big data in smart cities. *Sustainability*, 11, 987. doi:10.3390/su11040987.
- Narajewski, M., & Ziel, F. (2020). Econometric modelling and forecasting of intraday electricity prices. *Journal of Commodity Markets*, 19, 100107. doi:10.1016/j.jcomm.2019.100107.
- Nascimento, R. C., Souto, Y. M., Ogasawara, E. S., Porto, F., & Bezerra, E. (2021). STconvS2S: Spatiotemporal convolutional sequence to sequence network for weather forecasting. *Neurocomputing*, 426, 285–298. URL: <http://arxiv.org/abs/1912.00134>. arXiv:1912.00134.
- Nikolay Laptev, J. Y., & Rajagopal, R. (2018). Reconstruction and regression loss for time-series transfer learning. *Proceedings of the Special Interest Group of Knowledge Discovery and Data Mining (SIGKDD)*, . URL: [https://milets18.github.io/papers/milets18\\_paper\\_2.pdf](https://milets18.github.io/papers/milets18_paper_2.pdf).
- Nowotarski, J., Raviv, E., Trück, S., & Weron, R. (2014). An empirical comparison of alternative schemes for combining electricity spot price forecasts. *Energy Economics*, 46, 395–412. URL: <https://ideas.repec.org/a/eee/eneeco/v46y2014icp395-412.html>. doi:10.1016/j.eneco.2014.07.0.
- Nowotarski, J., & Weron, R. (2018). Recent advances in electricity price forecasting: A review of probabilistic forecasting. *Renewable and Sustainable Energy Reviews*, 81, 1548–1568. URL: <https://www.sciencedirect.com/science/article/pii/S1364032117308808>. doi:<https://doi.org/10.1016/j.rser.2017.05.234>.
- van den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A. W., & Kavukcuoglu, K. (2016). Wavenet: A generative model for raw audio. *CoRR*, abs/1609.03499. URL: <http://arxiv.org/abs/1609.03499>. arXiv:1609.03499.
- Oreshkin, B. N., Carpo, D., Chapados, N., & Bengio, Y. (2020). N-BEATS: neural basis expansion analysis for interpretable time series forecasting. In *8th International Conference on Learning Representations, ICLR 2020*. URL: <https://openreview.net/forum?id=r1ecqn4YwB>.

- Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Trans. on Knowl. and Data Eng.*, 22, 1345–1359. URL: <https://doi.org/10.1109/TKDE.2009.191>. doi:10.1109/TKDE.2009.191.
- Paszke et al. (2019). Pytorch: An imperative style, high-performance Deep Learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d Alché-Buc, E. Fox, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems 32* (pp. 8024–8035). Curran Associates, Inc. URL: <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- Ramachandran, P., Liu, P., & Le, Q. (2017). Unsupervised pretraining for sequence to sequence learning. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (pp. 383–391). Copenhagen, Denmark: Association for Computational Linguistics. URL: <https://www.aclweb.org/anthology/D17-1039>. doi:10.18653/v1/D17-1039.
- Rantonen, M., & Korpihalkola, J. (2020). Prediction of spot prices in nord pool’s day-ahead market using machine learning and deep learning. In N. G. et al. (Ed.), *International Conference on Machine Learning, Optimization, and Data Science. LOD 2020*. (pp. 676–687). Springer, Cham. doi:[https://doi.org/10.1007/978-3-030-64583-0\\_59](https://doi.org/10.1007/978-3-030-64583-0_59).
- Salinas, D., Flunkert, V., Gasthaus, J., & Januschowski, T. (2020). DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 36, 1181–1191. URL: <https://www.sciencedirect.com/science/article/pii/S0169207019301888>. doi:<https://doi.org/10.1016/j.ijforecast.2019.07.001>.
- Shishkin, J., Young, A., & Musgrave, J. (1967). The X-II variant of the census II method seasonal adjustment program. *Bureau of the Census, Technical Report, No. 15*. URL: <https://www.census.gov/library/working-papers/1967/adrm/shiskin-01.html>.
- Smyl, S. (2019). A hybrid method of exponential smoothing and recurrent neural networks for time series forecasting. *International Journal of Forecasting*, . doi:10.1016/j.ijforecast.2019.03.017.
- Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to Sequence learning with neural networks. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems*. Curran Associates, Inc. volume 27. URL: <https://proceedings.neurips.cc/paper/2014/file/a14ac55a4f27472c5d894ec1c3c743d2-Paper.pdf>.
- Ugurlu, U., Oksuz, I., & Tas, O. (2018). Electricity price forecasting using recurrent neural networks. *Energies*, 11. URL: <https://www.mdpi.com/1996-1073/11/5/1255>. doi:10.3390/en11051255.
- Uniejewski, B., Nowotarski, J., & Weron, R. (2016). Automated variable selection and shrinkage for day-ahead electricity price forecasting. *Energies*, 9. URL: <https://www.mdpi.com/1996-1073/9/8/621>. doi:10.3390/en9080621.
- Uniejewski, B., & Weron, R. (2021). Regularized quantile regression averaging for probabilistic electricity price forecasting. *Energy Economics*, 95, 105121. URL: <https://www.sciencedirect.com/science/article/pii/S0140988321000268>. doi:<https://doi.org/10.1016/j.eneco.2021.105121>.
- Uniejewski, B., Weron, R., & Ziel, F. (2018). Variance stabilizing transformations for electricity spot price forecasting. *IEEE Transactions on Power Systems*, 33, 2219–2229. doi:10.1109/TPWRS.2017.2734563.
- U.S. Census Bureau (2013). *X-13-ARIMA-SEATS Reference Manual*. U.S. Census Bureau. Washington, DC. URL: <http://www.census.gov/ts/x13as/docX13AS.pdf>.
- Wen, R., Torkkola, K., Narayanaswamy, B., & Madeka, D. (2017). A Multi-horizon Quantile Recurrent Forecaster. In *31st Conference on Neural Information Processing Systems NIPS 2017, Time Series Workshop*. URL: <https://arxiv.org/abs/1711.11053>. arXiv:1711.11053.
- Werbos, P. J. (1990). Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78, 1550–1560. URL: <https://ieeexplore.ieee.org/document/58337>.
- Weron, R. (2014). Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting*, 30, 1030–1081. URL: <https://www.sciencedirect.com/science/article/pii/S0169207014001083>. doi:<https://doi.org/10.1016/j.ijforecast.2014.08.008>.
- Weron, R., & Ziel, F. (2020). Electricity price forecasting. In U. Soytaş, & R. Sari (Eds.), *Routledge Handbook of Energy Economics* (pp. 506–521). Routledge. doi:10.4324/9781315459653-36.
- Yang, H., & Schell, K. R. (2020). HFNet: Forecasting real-time electricity price via novel GRU architectures. In *2020 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)* (pp. 1–6).

doi:10.1109/PMAFS47429.2020.9183697.

- Yao, Y., Rosasco, L., & Andrea, C. (2007). On early stopping in gradient descent learning. *Constructive Approximation*, 26(2), 289–315. URL: <https://doi.org/10.1007/s00365-006-0663-2>.
- Yosinski, J., Clune, J., Bengio, Y., & Lipson, H. (2014). How transferable are features in deep neural networks? In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems*. Curran Associates, Inc. volume 27. URL: <https://arxiv.org/abs/1411.1792>.
- Zhang, F., Fleyeh, H., & Bales, C. (2020). A hybrid model based on bidirectional long short-term memory neural network and catboost for short-term electricity spot price forecasting. *Journal of the Operational Research Society*, 0, 1–25. doi:10.1080/01605682.2020.1843976.
- Zhang, R., Li, G., & Ma, Z. (2020). A deep learning based hybrid framework for day-ahead electricity price forecasting. *IEEE Access*, 8, 143423–143436. doi:10.1109/ACCESS.2020.3014241.
- Zhou, S., Zhou, L., Mao, M., Tai, H., & Wan, Y. (2019). An optimized heterogeneous structure LSTM network for electricity price forecasting. *IEEE Access*, 7, 108161–108173. doi:10.1109/ACCESS.2019.2932999.
- Ziel, F., & Steinert, R. (2018). Probabilistic mid- and long-term electricity price forecasting. *Renewable and Sustainable Energy Reviews*, 94, 251–266. URL: <https://arxiv.org/abs/1703.10806>.

## Appendix A. Appendix

### *Appendix A.1. Best Single Models*

Table A1 shows that the best NBEATSx models yield improvements of 14.8% on average across all the evaluation metrics when compared to its NBEATS counterpart without exogenous covariates, and improvements of 23.9% when compared to ESRNN without time-dependent covariates. A perhaps more remarkable result is the statistically significant improvement of forecast accuracy over LEAR and DNN benchmarks, ranging from 0.75% to 7.2% across all metrics and markets, with the exception of BE. Compared to DNN, the RMSE improved on average 4.9%, the MAE improved 3.2%, the rMAE improved 3.0%, and sMAPE improved 1.7%.

When comparing the best NBEATSx models against the best DNN on individual markets, NBEATSx improved by 3.18% on the Nord Pool market (NP), 2.03% 2.65% on French (FR) and 5.24% on German (DE) power markets. The positive difference in performance for Belgian (BE) market of 0.53% was not statistically significant.

Figure A.1 provides a graphical representation of the GW test for the six best models, across the five markets for the MAE evaluation metric. The models included in the significance tests are the same as in Tables A1: LEAR, DNN, the ESRNN, NBEATS, and our proposed methods, the NBEATSx-G and NBEATSx-I. The p-value of each individual comparison shows if the improvement in performance (measured by MAE or RMSE) of the x-axis model over the y-axis model is statistically significant. Both the NBEATSx-G and NBEATSx-I model outperformed the LEAR and DNN models in all markets, with the exception of Belgium. Moreover, no benchmark model outperformed the NBEATSx-I and NBEATSx-G on any market.

Table A1: Forecast accuracy measures for day-ahead electricity prices for the best *single model* out of the four models described in the Subsection 4.3.5. The ESRNN and NBEATS, are the original implementations and do not include time dependent covariates. The reported metrics are *mean absolute error* (MAE), *relative mean absolute error* (rMAE), *symmetric mean absolute percentage error* (sMAPE) and *root mean squared error* (RMSE). The smallest errors in each row are highlighted in bold.

		AR1	ESRNN	NBEATS	ARx1	LEARx	DNN	NBEATSx-G	NBEATSx-I
NP	MAE	2.28	2.11	2.11	2.11	1.95	1.71	<b>1.65</b>	1.68
	rMAE	0.72	0.67	0.67	0.67	0.62	0.54	<b>0.52</b>	0.53
	sMAPE	6.51	6.09	6.06	6.1	5.62	4.97	<b>4.83</b>	4.89
	RMSE	4.08	3.92	3.98	3.84	3.60	3.36	<b>3.27</b>	3.33
PJM	MAE	3.88	3.63	3.48	3.68	3.09	3.07	3.02	<b>3.01</b>
	rMAE	0.8	0.75	0.72	0.76	0.64	0.63	<b>0.62</b>	<b>0.62</b>
	sMAPE	14.66	14.26	13.56	14.09	12.54	12.00	11.97	<b>11.91</b>
	RMSE	6.26	5.87	5.59	5.94	5.14	5.20	5.06	<b>5.00</b>
EPEX-BE	MAE	7.04	7.01	6.83	7.05	6.59	<b>6.07</b>	6.14	6.17
	rMAE	0.86	0.86	0.83	0.86	0.80	<b>0.74</b>	0.75	0.75
	sMAPE	16.29	15.95	16.03	16.21	15.95	<b>14.11</b>	14.68	14.52
	RMSE	17.25	16.76	16.99	17.07	16.29	15.95	15.46	<b>15.43</b>
EPEX-FR	MAE	4.74	4.68	4.79	4.85	4.25	4.06	3.98	<b>3.97</b>
	rMAE	0.80	0.78	0.80	0.86	0.71	0.68	<b>0.67</b>	<b>0.67</b>
	sMAPE	13.49	13.25	13.62	16.21	13.25	11.49	<b>11.07</b>	11.29
	RMSE	13.68	11.89	12.09	17.07	<b>10.75</b>	11.77	11.61	11.08
EPEX-DE	MAE	5.73	5.64	5.37	4.58	4.11	3.59	3.46	<b>3.37</b>
	rMAE	0.71	0.70	0.67	0.57	0.51	0.45	0.43	<b>0.42</b>
	sMAPE	21.22	21.09	19.71	18.52	16.98	14.68	14.78	<b>14.34</b>
	RMSE	9.39	9.17	9.03	7.69	6.99	6.08	5.84	<b>5.64</b>

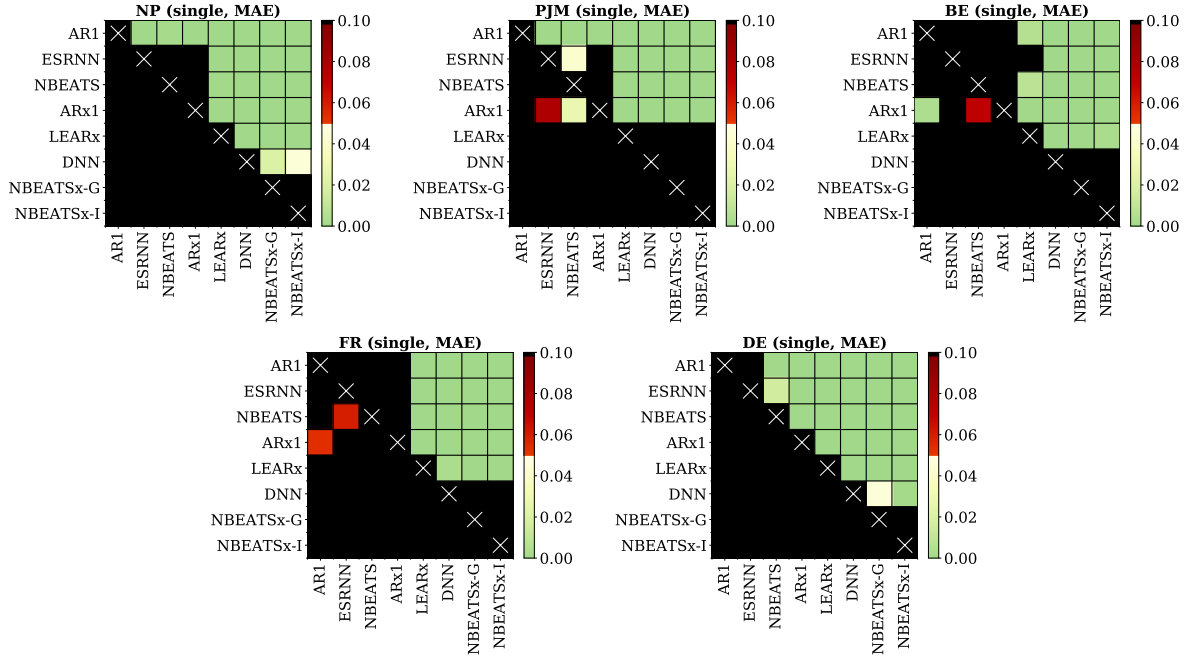


Figure A.1: Results of the Giacomini-White test for the day-ahead predictions with *mean absolute error* (MAE) applied to pairs of the single models on the five electricity markets datasets. Each grid represents one market. Each colored cell in a grid is plotted black, unless the predictions of the model corresponding to its column of the grid outperforms the predictions of the model corresponding to its row of the grid. The color scale reflects significance of the difference in MAE, with solid green representing the lowest  $p$ -values.