Electricity Price Forecasting: A methodological ANN-based Approach with special Consideration of Time Series Properties

Jan-Hendrik Meier[1][0000-0002-3080-2210], Stephan Schneider[1][0000-0003-1810-8813], Thies Schönfeldt[1], Philip Schüller[1], Bastian Wanke[1]

[1] Kiel University of Applied Sciences, Sokratesplatz 2, 24149 Kiel, Germany
jan-hendrik.meier|stephan.schneider@fh-kiel.de

Abstract. If one examines the spot price series of electrical power over the course of time, it is striking that the electricity price across the day takes a course that is determined by power consumption following a day and night rhythm. This daily course changes in its height and temporal extent in both, the course of the week, as well as with the course of the year. This study deals methodologically with non-linear correlative and autocorrelative time series properties of the electricity spot price. We contribute the usage of non-fully connectionist networks in relation to fully connectionist networks to decompose non-linear correlative time series properties. Additionally, we contribute the usage of long short-term-memory network (LSTM) to discover and to deal with autocorrelation effects.

Keywords: Electricity Prices, Artificial Neural Network, LSTM, ARIMAX.

1 Introduction

Despite all criticism of this approach, the random walk process has established itself for the modeling of stock prices. Pricing on electricity markets deviates significantly from the pricing on stock markets, as the underlying Markov property cannot be assumed for electricity markets as well. Produced electricity cannot be stored without significant losses and, accordingly, temporal arbitrage turns out to be highly inefficient. If one examines the spot price series of electrical power over the course of time, it is striking that the electricity price across the day takes a course that is determined by power consumption following a day and night rhythm. This daily course changes in its height and temporal extent in both, the course of the week, as well as with the course of the year. Accordingly, it can be concluded that the univariate time series shows non-linear correlative effects between daily, weekly, and yearly seasonal patterns as well as autocorrelative effects even without taking other explanatory variables into account.

The present study deals methodologically with non-linear correlative and autocorrelative time series properties of the electricity spot price. Correlation effects are adequately represented in classical fully connectionist networks but they cannot be meaningfully analyzed due to the high complexity of these networks. The research questions are, if the forecasting accuracy can be improved by (i) using different and complementary ANN-architectures to better reflect correlation effects and (ii) using a recurrent ANN-architecture to better reflect autocorrelation effects. To answer these questions, we contribute the usage of non-fully connectionist networks in relation to fully connectionist networks to decompose non-linear correlative time series properties.

Hence, we use (i) different ANN architectures with non-fully and fully connectionist networks to discover and to deal with correlation effects on exogenous side / input layer, (ii) using a long short-term-memory network (LSTM) to discover and to deal with autocorrelation effects, and (iii) an ARIMAX model with daily, weekly, and yearly seasonal patterns reflected as binary coded variables as a benchmark for the aforementioned models.

The paper is organized as follows: In section two, the current state of the literature is presented, and the research gap is identified. In section three, sample and methodology are introduced. In section four, the results are presented and discussed. The study closes with a conclusion.

2 Literature Review and Research Gap

The number of electricity price forecasting articles has increased significantly in recent years. A particularly good overview can be obtained by Weron (2014). The author could identify 30 publications with a focus on ARIMA and its extensions. We could not identify further more recent articles in this special field of ARIMA-modelling of electricity prices. More recent electricity price forecasting literature is focused mainly on probabilistic forecasting and artificial intelligence. With regard to ANN, Weron could identify 56 publications. Subsequently, two further articles were published on electricity price forecasting using ANN that were not included in Weron's review (Dudek, 2016; Marcjasz, Uniejewski, & Weron, 2018).

Comparing ARIMA(X) models of the Spanish and the Californian market with and without additional explanatory variables, Contreras, Espinola, Nogales, & Conejo (2003) recognize that additional explanatory variables, such as hydropower, are only required in months of a high correlation between the explanatory variable and the price, while in months of low correlations these variables do not show significant predictive power. The authors were able to show average daily mean errors between 5% and 10% with and without explanatory variables.

When forecasting with ARIMA, Conejo, Contreras, Espínola, & Plazas (2005) argue that it could be necessary to use a different notation of the model for nearly every week. Accordingly, ARIMA-models turn out to be very unstable in their predictive power over time. Especially in spring and summer where the volatility was very high the ARIMA forecast provided poor results. The authors also introduce several other techniques, e.g. an ANN with a multilayer perceptron and one hidden layer. The ARIMA model outperforms the ANN in every period except for the September. The mean week errors with ARIMA are between 6% and 27% whereas the ANN shows errors between 8% and 32%.

Garcia, Contreras, van Akkeren, & Garcia (2005) claimed that ARIMA-GARCH models show a better accuracy than seasonal ARIMA models. The authors present mean weekly errors of around 10% for relatively calm weeks. Misiorek, Trueck, & Weron (2006) compare some linear and non-linear time series models. In contrast to the aforementioned authors, the simple ARX model - the exogenous variable is the day-ahead load forecast - shows a better result than a model with an additional GARCH component.

Conejo, Plazas, Espinola, & Molina (2005) contributed a specified ARIMA model including wavelet transformation which was more accurate than the simple one. The wavelet transformation is applied to decompose the time series before predicting the electricity prices with ARIMA. This model outperforms the benchmark with a weekly error of 5% in winter and spring and 11% in summer and fall.

Applying a seasonal ARMA(X) process with three different explanatory variables of the temperature, Knittel & Roberts (2005) identified an inverse leverage effect with positive price reactions increasing the volatility more than negative ones. The authors further show that a higher order autocorrelation in the models is important to improve the results. The authors were able to show root mean squared errors for the out-of-sample week between 25.5 and 49.4 in the pre-crisis period and between 66.6 and 88.6 during the crises period. It is mentioned in the article that the data has a high frequency of large price deviations, which leads to these high forecast errors.

Zareipour, Canizares, Bhattacharya, & Thomson (2006) built an ARMAX and an ARX model with an average error in the 24-hour-ahead forecast of 8.1 and of 8.4 respectively, which is slightly better than the basic ARIMA model with an average of 8.8. With these models, it could be shown that market information in low-demand periods is not as useful as during high-demand periods. In general, the results have confirmed the contribution of the authors that market data improves the forecast results. Nevertheless, none of the models could forecast the extreme prices which increasingly occur in times of high-demand periods adequately.

Zhou, Yan, Ni, Li, & Nie (2006) suggested that including error correction will lead to a more accurate result in forecasting with ARIMA. Therefore, they developed an ARIMA approach which is extended by an error series. This novel method turned out to show quite good forecasting accuracies with an average error of 2% and lower despite of periods with a high price volatility.

Koopman, Ooms, & Carnero (2007) were using an ARFIMA model, which is an ARIMA model with seasonal periodic regressions, and combined it with a GARCH analysis. The authors pointed out the importance of day-of-the-week periodicity in the autocovariance function when forecasting electricity prices. Beneath the implementation of the day of the week, binaries were included for the holiday effect to consider demand variations.

With the increase in available computational power in recent years, ANN became more and more popular in forecasting and forecasting research. Both, classical multilayer perceptron (MLP) and recurrent networks (Hopfield, 1982; Haykin, 2009), especially long short-term memory (LSTM) (Hochreiter & Schmidhuber, 1997) networks, are used for forecasting purposes of time series data. Typically, all ANN architectures are composed of an input layer, a hidden layer with differing number of units, and an output layer. In fully connectionist networks, typically, lags and partially residuals are passed into the propagation function (Zhang, Patuwo, & Hu, 1998; Adebiyi, Adewumi, & Ayo, 2014). Each node of a layer is usually fully connected to the units of the subsequent layer. MLP as well as LSTM networks are fed with differentiated time series data. The reason is the underlying characteristics of a time series itself. If time is the explanatory factor for the values of the endogenous variable, in our case the electricity price, the time series must be made stationary by differentiation to avoid spurious correlations.

Recurrent ANN have the possibility to incorporate the output of latter layer units again into earlier layer units, which is not possible in MLP networks. Commonly, the units of all hidden layers of recurrent networks are in a chain-like informational loop. A hidden unit can use its output as input (direct feedback), or it is connected to a hidden unit of the preceding layer (indirect feedback), or it is connected to an unit of the same layer (lateral feedback), or it is connected to all other hidden units (fully recurrent). The recurrent type of LSTM is typically direct feedback (Malhotra, Vig, Shroff, & Agarwal, 2015). The LSTM network, with regard to its inherent properties of "[…] maintaining its state over time in a memory cell […]" (Greff, Srivastava, Koutník, Steunebrink, & Schmidhuber, 2017), is predestined for usage in time series analysis. In opposite to other recurrent network types a LSTM network solves the vanishing gradient problem (Hochreiter & Schmidhuber, 1997).

Fully connectionist MLP are the most used type of ANN for electricity price forecasting (Weron, 2014; Dudek, 2016). They differ in usage of different explanatory variables, e.g. power consumption, weather, wind conditions, in addition to lag variables. Furthermore, the results of a MLP serve as benchmarks in comparison with the results of other forecasting models like ARIMA. Additionally, MLP is often used as the nonlinear part within a hybrid model, e.g. in combination with ARIMA. A further type of ANN, occasionally seen in the extent literature, is a recurrent network (Weron, 2014), especially a nonlinear autoregressive exogenous model (NARX), a descendant of a recurrent network (Marcjasz, Uniejewski, & Weron, 2018).

Weron (2014) concluded that forecasting with univariate time series models is well known in the extent literature. Accordingly, including the right external input factors into the models, as well as dealing with nonlinear dependencies between endogenous and exogenous variables and among exogenous variables will become more important. In contrast to the author, we do not see that the time for univariate time series analysis of electricity prices is already over, as we still cannot see a satisfactory approach to meaningfully deal with the time-series characteristics of electricity prices. Although the Bayes-approach offers possibilities, it is rather unsuitable for practical use due to the high load of computer capacities during simulation operations. Hence, we see a research gap in handling the non-linear correlative effects between the exogenously modelled daily, weekly, and yearly seasonal patterns as well as autocorrelative effects within the time series and among the exogenously modelled variables. Our contribution is to close this research gap by using an ANN-based methodology. We perform a time series analysis for the German EEX "Phelix" Data using (i) different

ANN architectures with non-fully and fully connectionist networks to discover and to deal with correlation effects on exogenous side / input layer, (ii) using a long short-term-memory network (LSTM) to discover and to deal with autocorrelation effects, and (iii) an ARIMAX model with time series features as binary coded variables as a benchmark for the aforementioned models.

3 Sample and Methodology

3.1 Sample Selection

At the European Energy Exchange (EEX), electricity spot prices (EPEX Spot), as well as future contracts are traded. The vast number of German municipal utility companies, but also large industrial consumers on the demand side, and European electricity suppliers on the supply side take part at the electricity trading at the EEX. The electricity volumes can be traded on the same day (intraday) or for the following day (day-ahead). Purchase and sale orders can be placed on an hourly basis as well as for time blocks. The blocks are “baseload” (0.00am - 12.00pm) or “peakload” (8.00am - 8.00pm). These orders can be placed until 12.00pm of every trading day for the next calendar day and will be processed primarily over the internet. A computer system ensures the automatic settlement of the purchase and sale orders and the fixing of the exchange price. Finally, around 12.40pm the prices for the next day will be published via the internet and other data agencies.

The sample data used for this analysis is the EEX Phelix-DE day-ahead spot rate. It has established itself as a benchmark contract for European electricity. We considered time series data from January 1st 2015 until January 1st 2018 (Fig. 1). Each individual day has got 24 hourly price observations. The data underlying this analysis is complete

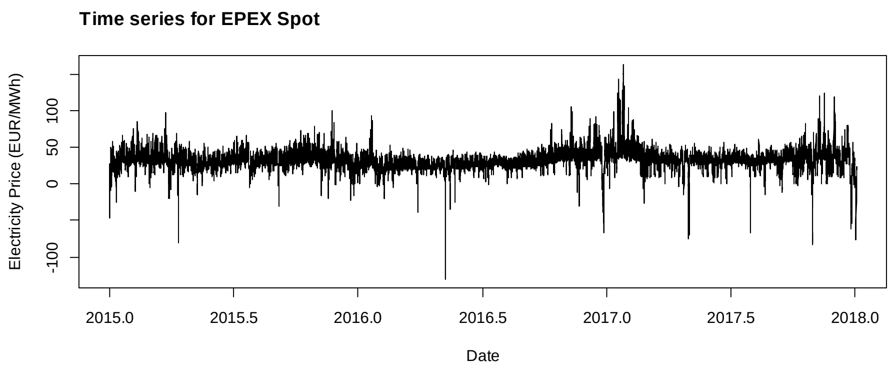

Fig. 1. Time Series for EPEX Spot

Since the storage of electrical power is not possible without significant efficiency losses, the price shows daily, weekly, and yearly seasonality patterns. The seasonality of the time series certainly has its origin in the electricity demand over a day and night rhythm (Fig. 2). Due to the daily, weekly and yearly seasonality patterns binary variables (“dummies”) for these categories were introduced. To capture the seasonality, our models contain 23 hour-dummies for the daily seasonality, 6 weekday-dummies for the weekly seasonality and 11 month-dummies for the yearly seasonality.

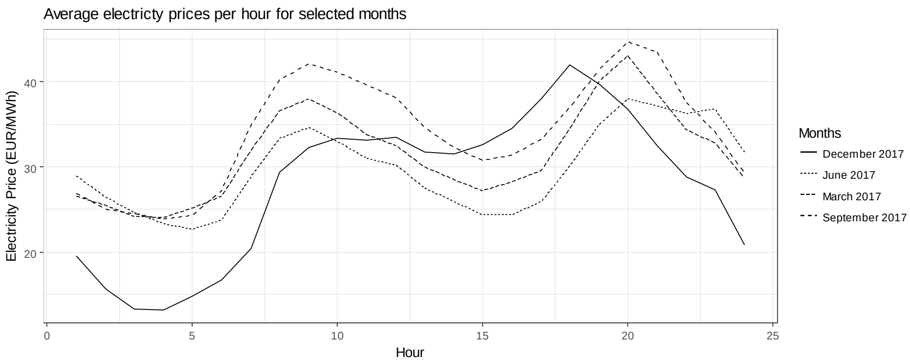

Fig. 2. Average electricity prices per hour for selected month

Beneath seasonal and calendar day effects, the effects of wind power and solar energy increase the volatility of the time series which is particularly challenging in the prediction of the spot prices (Bierbrauer, Menn, Rachev, & Trück, 2007). More and more often, even negative electricity prices are documented at the EEX, which is mainly observable in times of weak demand combined with sunlight or strong wind. Since the present study focusses on seasonality patterns, other explanatory variables (e.g. wind or temperature) were not included into the models.

In this study, our models are trained on a training data set of two years prior the predicted months. We predict the months March, June, September, and December 2017. The root mean squared error (RMSE) is selected to assess and compare the different models. In most of the extent papers, this is the standard forecasting accuracy measure (Weron, 2014).

3.2 ARIMAX-Model with Seasonality

The ARIMAX-model used in this study is an extension of the classical ARIMA-model, introduced by Box & Jenkins (1971). To include seasonality into the model, the binary variables for hour, weekday, and month are applied in the X-term of the model, which means, that these variables are supplemented as additional regressor in the AR-Term. We used the Hyndman-Khandakar algorithm to find the best notation for the ARIMAX model (Hyndman & Khandakar, 2008). The algorithm is using the KPSS tests to determine the number of differences (d) for the training dataset first. In a second step, the values of (p) and (q) are chosen for the training time series by minimizing the Akaike Information Criterion (AIC) out of every probable combination of these two parameters. As a result of this procedure, an ARIMAX(3,1,3) model with 40 dummy variables is used for the analysis.

3.3 ANN-Models

As described above, usually ANN are designed as fully connectionist networks. We suggest a slightly different approach to discover information about correlation of exogenous variables. Therefore, a four step approach is introduced: (i) Single Layer Perceptron (SLP), (ii) Multilayer Perceptron (MLP) with hidden layer and particular mapping (non-fully connectionist network), (iii) Multilayer Perceptron (MLP) with hidden layer without particular mapping (fully connectionist network), and (iv) Long Short-Term Memory (LSTM) with hidden layer, without particular mapping and direct feedback. A synopsis of the used architectures is illustrated in Fig. 3.

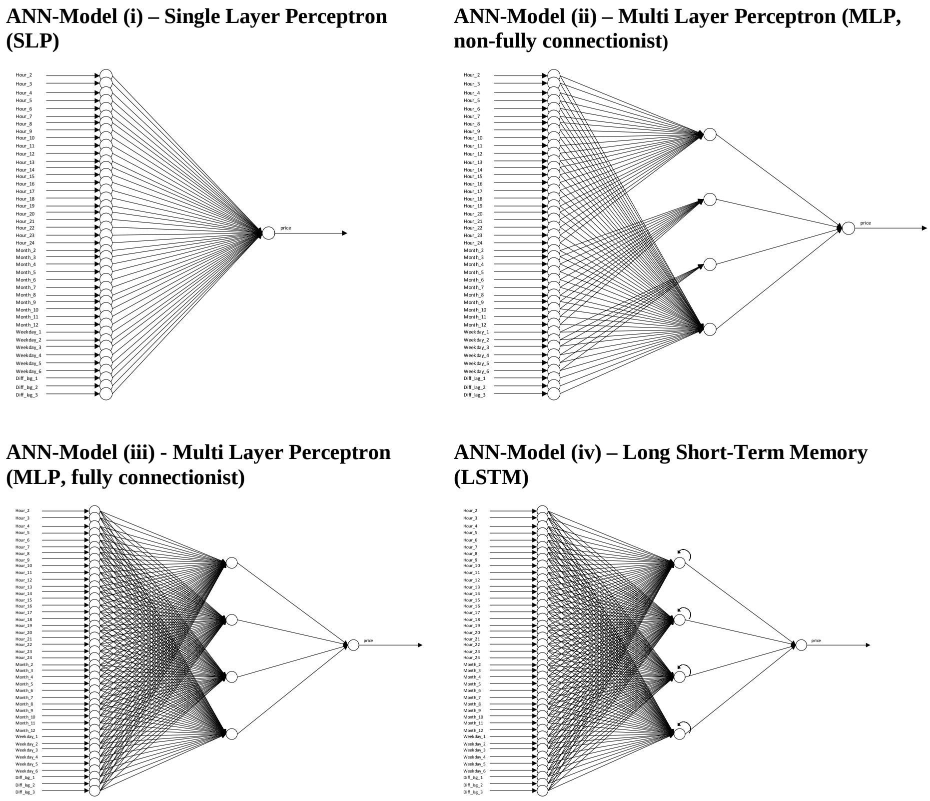

Fig. 3. Synopsis of used ANN-architectures

The input layer is composed of units for hours, weekdays, and months as well as of units for lags. The lag units are lag 1, lag 2 and lag 3 hours to be in line with the ARIMAX-model later. The hidden layers in model (ii) to (iv) consist of four units. The first three units in the hidden layer in (ii) are aggregated units. The first unit stands for an aggregated hourly information and gets its information only from the corresponding hourly input units. The second hidden unit is an aggregated weekday hidden unit and gets its information from all weekday input units. The third hidden unit represents an aggregated month hidden unit and gets its information only from all month input units. The fourth hidden unit can be seen as an all-unit which gets its information from all input units including the lagged variables. The hidden layer in (iii) differs from (ii). Each unit in the hidden layer gets its information from all units of the input layer. There is not such a particular mapping like in (ii). The LSTM model in (iv) is equally designed as the MLP model in (iii), despite the direct feedback for each hidden unit.

The ARIMAX model, which is used as a benchmark model at this point, is able to recognize endogenous autocorrelation of the time series using the lag variables. The binary-coded seasonal variables control the seasonality as additive constants for certain hours, certain days of the week, and certain months via the ARX-term. Relationships between these seasonal components cannot be recognized by this type of model.

ANN-model (i) is essentially equivalent to OLS regression but additionally, it is able to adopt to non-linearities. It is therefore comparable to the ARX term of the ARIMAX model. If ANN model (i) should yield better results than the ARIMAX model, this is due to the ability to map nonlinear relationships as well.

The ANN model (ii) allows the explicit modeling of daily, weekly and annual seasonality through its aggregated units, but prevents the consideration of exogenous correlative effects among these seasonalities. In this respect, it serves as a benchmark for model (iii), which explicitly allows the consideration of exogenous correlative effects. If ANN-model (iii) now delivers better results than ANN-model (ii) we can assume that one reason must be unconsidered exogenous correlation, but its' nature cannot be determined.

Expected better results of ANN-model (iv) would emerge the conclusion that the time series must have autocorrelative effects between seasonal binary variables, too.

4 Results

The model forecasting accuracies in terms of the root mean squared error (RMSE) for the five models as well as for all tested cases can be found in Table 1. In the light of a monthly forecast horizon, the forecasting accuracies are in line with the expectation. In the extent literature, winter season is known to be more volatile and the price is more influenced by exogenous correlation effects, e.g. wind power. Accordingly, the comparatively poorer forecasting result in December is also in line with expectations from literature. It is striking, that the ARIMAX benchmark outperforms the SLP as well as both MLP in three out of four tested periods. Only in the December period all ANN models show better results than the ARIMAX model. Overall, however, the LSTM network provides the best results compared to all other models.

Table 1. Model Forecasting Accuracy Results

Model	Type	RMSE			
		Mar	***Jun***	***Sep***	***Dec***
Benchmark	ARIMAX	7.940	8.842	8.429	28.554
ANN (i)	SLP	10.736	10.879	11.523	22.715
ANN (ii)	MLP (non-fully connectionist)	9.759	11.051	11.078	20.171
ANN (iii)	MLP (fully connectionist)	8.914	11.433	11.643	19.462
ANN (iv)	LSTM	6.947	8.189	9.551	19.273

The coefficients and model statistics of the ARIMAX model are given in Table 2 in the appendix. The coefficients show well pronounced daily, weekly and yearly seasonalities. Although the results are in line with expectation, the annual seasonality is based on only a few observations, resulting in high standard errors. The daily and weekly seasonality, on the other hand, can be described as stable, as inference is based on a large number of observations. It is evident that the daily and weekly seasonalities are far less exposed to structural changes than the annual ones. Nevertheless, the drift term of the ARIMAX model shows only low coefficients. By visual examination of the forecast, it can be seen that the ARIMAX model predicts repetitive daily patterns that oscillate slightly across the course of the week. Examples of this behavior for the months of March and December can be found in the charts in appendix 2.

The daily, weekly and annual seasonality already evident in the ARIMAX model coefficients are reflected in the weights of the SLP as well. Accordingly, it can be assumed, that due to the similarity to the ARX term of the ARIMAX, the poorer forecasting accuracy of the SLP is due to the nonexistence of the MA term in ANN or due to poorer adoption to the data as a consequence of the nonlinear activation function in the ANN. In the visual inspection, the model shows a less predictable behavior, in which daily patterns are recognizable but in a clearly distorted manner. The weaker forecasting quality is not surprising in this respect, although not in line with expectations. Which factor determines these distortions is not recognizable.

The importance of individual hidden units can be determined in ANN model (ii) by their weights to the output layer. It is striking that both, the hour unit, the day unit, and the month unit receive almost no weight and are therefore almost irrelevant to the model result. Only through the all-unit the electricity price forecast is achieved. Since the all-unit is comparable to the SLP, ANN model (ii) does not lead to a much better result than the SLP model. A possible explanation for this behavior is that only the interaction of the seasonal components and the lags provide a sufficient basis for the forecast. By visual examination of the forecast, it can be seen that the model shows a more repetitive result than the SLP, although here also unforeseeable distortions characterize the forecast.

In contrast to ANN model (ii), the weights in ANN model (iii) do no longer show seasonal structures. An interpretation of the individual weights is no longer possible. Even if there is still a strong weighting of a single hidden unit, the strongly correlative influence of the other hidden units on the forecast is clearly given. As already seen in the other MLP models, unpredictable distortions also shape the visual image of this model.

Comparable to ANN model (iii), the LSTM of ANN model (iv) shows strongly correlative influence, but strictly divided in two hidden units, whereof one hidden unit shows an excitatory and another hidden unit an inhibitory behavior. Two further hidden units show low weights, so that their influence is very limited. Due to the comparable architecture with ANN model (iii) the significantly higher forecasting accuracy is due to the ability to store the output of each unit and to feed it into this unit again (direct feedback). In other words: Not only the three lag variables fed into our model reflect autocorrelative effects but also the values stored inside the nets' units to deal with long-term dependencies. Accordingly, the LSTM achieves significantly smoother daily patterns, similar to the ARIMAX model. Like all other models, the LSTM is also unable to predict exogenous shocks, leading to some distortions.

The superiority of the ARIMAX model and the LSTM network in comparison to the other ANN-architectures clearly shows that an additive consideration of seasonal effects for electricity prices is entirely sufficient. An alternative consideration of correlation effects does not provide improved forecasting accuracy. Thus, the problem of electricity price prediction focuses on autocorrelation effects, which can be better considered in the LSTM network than in ARIMAX.

Due to the fact, that all models are fed with the same data – including lagged variables – it is surprising, that the SLP and the MLP models are not able to smoothen the forecast.

5 Conclusion

The electricity price at the electricity exchange EEX shows daily, weekly, and annual seasonality patterns. Due to the cyclicality of the considered seasonal components there are non-linear correlative relationships between them. Thus, the present study deals methodologically with non-linear correlative and autocorrelative time series properties of the electricity spot price. We propose a systematic ANN-based approach to address this problem. The usage of different architectures sheds light on the strength of these relationships and their influence on electricity price prediction.

A single layer perceptron shows lower forecasting accuracy than a standard ARIMAX model with binary coded seasonalities used as a benchmark. Possible reasons for the poorer predictive quality can be specified: The non-linear activation function of the SLP and, above all, the missing MA term, which smooths the results in the ARIMAX model.

A non-fully connectionist multi-layer perceptron (MLP) with seasonally specified aggregated units in the hidden layer is able to improve the forecasting accuracy only slightly, as correlative relationships of the components are taken into consideration individually. The non-fully connectionist MLP shows only low correlations and a specialization of one unit considering all information. Accordingly, the forecasting accuracy cannot be better than in the single layer perceptron by large extent. This gap is closed by the fully-connectionist MLP, where all interactive relationships between these components find their way into the forecasting model. Last but not least, the long short-term memory (LSTM) model provides the most accurate forecast, which, in addition to the correlative relationships already mentioned, also included autocorrelative relationships on the endogenous side over several periods into the forecast.

Appendix

Table 2. Results of the ARIMAX-Model

	March		June		September		December	
	value	**s.e.**	**value**	**s.e.**	**value**	**s.e.**	**value**	**s.e.**
ar1	16,785	0.0818	17,605	0.0824	16,935	0.0894	0.7862	0.1255
ar2	-11,043	0.1066	-12,112	0.1092	-11,114	0.1202	0.1662	0.1668
ar3	0.3519	0.0380	0.3903	0.0381	0.3529	0.0425	-0.0721	0.0563
ma1	-15,174	0.0814	-16,130	0.0824	-15,439	0.0894	-0.6353	0.1257
ma2	0.8308	0.0931	0.9595	0.0970	0.8666	0.1068	-0.2651	0.1491
ma3	-0.3018	0.0280	-0.3372	0.0285	-0.3128	0.0306	-0.0772	0.0371
Drift	0.0005	0.0044	0.0001	0.0042	0.0002	0.0042	0.0010	0.0054
hour_2	-18,453	0.1422	-17,613	0.1392	-17,176	0.1398	-17,153	0.1446
hour_3	-30,859	0.2203	-29,943	0.2139	-29,378	0.2151	-28,786	0.2218
hour_4	-37,111	0.2784	-36,319	0.2701	-35,291	0.2719	-34,848	0.2807
hour_5	-32,652	0.3187	-3,191	0.310	-31,394	0.3122	-30,797	0.3232
hour_6	-14,848	0.3459	-14,952	0.3372	-15,193	0.3402	-13,675	0.3559
hour_7	41,501	0.3649	38,950	0.3562	37,464	0.3600	39,900	0.3811
hour_8	109,063	0.3791	105,241	0.3704	101,211	0.3750	103,164	0.4006
hour_9	129,781	0.3905	126,016	0.3816	121,805	0.3868	124,477	0.4156
hour_10	116,282	0.3994	115,104	0.3906	111,091	0.3960	112,707	0.4266
hour_11	97,316	0.4059	95,963	0.3972	92,259	0.4027	94,941	0.4342
hour_12	90,601	0.4097	89,239	0.4012	85,333	0.4066	88,454	0.4387
hour_13	6,581	0.411	65,081	0.4026	61,004	0.4080	62,931	0.4401
hour_14	50,199	0.4097	49,358	0.4012	44,923	0.4066	46,252	0.4387
hour_15	42,780	0.4059	41,889	0.3972	37,259	0.4027	38,261	0.4342
hour_16	53,135	0.3994	52,199	0.3906	48,610	0.3961	50,092	0.4266
hour_17	70,571	0.3905	69,398	0.3816	66,280	0.3868	69,080	0.4156
hour_18	123,009	0.3792	121,728	0.3704	117,915	0.3750	121,331	0.4007
hour_19	152,748	0.3649	149,034	0.3562	145,497	0.3600	148,915	0.3811
hour_20	158,221	0.3460	152,503	0.3373	149,039	0.3402	151,111	0.3560
hour_21	121,306	0.3188	117,335	0.3101	114,028	0.3123	116,188	0.3233
hour_22	83,177	0.2784	80,509	0.2702	77,582	0.2719	80,662	0.2808
hour_23	66,735	0.2204	63,935	0.2140	61,303	0.2152	63,509	0.2219
hour_24	18,964	0.1423	17,294	0.1393	17,434	0.1400	19,562	0.1448
month_2	22,998	22,506	0.0658	22,414	-0.1652	22,475	0.7607	23,861
month_3	46,908	35,353	0.9387	30,205	0.3728	30,296	19,532	32,245

month_4	147,643	39,119	-19,747	35,096	-27,742	35,208	-11,062	37,491
month_5	103,643	41,226	180,039	39,284	168,511	39,372	198,808	41,808
month_6	93,538	42,189	147,728	42,176	156,443	41,431	192,500	44,097
month_7	87,065	42,244	137,409	42,218	149,274	42,407	189,628	45,211
month_8	53,909	41,181	95,537	41,188	118,970	42,519	160,880	45,387
month_9	61,107	39,078	92,912	39,026	100,672	39,184	161,780	44,395
month_10	73,183	35,688	97,470	35,608	104,586	35,773	89,663	42,306
month_11	23,042	30,485	36,872	30,405	42,197	30,533	59,004	37,933
month_12	47,579	23,006	52,791	22,842	55,008	22,983	57,146	24,238
wd 1	0.9330	0.3424	12,777	0.3340	11,154	0.3353	11,521	0.3483
wd 2	2,440	0.436	26,445	0.4270	24,740	0.4301	25,381	0.4457
wd 3	29,576	0.4654	29,748	0.4607	29,731	0.4627	29,139	0.4810
wd 4	27,587	0.4625	27,150	0.4577	28,138	0.4592	25,464	0.4775
wd 5	23,663	0.4211	25,284	0.4171	25,871	0.4173	22,658	0.4356
wd 6	20,898	0.3233	21,893	0.3221	20,446	0.3230	17,206	0.3371
sigma^2	13.51		13.17		13.28		14.38	
log likeli	-47705.96		-47481.28		-47557.49		-48251.5	
AIC	95507.92		95058.56		95210.98		96599.01	
AICc	95508.19		95058.83		95211.25		96599.27	
BIC	95881.00		95431.63		95584.05		96972.08	

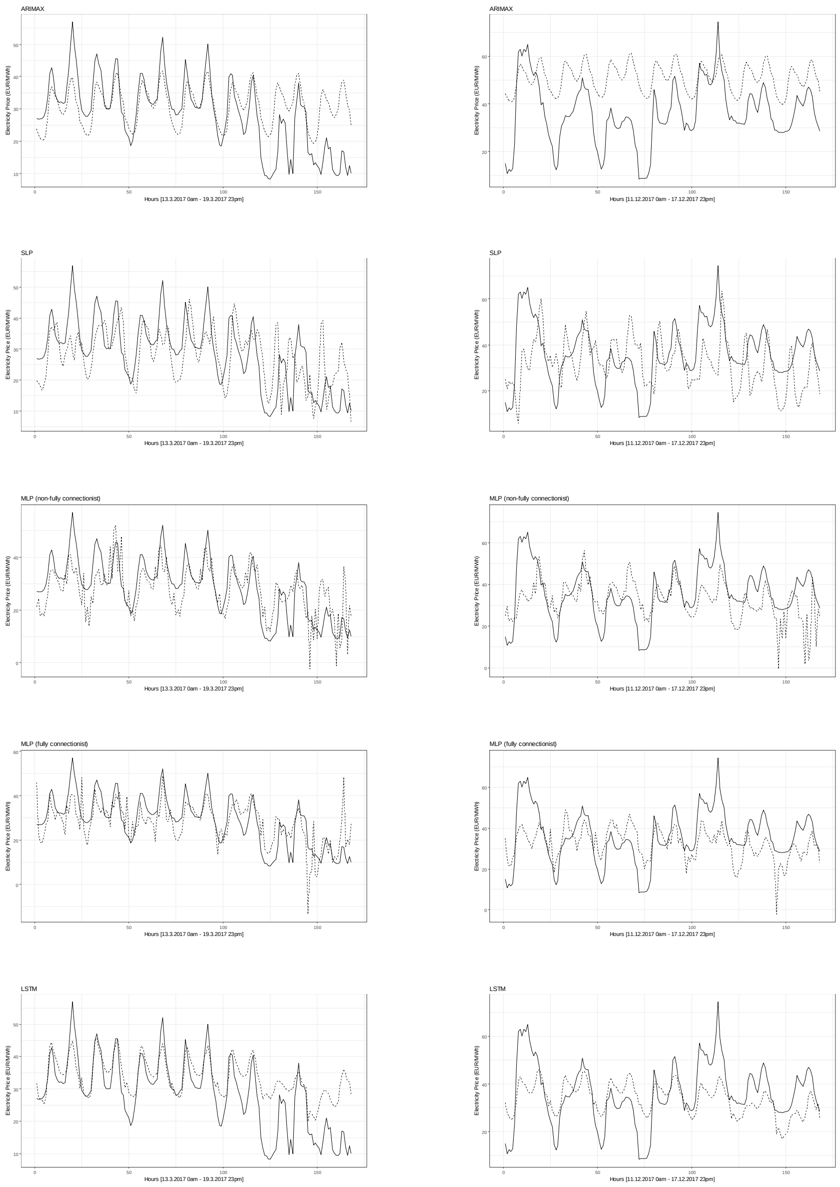

Fig. 4. Out of Sample Excerpt: Forecast vs Real Time Series on March (l) and December (r).

References

1. Adebiyi, A. A., Adewumi, A. O., Ayo, C. K.: Comparison of ARIMA and Artificial Neural Networks Models for Stock Price Prediction. Journal of Applied Mathematics, 1-7 (2014).
2. Bierbrauer, M., Menn, C., Rachev, S. T., Trück, S.: Spot and derivative pricing in the EEX power market. Journal of Banking & Finance, 31(11), 3462-3485 (2007).
3. Box, G. E. P., Jenkins, G. M.: Time series analysis: forecasting and control. 2nd edn. Holden-Day, San Francisco (1971).
4. Conejo, A.J., Plazas, M. A., Espinola, R., Molina, A. B.: Day-Ahead Electricity Price Forecasting Using the Wavelet Transform and ARIMA Models. IEEE Transactions on Power Systems, 20(2), 1035-1042 (2005).
5. Conejo, Antonio J., Contreras, J., Espínola, R., Plazas, M. A.: Forecasting electricity prices for a day-ahead pool-based electric energy market. International Journal of Forecasting, 21(3), 43-–462 (2005).
6. Contreras, J., Espinola, R., Nogales, F. J., Conejo, A. J.: ARIMA models to predict next-day electricity prices. IEEE Transactions on Power Systems, 18(3), 1014-1020 (2003).
7. Dudek, G.: Multilayer perceptron for GEFCom2014 probabilistic electricity price forecasting. International Journal of Forecasting, 32(3), 1057-1060 (2016).
8. Garcia, R. C., Contreras, J., van Akkeren, M., Garcia, J. B. C.: A GARCH Forecasting Model to Predict Day-Ahead Electricity Prices. IEEE Transactions on Power Systems, 20(2), 867-874 (2005).
9. Greff, K., Srivastava, R. K., Koutník, J., Steunebrink, B. R., Schmidhuber, J.: LSTM: A Search Space Odyssey. IEEE Transactions on Neural Networks and Learning Systems, 28(10), 2222-2232 (2017).
10. Haykin, S. S.: Neural Networks and Learning Machines. 3rd edn. Pearson Education, Upper Saddle River (2009).
11. Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural Computation, 9(8), 1735-1780 (1997).
12. Hopfield, J. J.: Neural Networks and Physical Systems with Emergent Colelctive Computational Abilities. Proceedings of the National Academy of Sciences, 79(8), 2554-2558 (1982).
13. Hyndman, R. J., Khandakar, Y.: Automatic Time Series Forecasting: The forecast package for R. Journal of Statistical Software, 27(3) (2008).
14. Knittel, C. R., Roberts, M. R.: An empirical examination of restructured electricity prices. Energy Economics, 27(5), 791-817 (2005).
15. Koopman, S. J., Ooms, M., Carnero, M. A.: Periodic Seasonal Reg-ARFIMA–GARCH Models for Daily Electricity Spot Prices. Journal of the American Statistical Association, 102(477), 16-27 (2007).
16. Malhotra, P., Vig, L., Shroff, G., Agarwal, P.: Long Short Term Memory Networks for Anomaly Detection in Time Series. (2015).
17. Marcjasz, G., Uniejewski, B., Weron, R.: On the importance of the long-term seasonal component in day-ahead electricity price forecasting with NARX neural networks. International Journal of Forecasting, 34(1), (2018a).
18. Marcjasz, G., Uniejewski, B., Weron, R.: On the importance of the long-term seasonal component in day-ahead electricity price forecasting with NARX neural networks. International Journal of Forecasting (2018b).

19. Misiorek, A., Trueck, S., Weron, R.: Point and Interval Forecasting of Spot Electricity Prices: Linear vs. Non-Linear Time Series Models. Studies in Nonlinear Dynamics & Econometrics, 10(3), (2006).
20. Weron, R.: Electricity price forecasting: A review of the state-of-the-art with a look into the future. International Journal of Forecasting, 30(4), 1030-1081 (2014).
21. Zareipour, H., Canizares, C. A., Bhattacharya, K., Thomson, J.: Application of Public-Domain Market Information to Forecast Ontario's Wholesale Electricity Prices. IEEE Transactions on Power Systems, 21(4), 1707-1717 (2006).
22. Zhang, G., Patuwo, E. B., Hu, M. Y.: Forecasting with artificial neural networks: The state of the art. International Journal of Forecasting, 14(1), 35-62 (1998).
23. Zhou, M., Yan, Z., Ni, Y. X., Li, G., Nie, Y.: Electricity price forecasting with confidence-interval estimation through an extended ARIMA approach. IEE Proceedings - Generation, Transmission and Distribution, 153(2), 187 (2006).