Jonas Tøndevoldshagen
Ola Veland Husøy
Marit Halleraker

Predicting electricity prices in Trøndelag using different statistical models

Bachelor thesis

NTNU
Norwegian university of science and technology
Faculty for economics
NTNU Business school

Bachelor thesis in Economics and managment
Supervisor: Denis Becker
Spring 2022

NTNU
Norwegian University of Science and Technology

Preface

This thesis marks the ending to our bachelor's degrees in Economics and Management. The thesis is written in our specialization which is Business analytics. This paper is written in the spring semester of 2022 and is worth 7,5 credits.

Working on this paper has been exciting, educational and challenging. In the process of writing this paper we have learned to build a full model, worked on improving the model, written a longer thesis and cooperated together. We hope this paper can be a helpful tool for anyone wishing to understand the elements behind a prediction model, or useful for anyone who plan to build their own model.

We would like to express our sincere gratitude to our advisor Denis Becker. He has given us useful advice, guidance, and helpful dialog when we have been stuck.

The contents of this paper are the sole responsibility of the authors.
Innholdet i denne oppgaven står for forfatternes regning.

Jonas Tøndevoldshagen Ola Veland Husøy Marit Halleraker

Trondheim, April 29, 2022

Abstract

In this thesis, we have built models for predicting electricity prices in Trøndelag, including a linear regression model (OLS), simple recurrent neural network model (SRNN), and a Long Short term memory network model (LSTM). Our dataset used the five-year period between 2017-2022, and it consisted of historical electricity spot prices from Nord Pool, together with weather data from the Frost API through the Norwegian Meteorological Institute. The paper first explains important concepts, methods, and theory to give some background to our analysis. We then described the methodology behind our models. We built a benchmark to measure our model that set tomorrows price as the same as today. We then built a linear regression model that attempts to predict spot prices by using linear relationships. We also built a SRNN model and a LSTM model.

We found that it is difficult to make good predictions outside the data that we trained the model on, and we were not able to beat our benchmark.

Sammendrag

I denne oppgaven har vi bygget modeller for å predikere strømpriser i Trøndelag, inkludert en lineær regresjonsmodell (OLS), en simple recurrent neural network (SRNN), og en Long short term memory nettverksmodell (LSTM). Vårt datasett går over en 5 års periode fra 2017-2022, og det består av historiske spotpriser for strøm fra Nord Pool, sammen med værdata fra Frost API hentet fra meteorologisk institutt. Oppgaven forklarer først viktige konsepter, metoder og teori for å gi et grunnlag for analysen vår. Vi beskriver så metodologien bak modellene våre. Deretter bygger vi en benchmark for å måle modellen vår som setter morgendagens pris til det samme som i dag. Så har vi bygget en lineær regressionsmodell som prøver å predikere spot priser ved å bruke lineære forhold. Vi har også bygget en SRNN modell, og en LSTM modell. Vi fant ut at det er vanskelig å lage gode prediksjoner utenfor dataen som vi trente modellen med, og vi klarte ikke å slå benchmarken.

Table of contents

1-Introduction **1**
1.1 Topicality *1*
1.2 Delimitation *2*
1.3 Structure *2*

2-Theory **3**
2.1 Analysis tools *3*
2.1.1 Artificial neural networks 3
2.1.2 LSTM 4
2.1.3 Overfitting vs. Underfitting 5
2.1.4 AIC and BIC 5
2.1.5 Activation functions 6
2.2 The power market *7*
2.2.1 Nord Pool 7
2.2.2 Spot prices 8

3-Method **11**
3.1 Retrieving data *11*
3.2 Data preparation *11*
3.3 Creating a benchmark model for predicting electricity prices *13*
3.4 Constructing and fitting the models *13*
3.5 Reliability and validity *13*

4-Analyzis **14**
4.1 Descriptive statistics for electricity prices *14*
4.2 Descriptive statistics for weather data *16*
4.3 The performance of the benchmark model. *20*
4.4 The performance of the OLS model *21*
4.6 The performance of the SRNN model *22*
4.8 The performance of a deliberately overfit model *24*

5-Conclusion **25**
5.1 Evaluating our models *25*
5.2 How the models can be improved *25*

List of references: **27**

Figures and Tables

Figure 1: Basic structure ANN---p.3
Figure 2: Sigmoid and ReLu---p.7
Figure 3: Price variation Norway--p.9
Figure 4: Histograms electricity prices---p.14
Figure 5: Correlation matrix electricity prices---p.15
Figure 6: Correlation plots electricity prices--p.16
Figure 7: Correlation matrix- weather and NO3--p.17
Figure 8: Heat map--p.18
Figure 9: Normalized data plot--p.19
Figure 10: Benchmark---p.20
Figure 11: OLS model---p.21
Figure 12: SRNN model---p.22
Figure 13: LSTM model---p.23
Figure 14: Overfit model--p.24

Table 1: Benchmark ---p.20
Table 2: OLS model---p.21
Table 3: Coefficient and intercept for the OLS model--p.21
Table 4: SRNN model results--p.22
Table 5: SRNN model structure---p.22
Table 6: LSTM model results---p.23
Table 7: LSTM model structure--p.23
Table 8: Overfit model results--p.24
Table 9: Overfit model structure---p.24

1-Introduction

As a part of our course of study we have chosen to specialise our bachelor thesis in Business Analytics. With Business Analytics as a background, we have chosen a topic for this paper that we find highly interesting and relevant for the time being, namely electricity pricing. The introductory chapter contains an outline for why we find the topic interesting, followed by our delimitations and a plan for our model. As a last part of the introduction, we will establish the further structure of our paper. The purpose of this is to give the reader an overview to better comprehend the paper throughout.

1.1 Topicality

What might be considered the number one topic for the Norwegian newspapers throughout the last months of 2021 and the starting months of 2022, has been electricity prices. It has been highly discussed because of the increase we have seen in the prices all over the country. The enormous price jump has had a significant impact both on businesses and on households' economy. It has been so impactful that the Norwegian government has seen it necessary to implement an aid scheme, in order to help households through the period of abnormally high prices. (Regjeringen.no, 2021).

What makes the market for electricity different from other financial markets, and why it is especially interesting to attempt to predict the prices, is the very nature of electricity. Electricity is non-storable, but the supply and demand does not always correlate and is affected by a range of variables. Despite this inconvenience with electricity not being able to be stored, it will always be necessary to produce more. Electricity can be considered a basic human need in the modern industrial world, as much of our way of life is dependent on it. We use electricity to heat our houses, to produce and store large quantum of food, in the production of other goods, for lighting, and in an increasing degree for transport. So much of our lives are dependent on electricity and therefore also affected by the prices. Additionally with population growth and income growth the demand is increasing. We find electricity pricing to be a highly current and interesting topic that we wanted to explore. As it is both relevant to our society, and to ourselves with our own electricity bills. We saw an opportunity to use the abilities we have learned through our Business Analytics courses in building models and that was our entrance to this project.

1.2 Delimitation

Within the subject of electricity prices, we have chosen to build models for electricity price forecasting including use of artificial neural networks. We have chosen to limit our model to looking at electricity prices from Nord Pools database, and then we further limited it to only focus on region NO3 *Trøndelag*. The reason for these limitations was to limit the extend of the models reach in order to be able to build a sounder model, and we found it most interesting to look at the electricity prices here in our own region. We choose to look at daily data for electricity prices from the past five years and our goal was to build a model that could predict the spot price for the following day.

1.3 Structure

The remainder of this paper is structured as followed: An introduction which we are now at the end of. Then followed by a theory chapter, where we will present the concepts and terms that we find relevant to present to the reader as background information before presenting our own model. In chapter 3 we showcase and present the method, we will introduce the data and our approach to building the model, as well as the challenges we have faced. Furthermore, chapter 4 which will be an analysis. Here we will present the model and look at how it works, and the aspects of it. At last in chapter 5 we will draw some conclusions as to the validity of our model, and our suggestions for how it could be further improved.

2-Theory

Before we get into our model for predicting electricity prices, we wish to cover some relevant theories to our paper and explain terms that will be commonly used. We do this to make and ensure a foundation of understanding that will be essential to follow our further analysis.
We will explain what an artificial neural network is, how the power market works including what is especially relevant to the electricity prices in Norway, since our model will only cover Trøndelag, Norway. Then we would also like to give an explanation as to the basic functions of the electricity market. Finally, we will look at some previous research on the area.

2.1 Analysis tools

2.1.1 Artificial neural networks

In order to understand neural networks, it is important to know that it is a branch of many other terms. We have artificial intelligence (AI), and then we have machine learning (enabling computers to act without being explicitly programmed for every action), and then a branch from machine learning is deep learning (which is machine learning with need for less human input, where the program is able to process data on its own), and on a branch off of deep learning we find neural networks (IBM, 2020).

Neural networks get their name 'neural' because they are meant to imitate the human brain, and how the neurons in our brains send signals to one another. It is called an artificial neural network because we are artificially imitating those connections. An artificial neural network (hereafter ANN) is built up of an input layer, an output layer and (often many) hidden layers.

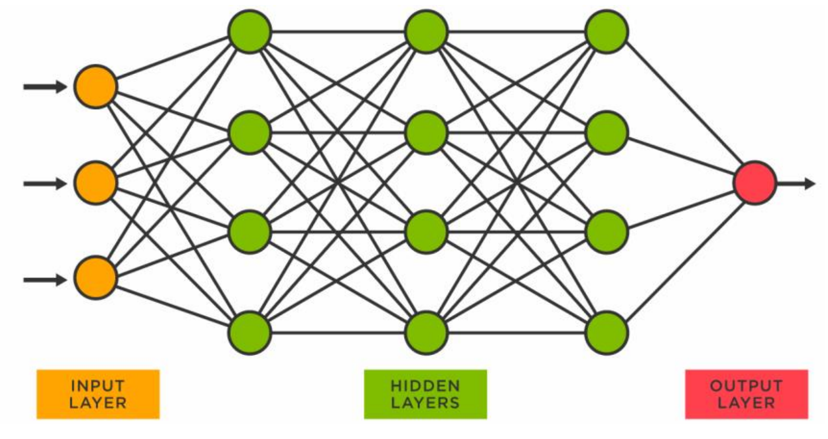

Figure 1: Basic structure ANN

The basic structure of an ANN is shown above. Each node is connected to another and sends data on to the next layer. Weights are assigned to each node to determine relevance of each variable; a larger weight means it will have more significance in determining the final output layer.
A neural network that consists of more than three layers is considered a deep network and is referred to as deep learning. (IBM, 2020).
ANNs can be divided into two main categories on whether it is feed-forward (meaning that the network has no loops) or recurrent (feedback) networks (where there are loops due to feedback connections) (Weron, 2014).

When it comes to the application of artificial neural networks, we have looked at Abiodun`s paper *State-of-the-art in artificial neural network application: a survey (*Abiodun, 2018). Artificial neural networks have become increasingly popular in solving various problems. ANNs are easy to use and useful in problem solving. The paper goes on to talk about how machine learning is the most recent development in a digital revolution and how it is making computing processes both more reliable and more cost-effective. Furthermore, it lists another reason for ANNs being more useful in practical application because unlike a lot of statistical models, ANNs does not require any hypothesis. In addition, the application of ANNs is so varied that it has no boundary.

In order to review both sides of the argument we have also looked at Vivek Kumar`s *Why are neural networks not the answer to everything?* (Kumar, 2019). Kumar points out some of the disadvantages that also comes from the use of neural networks. That includes the black box nature of neural networks, how a lot of advantages, such as speed, is lost when the libraries do not fit, and that it requires large amounts of data. Most of the time, neural networks do not perform well with small amounts of data, and it would then be more beneficial to use a different type of network.

2.1.2 LSTM

When using a recurrent neural network (hereafter RNN) we face a problem with memory loss, called the vanishing gradient problem. This means that the network is only looking at the most recent input. In order to build a model that is capable of looking further back and remembering in order to see the connections, we need to implement a long-short term memory (hereafter LSTM). LSTM is therefore a more advanced version of RNN. There are many situations where we need to process not only the most recent input. For understanding the context in a book, you need to remember more than just the previous page, and this is the same when building a prediction model. However, the network cannot keep all the information, so it needs to sort out the irrelevant information.

For each stage the LSTM needs to first determine what of the current information it can forget, then it needs to analyse the new information it is getting in each stage and then send that on to the next stage, where this needs to be repeated (Shipra_saxena, 2021).

2.1.3 Overfitting vs. Underfitting

The terms overfitting vs. underfitting refers to how well the model`s prediction is correct, when looking at the test data compared to the training data. When building a model, it would be relevant to be correct also on yet unseen data, and how it works on the training data is an indication of this.

Underfitting (also referred to as error due to bias) can happen when the complexity of the model is too low, the neural network is then not able to learn the pattern from input variable to the output variable.

Overfitting (also referred to as error due to variance) can happen when there are some fluctuations in the training set that the model misinterprets as patterns, and so when the model sees new data the performance decreases. (Oppermann, 2021).

2.1.4 AIC and BIC

$$AIC = n \ln(L) + 2p$$

$$BIC = n \, ln(L) + p \ln(n)$$

If one were only to look at R-squared to select the best models, then the model selected would almost always be the most complex one, since R-squared does not consider the number of parameters. Selecting complex models can be a problem when it comes to the model’s performance outside what it has been trained on. In other words, more complex models tend to just remember the data and quickly overfit. AIC and BIC try to address this by punishing models with more parameters. If one adds a parameter that does not improve the performance off the model enough, then the AIC and BIC would get worse. We end up calculating AIC and BIC like the formulas above. The n is the length of the trained data, L is the model performance (MSE), and p is the number of parameters used by the model. When working with AIC and BIC, lower score indicates a better model (Zajic, 2019). The difference between AIC and BIC lies in how they calculate the punishment of additional parameters. BIC have a larger punishment and therefore prefers simpler models than AIC (Analyttica Datalab, 2019).

2.1.5 Activation functions

An activation function is the factor in a neural network that determines how 'active' a node should be, hence its name. It considers all the input and bias of the node and then determines what to send forward in the network. When it comes to choosing an activation function you must consider the function you are trying to approximate. Based on that you should choose an activation function that will approximate that function faster. This is because that will make the training faster. There is not one activation function that is superior in all instances, but they each have their different areas of use. The sigmoid function is good when using it as a classifier due to its shape. However, the sigmoid function struggles to change its output with changing input when it is close to the end part of the function. This is what's called a vanishing gradient problem, and it could potentially stop the neural network from continuing its training. This is because the gradient is so small that it (depending on the model) can stop it or significantly slow it down. Moreover, often you don't know the shape you are trying to approximate. Then a ReLu function can be more useful since it is more of a general approximator.

There are several factors that determine the choice of activation factor. The shape, the range of output values, whether it is a binary classifier, whether it is for a hidden layer or not, and how computationally expensive it is. Often different activation functions are used for different parts of the model, and it can also be necessary to try out different ones (Sharma, 2017).

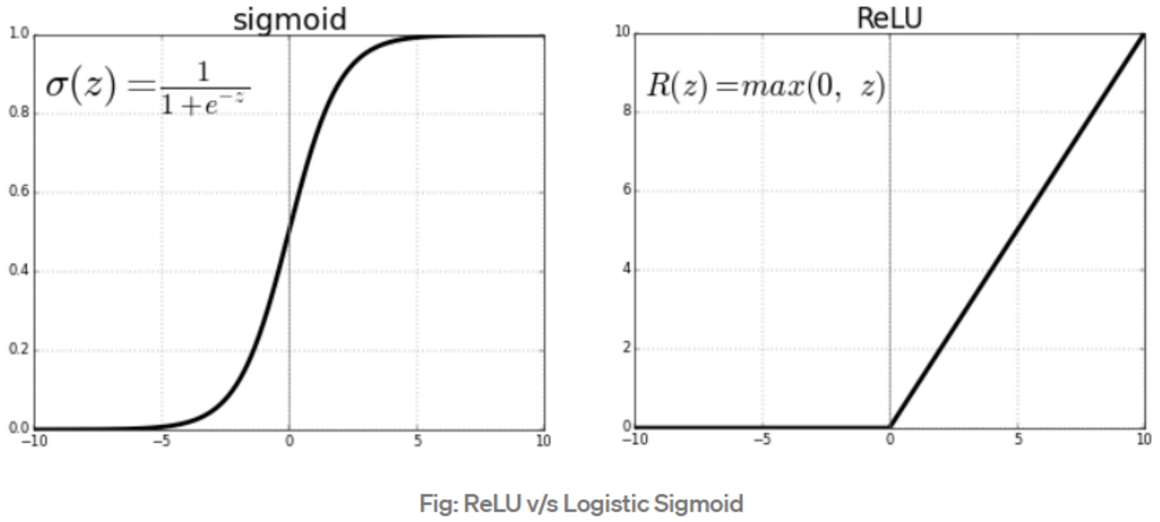

Figure 2: Sigmoid and ReLU (Sharma, 2017)

2.2 The power market

Since our model is built based on data from Trøndelag in Norway, we must look at how the power prices are specifically in Norway. Before 1991 there was regulation on the power market. That means that politicians decided the price of electricity. This was criticized based on many different factors. How much the electricity producing companies' supplies is among other things affected by weather (windmills are affected by wind amount, hydropower by rainfall and so on), cost of labour, cost of machines, development in technology, and political factors (subsidies). The supply is also naturally affected by demand, and vice versa. Some of the factors that determine demand is again weather (for example we use more electricity when it is cold and more air-condition when it is hot), economic growth (people rising out of poverty increase the demand), population growth (long term effect) and again technologic advances that may require more electricity. (Nord pool, no date).

2.2.1 Nord Pool

As a consequence of the deregulation in 1991 there was in 1992 established a Nordic commodity market for electricity, *Nord Pool*. Initially the market only consisted of Norway, but Sweden, Finland and Denmark all joined the market in turn. And from the year 2000, Nord Pool was the major power exchange for the Nordic countries. Nord Pool is an exchange that operates with spot prices which we will explain later under 2.2.2. Spot prices (Weron, 2006).

The idea behind the deregulation was to make access to electricity created through a multitude of sources that are not affected by the same factors, such as weather. This should make the power supply more secure. It is also supposed to be more efficient as we are using the types of power that in the moment are in large quantity, and trading with a large market. (Nord pool, no date). The intention with Nord pool is to even out risk factors. An example of a risk factor is how Norway has a lot of hydropower and is therefore susceptible as it is vulnerable to changes in rainfall. However, since we through Nord pool are a part of a power market our prices are also impacted by the power supply and demand in other countries in Nord pool and those that supply Nord pool.

Another important factor to note about Nord Pool is that electricity is traded in euros. Therefore, if the Norwegian kroner is weak, we pay more for power, and the other way around. It is also important to keep in mind that in Norway our power consumption is increasing more rapidly than our power production. This means that we continuously must import more and more of our power if we cannot improve our production or decrease our consumption. (NTE, no date).

2.2.2 Spot prices

Spot price is one way of paying for and getting your electricity in Norway. Spot prices are giving the consumers the price that the power companies are buying electricity for at Nord pool. The price for 1 Mwh will therefore be equal between the power companies, with an additional surcharge. The spot price for power will change on an hourly rate. The market for spot prices is a day-ahead market that does not allow for continuous trading since the prices are set for each hour of the day on the previous day.

The other option is through the contract market, which is agreeing on a fixed price with the company that supplies your power. With this power agreement the costumer will be committed to a given company over time. There will be certain benefits connected to each option, the fixed price provides safety in the way of not getting surprised by the electricity bill, but in “normal times” one can argue that the spot price will be the better (cheaper) option. The subject of this paper is chosen because of its actuality of the time being, the price for electricity has gone up all over Europe, and in Norway records has been broken multiple times starting from November of 2021. (Lier, 2021)

	Oslo	Kr.sand	Bergen	Molde	Tr.heim	Tromsø
01-02-2022	143,79	143,79	135,97	19,80	19,80	18,87
31-01-2022	162,90	162,90	142,86	27,44	27,44	27,40
30-01-2022	120,67	120,67	120,67	12,55	12,55	12,55
29-01-2022	107,61	107,61	107,61	15,04	15,04	15,04
28-01-2022	135,99	135,99	135,99	15,39	15,39	15,39
27-01-2022	126,41	126,41	126,41	27,26	27,26	27,26
26-01-2022	130,30	130,30	130,30	14,99	14,99	14,99
25-01-2022	154,84	154,84	154,84	14,27	14,27	14,27
24-01-2022	134,20	134,20	134,20	12,47	12,47	12,47
23-01-2022	136,14	136,14	136,14	12,14	12,14	12,14
22-01-2022	138,81	138,81	138,81	14,54	14,54	14,54
21-01-2022	137,84	137,74	137,84	18,01	18,01	18,01
20-01-2022	130,24	130,24	130,24	13,61	13,61	13,61
19-01-2022	135,65	135,65	135,65	15,23	15,23	15,23
18-01-2022	161,58	161,58	149,60	17,87	17,87	17,87
17-01-2022	132,33	132,33	132,33	17,20	17,20	17,20
16-01-2022	127,73	127,73	127,73	17,44	17,44	17,44
15-01-2022	155,90	155,90	147,38	19,39	19,39	19,39
14-01-2022	129,76	129,76	129,76	18,23	18,23	18,23

Figure 3: Price variation Norway from Nord Pool

The figure above shows the differences in prices throughout Norway from the middle of January 2022. We can see that the prices are showing huge differences across the country, and that the people living in the southern cities of Norway have suffered the most throughout this on-going crisis.

2.3 Electricity price forecasting

A variety of methods and ideas have been tested within the subject of electricity price forecasting (EPF) over the years, with varying degrees of success. "Electricity is an interesting commodity because it is "economically non-storable, and the power system stability demands a persistent equilibrium between production and consumption" (Weron, 2014).

In order to gain an understanding of what has been done before, we looked at the review article by Rafael Weron titled *Electricity price forecasting: A review of the state-of-the-art* *with a look into the future*. The nature of electricity includes it being non storable and the need for balance between production and consumption means that the need for price forecasting is perhaps greater on electricity than on any other market. Also, that with spot prices which is mostly used it is a significant difference from other financial markets that it is not a continuous market, but a day-ahead market. In many electricity markets negative prices are allowed, this is because it can be more expensive to shut down and later power up production than it is to accept negative prices. Negative prices can happen for example when the renewable electricity sources are producing especially much for example extra rainfall leading to extra hydropower. The paper discusses a lot of different models for electricity price forecasting. For example, similar-day method which predicts price by looking at a day that has the same characteristics such as which weekday, month, and weather with more (Weron, 2014).

3-Method

3.1 Retrieving data

In the process of retrieving data on electricity prices, we have as mentioned used data from Nord Pool. In chapter 2.2.1 we have described what Nord Pool is. We choose to download data from the past five years, with daily data. We downloaded this data to a csv file, and then we uploaded it into python. In python we have visualised the data and built our models. In the starting process of our assignment, we considered all the spot prices that we could collect from Nord Pool, but we ended up only using data from Trøndelag and connected power regions to try to achieve more specific results.

To predict electricity prices, we found it necessary to also downloaded the weather data for the same period. We did this through an API called Frost. The API delivers data directly from the Norwegian meteorological Institute. The data we first collected was from three weather stations, and consisted of measurements of temperature, precipitation, and wind data. After we looked at the data, we decided to drop the wind data since it was both often lacking data, and it had low correlation to the spot-price.

The locations of the used weather stations are (weather station id):

- Voll-Trondheim (SN68860)
- Tunnsjødal (SN74350)
- Nea (SN68290)

The decision on which stations to use was based on their proximity to population centres and powerplants, and we made sure not to have them too close to each other. We did not experiment further with different stations or data types.

3.2 Data preparation

To be able to create our different models, we first had to prepare the data. Most of this work was done through python and with use of the data preparation library called pandas[1]. Each

[1] Pandas website: https://pandas.pydata.org/

model needed the data prepared in different ways and here are some of the general data preparation we went through:

- Got all the data into a combined data frame.
- Removed the features that we decided not to include (windspeed and non-connected power regions)
- Shifted the weather data one day, so that the data for day 2 appears as data for day 1, day 3 for day 2 and so on[2].
- Removed the last row, since it was now missing weather data
- Made a dependent variable out of NO3 (Trøndelag)
- Made the rest of the data to be independent variables
- Sliced the data into a training set and a test set. The training set was the first 4 years, and the test set was the last year

Some of the more model specific preparation:

- In the Neural Network models, we Scaled all the data to be between 0 and 1. And this was done to improve model results. All the data from the models was scaled back before we evaluated the models
- For the LSTM model, we created an 3D array that we needed to fit the model. For this task we used the 'time series generator' by Keras[3]. The array got all the independent variables for the previous 3 days to make the predictions.
- In all models except the LSTM model, we shifted the dependent variable by one day, so that the dependent variable for day 2 is lined up with the independent variables for day 1. This was apparently automatically done in the Keras generator for the LSTM
- For the SRNN and the linear model we created 2 additional lagged variables for the spot price in Trondheim, so that the model could use the last 3 days in its predictions.

---

[2] If someone were to try to predict power prices in real time, they would not be able to use the exact weather data in the future, but they could use forecasts. The data we collect is real data (we were not able to collect historical weather forecasts). But we used the data as if it was a forecast for the future. Doing it this way will avoid some noise that we would have in real time prediction.

[3] Keras website: https://keras.io/

3.3 Creating a benchmark model for predicting electricity prices

The benchmark was constructed to give some sort of scope of how good our model was at predicting the prices. The benchmark takes the price of the given day and predicts that it's going to be the same for the day after. We think this is a suitable benchmark that a good model should be able to beat.

3.4 Constructing and fitting the models

The difficulty to construct and fit the different models, was varied. The OLS was straight forward, we used the 'Linear Regression' function from the python library sklearn[4]. For linear regression one should expect to receive the same results when repeating the construction and fitting with the same data. In other words, reliability should not be a problem (as long as the same data is being used).

The SRNN and LSTM was a bit trickier than the OLS. We used Keras to help us Construct and fit these Models. Our approach to find a good model structure that could fit well with the data, consisted mostly of trial-and-error. There are a lot of parameters in ANNs, and we needed to experiment to be able to adjust the parameters correctly. Our trial-and-error approach could probably be better if we used a more systematic approach. In the trial-and-error prosses we needed to make sure that our model was not too much overfit or underfit, and to do that we looked at the estimations on the test-set. This can potentially mean that our model will not do as good as the test set on out-of-sample data. This can be the case because the model is indirectly fitted to the test set. Using validation set could have helped in this case. We did look at AIC and BIC scores, but AIC and BIC can only help comparing models where the number of features are different. When we had our structure and tweaked the parameters, we tried to fit the model multiple times to get the best fit. When fitting ANNs you would ideally like to collect an average of multiple runs, but this was difficult in our case, since many of the attempted runs got stuck early on and could not be fitted more. Reliability in ANNs is difficult because of the random nature of the initial position when you start the fit.

3.5 Reliability and validity

The term's reliability and validity can tell us something about the quality of the method used. Reliability refers to whether the results you have obtained can be reproduced. Validity refers

[4] Sklearn website: https://scikit-learn.org/stable/

to whether we are measuring the things we are meant to measure, and not just seeing a false correlation. (Jacobsen, 2015). We have chosen to use data from Nord Pool and the Norwegian meteorological institute which are both accurate and reliable sources. We have previously listed how our models were built. We have split the dataset into a test and a training set in order to test the validity.

4-Analyzis

We are going to start the analysis by looking into each of our datasets separately. We will look at: descriptive statistics, some calculations that were made, as well as some plots, in order to better understand the data. Then we will merge our datasets and move on to look at the results of our models.

4.1 Descriptive statistics for electricity prices

As described in the method section we got data about electricity prices from Nord pool and used python to analyse and construct the large data sets to be interpretable. We wanted to start off with doing some descriptive statistics to get a better understanding of the data. Nord pool gave us data for 17 different regions in Scandinavia and the Baltic states. We could see that the average price varied from the smallest being 297,57 in Tromsø (NO 4) to the highest with an average of 509,09 in Lithuania (LT). The unit of measurement being euro per MWh. The difference was even more apparent when we looked at the minimum price in the five-year period. The one with the highest minimum price 47,76, that was the minimum in all the Baltic states, Lithuania, Estonia and Latvia. The Zone with the smallest minimum price was DK1 with -152,61. Negative electricity prices can happen when the supply greatly exceeds the demand, which can happen for power sources that cannot control the amount of energy they create, such as solar panels and windmills.

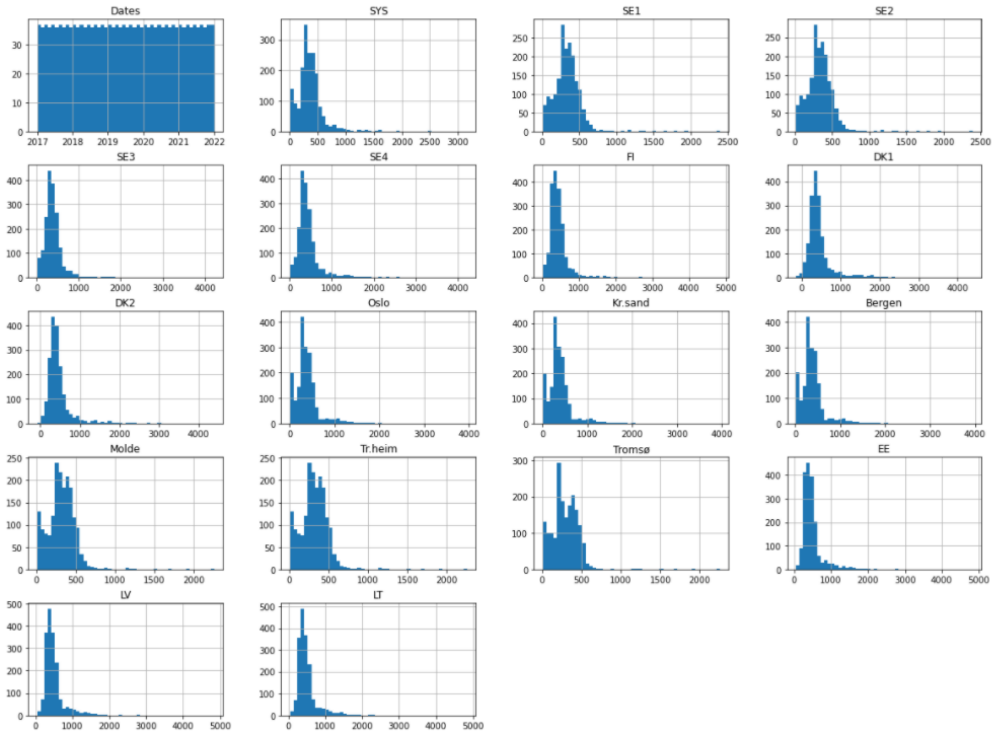

Figure 4: Histograms electricity prices

From the figure above you can see the different histograms for the weather zones. Even though we can see that they have different values as we also saw from the descriptive numbers, they all follow roughly the same shape. Which makes sense since they are all part of the same market, and they are affected by a lot of the same factors.

We then look at how the other zones compared correlated to our own Trøndelag zone.

The results are shown in the picture below:

```
Tr.heim     1.000000
Molde       1.000000
SE2         0.950324
SE1         0.949741
Tromsø      0.945802
SYS         0.776043
Oslo        0.648216
Bergen      0.644438
Kr.sand     0.641261
SE3         0.632384
FI          0.586083
SE4         0.552873
EE          0.507887
LV          0.499859
LT          0.489935
DK2         0.484762
DK1         0.465711
Name: Tr.heim, dtype: float64
```

Figure 5: Correlations for Nord Pool zones relative to NO3

We discovered that Molde is identical to Trondheim and has 100% correlation. The explanation for this is that Molde and Trondheim it a part of the same power network the NO3. Furthermore, we see that the correlation between Trondheim and the other zones vary greatly. Both within Norway with Kristiansand only having a 64,12% correlation but that Trondheim has the least correlation with Denmark which only has a 48,47% and a 46,57% correlation.

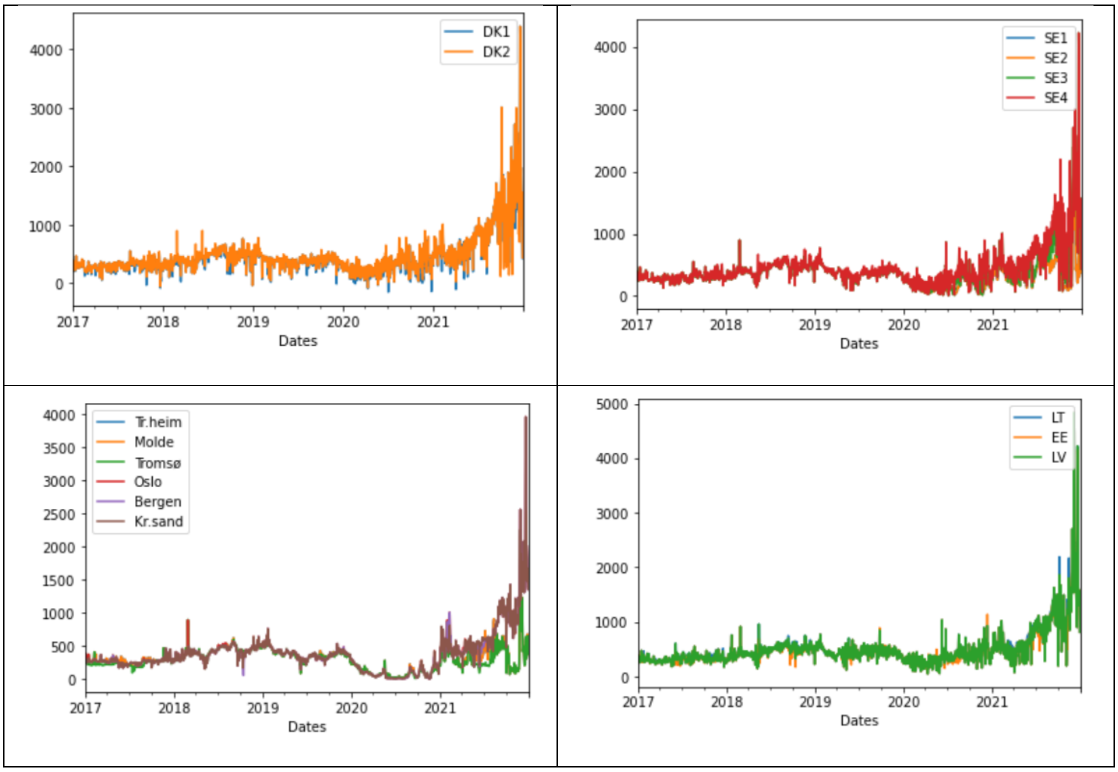

Figure 6: Correlation plots

We plotted the different zones separated by countries to look at what correlation it was within each country, in Norway, Sweden and Denmark and we also made one to compare the three countries of the Baltic states. As you can see there were very little differences within each country (and within the Baltic states). The plot with the biggest difference is Tromsø vs. Kristiansand, but these zone although both being in Norway have a great geographical distance.

4.2 Descriptive statistics for weather data

As explained in the method section we downloaded our weather data through Frost. We now want to do some descriptive statistics for the weather data along with the electricity prices and look at the correlation.

For each of our three weather stations we looked at the mean temperature, the total precipitation and the mean windspeed for each day. Including that for each station in addition the electricity prices, we now have 25 variables. We found that the three weather stations had similar mean values for the mean temperature, precipitation and windspeed over the five-year period.

We looked at the correlations (Figure 7 below) and found all the variables for weather had a negative correlation with the electricity price in NO3. This means that when these variables go down, the electricity prices go up and vice versa.

```
NO3                             1.000000
SN68860_sum_precipitation_mm   -0.069951
SN68290_sum_precipitation_mm   -0.071060
SN74350_mean_wind_speed_mps    -0.075961
SN74350_sum_precipitation_mm   -0.116751
SN68290_mean_wind_speed_mps    -0.129623
SN68860_mean_temp_degC         -0.175571
SN74350_mean_temp_degC         -0.177387
SN68290_mean_temp_degC         -0.185219
SN68860_mean_wind_speed_mps    -0.189403
Name: NO3, dtype: float64
```

Figure 7: Correlations weather and NO3

We then used the seaborn library[5] to generate a heat map. All the variables for weather had a negative correlation as we saw in the correlation matrix, and we can see in the heat map that the weather indicators that have the strongest negative correlations (shown with the colour black) is the mean temperature for all three weather stations, and the mean wind speed in weather station SN68860 (Voll-Trondheim). Another interesting takeaway from the heat map is that the temperatures from the different weather stations correlates a lot. This could mean that having temperature data from 3 different stations is unnecessary since the data is adequatley similar.

[5] Seaborn website: https://seaborn.pydata.org/

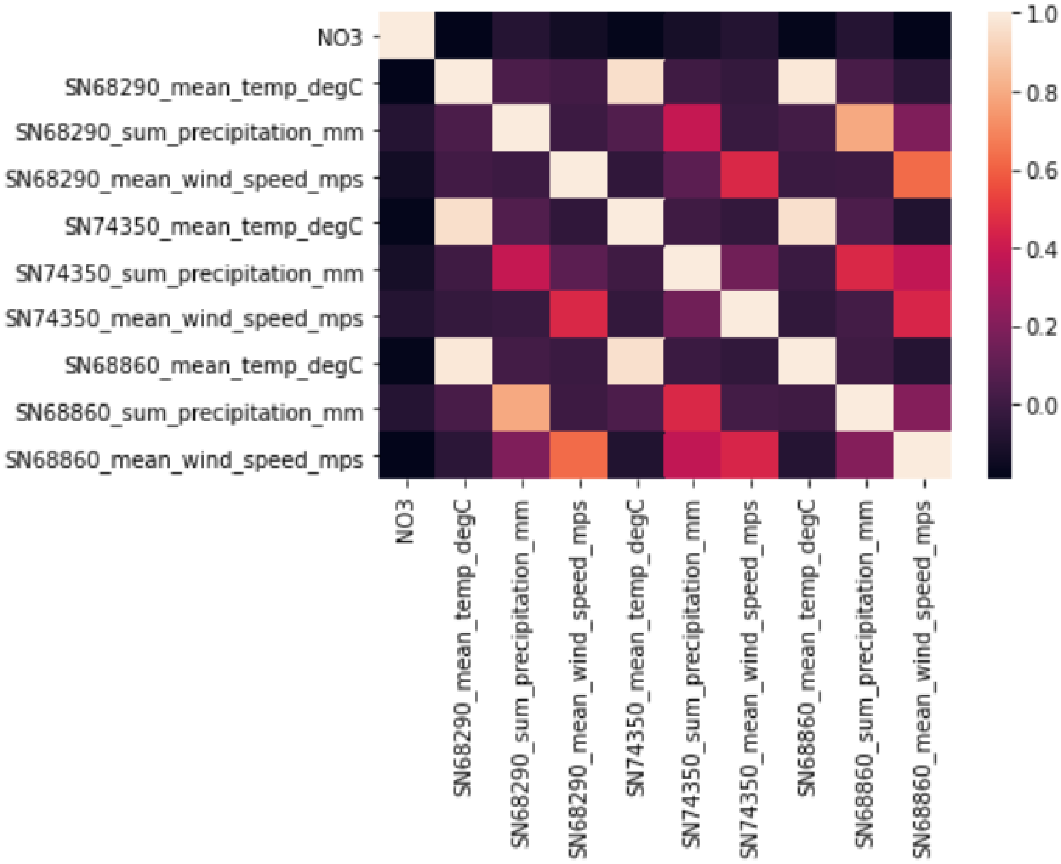

Figure 8: Heat map

Lastly, we wanted to look at how the electricity price in NO3 developed related to the weather factors. Since these factors are on very different ranges, we normalized the data to be able to compare. Since we saw in the heat map that the different stations were very similar, we only looked at one weather station to make the plot cleaner and easier to understand. In addition, to see how the weather factors change based on seasons we limited the outlook to a single year, 2017.

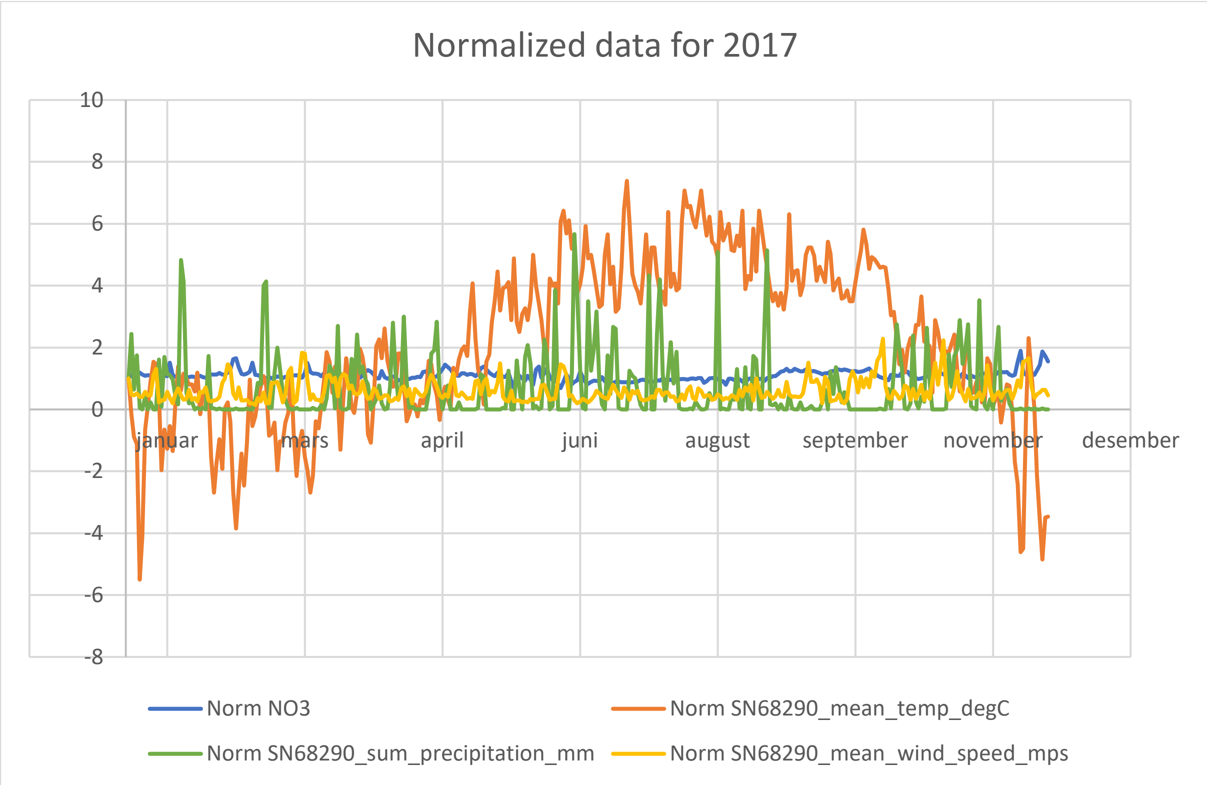

Figure 9: Normalized data plot

From this plot we can see that 2017 was a year with moderately even electricity prices. There were no great fluctuations like the ones we have seen in the past year. We see that changes in precipitation is relatively even throughout the year as precipitation both includes rain and snow. It is not a large impact, but we can see that the electricity prices seem to go slightly up when the mean temperature decreases. It is important to note that although we have normalized the data, the variables still move on very different scales, which makes the resaults difficult to compare. Mean temperature is the only variable which moves below zero, and it also has a much more significant change over time then the others.

4.3 The performance of the benchmark model.

The following results are from the benchmark. The benchmark is supposed to let us know if the predictions made by our other models are reasonable.

Table 1: Benchmark

	training	test
R^2	0.947	0.771
MSE	1068.436	14724.638
RMSE	32.687	121.345
BIC	10182.281	-
AIC	10176.995	-

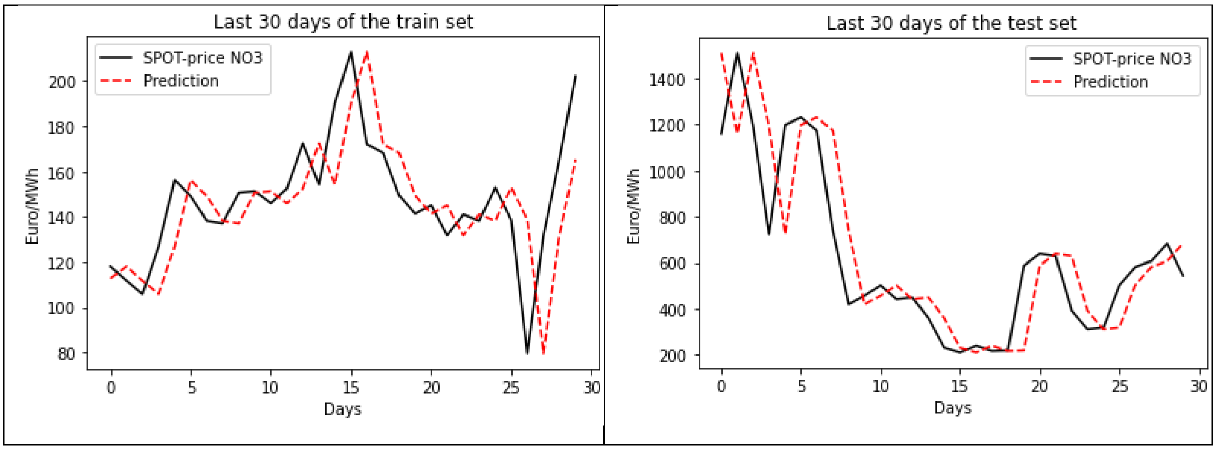

Figure 10: Benchmark

We can see from the results on the benchmark, that our training and test sets are very different. The mean square error on the test set is 14724, while the mean square error on the training set is only 1068. And one should also remember that the test set is only 1 year of data, while the training set is 4 years. This difference will most likely make it difficult to make good predictions on the test set when only fitting on the training set. It is probably not realistic to expect any model to predict something that is really different from what it has been shown previously. We might still expect to get somewhat better predictions than the benchmark on the test set.

4.4 The performance of the OLS model

The following results are from an OLS model. The thought behind this model was mainly to have something to compare the ANN models to.

Table 2: OLS model

	training	test
R^2	0.954	0.723
MSE	924.479	17787.045
RMSE	30.405	133.368
BIC	10059.451	-
AIC	9980.189	-

Table 3: Coefficient and intercept for the OLS model

Independent variables:	The coefficients:
SE2	-0.007
NO1	-0.155
NO5	0.266
NO3	0.703
NO4	0.120
SN68290_mean_temp_degC	-0.268
SN68290_sum_precipitation_mm	0.368
SN74350_mean_temp_degC	1.092
SN74350_sum_precipitation_mm	-0.491
SN68860_mean_temp_degC	-1.197
SN68860_sum_precipitation_mm	-0.026
NO3_lag1	-0.155
NO3_lag2	0.209

Intercept:	13.022

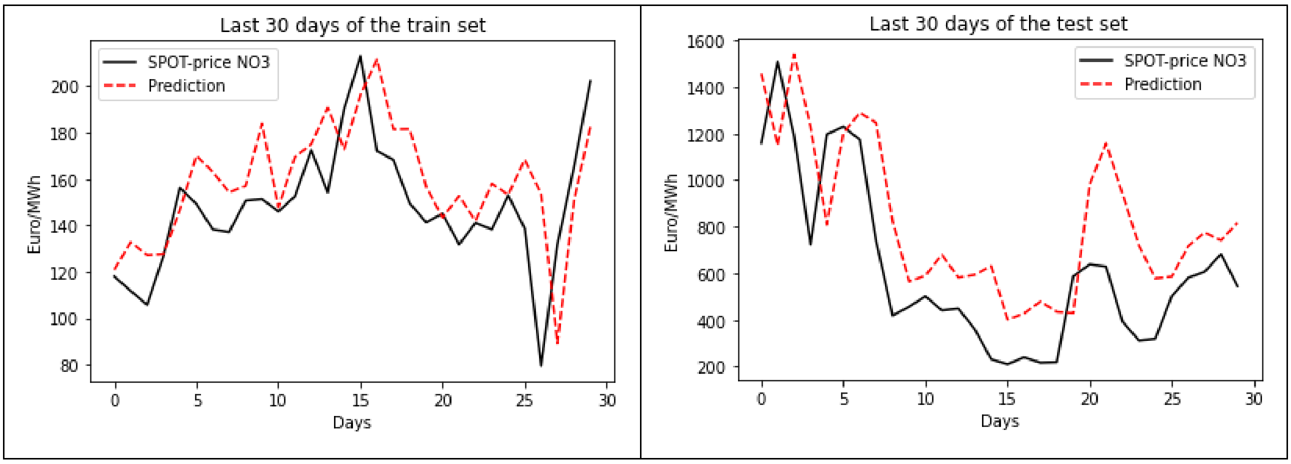

Figure 11: OLS model

The linear regression gave us not much better results than what only using the benchmark on the training data did. It performed also quite a lot worse than the benchmark on the test data. This may indicate that our dependent variable is not very linearly dependent on our independent variable. We can see from the plots that the predictions are heavily relying on the previous day price, especially on the volatile test-set.

4.6 The performance of the SRNN model

The following results are from an SRNN model. The specific results came from running the training model multiple times and choosing one of the results that fitted correctly.

Table 4: SRNN model results

	training	test
R^2	0.954	0.738
MSE	925.141	16842.456
RMSE	30.416	129.778
BIC	10067.778	-
AIC	9983.232	-

Total parameters: 16

Table 5: SRNN model structure

	Node amount:	Node type	Activation:
Input	1	dense	RELU
Output	1	dense	linear

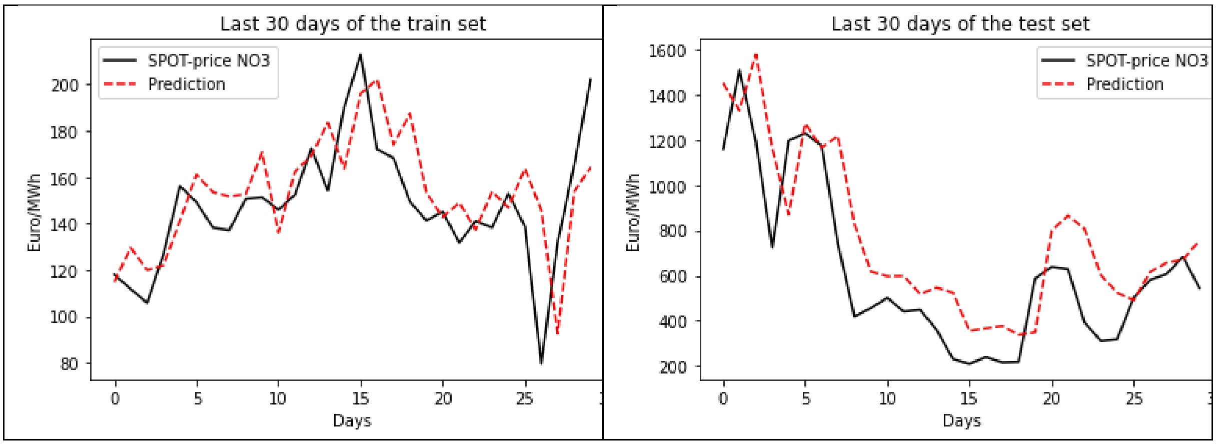

Figure 12: SRNN model

We ended up with a small model and we saw that the model got better fitted to the training data when we increased the size of the network with more nodes, as well and adding hidden layers. On the other side, the AIC, BIC and the performance on the test-set got worse. The model ended up performing quite similarly as the linear model. That might indicate that this model is probably not deep or complex enough to detect meaningful patterns in the data other than probably the linear relationships.

4.7 The performance of the LSTM model

The following results are from an LSTM model. The specific results came from training the model multiple times and choosing one of the results that fitted correctly.

Table 6: LSTM model results

	training	test
R^2	0.937	0.276
MSE	1271.037	46718.282
RMSE	35.652	216.144
BIC	11251.711	-
AIC	10644.036	-

Total parameters: 115

Table 7: LSTM model structure

	Node amount:	Node type	Activation:
Input	2	LSTM	RELU
Output	1	dense	linear

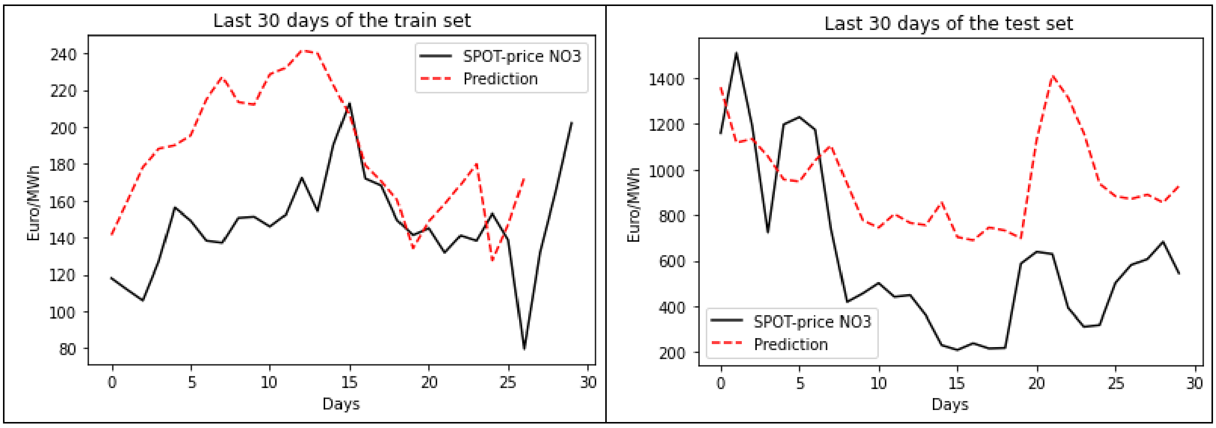

Figure 13: LSTM model

We had a lot of trouble getting the LSTM to make any reasonable predictions. We know from the literature that it is supposed to be able to improve timeseries predictions with the ability to forget. There is a chance that we could have set up the model differently, but we think that the advantage of using LSTM probably goes away when the data amount is too low.

4.8 The performance of a deliberately overfit model

The following results stems from a model created to show why we cannot just create overly big and complex models.

Table 8: Overfit model results

	Training	test
R^2	0.994	-0.553
MSE	130.298	99694.493
RMSE	11.415	315.744
BIC	165168.341	-
AIC	50497.332	-

Table 9: Overfit model structure

	Node amount	Node type	Activation
Input	100	dense	RELU
Hidden	100	dense	RELU
Hidden	100	dense	RELU
Output	1	dense	linear

Total parameters: 21701

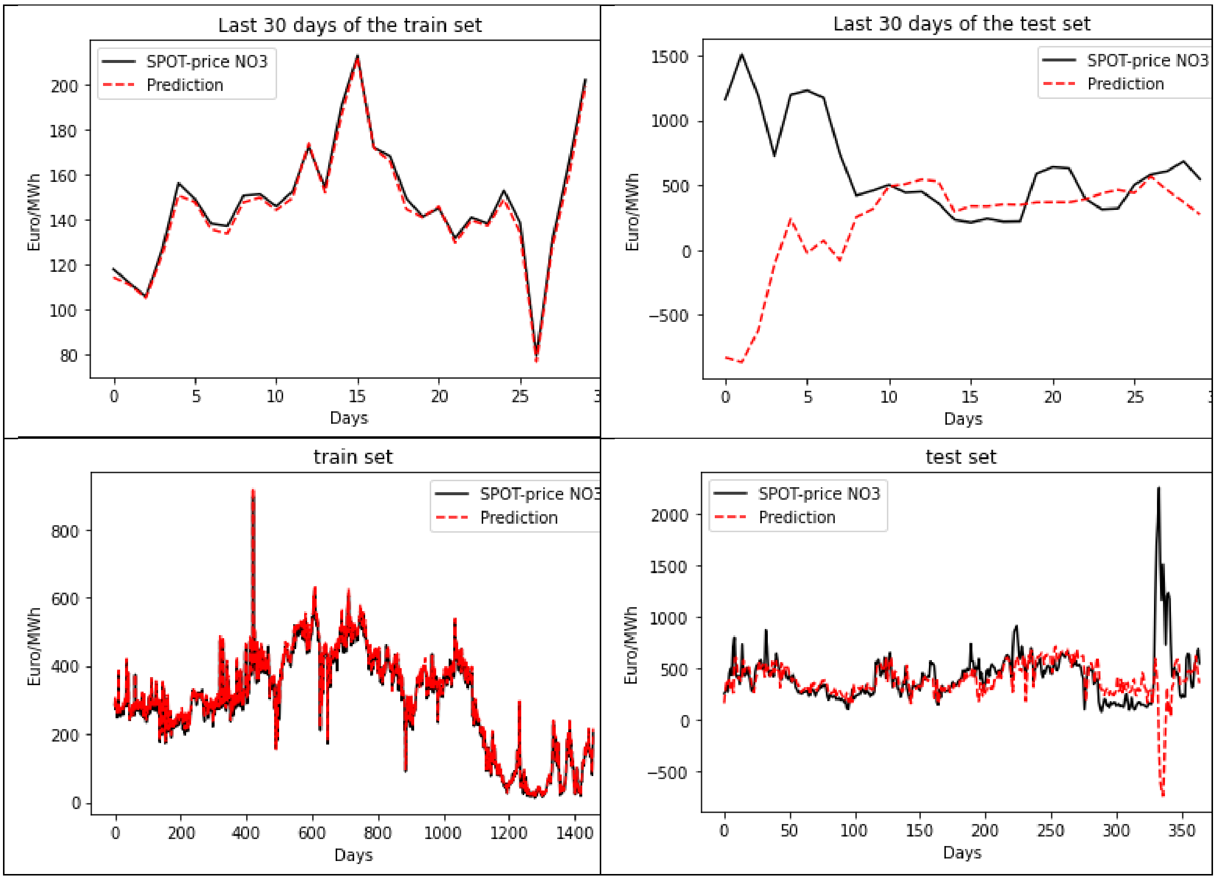

Figure 14: Overfit model

To show an example of an overfit model, we decided to create a large network that could easily fit the training set. The training-set is here memorized. However, this memorization does not show any obvious benefits on the test set.

5-Conclusion

5.1 Evaluating our models

The main task of this assignment has been to build statistical models that can predict up-coming electricity prices in Trøndelag, Norway. We can start by saying that all our models performed badly on the test-set we gave them, and none of them got better results than the benchmark did on the test-set. This indicates that we did not manage to predict the prices to a valid nor reliable point. We worked with a few different types of models that we thought could give accurate results. In addition to making a benchmark model to compare the ANN models, we built a model using linear regression. However, we found the linear regression model performed worse than the benchmark, which might suggest that there is not a strong linear relationship between the dependent and independent variables. We then made a SRNN model that performed about the same as the OLS model. We believe that this might indicate that the model was not big enough to discover the potential complex patterns in the data, but making the model bigger gave worse results indicated by AIC, BIC and results on test data. Bigger models usually need more data, and we believe that too little data prevented us from having bigger networks, which in turn prevented us from finding complex patterns.

We then moved on to building a LSTM model, but found that it did worse in both the training and the test set. This might also be because our dataset was too small for this type of model. Finally, we made an overfitted model, and as expected it did not improve performance on the test set from our previous models. In conclusion none of the models we made were fit to predict electricity prices, but we would like to present some reflections on how it might be improved.

5.2 How the models can be improved

First and foremost, a given way to improve the model would be to add more data. Some potential ways to do that could be through having a larger training set, either by using hourly data instead of the daily data, or by increasing the time horizon. The model could possibly be improved by implementing more weather stations. Or possibly using an average from multiple stations instead of looking at individual ones.

It could also be interesting to apply more data through different features. We believe it could be interesting to include power consumption prognosis, power production prognosis, and cross border energy trade/transfer. Because the market for electricity, like other markets, is affected by supply and demand. It would also probably improve the model to look further into the supply of energy, prices on energy producing materials (such as gas and oil), prices on energy dependent raw materials (such as aluminium and artificial fertilizer). We also would have liked to include seasonality in the model as the vastly different seasons in Norway has a significant impact on the consumption. This could for example be done by adding a number that changes between 0 and 1 from summer to winter. Another dummy variable that might be interesting to add is weekend or weekday. Finally, it could be interesting to add a variable dependent on the variation in the previous week power price since this might affect the price, since the electricity producers might demand a higher price if the spot price is unstable or difficult to predict.

Other than more data, we could have implemented a validation set, to improve validity. Experimented more with parameters and different number of nodes and layers. Gotten average results and represented the data in boxplots so that we could have shown reliability. Implemented feature selection to make sure we used good data.

There could always be different opinions or priorities made, regarding which variables to choose to produce the best results for showcasing the data used for this paper. Given the complexity of the topic of predicting electricity prices, we had to make decisions based on the variables we considered to be significant, and during the process of this assignment we might have realized, as mentioned in the previous section, which changes we might have done if we were to start our research all over. Our conclusion, given the results we have discovered and visualised throughout this paper, is that larger data sets and more information obtained would be necessary to improve our model and end-results.

List of references:

Aibodun, O. (2018) *State of the art in artificial neural network application: A survey*. Available at: https://www.sciencedirect.com/science/article/pii/S2405844018332067 (Accessed:30.03.2022)

Analyttica Datalab (2019) *What is Bayesian Information Criterion (BIC)?*. Available at: https://medium.com/@analyttica/what-is-bayesian-information-criterion-bic-b3396a894be6 (Accessed: 25.04.2022)

Brownlee, J. (2020) *How to Diagnose Overfitting and Underfitting of LSTM Models.* Available at: https://machinelearningmastery.com/diagnose-overfitting-underfitting-lstm-models/ (Accessed: 22.03.2022)

Brownlee, J. (2020) *How to Tune LSTM Hyperparameters with Keras for Time Series Forecasting.* Available at: https://machinelearningmastery.com/tune-lstm-hyperparameters-keras-time-series-forecasting/ (Accessed: 20.03.2022)

Brownlee, J. (2020) *How to Use the TimeseriesGenerator for Time Series Forecasting in Keras*. Available at: https://machinelearningmastery.com/how-to-use-the-timeseriesgenerator-for-time-series-forecasting-in-keras/ (Accessed: 30.03.2022)

Brownlee, J. (2020) *Time Series Prediction with LSTM Recurrent Neural Networks in Python with Keras*. Available at: https://machinelearningmastery.com/how-to-use-the-timeseriesgenerator-for-time-series-forecasting-in-keras/ (Accessed: 17.03.2022)

IBM (2020) *Neural networks*. Available at: https://www.ibm.com/cloud/learn/neural-networks (Accessed: 24.01.2022)

IBM (2020) *Machine learning*. Available at: https://www.ibm.com/cloud/learn/machine-learning (Accessed: 03.02.2022)

Jacobsen, D. I. (2015). Hvordan gjennomføre undersøkelser? Innføring i samfunnsvitenskapelig metode. Oslo: Cappelen Damm

Kumar, V. (2019) *Why are neural networks not the answer to everything?* Available at: https://www.analyticsinsight.net/neural-networks-not-answer-everything/ (Accessed: 25.03.2022)

Lier,T., Rafelsen, M. and Resvoll, A. (2021) *Ny strømprisrekord i Sør-Norge mandag* Available at: https://e24.no/privatoekonomi/i/Qy3jMx/ny-stroemprisrekord-i-soer-norge-mandag (Accessed: 26.04.2022)

Meteorologisk institutt (no date) *What is Frost?* Available at: https://frost.met.no/index.html (Accessed: 21.04.2022)

Nord Pool (no date) *The power market.* Available at: https://www.nordpoolgroup.com/the-power-market/ (Accessed: 24.01.2022)

NTE (no date) *Hva påvirker strømprisen?* Available at: https://nte.no/blogg/hva-pavirker-stromprisen/ (Accessed: 24.01.2022)

Oppermann, A. (2021) *Underfitting and overfitting in deep learning.* Available at: https://medium.com/mlearning-ai/underfitting-and-overfitting-in-deep-learning-687b1b7eb738 (Accessed: 21.04.2022)

Regjeringen.no (2021*) Regjeringen med milliardpakke i strømstøtte.* Available at: https://www.regjeringen.no/no/aktuelt/regjeringen-med-milliardpakke-i-stromstotte/id2891839/ (Accessed: 25.03.2022)

Sharma, A. (2017) *Understanding activation functions in Neural Networks.* Available at: https://medium.com/the-theory-of-everything/understanding-activation-functions-in-neural-networks-9491262884e0 (Accessed: 01.04.2022)

Sharma, S. (2017) *Activation functions in Neural Networks.* Available at: https://towardsdatascience.com/activation-functions-neural-networks-1cbd9f8d91d6 (Accessed:23.04.2022)

Shipra_saxena. (2021) *Introduction to long short term* Available at: https://www.analyticsvidhya.com/blog/2021/03/introduction-to-long-short-term-memory-lstm/ (Accessed: 28.03.2022)

Weron R. (2006) Modelling and forecasting electricity loads and prices: a statistical approach. West Sussex: Wiley Finance.

Weron, R. (2014) *Electricity price forecasting: A review of the state-of-the-art with a look into the future*. Available at: https://www.sciencedirect.com/science/article/pii/S0169207014001083 (Accessed: 21.03.2022)

Zajic, A. (2019) *Introduction to AIC*. Available at: https://towardsdatascience.com/introduction-to-aic-akaike-information-criterion-9c9ba1c96ced (Accessed: 28.03.2002)