Bacheloroppgave

NTNU
Norges teknisk-naturvitenskapelige universitet
Fakultet for økonomi
NTNU Handelshøyskolen

Huseby, Eivind Rytter
Grevstad, Oskar

Can LSTM-networks predict electricity prices in Trondheim?

Bacheloroppgave i Økonomi og Administrasjon
Veileder: Blomsø-Becker, Mike Denis
April 2022

NTNU
Kunnskap for en bedre verden

Huseby, Eivind Rytter
Grevstad, Oskar

Can LSTM-networks predict electricity prices in Trondheim?

Bacheloroppgave i Økonomi og Administrasjon
Veileder: Blomsø-Becker, Mike Denis
April 2022

Norges teknisk-naturvitenskapelige universitet
Fakultet for økonomi
NTNU Handelshøyskolen

NTNU
Kunnskap for en bedre verden

AF3035

Bacheloroppgave i Business Analytics

Vår 2022

Can LSTM-networks predict electricity prices in Trondheim?

Eivind Rytter Huseby & Oskar Grevstad

Supervisor: Mike Denis Blomsø-Becker

Norwegian University of Science and Technology

Fakultet for Økonomi: NTNU Handelshøyskolen

Preface

We are two students in our sixth semester of our degree at NTNU Handelshøyskolen in Trondheim. As this is our final semester, this thesis is the final assignment of our bachelors degree in economics and administration with specialization in business analytics.

As our specialization is business analytics, we wanted to utilize this opportunity to extend our understanding of analytical modelling. In particular we found great interest in the use of LSTM-networks for prediction purposes. For this type of modelling, we found the complex world of electricity pricing to be both instructive and relevant.

We would like to thank our supervisor Mike Denis Blomsø-Becker for inspiring us to pursue the study of business analytics, as well as providing great support and feedback during our work on this thesis.

The contents of thesis is at the expense of the authors.

Eivind Rytter Huseby

Oskar Grevstad

Abstract

In this bachelors thesis in economics and administration we are investigating whether long-short-term-memory neural networks can be used to predict the price of electricity in Trondheim one day in advance. To make these predictions we are using the price of natural gas, price of coal, precipitation measured in Bykle and Glomfjord, temperature measured at Værnes, the exchange rate between NOK and USD, and the exchange rate between NOK and EUR.

To predict tomorrow power price, our research includes predictions based on the prior 50 and 30 days, as well as predictions based only the day before. We are doing this with all the variables mentioned over as independent variables, as well as separate predictions only using the historical electricity price to predict tomorrows price of power. Our model is learning by minimizing the the mean squared error of the predictions in the learning phase, and then validating the learning process with a set of validation data.

Through our research and testing we have found signs indicating that our models predictions improve when we reduce the amount of variables and days included in the prediction. It is not likely for an LSTM-network to be able to provide better predictions with less variables and time steps as the input, unless the variables chosen only contain noise. In that case, the addition of more variables with more noise would only prevent the models ability to produce reliable predictions.

In conclusion we can conclude that our model is not capable of producing reliable predictions based on the chosen set of independent variables. Through our testing we have also seen tendencies indicating that the power price could be subject of random-walk theory.

Sammendrag

I denne bacheloroppgaven i økonomi og administrasjon skal vi undersøke hvor vidt long-short-term-memory neurale nettverk kan brukes til å forutsi moregndagens strømpris i Trondheim. Dette basert på prisen på naturgass, kullpriser, nedbør målt i Bykle og Glomfjord, temperaturen målt på Værnes, vekslingsraten mellom Norske kroner og Amerikanske dollar og vekslingsraten mellom Norske kroner og Euro

Vi forsøker å prediktere morgendagens strømpris i Trondheim ved å se på de foregående 30 og 50 dagene, samt en dag i forveien. Vi forsøker å gjøre dette ved å benytte alle variablene nevnt over som uavhengige variabler, og ved å kun bruke historisk strømpris som uavhengig variabel til å forutsi morgendagens strømpris. Modellen lærer ved å minimere det gjennomsnittlige kvadrerte avviket i treningsfasen, for så å validere læringen mot et sett med valideringsdata. Dette med mindre variablene kun inneholder støy og ikke utgjør noen nyttig funksjon for modellens evne til å produsere gode forutsigelser.

Gjennom undersøkelser og testing har vi funnet indikasjoner på at modellens evne til å predikere blir bedre ved å redusere antall variabler og dager som benyttes i beregningene. Det er for øvrig lite sansynlig at et LSTM-nettverks evne til å forutsi verdier bedres ved reduksjon av variabler og antall dager som inkluderes som input til modellen.

Basert på disse funnene kan vi konludere med modellen vår ikke er i stand til å produsere pålitelige resultater basert på de uavhengige variablene valgt i denne oppgaven. Vi ser også gjennom disse undersøkelsene tegn til at strømprisen kan betegnes som en random-walk tidsserie.

Contents

List of Figures

List of Tables

List of Abbreviations

RNN	Recurrent Neural Network
LSTM	Long Short-Term Memory
BP	Backpropagation
tanh	Hyperbolic tangent function
Adam	Adaptive Moment Estimation
MSE	Mean Squared Error
MWh	Mega Watt hour
USD	US dollars
NOK	Norwegian kroner

1 Introduction

The Norwegian electricity market has for a long time been characterized by its low prices and consistent nature. However, starting in the latter part of 2020, Norwegian electricity prices, specially in the southern Norway have skyrocketed(Statistics Norway, 2022) and become a big talking point for all consumer segments. With some geographical areas seeing prices sextuple in a years time, this impacts both consumers and industry as electricity pricing is a big input factor in multiple production processes.(Eika & Jørgensen, 2003). The prices increased so much that the Norwegian government implemented a electricity support package at the end of 2021 (Ministry of Finance and Office of the Prime Minister, 2021).

By being able to reliably predict the fluctuations in the electricity prices or spot price-trend tendencies before they happen, businesses as well as households could cut their bills by running heavy power consuming processes at times when there is an energy surplus and prices are low. In addition to this, knowing when to use less electricity could also give valuable insight for cutting costs. However, it is worth mentioning that a tool that with great certainty could predict prices being commercially available, would most likely in it self affect the prices. Thus making it an even harder task to accomplish.

As the world is more complex than what can fit into an equation, it is common to simplify such problems with models. To be able to accurately make predictions or forecasts on the electricity price, one would need a model that evaluates relevant variables and is able to learn from historic data to predict the market price some time in advance. For this, a model utilizing neural network technology could be the right place to start. In the case of electricity price, this would be the power price, traded at Nord Pool for the Nordic regions.

With hydropower plants being Norway's largest source of energy production, as well as one of the only somewhat store-able renewable energy sources, this could be valuable. Being able to know when to utilize filled magazines, and when to hold on to it could benefit European consumers by providing cheaper, cleaner energy when the market otherwise is tight and dependent on non-renewable sources. In the case of business analytics, forecasts of electricity pricing would be beneficial for budgeting, break even analysis, cost analysis and calculations regarding the ideal price of a product or a service. This in particular for power heavy business-models such as companies utilizing heavy industrial machinery in production and processing. All this could help businesses gain better control of their profit margins.

With this in mind, our thesis revolves around the question: *Whether LSTM-networks are useful to predict electricity prices in Trondheim, based on the chosen set of variables.*
In reality, all the data needed for a precise prediction of this scale would be a great feat to achieve, and would require computing power beyond what we have accessible. Instead we have chosen variables that according to NTE(Nord Trøndelag Elektrisitetsverk AS, 2022) has the biggest influence on the electricity prices. The variables available and in focus of this thesis are the price of coal, the natural gas price, temperature at a specific point in Trondheim, precipitation at two different stations in Norway, and the exchange rate from euro to NOK and from US dollars (USD) to NOK.

Due to the current situation regarding Ukraine and Russia, which impacts the electricity market both directly and indirectly, the data-set used in this paper is set to end on the 28th of February 2022. This to prevent uncertainty and noise to clutter up further analysis, and to maintain the focus on the predictability of the model.

2 Theory

2.1 Electricity Market and Pricing

Because of its inability to be efficiently stored(Ministry of Petroleum and Energy, 2021), the market for electricity differentiates itself from other markets. Electricity has to be distributed and consumed directly after production. This process both complicates and simplifies the market pricing by somewhat shifting the focus solely to the current supply and demand status.

Norway is part of a shared Nordic Market with Denmark, Sweden and Finland, which is connected to the European market through cable with The Netherlands, Baltics, Germany, Poland and Russia.

In short power is produced an mainly traded on the power exchange or directly to power suppliers and power heavy industry institutions. In the Nordics case this exchange is Nord Pool.(Ministry of Petroleum and Energy, 2021) As smaller scale businesses, households and other small to medium scale end users are unable to trade at the quantities provided on the power exchange, power suppliers trade on their behalf. In this market, end users are defined as those who purchase power for consumption. This system is illustrated in Fig 2.1.(Ministry of Petroleum and Energy, 2021)

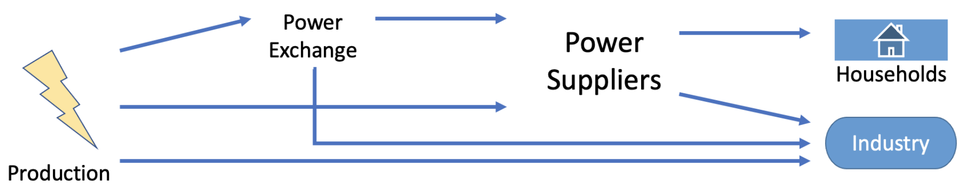

Figure 2.1: Illustration of the electricity market

2.1.1 Markets

Wholesale of power takes place in three different markets: day-ahead, intraday and balancing markets. Of these, the main part of electricity trade in Nord Pool is done by auction in the day-ahead market.(Ministry of Petroleum and Energy, 2021)In this market hourly bids are put for the next 24 hours as the name suggests. In Norway, the balancing market is run by Statnett which the transmission system operator(TSO).

Financial trade however, is done at the NASDAQ Stockholm exchange. This is where trading of financial instruments such as price securities, futures and derivatives take place.

2.1.2 Pricing

Breaking down the contributing factors for the power pricing can be a complex task, as there are many variables that have an impact on the end price. On the ground level, Nord Pool calculates and sets a system price as a reference point for the Nordic markets. Producers then place their bids for which quantities of power they are willing to produce at a specified price. This price is the somewhat most traditional part of the price formation, as it is mostly set by production cost at the power plant(Ministry of Petroleum and Energy, 2021). From here, end users and power suppliers place their bids for what quantities they are willing to purchase, and at what price. the market price is then determined at the equilibrium between supply and demand(Ministry of Petroleum and Energy, 2021). Equilibrium is defined as the point where the supplied quantities is equal to the demand as shown in Fig. 2.2.(Khan Academy, 2019)

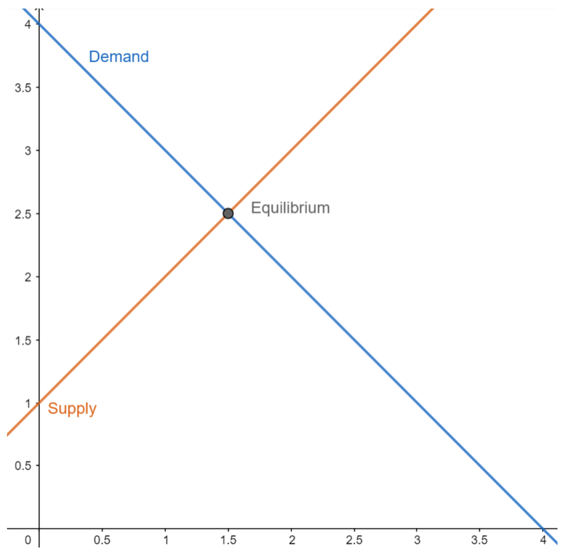

Figure 2.2: The market equilibrium is located where supply meets demand.

At the market equilibrium, it is the marginal cost that determines the market price. The marginal cost is the production cost of supplying one more unit, and as long as this is profitable, one should keep producing. The market equilibrium coincides with the point where the marginal cost equals the marginal willingness to pay in a perfect competition scenario.

By the price being determined in the market equilibrium, it ensures that power is sold and produced at the lowest possible cost to society, while also generating the maximum amount of social income. Being a necessity good, society benefits in total from the price being determined in this manner.

Because of this structure, the main foundations of the power price are the power plants costs and therefore the price of coal, natural gas and emission allowances(Ministry of Petroleum and Energy, 2021). In addition to this, the production of renewable energy in the network also has an influence on the pricing (Ministry of Petroleum and Energy, 2021).

Norway in particular is divided into five power areas with separate bidding prices set by Nord Pool(Skånlund & et al., 2013). This is mainly because of the congestion that arises when the power situation varies between the areas and there is insufficient grid capacity to transport electricity between the areas. As they are both geographically and demographically different from each other, factors like seasons, weather, magazine filling and many more contribute to deviations between each area. The balancing of transporting power from areas with a surplus to areas with a deficit is also cause for the different pricing across the bidding areas.(Ministry of Petroleum and Energy, 2021) In this paper, the area NO3 is selected as the price variable, as it contains Trondheim which is the area of interest. N01 is also used in analysis, but only as an independent variable, with the assumption that it is the area most affected by the European market. This in the hopes that it will capture intangible factors in the market, to help the learning process of the model

2.1.3 End user Price

The end user price consists of multiple factors other than just the energy price traded at Nord Pool. In Norway the additions of grid tariff, VAT, Enova levy, electricity tax and electricity certificate contribute to the total price paid for electricity by the end consumer.
In this paper, analysis is based on the power price in the day-ahead market, and not the end user price. For further reading on the subject of the Norwegian power market, we recommend the article "The Power Market" by the Norwegian Ministry of Petroleum and Energy(Ministry of Petroleum and Energy, 2021)

2.2 Neural Network

2.2.1 Recurrent Neural Network

A Recurrent Neural Network (RNN) is a neural network that utilizes recurrence in the hidden layer to feedback new data into the model (Hochreiter & Schmidhuber, 1997). By being able to consider both the current input as well as what the model has learned form previous inputs, RNN's differ from a traditional feed-forward neural network. Because of this RNN's are very useful when working with sequences.(Goodfellow, Bengio, & Courville, 2016) This is illustrated in Fig 2.3. (Donges, 2021) Because of this structure, a recurrent neural network has multiple inputs, both the current and the immediate past input. This proves advantageous when modelling text analysis, time series analysis and forecasting.(Donges, 2021) The model is learning by adjusting its weights employing a backpropagation(BP)-algorithm to identify the derivatives of the error found in the output. Normally a gradient descent algorithm is used in BP to find the local minimum to adjust the weight. This is done in every step of the models training. Even though recurrent neural networks can provide great results, it is haunted by two major problems: vanishing and exploding gradients(Li & et al., 2018).

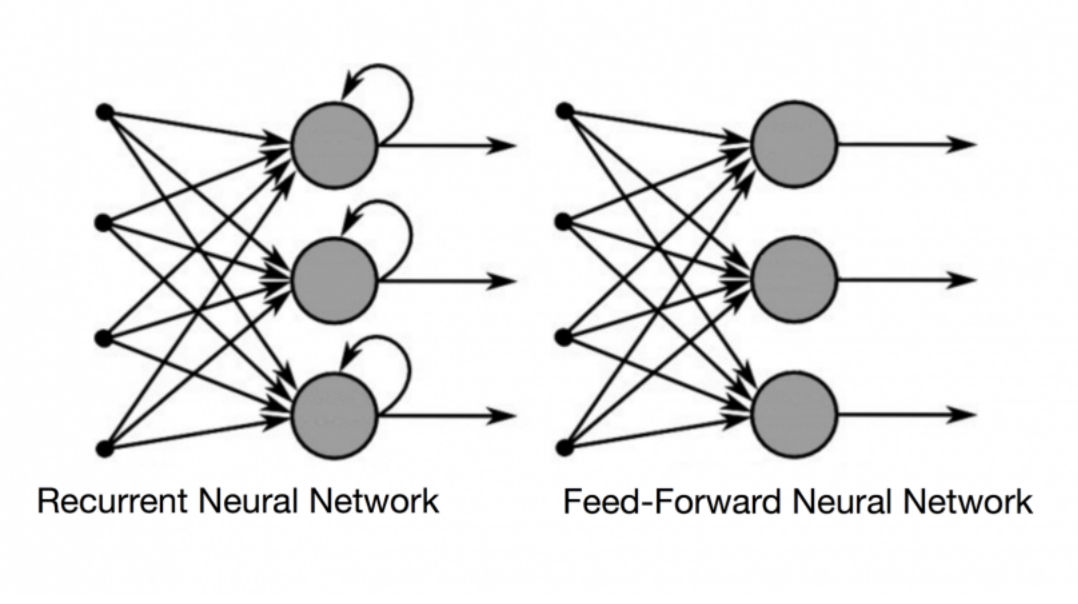

Figure 2.3: Illustration of the differences between a recurrent neural network, and a feed-forward neural network(Donges, 2021)

2.2.2 Long Short-Term Memory

Long Short-Term Memory (LSTM) is a RNN variant first presented in 1997 by Sepp Hochreiter and Jürgen Schmidhuber(Hochreiter & Schmidhuber, 1997). LSTM was created to solve the vanishing gradient problem of the RNN. (Nielsen, 2018) LSTM is a constitution of a chain of modules and the storage space of these modules are referred to as the cell state (Olah, 2015). The cell state is where the information flows through the network. The modifications to the network is done in three "gates"; the input gate, the forget gate and the output gate (Olah, 2015) These gates uses a Sigmoid function and/or a hyperbolic tangent function(tanh(x)) to decide what values to continue using and what values to change and the values of these new values.
As shown below, in section 2.2.2.1, the Sigmoid function will give outputs between 0 and 1. In the forget gate this is used to determine how much of the combination of the input and the previous output to carry on to the cell state. This is done by the Sigmoid funtion weighting each part of the inputs(Olah, 2015). In the input gate the Sigmoid function is used to determine which values need a update, and then the tanh is used to determine a vector for the updated values (Olah, 2015). The values created here are then added to the cell state, so that the new cell state will be the previous cell state updated by the Sigmoid function summed with the post input gate vector updated with the Sigmoid function(Olah, 2015).
The final part of the module is the output gate where the output to carry on to the next module is decided. The output gate will use the Sigmoid function to decide what to carry on to the cell state. Then this is multiplied with the tanh to make sure that the output will be in the desired range (-1 $\rightarrow$ 1)(Olah, 2015), this is proven in section 2.2.2.2.

2.2.2.1 Sigmoid function

The Sigmoid function is given as follows(Moolayil, 2019):

$$S(x) = \frac{e^x}{e^x + 1}$$

Since:

$$\lim_{x\to\infty} S(x) \rightarrow S(x) = \frac{\infty}{\infty + 1} \simeq \frac{\infty}{\infty} = 1$$

$$\lim_{x \to -\infty} S(x) \to S(x) = \frac{0}{0+1} = \frac{0}{1} = 0$$

As a result of this, the Sigmoid function will always return a value between 0 and 1.

2.2.2.2 Hyperbolic tangent function

The hyperbolic tangent function (tanh) is given as

$$tanh(x) = \frac{sinh(x)}{cosh(x)} = \frac{e^x - e^{-x}}{e^x + e^{-x}}$$

Where sinh and cosh are the hyperbolic sine- and hyperbolic cosine-function. Proof:

$$\lim_{x \to \infty} tanh(x) \to tanh(x) = \frac{\infty - 0}{\infty + 0} = \frac{\infty}{\infty} = 1$$

$$\lim_{x \to -\infty} tanh(x) \to tanh(x) = \frac{0 - \infty}{0 + \infty} = \frac{-\infty}{\infty} = -1$$

This proves that tanh(x) always will produce an output value between -1 and 1.

2.3 Keras

Keras is a library built in in python, running on top of TensorFlow. This library have different built in function for deep learning(Chollet et al., 2015).
The advantage we have using Keras is that the models we want to create are easily created by downloading the Keras package and using its built in functions.

2.4 Adaptive Moment Estimation

To implement the gradient descent method we use the optimizer known as Adaptive Moment Estimation (Adam). Adam will work by minimizing the loss of the model(Singarimbun, Nababan, & Sitompul, 2019). The model will learn by adapting first-order-gradients, which means that it requires little memory(Kingma & Ba, 2017). Adam is often used for LSTM because it uses the advantages of the AdaGrad and the RMSProp(Kingma & Ba, 2017). Which are two different ways to minimize the loss.

2.5 Mean Squared Error

The Mean Squared Error (MSE) is used to measure the loss of the model. The aim is to minimize the MSE. The advantage of the MSE is that it shows us how far from the real value the models' predicted value is.

$$MSE = \frac{1}{N}\sum_{i=1}^{N}(x_i - y_i)^2$$

In the equation above, x is the observed values and y is the predicted values.
The function will hence penalize the biggest error much harder than the smaller errors.

2.6 Random Walk

Random Walk is defined as a process for determining the likely value of a point subject to random motions, based on the probability of it, at each point in time, moving some distance and direction(The Editors of Encyclopaedia , 2008b). In the stock market, as well as the power market, this means that prices are subject to a Markov Process(The Editors of Encyclopaedia , 2008a), meaning that future prices are independent of historic prices and movements. By means of this theory, one should not be able to predict the price of some stock or power based on historic data in the short term. If the accuracy of a predictive LSTM model is getting better, the fewer time steps it uses, it could be a sign of the time series being a random walk.(Flovik, 2018)

3 Methodology

3.1 Data Gathering

Data used in further analysis is gathered from the Thomson Reuters Eikon[1], and The Norwegian Meteorological Institute[2]. The data used is from from 27.11.2014 to 28.02.2022, and the following variables has been taken into account; the average daily electricity price in Trondheim measured in NOK/MWH, the average daily electricity price in Oslo measured in NOK/MWh, the daily closing price for coal measured in dollars per metric ton, the daily closing price for natural gas measured in Euro per Mega Watt hour , the daily precipitation from the station SN40420 - Bykle, the daily rainfall from the station SN80705 - Glomfjord and the temperature from the station SN69100 - Værnes.

The reason for using the Station at Bykle and the station at Glomfjord are that they are the closest weather stations to the two biggest water deposits producing power in Norway. The station in Bykle is relatively close to Blåsjø and the station at Glomfjord is relatively close to Storglomvatnet(Rosvold & Halleraker, 2021) We implemented the daily exchange rate of US dollars (USD) to Norwegian kroner(NOK) and for Euros to NOK downloaded from Norges Bank (Norges Bank, 2022).

The temperature is used to account for the effects of supply and demand. During days with low temperature the demand will be higher since more people are using electricity for e.g. heating, while on warmer days less heating is needed. We were unable to find reliable data regarding industrial power usage, valid for this model. The data used in the model as well as the code file is found in the appendix A.

3.2 Creating the data set

The data set is imported to and modified in Excel where we have changed the data series into tables and then combine them using the combine tables feature in excel. We then chose to combine them by the dates by creating Power Queries for each table, and merging these queries using excels built in merge function. By doing this we were able to merge the different tables to ensure that they all include the same dates. Removing dates where the weather station was not

[1]Thomson Reuters Eikon: `https://eikon.thomsonreuters.com/index.html`

[2]The Norwegian Meteorological Institute: `https://seklima.met.no/observations/`

operating from the Nord Pool and resource price data, as well as weekends and other holidays from the weather station data.

The implementation of the exchange rates are done using a built in merge function in Pandas, this allows us to implement based on a chosen parameter, in our case on the date. The advantage is that Pandas merge will remove the columns in which the date does not correspond with the different data sets. By creating a correlation matrix of this data set as shown in Fig.3.1, we can clearly see some indications of correlations between the independent variables and the electricity price in Trondheim. Mainly that the Coal Price, and the EEX EGIX Index(the price of natural gas) correlates in some regard with the power prices in Trondheim. Naturally, the electricity average in Oslo and Trondheim correlates rather well.

[11]:

	Nedbør (døgn)	Nedbør (døgn)2	Coal ICE API2 CIF ARA Nr Mth $/MT - SETT. PRICE	Electricity.Avg.Trondheim	Nordpool-Electricity Avg Oslo	Middeltemperatur (døgn)	OBS_VALUE_x	OBS_VALUE_y	EEX EGIX THE Index Eur/MWh
Nedbør (døgn)	1.000000	-0.079314	0.074092	-0.041656	0.020637	0.003605	0.034651	0.006170	0.052567
Nedbør (døgn)2	-0.079314	1.000000	0.055183	-0.096572	-0.020021	0.092228	-0.029304	-0.005417	0.075171
Coal ICE API2 CIF ARA Nr Mth $/MT - SETT. PRICE	0.074092	0.055183	1.000000	0.388421	0.816723	0.045324	-0.171941	0.009549	0.906691
Electricity.Avg.Trondheim	-0.041656	-0.096572	0.388421	1.000000	0.636262	-0.150519	-0.278073	-0.125837	0.240508
Nordpool-Electricity Avg Oslo	0.020637	-0.020021	0.816723	0.636262	1.000000	-0.161223	-0.118148	0.059300	0.808565
Middeltemperatur (døgn)	0.003605	0.092228	0.045324	-0.150519	-0.161223	1.000000	0.024371	0.037553	-0.004438
OBS_VALUE_x	0.034651	-0.029304	-0.171941	-0.278073	-0.118148	0.024371	1.000000	0.809343	-0.150408
OBS_VALUE_y	0.006170	-0.005417	0.009549	-0.125837	0.059300	0.037553	0.809343	1.000000	-0.032897
EEX EGIX THE Index Eur/MWh	0.052567	0.075171	0.906691	0.240508	0.808565	-0.004438	-0.150408	-0.032897	1.000000

Figure 3.1: Correlation Matrix

3.3 Source of error

In the creation of the data set many days have been removed from the weather stations in Bykle and Glomfjord, which means that we are removing potential days with weather that could increase the water reserves or days that decrease the water reserves. The precipitation is not encoded according to temperature, giving some uncertianty as snow, rain, drizzle and hail all have different effects on the magazine filling. The weather data from Værnes is also incomplete in the matter of available dates and this results in potential huge temperature differences. The temperature is measured as a daily middle temperature and will not take into account that in many cases the temperature is lower at the earlier and later stages of the day as well as higher temperatures usually occur during the middle of the day. Our model is not taking to account the consumption and in an economic market where supply and demand often affect the price this might be a source of error. Another source of error worth mentioning, is that there are as many factors contributing to the pricing of power as there are people in the region. By sampling only a small set of variables, as well as a defined time-span for the data sets you always run the risk of missing out on factors that could shift the analysis one way or the other.

3.4 Creating the model

The model is created in Jupyter Lab using a Python body. The packages used are Pandas, Numpy, Sklearn, Matplotlib and Keras.

In the model the 70% first variables rounded to the nearest whole number are used as the training set. The 20% next variables rounded to the nearest whole number are used as a validation set and the remaining variables rounded to the nearest whole number are used as the test set. The model is scaled using the MinMaxScaler changing the values of the data set to be, in our study between 0 and 1.

The model we created is a stacked LSTM with two hidden LSTM layers, two dropout layers, using Adam as the optimizer and MSE to calculate the loss at each epoch. To avoid that the model runs for ever without much more learning and to avoid overfitting, early stopping is implemented with a patience of five.

3.5 Predicting the price using the price

In this section we are only using the past average daily price in Trondheim to predict what the future daily average electricity price will be in Trondheim, giving us a reference to see whether or not our model can outperform the price at predicting itself. In this model we first need to reshape our data to be three dimensional since LSTM operates with three dimensional data as the input. The dates are deleted from the data set, and exchanged with row-numbers. This section of the study was done in three cycles starting with using the previous 30 inputs, in section 4.1, 50 inputs, in section 4.2 and then only the previous input, in section 4.3. The reasoning for changing this will be discussed in discussion. The batch size used to create this model is 32.

3.6 Adding the variables

However, the electricity price is affected by than just the historical electricity price. To build on our thesis, we created new models were we implemented the price of coal, natural gas, the temperature data from Værnes, the precipitation in Bykle and Glomfjord, the electricity price in Oslo and the Exchange rates from Euro to NOk and USD to NOK. This to see how these variables affected the price and how this compared to the already created model. We once again delete the date variables from the data set before running the model to not let the dates affect the prediction. The hope is that this will improve the model when adding more historical variables. The model is looking a lot like in the previous section, but there is no need to reshape the data in this one since the input is already in three dimensions. The graphs produced within this section is created using the previous 30 inputs, in section 4.4, 50 inputs, in section 4.5 and then only the previous input, seen in section 4.6. These three sizes of time steps are used to keep the analysis coherent to the reference point constructed in the former section.

4 Results

4.1 Predicting average electricity price with 30 time steps

In this section the prices are predicted by the model having a look at the 30 previous inputs to try and predict the next step. This model will be referred to as the 30-step model. The models in this section using was initially made to create a point of reference for comparison when introducing the models with all our variables.

4.1.1 The loss

The graph under, Fig 4.1a shows how the 30-step model evolves and improves by calculating the loss in form of MSE at every epoch. The graph shows that the model has a steep learning curve for the first epochs before it flattens out as the model gets better at predicting the values, or in other words, as the MSE is decreasing. The loss is the value for the training set and the val_loss is the loss for the validation set. This will be the case for all the loss graphs in this paper.

Because the early stopping feature is implemented with a patience of five in this model, the graph shows some spikes in loss for both the validation and the training set.

4.1.2 The training set

The Training set figure, Fig 4.1b, shows that the 30-step model is able to forecast some of the trends in the actual price fluctuations, but it is not perfectly predicting the prices. However, the training data shows the predicted values on data that the model has already seen, and the values in this set are the ones that are used by the model to create weights for the independent variables. The model is not capable to predict the highest prices, but it does predict the spikes on the training set.

4.1.3 The validation set

Looking at the validation set, Fig 4.1c, it becomes even clearer that the weights are not perfect for this model. The validation set is also a set the model has seen before but the weights that are calculated does not suit the data in this set. The graph shows clearly that the model is picking up on some trends, but in the first 100 values the model does not predict the spikes and the model does not manage to predict the volatility seen in the last 50 variables. The graph is also much rounder on the local maximums meaning that the model seems to struggle with predicting sudden day to day changes.

4.1.4 The test set

The test set is showing the 30-step model perform on data it has not been introduced to before. Looking at the graph in this section, Fig 4.1d, shows that the model is performing very poorly on the test data set. It does not match the observed values. The positive take from this is that the graph seems to have spotted some few trends despite them in most cases being predicted to early. The changes in price also seems to be predicted to be much smoother over time than they actually appear in the observations.

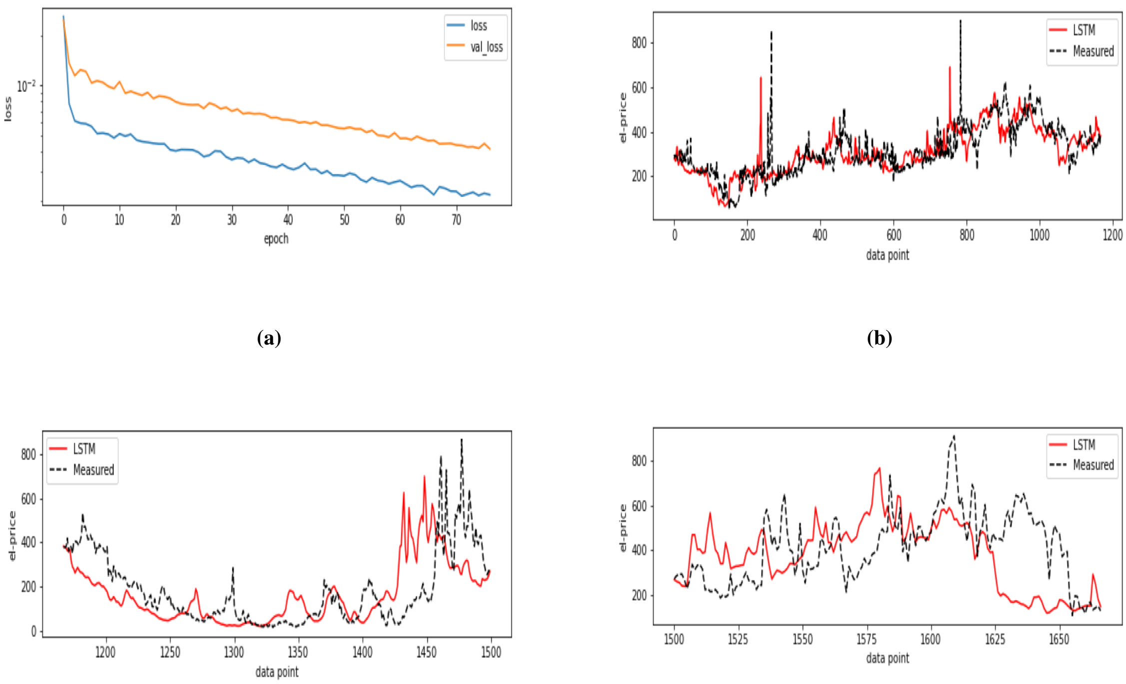

(a) (b)

(c) (d)

Figure 4.1: Illustration of the 30-step model predicting the daily average electricity price in Trondheim based on the daily average electricity price in Trondheim using 30 time steps with the (a) loss , (b) training set, (c) validation set, and (d) test set.

4.2 Predicting average electricity price with 50 time steps

In this section we are looking to predict the price using the 50 previous time steps. The increase of input is done in hopes that the the model will work better than what it did in the last section. An increase in the amount of data used to calculate the next step might help the model catch more trends during training, and hopefully create a more accurate series of outputs. This model will be referred to as the 50-step model.

4.2.1 The loss

The loss function, illustrated in Fig 4.2a, performed relatively similar to the 30-step loss function, observing a steep learning curve in the beginning before flattening out towards the 70th epoch. The model still runs quickly, and its training is taking a little shorter time than in the previous section. This is not to be expected when a larger amount of data is fed into the model at each epoch. The validation set and the training set is showing there is not too big of a difference in the loss. There is actually an deterioration for both the set, however not a big one, we see this more clearly in the figure 5.1a

4.2.2 The training set

As for the training data, shown in Fig 4.2b it is difficult to separate it from the training set produced by the 30-step model. It predicts the same trends, but it seems that the trends are predicted somewhat earlier in the 50-step model than what we observed in the previous section. The model is also better at predicting the peaks as we see the maximums are larger than those produced by the 30-step model. However, still predicted to happen way earlier than they actually appear in real life.

4.2.3 The validation set

The validation set for this model, as seen in Fig 4.2c, is better at predicting peaks and shows to be slightly more volatile than the 30-step. We also see that this validation still is not predicting the absolute maximum price. However, the results of the validation set shows that the model is predicting the values and trends to come at an much earlier stage that they actually do.

4.2.4 The test set

The test set for the 50-step model, illustrated in Fig 4.2d, is also showing a tendency to be slightly more volatile and is picking up on more of the trends that are observed in the electricity price, but still not nearly close enough to be reliable. At the same time the 50-step models ability to predict is poor throughout the whole set, considering that its predictions are incorrect at almost every single data point. The model also has reoccurring tendency of not being able to predict the maximums or the minimum values of the price. Yet, in the final data points it is able to predict a drop in price despite being far off the actual observed price.

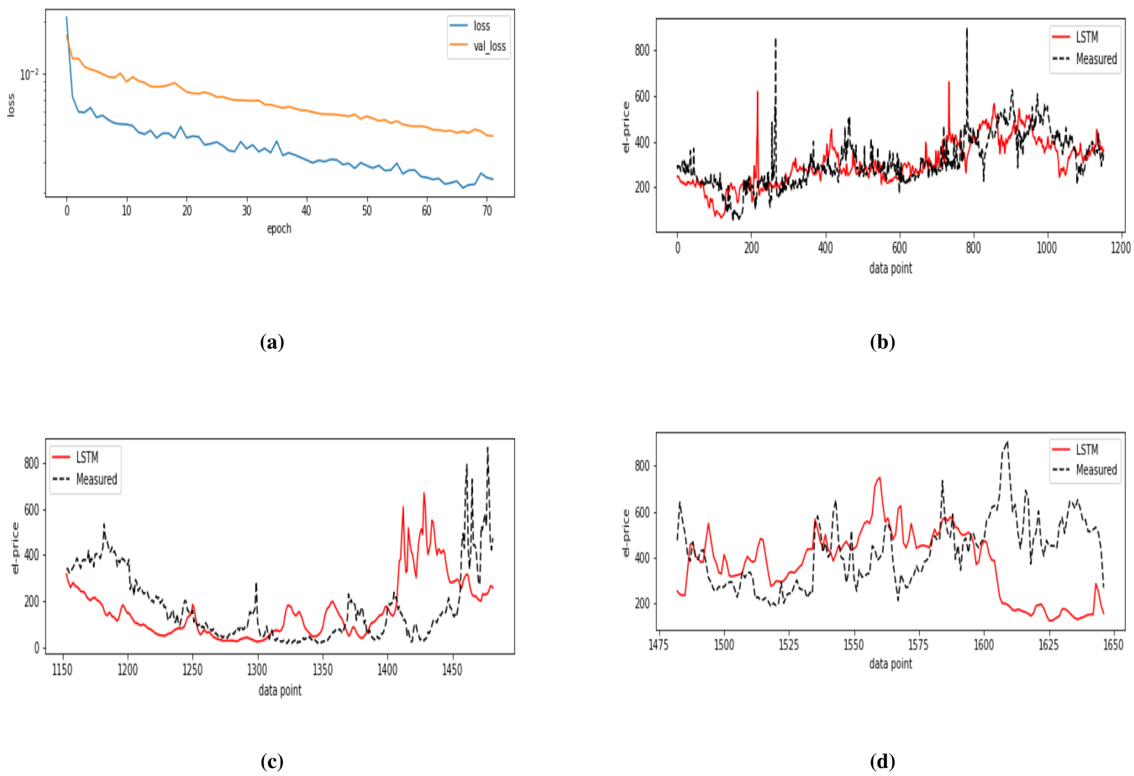

Figure 4.2: Illustration of the 50-step model predicting the daily average electricity price in Trondheim based on the daily average electricity price in Trondheim using 50 time steps with the (a) loss, (b) training set, (c) validation set, and (d) test set.

4.3 Predicting average electricity price using the price with 1 time steps

After analyzing the trends of the previous models they both seem to predict all the values too early in the training set for the 30-step and the 50-step models, we decided to try using only one time step to see how this affects the models ability to make accurate predictions. This will be referred to as the one-step model.

4.3.1 The loss

The loss graph, illustrated in Fig 4.3a, shows that the model is performing better now, having reduced the MSE even more than it did before. At this point we see that the val_loss is actually getting worse in the beginning before it drops and flattens out at as it approaches 0. The main difference between this case and earlier models is that the graph for loss appears to be much smoother with less spikes.

4.3.2 The training set

Using only the previous input to predict the next variable the graph for the training set, illustrated in Fig 4.3b, gives a almost perfect match with the observed values. It predicts all the trends, yet the model still seems to have some difficulties with the predictions the maximum and minimum values.

4.3.3 The validation set

The validation set, illustrated in Fig 4.3c, is showing the same as we saw in the training set. The model is hitting the trends perfectly, despite it missing on the lower prices, where our model is predicting them to be a bit higher than observed. The model is not hitting the highest values perfectly but it is closer in its predictions than in the two earlier models.

4.3.4 The test set

The test set, illustrated in Fig 4.3d, is very promising when using only one time step. The one-step model is predicting all the trends and is pretty much spot on when the price is increasing or decreasing. It is at first eyesight just off in those situations where the price is reaching a local maximum. This is a curious trend as the model returns to hit close to the measured values again once the price decreases after a local maximum.

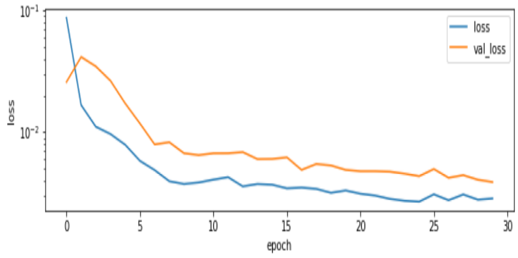

(a)

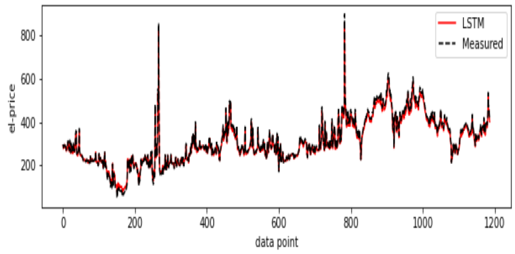

(b)

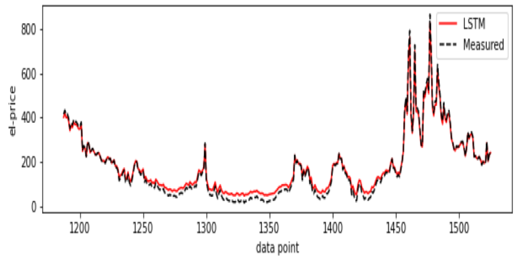

(c)

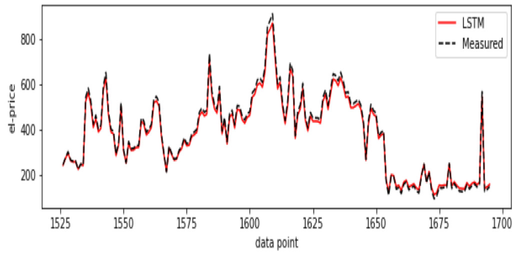

(d)

Figure 4.3: Illustration of the one-step model predicting the daily average electricity price in Trondheim based on the daily average electricity price in Trondheim using one time steps with the (a) loss, (b) training set, (c) validation set, and (d) test set.

4.4 Predicting average electricity price using all our variables and a time step of 30

We will now look at the results produced when implementing all independent variables in the model, introduced in section 3.1.

4.4.1 The loss

The loss graph, illustrated in Fig 4.4a, using the 30 previous variable shows us clear signs that the training set is learning fast and managing to decrease the loss over all. The validation set is however showing signs of bad learning. The graph is decreasing slowly, but there are many spikes and every point of a bigger decrease in the loss is matched by a resulting increase. This is also observed in the diagrams and table in figure 5.1.

4.4.2 The training set

The training set for the 30-step model, illustrated in Fig 4.4b, is proving that the model is able to adapt to the data it is fed. It shows us that the model does not predict the peaks neither the lows, but it shows the trends. The trends shown in the model are very close to the actual observed trends in the market. To add on to this, it also seems to be able to predict the volatility of the market despite its predictions being a bit off the exact measured value.

4.4.3 The validation set

In the validation set, illustrated in Fig 4.4c, the 30-step models values are in contrary to the training set, quite far of the measured values. The graph shows that it is able to notice general trends in the market, it shows that it has the ability to predict some volatility, but mostly it builds upon the general observation this far in the study. This being that LSTM models seem to struggle with predicting the peaks and the lows of the price. The model is also in this section seemingly predicting the larger changes a couple of days before they actually happen, which perhaps could be a useful trait.

4.4.4 The test set

In the test set, illustrated in Fig 4.4d, the 30-step model proves that it is absolutely not able to be used to predict the values of the average daily electricity price in Trondheim. As well as in the other sets it is not able to predict the peaks or the lows, but in the training set it is almost not able to produce a correct prediction at any time throughout the whole set. It is still is able to catch and predict market tendencies at some points along the set, but the predictions are to inaccurate to be of any reliable use.

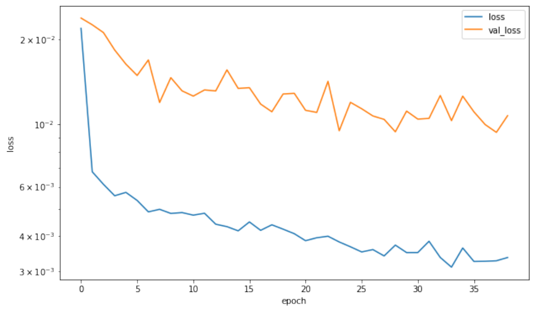

(a)

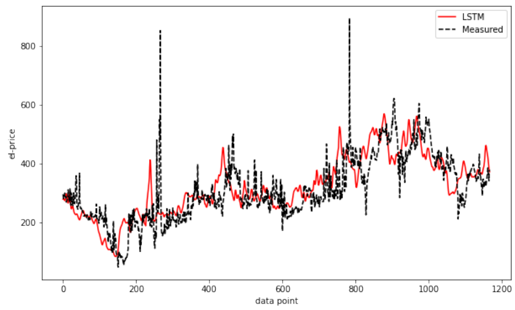

(b)

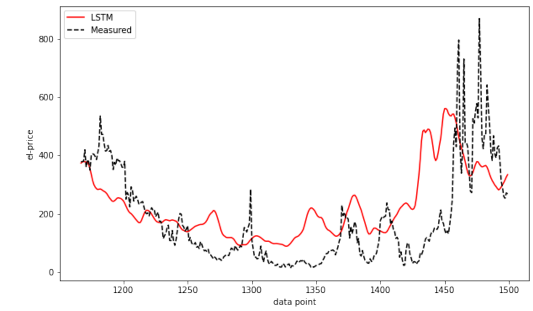

(c)

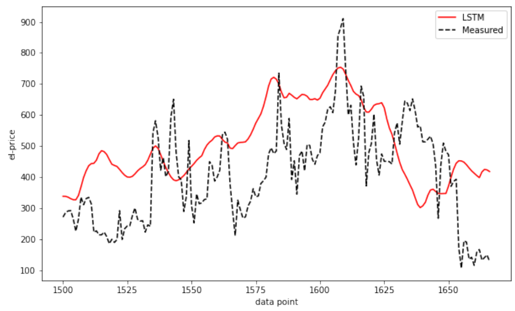

(d)

Figure 4.4: Illustration of the model predicting the daily average electricity price in Trondheim based on all our variables using 30 time steps with the (a) loss, (b) training set, (c) validation set, and (d) test set.

4.5 Predicting average electricity price using all our variables and a time step of 50

Again, to try to make the predictions better we implemented more inputs in the model. In this occasion we use 50 time steps to predict the next value.

4.5.1 The loss

The loss graph, illustrated in Fig 4.5a, for using 50 time steps shows that on the training set, the model is yet again learning well. It has a steep learning curve to begin with before the learning curves flattens out towards the final epochs. For the validation set, the learning curve also starts of steep, but it does not flattens out towards the final epochs. Instead, it is constantly changing from improving to getting worse. The total decrease in MSE in the validation set shows that it has learned something, but not to the extent, or in the smooth way, we were hoping.

4.5.2 The training set

The 50-step model used on the training set, illustrated in Fig 4.5b, is showing to fail on predicting the prices more often than the 30-step model. On the first variables the model is not picking up the increased prices and still has big problems predicting the highest and the lowest prices. The market trends are somewhat predicted, but they seems to be predicted at an much earlier state than observed. Around the 800th and the 1050th data point we also observe that the 50-step models predictions are far off the measured values. This is an alarming sign going into the validation and test set, as these are data points already familiar to the model. Based on this we can already predict that the 50-step model will struggle with the unfamiliar data in the test set.

4.5.3 The validation set

The graph produced by the 50-step validation set, illustrated in Fig 4.5c, explains the spiky and slow learning loss curve. For the first values it completely ignores the trend where observed data shows an increase in the average electricity price in Trondheim before a longer period of decreasing price. The model picked up on the trend of the price decrease, but expected the price to be much lower than observed, and then a period of it being much higher than observed. Once again the issue with maximum and minimum values prevail, and these values still prove difficult for an LSTM-model to predict. In the validation set the model seems to have a hard time with picking up on short term volatility, and mostly the market trends are predicted to happen long before they are measured.

4.5.4 The test set

The prediction on the test set, illustrated in Fig 4.5d, proves once again that the 50-step model has accuracy and prediction issues. The highs and lows are once again not predicted correctly and on the test set most of the actual trends are missing. At some points through the data set it is matching some of the market trends, and around the 1600th data point, and around the 1620th data point it is actually almost predicting the correct value. This is however the only time it predicts close to correctly, and could be a coincidence given the inaccuracy on the rest of the test set. Once again the model shows to not be a reliable tool of predicting power prices based on these variables. The test set also proves that the volatility of the real life market is not something this 50-step model is able to reliably predict. Despite the fact that the 50-step model at some points is able to catch some trends, the accuracy is yet again too poor to be useful in any way.

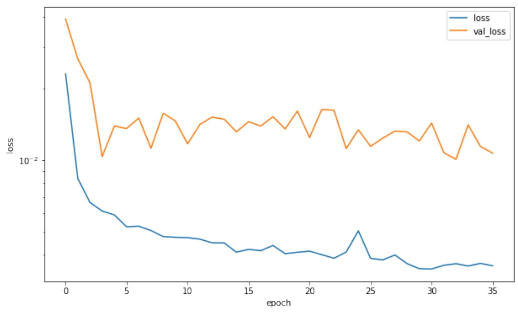

(a)

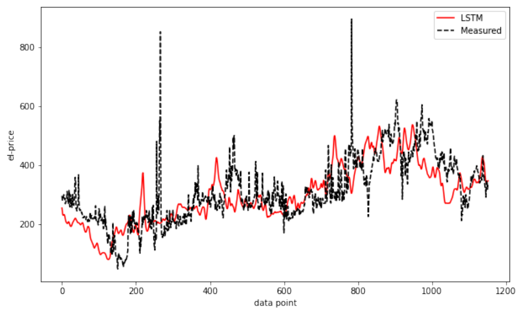

(b)

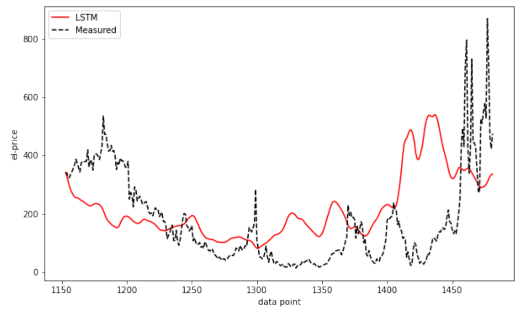

(c)

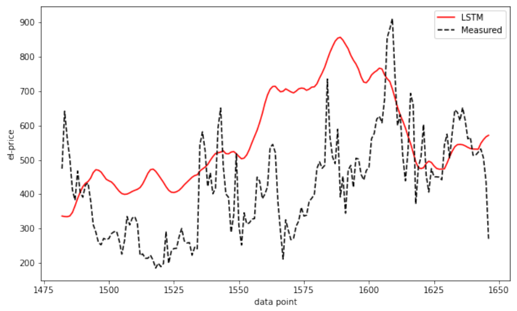

(d)

Figure 4.5: Illustration of the model predicting the daily average electricity price in Trondheim based on all our variables using 50 time steps with the (a) loss, (b) training set, (c) validation set, and (d) test set.

4.6 Predicting average electricity price using all our variables and a time step of 1

We had some success in the previous one-step model where we only used the average daily electricity price in Trondheim to predict the average daily electricity price in Trondheim. As that one-step model showed great improvements from the 30- and 50-step models we wanted examine if the same trend applied with all variables. In this section we then use a one-step model with all variables to try and predict the average daily power price in Trondheim.

4.6.1 The loss

The loss for the training set, illustrated in Fig4.6a is showing to be efficiently decreasing at an appealing rate, proving that the training set for the one-step model is once again learning quite well. The validation set is getting worse for the first epochs before also this shows to be learning at a good rate. The validation loss is less volatile than we have seen for the two earlier models, with a smoother, steeper learning curve over all. But we observe that for the final epoch the MSE is increasing.

4.6.2 The training set

The model is predicting very well om the training set, illustrated in Fig4.6b and is on most occasions close to perfectly predicting the price. Once again the lowest and the highest average daily electricity prices are not predicted correctly by the model. On the side of trends, it is spot on, showing that the model has been able to create valid weights in its learning phase. The values where we see a clear miss on the measured values is right before the 200th, around the 400th data point, as well as for the peaks.

4.6.3 The validation set

The models performance on the validation set, illustrated in Fig4.6c, is proving that on this set as well, the model is able to predict many of the trends seen in the market. As has become the norm for these models, it cannot predict the extreme values. The minimum price predicted is still above the measured minimums, and the predicted maximum is still lower than the observations. The best sign shown in the validation set is that market volatility actually is picked up well in the model.

4.6.4 The test set

On the test set, illustrated in Fig4.6d, we observe a model that clearly is good at picking up on trends. The model starts off by predicting the values to be to high, before correcting itself. The model is once again showing that the maximum and minimum price is something it is not able to predict. The final 75 data points prove to be a bit alarming, as the model is predicting the price to be much higher than the observed values. However, the model is picking up on all the trends in the market. The model is then also predicting the volatility in the market quite well.

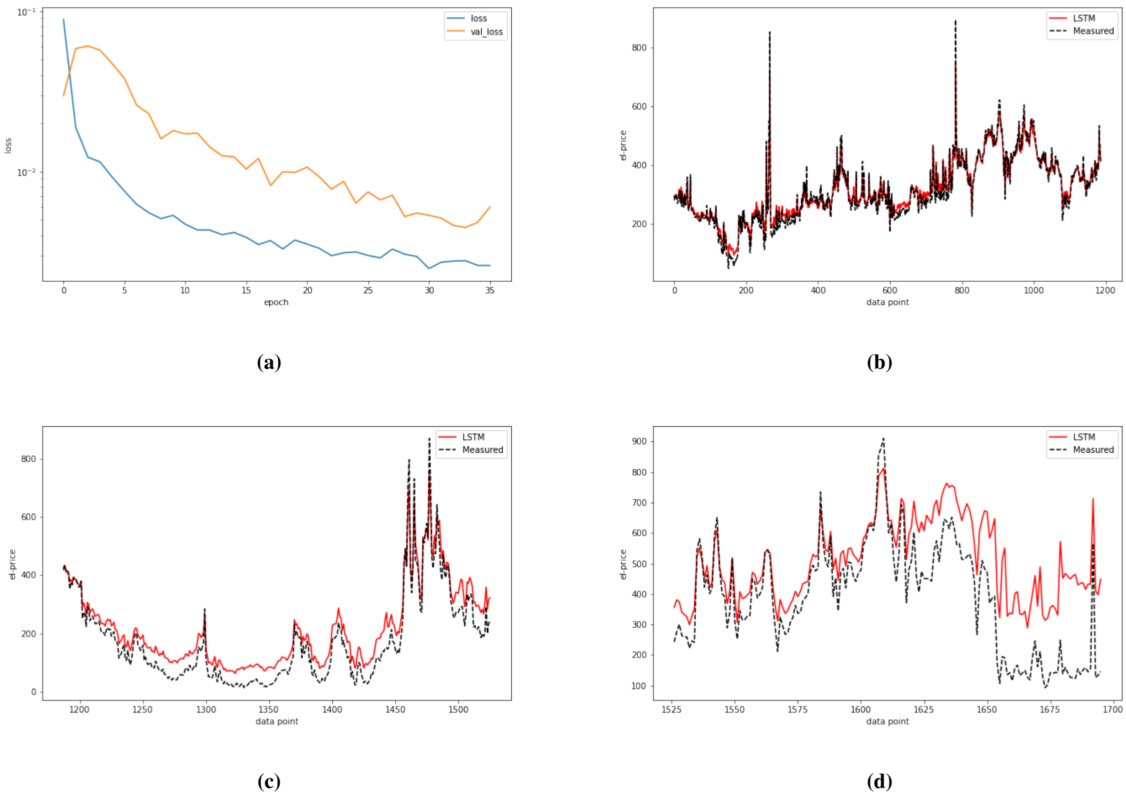

(a) (b) (c) (d)

Figure 4.6: Illustration of the model predicting the daily average electricity price in Trondheim based on all our variables using 1 time steps with the (a) loss, (b) training set, (c) validation set, and (d) test set.

5 Discussion

5.1 Analyzing the loss

The model is showing clear trends of learning and as described in the former chapter, the loss in form of the MSE is increased for all the models when adding more variables. All in all, all the models achieve a remarkably low MSE given their poor outputs as seen in fig 5.2a. This however, does not mean that the model is good, and as we could see in the last chapter most of the models proved not to be able to produce any reliable predictions. The low loss is a result of the models relatively good prediction on the training data, but as the validation loss showed, in most occasions the models weights are working poorly when new data is introduced to the model for validation and testing. As seen in fig. 5.2b and fig. 5.2c the models only using the electricity price to make predictions in almost all cases results in the lowest loss.

	all independent var			only el price in Trondheim		
timesteps	1	30	50	1	30	50
loss	0,0025949	0,0033627	0,0035949	0,0028743	0,0021674	0,0023865
val_loss	0,0060062	0,0107442	0,0107054	0,0039163	0,0041154	0,0042794

(a) Table comparing the MSE for all models

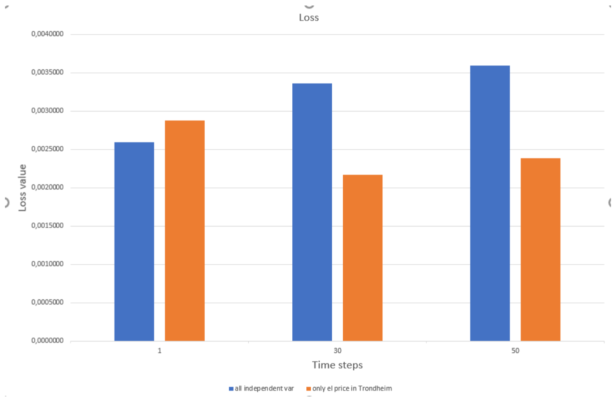

(b) Direct comparison of the loss in the models containing all variables(blue) and models only containing the electricity price in Trondheim(orange)

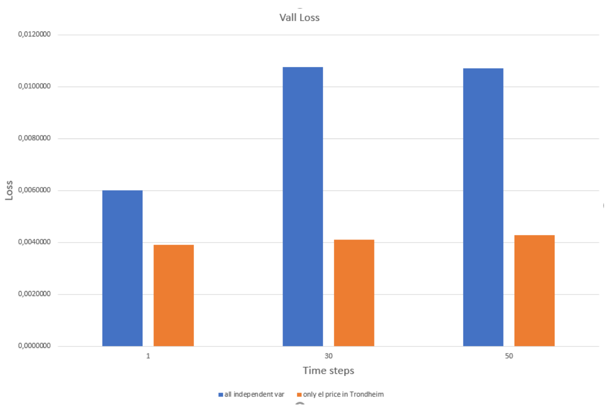

(c) Direct comparison of the loss in the validation sets for models containing all variables(blue) and models only containing the electricity price in Trondheim(orange)

Figure 5.1

5.2 Is the model getting better?

From fig 5.2 we observe that the model in most cases seem to be better using fewer time steps. Why does this happen?

Using fewer time steps means that the model looks at fewer days to predict the next. With one time step the models input is only using the data from the day before, and because of this the model will calculate tomorrow's value based on only the value from the day before. As a result of the way the network works, it is fed the price from yesterday as the input and the price today as the output to predict the price today. This then results in the model seemingly hitting all the trends, but the model is still not actually managing to hit the values perfectly on the training set, as we see in the figure 5.2.

This can also be used to explain why the model is working poorly for the 30 time step and the

50 time step prediction. In this case all the other variables that are fed into the model appear to simply just be noise that forces the model to weight and utilize non-relevant data to make predictions. Because of this the model is not improving.

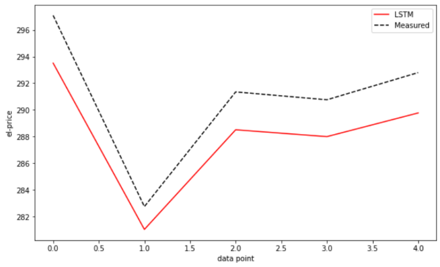

Figure 5.2: Illustration of the misfit using 1 time step

5.3 Why does the model struggle with minimum and maximum values?

In the maximum and minimum points of the dependent values, most of the input data taken into consideration by the model is yet again just noise. Because of this, LSTM will not in our situation, be valid for predicting extreme cases.

5.4 Is our time series a random walk?

The model has a tendency to improve as a predictor as you reduce the amount of time steps utilized to make the prediction. As seen in chapter 4, the model with one time step in both occasions provided the best results. If this model was actually able to make valid predictions this would not be the case, and increasing the amount of time steps should result in a better prediction. This could be an indication that the time series we try to make predictions on, is in fact a random walk, and by that a Markov process. This could be an explanation for the models strange behaviour. All the extra variables are clearly affecting the predicting power of the model, and we see it especially in fig 4.6d where the models prediction on the test set is very different from what it manages without the extra variables in figure 4.3d. By means of this, our research could be an indication of the electricity price being somewhat of a Markov process, but more likely this is a result of the input variables not being sufficient to capture the complexity of the power market. Because there are immense amounts of geographical, political, natural, and psychological factors influencing the power price both directly and indirectly, proving the power price to be or not to be a random walk is a difficult and possibly controversial task not suited for this thesis. We could be receiving these results simply because of the scarcity of the variables chosen for this data set.

6 Conclusion

In this thesis we wanted to explore how accurate LSTM-networks can be used to predict short term electricity prices in Trondheim, based on our set of chosen variables.
Throughout this study we have uncovered that an LSTM network is not able to make valid predictions for the variables we chose. As we discovered throughout chapter 4 the model proved to provide the best results with fewer variables and fewer time steps. Despite the accuracy of the model appearing to be rather good using only one time step, it would most likely never be of any valid use for actual prediction in a business scenario.

The fact that the models predictability increases with fewer time steps proves that there is a problem with the models credibility. A possible explanation, as we introduced in section 5.4, is that the time series the model is making predictions on is subject to the random walk theory. What is even more likely however, is that the chosen set of variables are not suited for making these predictions, and by that proving that an LSTM network using our chosen set of variables is unable to make valid short term predictions for the electricity prices in Trondheim. This, despite the high accuracy achieved with the one-step model, makes the it unusable for the suggested areas for budgeting, power-usage planning and cost calculations and estimations.

6.1 Further studies

There are many ways one could create a better model for this task. Experimenting with more LSTM layers and with different values than the ones used in this study could prove different results than what we achieved. However, experimenting with adding layers to our model during this study did not yield any noticeable results. Another approach would be to implement different variables and more values, getting even more precise weather data by creating a difference between rain, snow, hail and other forms of precipitation. Implementing the export and import as well as the supply and demand inn the different electricity zones in Norway and the other regions using Nord Pool.

As seen in the correlation matrix in fig. 3.1, some of our selected variables show a higher correlation with the price of electricity in Oslo, than in Trondheim. Because of this, the electricity prices in the Oslo area seem to be better described by these measurable variables, and an LSTM model could provide better predictability in this area.

In addition to this, it would be interesting to investigate if the power market is in fact subject to the random walk theory and by that is a Markov process.

References

The Editors of Encyclopaedia . (2008a). *Markov process.* Retrieved from https://www.britannica.com/science/Markov-process

The Editors of Encyclopaedia . (2008b). *Random walk.* Retrieved from https://www.britannica.com/science/random-walk

Chollet, F., et al. (2015). *Keras.* https://keras.io.

Donges, N. (2021). *A guide to rnn: Understanding recurrent neural networks and lstm networks.* Retrieved from https://builtin.com/data-science/recurrent-neural-networks-and-lstm

Eika, T., & Jørgensen, J.-A. (2003). Makroøkonomiske virkninger av høye strømpriser i 2003.

Flovik, V. (2018). *Markov process.* Retrieved from https://towardsdatascience.com/how-not-to-use-machine-learning-for-time-series-forecasting-avoiding-the-pitfalls-19f9d7adf424

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press. (http://www.deeplearningbook.org)

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, *9*(8), 1735–1780. doi: 10.1162/neco.1997.9.8.1735

Khan Academy. (2019). *Lesson summary: Market equilibrium, disequilibrium, and changes in equilibrium.* Retrieved from https://www.khanacademy.org/economics-finance-domain/ap-macroeconomics/basic-economics-concepts-macro/market-equilibrium-disequilibrium-and-changes-in-equilibrium/a/lesson-summary-market-equilibrium-disequilibrium-and-changes-in-equilibrium

Kingma, D. P., & Ba, J. (2017). *Adam: A method for stochastic optimization.*

Li, S., & et al. (2018). Independently recurrent neural network (indrnn): Building a longer and deeper RNN. *CoRR*. Retrieved from http://arxiv.org/abs/1803.04831

Ministry of Finance and Office of the Prime Minister. (2021). *Government launches electricity support package worth billions.* Retrieved from https://www.regjeringen.no/en/aktuelt/government-launches-electricity-support-package-worth-billions/id2891839/

Ministry of Petroleum and Energy. (2021). *Kraftmarkedet.* Retrieved from https://

energifaktanorge.no/norsk-energiforsyning/kraftmarkedet/

Moolayil, J. (2019). *Learn keras for deep neural networks: A fast-track approach to modern deep learning with python*. Apress. doi: 10.1007/978-1-4842-4240-7

Nielsen, M. A. (2018). *Neural networks and deep learning.* Determination Press. Retrieved from http://neuralnetworksanddeeplearning.com/

Nord Trøndelag Elektrisitetsverk AS. (2022). *Hva påvirker strømprisen?* Retrieved from https://nte.no/blogg/hva-pavirker-stromprisen/

Norges Bank. (2022). *Exchange rates.* Retrieved from https://www.norges-bank.no/en/topics/Statistics/exchange_rates/?tab=currency&id=USD

Olah, C. (2015). *Understanding lstm networks.* Retrieved from http://colah.github.io/posts/2015-08-Understanding-LSTMs/

Rosvold, K. A., & Halleraker, J. H. (2021). *vannmagasin (kraftverk).* Retrieved from https://snl.no/vannmagasin_-_kraftverk

Singarimbun, R. N., Nababan, E. B., & Sitompul, O. S. (2019). Adaptive moment estimation to minimize square error in backpropagation algorithm. , 1-7. doi: 10.1109/ICoSNIKOM48755.2019.9111563

Skånlund, A. B., & et al. (2013). Nordic bidding zones. *THEMA Consulting Group*. Retrieved from https://old.thema.no/wp-content/uploads/2013/10/THEMA-report-2013-27-Nordic_Bidding_Zones_FINAL.pdf

Statistics Norway. (2022). *09364: Kraftpriser i sluttbrukermarkedet, etter kontraktstype, statistikkvariabel og kvartal.* Retrieved from https://www.ssb.no/statbank/sq/10065629

A The Code

Please see attached PDF file for all code used in this thesis.

NTNU

Kunnskap for en bedre verden