


Eivind Hagemann Brataas

Statistical Machine Learning on Covid-19 Time Series using Econometrics

Master's thesis in MSMNFMA
Supervisor: Gunnar Taraldsen
Co-supervisor: André Voigt
June 2022



Eivind Hagemann Brataas

Statistical Machine Learning on Covid-19 Time Series using Econometrics

Master's thesis in MSMNFMA
Supervisor: Gunnar Taraldsen
Co-supervisor: André Voigt
June 2022

Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Mathematical Sciences



Kunnskap for en bedre verden

DEPARTMENT OF MATHEMATICAL SCIENCES

MA3911 - MASTER'S THESIS IN MATHEMATICAL SCIENCES

Statistical Machine Learning on Covid-19 Time Series using Econometrics

Author:

Eivind Hagemann Brataas

Primary supervisor:

Gunnar Taraldsen

Secondary supervisor:

André Voigt

June, 2022

Abstract

This thesis compares three models for forecasting daily new cases for Covid-19. The first model is a class of machine learning models, called an CNN-LSTM model, and was the state-of-the-art model for the stated task on the global daily new cases data set during the summer of 2021. Some months earlier, Taraldsen published two so-called toy models for forecasting the daily new cases in Norway. Both these models are SARIMA models, but one of them assumes Gaussian white noise, while the other assumes that the noise is conditionally heteroscedastic and is modeled by a GARCH model. They both gave accurate predictions and come with a prediction interval, as opposed to the CNN-LSTM model. Additionally, they were much less computer intensive, with only three and four parameters, respectively. On the other hand, the CNN-LSTM model must fit more than 300000 parameters. The models had not yet been applied to other Covid-19 data sets than what was used in their respective articles. This thesis compared the performance of the three models on the global time series data, as well as the Norwegian data. A lot more data have become available since the articles were published. This makes it possible to compare the models on other partitions of both the data sets, and to experiment with reduced sample sizes across all these partitions. Finally, a parametric bootstrap experiment was conducted to get a grasp of the uncertainty in the forecast from the CNN-LSTM model. While the CNN-LSTM model achieved some accurate forecasts, the results of this thesis suggest that the SARIMA model with GARCH noise may be the model of choice for the earliest parts of the data sets, while the SARIMA model with Gaussian white noise would be the best choice on the rest of the data sets, where its predictions are more accurate and has about the same spread. These results are mostly explained by the varying heteroscedasticity of the two time series.

Preface

I would like to preface this by thanking the people that made the creation of this thesis possible.

I would first like to express my deepest gratitude to my supervisor Gunnar Taraldsen. Gunnar introduced me to the interesting field of time series two years ago. We have since then shared many good discussions, and he has kept me motivated. With his great theoretical knowledge, he has always kept the big picture in mind. I would like to thank him for devoting a lot of his time and for being patient with me. Without him, the thesis would not have existed.

I would also like to thank my secondary supervisor Andre Voigt for sharing his insights about the application of the models in the real world.

Finally, I want to thank all my friends, my family, and my partner Elin for supporting me while writing this thesis.

The thesis was exiting to work on, as it includes topics from applied and theoretical statistics, to analysis, to machine learning. It was therefore the perfect ending to my time as a master student.

Table of Contents

List of Tables	v
List of Figures	vi
1 Introduction	1
2 Theory	3
2.1 Time Series	3
2.2 Defining the ARMA process and its extensions	3
2.3 The GARCH process	6
2.3.1 The existence of the GARCH(1,1) model	6
2.4 Parameter estimation for ARMA models	11
2.5 Prediction with ARMA models	13
2.6 Simulation with ARMA models	14
2.7 Model Diagnostics for ARMA models	15
2.8 The Machine Learning Model	17
2.8.1 Convolutional Neural Networks	17
2.8.2 Long Short Term Memory models	19
2.9 Measures of accuracy	23
3 Methods	24
3.1 Defining the three models	24
3.1.1 The SARIMA model	24
3.1.2 The Gandalf model	24
3.1.3 The CNN-LSTM model	25
3.2 On the different ways of forecasting	26
3.2.1 Prediction scheme 1	26
3.2.2 Prediction scheme 2	27
3.3 Programming languages	28
4 Applying the models on Covid-19 data	29
4.1 Data	29
4.2 Exploring the fit of the SARIMA models on the Taraldsen's data set	29
4.3 Exploring the fit of the CNN-LSTM model on Zain's data set	32
4.4 Comparing the three models	34
4.4.1 Comparisons on Taraldsen's data set	34

4.4.2	Comparisons on Zain's data set	37
4.4.3	Comparisons with less training data and at different partitions	40
4.5	The models on simulated realizations	45
5	Discussion	50
5.1	The performance of the three models	50
5.1.1	Accuracy	50
5.1.2	The changing volatility and the effect of sample reduction	51
5.1.3	Results from the simulation study	52
5.2	Possible improvements	52
5.2.1	The machine learning model	52
5.2.2	The SARIMA models	53
5.3	A note on the application of the models in the real world	54
6	Conclusion	55
	References	56
	Appendix	58
A	Figures and tables	58
B	Code	61

List of Tables

1	Relevant computer specifications.	28
2	Time used to fit and forecast with each of the models, trained on Taraldsen's data set.	36
3	All forecast results from the SARIMA model and the Gandalf model, and the mean CNN-LSTM prediction. Note that the RRMSE and MAPE from the CNN-LSTM model is based on the mean prediction of ten models. All 28-day predictions used prediction scheme 1, while prediction scheme 2 was used for all seven-day predictions.	43
4	Relative standard deviation of accuracy measures of 28-day forecasts on 1000 simulated realizations from the Gandalf model. The models were initially trained on Zain's data set.	49
5	All forecast results from the mean prediction of ten CNN-LSTM models and the mean RRMSE of the ten (individual) models. All 28 day predictions used prediction scheme 1, while prediction scheme 2 was used for all seven day predictions.	59

List of Figures

1	The Norwegian and global Covid-19 new cases time series.	2
2	An illustration of a one-dimensional kernel operation with kernel size is 3 (TowardsDataScience, n.d.).	17
3	An illustration of a one-dimensional convolutional neural network (Kiranyaz et al., 2021).	18
4	An illustration of a one-dimensional pooling operation, with size 2 and stride 2 (Peltarion, 2022).	18
5	Construction of samples from time series data.	19
6	An illustration of a RNN (W. Zhang et al., 2019).	19
7	An illustration of a LSTM cell (Varsamopoulos et al., 2018).	20
8	The architecture of the specific CNN-LSTM model from Zain et al. (2021).	25
9	An illustration of prediction scheme 1 for the CNN-LSTM model.	27
10	Illustration of prediction scheme 1 for the SARIMA models.	27
11	Illustration of prediction scheme 2 for the CNN-LSTM model.	28
12	The Norwegian and the global new cases time series with transformations.	30
13	Replication of results from Taraldsen (2020b) with the SARIMA model and the Gandalf model.	31
14	ACF plot of residuals and squared residuals for SARIMA model on Taraldsen’s data set.	32
15	ACF plot of residuals and squared residuals for Gandalf model on Taraldsen’s data set.	32
16	Mean loss function of ten equally specified CNN-LSTM models, trained on Zain’s data set.	33
17	Forecast with ten CNN-LSTM models from July 18th to August 14th 2020 on the global new-cases data with prediction scheme 1. The models were trained on Zain’s data set.	33
18	Forecast with ten CNN-LSTM models from November 11th to November 18th 2020 on the Norwegian new-cases data with prediction scheme 1. Taraldsen’s data set was used for training.	34
19	Forecast from the three models with prediction scheme 2, from November 11th to November 18th 2020 on the Norwegian new-cases data. Taraldsen’s data set was used for training.	35
20	Forecast from the three models with prediction scheme 1, from November 11th to December 8th 2020 on the Norwegian new-cases data. Taraldsen’s data set was used for training.	35
21	The parameter estimates from the Gandalf models used to predict the 28 days shown in Figure 20. The dates along the x-axis is the last day used to train the model used to predict the day after. The colored dotted lines are the parameter estimates of the Gandalf model fitted on Taraldsen’s data set.	36
22	The parameter estimates from the SARIMA models used to predict the 28 days shown in Figure 20. The dates along the x-axis is the last day used to train the model used to predict the day after. The colored dotted lines are the parameter estimates of the SARIMA model fitted on Taraldsen’s data set.	37
23	Forecast from the SARIMA models with prediction scheme 2 from July 18th to July 24th on the global data set. Both models were trained on Zain’s data set.	37
24	ACF plot of residuals and squared residuals for the SARIMA model and the Gandalf model, trained on Zain’s data set.	38

25	Forecast from the three models with scheme 2, from July 18th to July 25th 2020 on the global new-cases data. All models were trained on Zain's data set.	39
26	Forecast from the three models with scheme 1, from July 18th to August 14th 2020 on the global new-cases data. All models were trained on Zain's data set.	39
27	The parameter estimates for the two SARIMA models used to predict the 28 days shown in Figure 26. The dates along the x-axis is the last day used to train the model used to predict the day after. The colored dotted lines are the parameter estimates of the SARIMA model fitted on Zain's data set.	40
28	Forecast from the three models with scheme 2, from July 17th 2020 on the global data set. The models were trained on Zain's data set.	40
29	Forecasts using the 100 previous Norwegian data training data from November 10th 2020.	41
30	Forecasts using the 50 previous Norwegian data training data from November 10th 2020. .	41
31	The parameter estimates for the two SARIMA models used to predict the 28 days shown in Figure 30. The dates along the x-axis is the last day used to train the model used to predict the day after. The colored dotted lines are the parameter estimates of the SARIMA model fitted on Zain's data set.	42
32	MAPE scores for each model from Table 3.	44
33	Comparisons of models based on MAPE scores from Table 3	44
34	Three simulations from the Gandalf model. The model was trained on the original global training data. The 28 ensuing days were then simulated.	45
35	Predictions on three simulated realizations from the Gandalf model with prediction scheme 1.	46
36	Distribution of RRMSE and MAPE based on predictions from the Gandalf model on 1000 simulated realizations from the Gandalf model. Prediction scheme 1 was used to generate the predictions, once for each simulation.	46
37	The deviation of the four Gandalf parameters on the 200 simulated realizations. For each simulated realization, prediction scheme 1 was used to fit models and predict one day at a time. This gives 28 fitted models and parameter estimates for each simulation.	47
38	Distribution of RRMSE and MAPE based on log-scaled predictions from the Gandalf model using simulated realizations from the Gandalf model on log scale. Prediction scheme 1 was used to generate the predictions, once for each of the 1000 realizations.	47
39	Distribution of RRMSE and MAPE based on predictions from the SARIMA model on 1000 simulated realizations from the Gandalf model. Prediction scheme 1 was used to generate the predictions, once for each realization.	48
40	Distribution of RRMSE and MAPE based on predictions from the CNN-LSTM model on 1000 simulated realizations from the Gandalf model. Prediction scheme 1 was used to generate the predictions, once for each realization. The mean CNN-LSTM prediction was used to generate the forecasts.	48
41	Forecast with ten CNN-LSTM models from November 11th to December 8th 2020 on the Norwegian new-cases data. The previous 264 days there used as training data.	58
42	RRMSE scores for each model from Table 3.	60
43	Comparison of models based on RRMSE scores in Table 3	60

1 Introduction

The novel corona virus 2019 (Covid-19) took the world by storm, killing over 6 million people in two years. Unlike previous pandemics, lots of data have been gathered about the disease. This enables accurate forecasts of the spread of the virus, which have been able to function as an early warning system for authorities to react with appropriate responses (Dehning et al., 2020). Even though the current pandemic might be under control, globalisation will likely cause future viruses with the same potency. Thus, the development of better forecasting tools for disease spread should continue, to hopefully prevent another disaster. There are at least two schools of thought when it comes to modeling the spread of diseases. The epidemiological approach attempts to understand the inner workings of the virus, which can then be used to model its spread, while a data-driven approach ignores where the data comes from and purely uses the data at hand. In its simplest form, the latter approach only considers the new Covid-19 cases X_t at day t . After collecting the observations x_1, \dots, x_n , the goal is to predict the amount of future new cases X_{n+1}, \dots, X_{n+h} as accurately as possible. This thesis focuses on solving this exact problem.

In an article published by Zain et al. (2021), several predictive models were trained to solve the above problem. It included a proposed machine learning model that combined the abstraction capabilities of a Convolutional Neural Network (CNN) with a Long-Short Term Memory (LSTM) model's ability to learn long term dependencies. The current state-of-the-art model at the time of publishing was also among the compared models. From their empirical results on the global daily new cases data set from January 4th to August 14th 2020, they concluded that the proposed model was the best model for the task at hand. This conclusion was based on the accuracy of 28-day forecasts from each model. Also among the most accurate models in this study was an Auto Regressive Moving Average (ARMA) model. The ARMA model is a large class of models, with many possible alterations. None of these alterations were explored in their article. Some months earlier, Taraldsen (2020b) published two so-called toy models for the Norwegian daily new cases from February 21st to November 10th 2020. They were both Seasonal Auto Regressive Integrated Moving Average (SARIMA) models, but one of them assume that the underlying noise is Gaussian and independent of time, while the other one assume GARCH(1, 1) noise; a more general noise model that allows conditional heteroscedasticity. On the data set used by Taraldsen, both these models delivered accurate forecasts. For the daily new cases of Covid-19 in some of the key countries, SARIMA models and ARMA-GARCH models both have been shown to give accurate forecasts (Kumar et al., 2021)(Ekinci, 2021). However, at the time of writing, a combined SARIMA-GARCH model for daily new cases have yet to be tested on the global new cases data. In this thesis, the three models were replicated as precisely as possible, and their performances were compared on both the global and the Norwegian data set, seen in Figure 1. The articles were written based on data from 2020. Since then, a lot more data has become available. The two time series therefore include the daily new cases up to February 20th 2022. A lot has happened to the handling of the virus over the course of the two years: new variants have emerged, vaccines have been developed, and rulings and attitudes surrounding the disease have evolved. How will these changes affect the models fit and prediction capabilities? And what happens to the results if the sample size is reduced? This thesis attempts to answer these questions.




Figure 1: The Norwegian and global Covid-19 new cases time series.

The replicated models were initially tested on the two data sets from the articles from Zain et. al. and Taraldsen. This made the results from the models in the thesis comparable to the results from the articles. The models were later tested on newer and older partitions of the data sets, by regulating the last date included in the training set, and the number of samples to include preceding said date. The primary goal was to compare the accuracy of the models across these partitions of the data sets. However, this was not the only task. If the models were to be used in practice, the uncertainty of the forecasts is almost as important as the forecast itself. When using Taraldsen’s SARIMA models, this uncertainty can easily be approximated with theoretical 95% prediction intervals. The CNN-LSTM model has no way to generate such an interval. To get a grasp of the spread of this model’s forecast, a small parametric bootstrap study was conducted. The secondary goal of the thesis was to describe and compare the spread in the forecasts of each model. The two SARIMA models only contain 3 and 4 parameters, respectively. Some of these specifies the structure of the auto regressive components and the moving average components, while others describe the noise component of the models. The third goal of the thesis was to investigate the behaviour of these parameters, especially the ones connected to the noise of the models. The CNN-LSTM model contains more than 300000 parameters. Such an investigation was therefore not feasible for this model.

The time series analysis made in this thesis require some baseline knowledge about ARMA models and machine learning models. The necessary theory is given in Section 2. Details about the specific models and how the forecasts in this thesis were constructed are given in Section 3. The models were then applied to the two data sets, as described above. These results are presented in Section 4 and discussed in Section 5. The final thoughts and conclusions in Section 6 finalize the thesis.

2 Theory

This section introduces the necessary concepts and theoretical results needed to understand the thesis and demonstrates how the later results were generated.

2.1 Time Series

The **random variables** lay the foundation of statistical theory. The exact definition is given.

Definition 2.1 (Random variable). *Let (Ω, \mathcal{F}) and $(\tilde{\Omega}, \tilde{\mathcal{F}})$ be measurable spaces.*

A map $X : \Omega \rightarrow \tilde{\Omega}$ is called a random variable if

$$X^{-1}(\tilde{F}) \in \mathcal{F} \quad \text{for all } \tilde{F} \in \tilde{\mathcal{F}}.$$

In probability theory, the measurable spaces (Ω, \mathcal{F}) and $(\tilde{\Omega}, \tilde{\mathcal{F}})$ are called event spaces. In this thesis, $\tilde{\Omega} = \mathbb{R}$ and $\tilde{\mathcal{F}} = \mathcal{B}(\mathbb{R})$, the Borel σ -algebra. The associated probability space (measure space) is (Ω, \mathcal{F}, P) , where P is the cumulative distribution function of X .

The Covid-19 new-cases data are viewed as a series of events from random variables X_t , one for each passing day t . A realization of these random variables is x_1, \dots, x_n . Figure 1, shows one such realization for each of the two time series.

Definition 2.2 (Time series). *A time series is a sequence of random variables $\{X_t\} = \{X_t\}_{t \in \mathbb{R}}$ indexed by time. In this thesis, a realization of X , denoted by (x_1, \dots, x_n) , is also called a time series.*

More general definitions of time series exist. However, for the purposes of this thesis, the above definition will suffice.

The remainder of this section outlines how to model time series. For the statistical models, a desired property for time series when modeling is *weak stationary*. Essentially, the elements of time series with this property maintains the same relationship to its neighbouring data points across the whole time series.

Definition 2.3 (Weak Stationarity). *A time series X is weakly stationary if it is shift-invariant. That is, X_t has a constant mean for all t and $\text{cov}(X_i, X_j) = h(|i, j|)$ for some unsigned function h .*

In other words, the correlation between two variables of a weakly stationary time series is only dependent on the distance between their respective points in time, not time itself. On the other hand, some time series has no correlation between the neighbouring variables. These are called white noise processes.

Definition 2.4 (White noise). *If $Z \sim \text{WN}(0, \sigma^2)$, then $\{Z_t\}$ is a sequence of uncorrelated random variables that all have mean zero and the same variance σ^2 . This sequence is called a white noise process.*

2.2 Defining the ARMA process and its extensions

The **Auto Regressive Moving Average** (ARMA) process is widely used to model weakly stationary time series. It is a general class of processes that can be specified by the hyper parameters p and q .

Definition 2.5 (ARMA(p, q) process). *Let $\{X_t\}$ be a zero-mean weakly stationary time series. $\{X_t\}$ is an ARMA(p, q) process if it is a weakly stationary solution to the equations*

$$\phi(B)X_t = \theta(B)Z_t, \tag{1}$$

where $\phi(B) = 1 - \phi_1 B - \phi_2 B^2 - \dots - \phi_p B^p$, $\theta(B) = 1 + \theta_1 B + \theta_2 B^2 + \dots + \theta_q B^q$ and B is the back-shift operator defined by $BX_t = X_{t-1}$. Additionally, $Z \sim \text{WN}(0, \sigma^2)$ and $\phi(B)$ and $\theta(B)$ have no common factors.

The special cases ARMA($p, 0$) and ARMA($0, q$) are the AR(p) and MA(q) process, respectively.

For most time series, $E[X_t] = \mu \neq 0$. However, the process $\{X_t - \mu\} \sim \text{ARMA}(p, q)$. These processes are called an ARMA process with mean μ . The remaining theoretical results only considers zero-mean time series without any loss of generality.

The stationarity assumption of the solution of the ARMA model implies some restriction of the possible solutions of Equation (1). If $\phi(z) \neq 0$, a candidate solution is

$$X = \frac{\theta(B)}{\phi(B)} Z_t \quad (2)$$

Let $\psi(z) = \frac{\theta(z)}{\phi(z)} = \psi_0 + \psi_1 z + \psi_2 z^2 + \dots$. The solution $X_t = \psi_0 + \psi_1 Z_{t-1} + \psi_2 Z_{t-2} + \dots$ is zero-mean since it is a linear combination of zero-mean random variables. To satisfy the definition of weak stationarity, the variance of the solution X_t needs to be finite. However, for $h \geq 1$,

$$\begin{aligned} E[X_t X_{t-h}] &= \psi_0^2 + \psi_0 \psi_1 (E[Z_t] + E[Z_{t+h}]) + \dots + \psi_h \psi_0 E[Z_{t-h} Z_{t-h}] + \dots \\ &= \psi_0^2 + \sum_{i=h}^{\infty} \psi_i \psi_{i-h} E[Z_{t-i}^2] = \psi_0^2 + \sigma^2 \sum_{i=h}^{\infty} \psi_i \psi_{i-h}, \end{aligned}$$

which may diverge if $\psi(z)$ is not absolutely summable. Assuming absolute summability of $\psi(z)$ is sufficient to obtain a stationary solution of the ARMA equations. This would also imply uniqueness of the solution. However, to verify that an ARMA process satisfies this condition can take a lot of time, which makes this criterion impractical. An equivalent, and more practical condition is the following:

$$\phi(z) \neq 0, \quad \text{for all } z \in \mathbb{N}$$

i.e., $\phi(z)$ has no zeros on the unit circle. To understand why this is equivalent, consider the Taylor expansion of $\frac{1}{\phi(z)}$ around zero. For $\delta > 0$,

$$\frac{1}{\phi(z)} = \sum_{i=-\infty}^{\infty} a_i z^i < \infty, \quad \text{for } 1 - \delta < z < 1 + \delta.$$

Since z can be made arbitrarily close to 1, the sum $\sum_{i=-\infty}^{\infty} |a_i|$ converges. Applying $\frac{1}{\phi(B)}$ to both sides of the ARMA equations shows that

$$X_t = \frac{1}{\phi(B)} \phi(B) X_t = \frac{1}{\phi(B)} \theta(B) Z_t = \psi(B) Z = \sum_{i=-\infty}^{\infty} \psi_i Z_{t-i}.$$

Thus, $\psi(B)Z$ solves Equation (1) if and only if $\phi(z) \neq 0$ for all $z \in \mathbb{N}$. The above argument shows that in most cases, the ARMA process can be expressed as a function of all the past, present and future noise components. In certain cases, X_t can be expressed uniquely using only Z_{t-h} for $h \leq 1$.

Definition 2.6 (Causal process). *The process $\{X_t\}$ is causal if the polynomials $\theta(z)$ and $\phi(z)$ have no common zeros and there exists a sequence $\{\psi_j\}_{j=0}^{\infty}$ with $\sum_{j=0}^{\infty} |\psi_j| < \infty$ such that*

$$X_t = \sum_{j=0}^{\infty} \psi_j Z_{t-j}, \quad \text{for all } t.$$

In other words, a causal process is independent of future noise. For an ARMA(p, q) process, this is equivalent to $\phi(z)$ having no zeros inside the unit circle (Brockwell et al., 2016, p. 75). Note that $\sum_{j=0}^{\infty} \psi_j Z_{t-j}$ is an MA(∞) process. Thus, any causal ARMA(p, q) process can be expressed as a MA(∞) process.

The dual property for $\{Z_t\}$ is called invertibility.

Definition 2.7 (Invertible process). *The process $\{X_t\}$ is invertible if the polynomials $\theta(z)$ and $\phi(z)$ have no common zeros and there exists a sequence $\{\pi_j\}_{j=0}^{\infty}$ with $\sum_{j=0}^{\infty} |\pi_j| < \infty$ such that*

$$Z_t = \sum_{j=0}^{\infty} \pi_j X_{t-j}, \quad \text{for all } t.$$

Equivalently, an invertible process is a process where the noise only depends on past observations. The invertibility property of an ARMA process can be shown to be equivalent to all zeros of $\theta(z)$ being outside the unit circle (Brockwell et al., 2016, p. 76).

Both these properties are highly desirable, as they simplify some calculations. One of them is the derivations surrounding large-sample prediction with ARMA models. Luckily, it can be shown that a non-causal or non-invertible ARMA process can be transformed into a causal and invertible ARMA process by finding a new white noise sequence $\{W_t\}$. Hence, without loss of generality, all future ARMA processes are assumed to be causal and invertible (Brockwell et al., 2016, p. 50). Note that for an MA(q) process, $\psi(z) = \theta(z)$, whereas for an AR(p) process, $\pi(z) = \phi(z)$. In other cases, the coefficients ψ_i and π_i can be found by comparing the sides of the following equations:

$$\begin{aligned} \psi(z)\phi(z) &= \theta(z) \\ \pi(z)\theta(z) &= \phi(z). \end{aligned}$$

The assumption that $\{X_t\}$ is weakly stationary, is important to note. In the real world, this is not always the case. For example, both time series seen in Figure 1 have upward trends, implying that the mean of these series are not constant. Even if $\{X_t\}$ is non-stationary, it is often possible to transform it to a stationary time series. The remainder of this section describe how small alterations to the ARMA process can be made to address time series with trend and seasonality.

Definition 2.8 (The ARIMA(p, d, q) process). *The time series $\{X_t\}$ is a ARIMA(p, d, q) process if*

$$(1 - B)^d X_t$$

is a causal ARMA(p, q) process.

The ARIMA(p, d, q) model can be used to transform a series with a linear trend to a stationary process, which can then be modeled by an ARMA(p, q) process. The time series for the new Covid-19 cases seems to both have a trend and a seasonal component (of period 7). This time series would not be stationary by applying a single difference operation. The **Seasonal Auto Regressive Integrated Moving Average** (SARIMA) model aims to remove both the trend and the seasonal component of $\{X_t\}$ by differencing not only in the first order, but also for order s , where s is the period of the time series.

Definition 2.9 (The SARIMA process). *The time series $\{X_t\}$ is a SARIMA(p, d, q) \times (P, D, Q) $_s$ process if the differenced series $Y_t = (1 - B)^d (1 - B^s)^D X_t$ is an ARMA process of the of the following form*

$$\phi(B)\Phi(B^s)X_t = \theta(B)\Theta(B^s)Z_t, \tag{3}$$

where $\phi(z) = 1 - \phi_1 z - \dots - \phi_p z^p$, $\Phi(z) = 1 - \Theta_1 z - \dots - \Theta_P z^P$, $\theta(z) = 1 + \theta_1 z + \dots + \theta_q z^q$ and $\Theta(z) = 1 + \Theta_1 z + \dots + \Theta_Q z^Q$.

Note that Equation (3) can be rewritten as an ARMA(pP, qQ) process (possibly with some of the coefficients being fixed to zero and others as products of previous coefficients). Future results will only consider ARMA processes, as these are notionally simpler, but these also apply to SARIMA models.

2.3 The GARCH process

For all processes defined up to this point, the mean of X_t has been conditionally dependent on its previous iterations, while the variance σ^2 has stayed constant. Sometimes it is reasonable to assume that not only the mean, but also the variance of $\{X_t\}$ changes over time. The noise component can be modeled by its own process. A popular choice is the **Generalized Auto Regressive Conditional Heteroscedasticity** (GARCH) process, developed by Bollerslev (1986). It allows the variance at the current time step to depend on previous variances and the previous iterations of Z_t .

Definition 2.10 (GARCH(p, q) process). *The GARCH(p, q) process is the weekly stationary solution to the equations*

$$Z_t = \sigma_t e_t, \quad (4)$$

for $e_t \sim \text{IID}(0, 1)$, where σ_t is defined by the recursions

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^p \alpha_i Z_{t-i}^2 + \sum_{j=1}^q \beta_j \sigma_{t-j}^2, \quad \forall t, \quad (5)$$

for $\alpha_0 > 0$ and $\alpha_i, \beta_j \geq 0$ for all i, j .

Under these assumptions, $\sigma_t^2 > 0$ for all t . When modelling, one must assume a particular distribution of e_t . This is often the standard normal distribution or Student's t-distribution. The quantity σ_t^2 is commonly referred to as the volatility of the time series. A time series where the fluctuations are followed by fluctuations of similar magnitude is said to exhibit **persistence of volatility**. In these cases, it might be wise to add a GARCH model to the underlying noise.

2.3.1 The existence of the GARCH(1,1) model

The GARCH(1, 1) model, specified in Equation (6), has been particularly popular in time series forecasting.

$$Z_t = \sigma_t^2 e_t, \quad \sigma_t^2 = \alpha_0 + \alpha_1 Z_{t-1}^2 + \beta_1 \sigma_{t-1}^2 \quad (6)$$

However, a weekly stationary solution to Equation (6) only exists for certain values of α_0, α_1 and β_1 . Observe that

$$\begin{aligned} Z_t^2 &= \sigma_t^2 e_t^2 = (\alpha_0 + \alpha_1 Z_{t-1}^2 + \beta_1 \sigma_{t-1}^2) e_t^2 = \alpha_0 e_t^2 + e_t^2 (\alpha_1 e_{t-1}^2 + \beta_1) \sigma_{t-1}^2 \\ &= \alpha_0 e_t^2 + e_t^2 (\alpha_1 e_{t-1}^2 + \beta_1) (\alpha_0 + \alpha_1 Z_{t-2}^2 + \beta_1 \sigma_{t-2}^2) \\ &= \alpha_0 e_t^2 + \alpha_0 e_t^2 (\alpha_1 e_{t-1}^2 + \beta_1) + e_t^2 (\alpha_1 e_{t-1}^2 + \beta_1) (\alpha_1 e_{t-2}^2 + \beta_1) \sigma_{t-2}^2 \\ &= \dots = \alpha_0 e_t^2 \left(1 + \sum_{k=1}^n \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right) + e_t^2 \sigma_{t-(n+1)}^2 \prod_{i=1}^{n+1} (\alpha_1 e_{t-i}^2 + \beta_1). \end{aligned}$$

This gives a candidate solution to Equation (6), namely

$$\lim_{n \rightarrow \infty} Z_{t,n} = e_t \sqrt{\alpha_0 \left(1 + \sum_{k=1}^{\infty} \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right)},$$

where

$$Z_{t,n}^2 = \alpha_0 e_t^2 \left(1 + \sum_{k=1}^n \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right).$$

However, it must be shown that this limit exists. As will be clear, this is not always the case.

Theorem 2.11. $Z_{t,n}$ is a Cauchy sequence and converges to

$$\tilde{Z}_t = e_t \sqrt{\alpha_0 \left(1 + \sum_{k=1}^{\infty} \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right)}$$

in L^2 if and only if

$$\alpha_1 + \beta_1 < 1.$$

Proof. To show that $Z_{t,n}$ is Cauchy, one needs to show that $\forall \epsilon \exists N \in \mathbb{N}$ s.t.

$$E[(Z_{t,n} - Z_{t,m})^2] < \epsilon$$

for $n, m \geq N$. Expanding this expression gives that

$$E[(Z_{t,n} - Z_{t,m})^2] = E[Z_{t,n}^2 - 2Z_{t,n}Z_{t,m} + Z_{t,m}^2] = E[Z_{t,n}^2] + E[Z_{t,m}^2] - 2E[Z_{t,n}Z_{t,m}]. \quad (7)$$

Note that

$$\begin{aligned} Z_{t,n}Z_{t,m} &= e_t \sqrt{\alpha_0 \left(1 + \sum_{k=1}^n \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right)} e_t \sqrt{\alpha_0 \left(1 + \sum_{k=1}^m \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right)} \\ &= e_t^2 \alpha_0 \sqrt{1 + \sum_{k=1}^n \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1)} \sqrt{1 + \sum_{k=1}^m \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1)} \end{aligned}$$

Without loss of generality, assume $n \geq m$. Since all terms in the above equation are positive, $Z_{t,n}Z_{t,m} \geq Z_{t,m}^2$. By monotonicity of the expected value operator, it is clear that

$$E[Z_{t,n}Z_{t,m}] \geq E[Z_{t,m}^2].$$

Plugging this result into Equation (7) gives

$$\begin{aligned} E[(Z_{t,n} - Z_{t,m})^2] &\leq E[Z_{t,n}^2] - E[Z_{t,m}^2] \\ &= E \left[\alpha_0 e_t^2 \left(1 + \sum_{k=1}^n \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right) \right] - E \left[\alpha_0 e_t^2 \left(1 + \sum_{k=1}^m \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right) \right] \\ &= \alpha_0 \sum_{k=1}^n E \left[e_t^2 \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right] - \alpha_0 \sum_{k=1}^m E \left[e_t^2 \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right] \\ &= \alpha_0 \sum_{k=m}^n E \left[e_t^2 \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right] \end{aligned}$$

The following lemma can be utilized to reduce the above expression even further.

Lemma 2.12. $E \left[e_t^2 \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right] = (\alpha_1 + \beta_1)^k$

Proof. Proof by induction:

$k = 1$:

$$E[e_t^2(\alpha_1 e_{t-1}^2 + \beta_1)] = \alpha_1 E[e_t^2 e_{t-1}^2] + \beta_1 E[e_t^2] = \alpha_1 E[e_t^2] E[e_{t-1}^2] + \beta_1 E[e_t^2] = \alpha_1 + \beta_1$$

Assume it holds for $k = m$. Then

$$\begin{aligned} E \left[e_t^2 \prod_{i=1}^{m+1} (\alpha_1 e_{t-i}^2 + \beta_1) \right] &= E \left[(\alpha_1 e_{t-m-1}^2 + \beta_1) e_t^2 \prod_{i=1}^m (\alpha_1 e_{t-i}^2 + \beta_1) \right] \\ &= E \left[(\alpha_1 e_{t-m-1}^2 + \beta_1) E \left[e_t^2 \prod_{i=1}^m (\alpha_1 e_{t-i}^2 + \beta_1) \mid e_{t-m-1} \right] \right] \\ &= E \left[(\alpha_1 e_{t-m-1}^2 + \beta_1) E \left[e_t^2 \prod_{i=1}^m (\alpha_1 e_{t-i}^2 + \beta_1) \right] \right] \\ &= E[(\alpha_1 e_{t-m-1}^2 + \beta_1)(\alpha_1 + \beta_1)^m] = (\alpha_1 + \beta_1)^m E[\alpha_1 e_{t-m-1}^2 + \beta_1] \\ &= (\alpha_1 + \beta_1)^{m+1}, \end{aligned}$$

where the second line equations used the law of iterated expectation and the fact that $e_t \sim \text{IID}(0, 1)$. ■

The above lemma implies that

$$E[(Z_{t,n} - Z_{t,m})^2] \leq \alpha_0 \sum_{k=m}^n (\alpha_1 + \beta_1)^k.$$

Since α_1 and β_1 are assumed positive, the series converges if and only if $\alpha_1 + \beta_1 < 1$. In this case, the sum can be made arbitrarily small by increasing n and m . Thus, $Z_{t,n}$ is Cauchy and the limit

$$\tilde{Z}_t = e_t \sqrt{\alpha_0 \left(1 + \sum_{k=1}^{\infty} \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right)}$$

exists. ■

The following argument shows that \tilde{Z}_t satisfies Equation (6). Let $\tilde{\sigma}_t^2 = \tilde{\sigma}_t^2 e_t$. Then

$$\tilde{\sigma}_t^2 = \alpha_0 \left(1 + \sum_{k=1}^{\infty} \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right) = \alpha_0 \left(1 + \alpha_1 e_{t-1}^2 + \beta_1 + \sum_{k=2}^{\infty} \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right).$$

Note that

$$\sum_{k=2}^{\infty} \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) = (\alpha_1 e_{t-1}^2 + \beta_1) \sum_{k=1}^{\infty} \prod_{i=1}^k (\alpha_1 e_{t-1-i}^2 + \beta_1)$$

and that

$$\tilde{\sigma}_{t-1}^2 = \alpha_0 \left(1 + \sum_{k=1}^{\infty} \prod_{i=1}^k (\alpha_1 e_{t-1-i}^2 + \beta_1) \right).$$

Shuffling the terms around in the above equation gives

$$\sum_{k=1}^{\infty} \prod_{i=1}^k (\alpha_1 e_{t-1-i}^2 + \beta_1) = \frac{\tilde{\sigma}_{t-1}^2}{\alpha_0} - 1.$$

Combining the above equations results in the following:

$$\tilde{\sigma}_t^2 = \alpha_0 \left(1 + \alpha_1 e_{t-1}^2 + \beta_1 + (\alpha_1 e_{t-1}^2 + \beta_1) \left(\frac{\tilde{\sigma}_{t-1}^2}{\alpha_0} - 1 \right) \right) = \alpha_0 + \alpha_1 e_{t-1}^2 \tilde{\sigma}_{t-1}^2 + \beta_1 \tilde{\sigma}_{t-1}^2 = \alpha_0 + \alpha_1 \tilde{Z}_{t-1}^2 + \beta_1 \tilde{\sigma}_{t-1}^2.$$

Hence $\tilde{\sigma}_t^2$ can be written as in Equation (6), and \tilde{Z}_t is a solution to the GARCH(1, 1) equations. For \tilde{Z}_t to be a valid solution, it needs to be weekly stationary. When $\alpha_1 + \beta_1 < 1$, the unconditional variance

$$\begin{aligned} E[Z_t^2] &= E \left[e_t^2 \alpha_0 \left(1 + \sum_{k=1}^{\infty} \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right) \right] = \alpha_0 + \alpha_0 \sum_{k=1}^{\infty} E \left[e_t^2 \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right] \\ &= \alpha_0 \left(1 + \sum_{k=1}^{\infty} (\alpha_1 + \beta_1)^k \right) = \alpha_0 \sum_{k=0}^{\infty} (\alpha_1 + \beta_1)^k = \frac{\alpha_0}{1 - \alpha_1 - \beta_1} < \infty. \end{aligned}$$

Further, by using the law of total expectation

$$E[Z_t] = E[Ae_t] = E[AE[e_t | e_s, s < t]] = E[AE[e_t]] = 0$$

where $A = \sqrt{\alpha_0 \left(1 + \sum_{k=1}^{\infty} \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right)}$. By the same argument,

$$E[Z_{t+h}Z_t] = E[Be_{t+h}e_t] = E[BE[e_{t+h}e_t | e_s, s < t+h]] = E[BE[e_{t+h}e_t]] = 0,$$

for $h < 0$ and with $B = \sqrt{\alpha_0 \left(1 + \sum_{k=1}^{\infty} \prod_{i=1}^k (\alpha_1 e_{t+h-i}^2 + \beta_1) \right)}$. These three qualities imply that $Z_t \sim WN(0, \frac{\alpha_0}{1 - \alpha_1 - \beta_1})$, which is a weekly stationary process.

To demonstrate that the solution is unique, denote $\tilde{\tilde{Z}}_t = \tilde{\sigma}_t e_t$ as another solution. Once again, iterating thought Equation (6) gives

$$\tilde{\tilde{Z}}_t^2 = \dots = \alpha_0 e_t^2 \left(1 + \sum_{k=1}^n \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right) + e_t^2 \tilde{\sigma}_{t-(n+1)}^2 \prod_{i=1}^{n+1} (\alpha_1 e_{t-i}^2 + \beta_1).$$

When $\alpha_1 + \beta_1 < 1$, the mean square limit

$$\begin{aligned} E \left[(\tilde{\tilde{Z}}_t - \tilde{Z}_t)^2 \right] &= E \left[e_t^2 \tilde{\sigma}_{t-(n+1)}^2 \prod_{i=1}^{n+1} (\alpha_1 e_{t-i}^2 + \beta_1) \right] \\ &= E \left[\tilde{\sigma}_{t-(n+1)}^2 E \left[e_t^2 \prod_{i=1}^{n+1} (\alpha_1 e_{t-i}^2 + \beta_1) \mid \tilde{\sigma}_{t-(n+1)}^2 \right] \right] \\ &= (\alpha_1 + \beta_1)^{n+1} E[\tilde{\sigma}_{t-(n+1)}^2] \rightarrow 0 \end{aligned}$$

as $n \rightarrow \infty$. The above derivation used the fact that $E[\tilde{\sigma}_{t-(n+1)}^2] < \infty$ since any solution to Equation (6) has to be weekly stationary. Hence, \tilde{Z}_t is the unique weekly stationary solution to the GARCH(1, 1) equations. Here is a summary of the results:

Solution of GARCH(1,1) Equations:

If $(\alpha_1 + \beta_1) < 1$, a weekly stationary solution of the GARCH(1,1) equations is

$$Z_t = e_t \sqrt{\alpha_0 \left(1 + \sum_{k=1}^{\infty} \prod_{i=1}^k (\alpha_1 e_{t-i}^2 + \beta_1) \right)}.$$

It has the properties

$$E[Z_t] = 0,$$

$$Var(Z_t) = E[Z_t^2] = \frac{\alpha_0}{1 - \alpha_1 - \beta_1}$$

and

$$E[Z_{t+h}Z_t] = 0, \text{ for } h > 0.$$

It can be shown that, in general, a weekly stationary solution to the Equation (4) and Equation (5) exists if and only if $\sum_{i=1}^p \alpha_i + \sum_{j=1}^q \beta_j < 1$ (Bollerslev, 1986).

Remark: The above derivations only showed that Z_t is unconditionally independent of Z_{t+h} for $|h| \geq 1$. This does not imply that, for example, $E[Z_t Z_{t+h} | Z_t] = 0$. After all, the whole point of the GARCH process is to allow such quantities to be nontrivial.

2.4 Parameter estimation for ARMA models

There are several ways to construct the likelihood function $L(\boldsymbol{\theta} \mid x_1, \dots, x_n)$ for an ARMA model with $\boldsymbol{\theta} = (\phi_1, \dots, \phi_p, \theta_1, \dots, \theta_q, \sigma^2)$. This section considers two alternatives: An easier (but less precise) way and a harder (but exact) way.

The **residuals** (often called innovations),

$$\epsilon_t := x_t - \phi_1 x_{t-1} - \dots - \phi_p x_{t-p} - \theta_1 \epsilon_{t-1} - \dots - \theta_q \epsilon_{t-q},$$

for $t = q + 1, \dots, n$, can be calculated from the observations. To do the calculation, $\epsilon_1, \dots, \epsilon_q$ have to be assumed, and are usually set to zero, which is their unconditional expected value. The residuals serve as estimates for the Z_{q+1}, \dots, Z_n , which are not observable. Let

$$\xi = (x_1, \dots, x_p, \epsilon_1, \dots, \epsilon_q).$$

The **conditional likelihood function** for an ARMA(p, q) model is derived by conditioning the density function on the first p observations and the first q residuals.

$$\begin{aligned} L(\boldsymbol{\theta} \mid \xi, x_{p+1}, \dots, x_n) &= f(x_{p+1}, \dots, x_n \mid \boldsymbol{\theta}, \xi) \\ &= f(x_{p+2}, \dots, x_n \mid \boldsymbol{\theta}, \xi, x_{p+1}) f(x_{p+1} \mid \boldsymbol{\theta}, \xi) \\ &\quad \vdots \\ &= f(x_n \mid \boldsymbol{\theta}, \xi, x_{p+1}, \dots, x_n) \cdots f(x_{p+1} \mid \boldsymbol{\theta}, \xi). \end{aligned}$$

Assuming $\{Z_t\}$ Gaussian, all these terms are univariate Gaussian CDFs with mean zero (by an earlier assumption of ARMA processes) and variance σ^2 . The conditional log likelihood can then be expressed as

$$l(\boldsymbol{\theta} \mid \xi, x_{p+1}, \dots, x_n) = \log(L(\boldsymbol{\theta} \mid \xi, x_{p+1}, \dots, x_n)) = \frac{n-p-1}{2} \log(2\pi) - \frac{n-p-1}{2} \log(\sigma^2) - \sum_{t=p+1}^n \frac{\epsilon_t^2}{2\sigma^2}.$$

This expression can be estimated numerically w.r.t. $\boldsymbol{\theta}$ to obtain the maximum likelihood estimates (Hamilton, 1994).

Remark: The conditional likelihood is a good measure of goodness of fit even for non-Gaussian time series (Brockwell et al., 2016, p. 140).

Assuming $Z_t \sim \text{GARCH}(1, 1)$ and $e_t \sim N(0, 1)$ gives the following distribution of X_t :

$$X_t \mid \xi, \sigma_t^2 \sim N(0, \sigma_t^2). \quad (8)$$

Let $\boldsymbol{\theta} = (\phi_1, \dots, \phi_p, \theta_1, \dots, \theta_q, \alpha_0, \alpha_1, \beta_1)$ and let $\xi = (x_1, \dots, x_p, \epsilon_1, \dots, \epsilon_m)$, with $m = \max(p, q)$. Given ξ and σ_p^2 , the recursions

$$\sigma_t^2 = \alpha_0 + \alpha_1 \epsilon_{t-1} + \beta_1 \sigma_{t-1}^2 \quad (9)$$

give σ_t for $t = p + 1, \dots, n$, and the likelihood of the above model can be expressed in the following way:

$$\begin{aligned}
L(\boldsymbol{\theta} \mid \xi, \sigma_p^2, x_{p+1}, \dots, x_n) &= f(x_{p+1}, \dots, x_n \mid \boldsymbol{\theta}, \xi, \sigma_p^2) \\
&= f(x_1, \dots, x_n \mid \xi, \boldsymbol{\theta} \sigma_p^2) \\
&= \prod_{t=p+1}^n \frac{1}{\sigma_t \sqrt{2\pi}} e^{-\frac{\epsilon_t^2}{2\sigma_t^2}},
\end{aligned}$$

Natural choices for $\epsilon_1, \dots, \epsilon_m$ and σ_p^2 are their unconditional expected values, i.e., zero and $\frac{\alpha_0}{1-\alpha_1-\beta_1}$, respectively. The log of the above likelihood can be maximized numerically w.r.t. $\boldsymbol{\theta}$. Another popular choice of distribution of e_t is Student's t-distribution. When Matlab estimates the parameters of an ARMA model, it maximizes the above likelihoods (The MathWorks, 2022a). However, by default, the conditional values are $\xi = x_0, \dots, x_{1-p}, \epsilon_0, \dots, \epsilon_{1-q}$, which are all set to zero. To obtain confidence intervals for the estimates, Matlab uses a method called the **outer product of gradients** for ARMA models (The MathWorks, 2022c).

While the conditional likelihoods were easy to derive, one must make assumptions on the first residuals and the first p observations have to be used as presamples, leading to a slight reduction in the sample size. This should not be a problem if the sample size is large, as the first couple of terms only have a small contribution to the total log likelihood. It is, however, possible to calculate an **exact likelihood function** for an ARMA(p, q) model. Let Γ be the covariance matrix for X_1, \dots, X_n , defined by $\Gamma_{ij} = \text{cov}(X_i, X_j) = E(X_i X_j)$, for $\{X_t\} \sim \text{ARMA}(p, q)$. Let $x = (x_1, \dots, x_n)^T$ be a vector containing the n zero-mean observations of $\{X_t\}$ and let $\boldsymbol{\theta} = (\phi_1, \dots, \phi_p, \theta_1, \dots, \theta_q, \sigma^2)$. Assuming $Z_t \sim N(0, \sigma^2)$, the exact likelihood is

$$L(\boldsymbol{\theta} \mid x_1, \dots, x_n) = (2\pi)^{-\frac{n}{2}} \det(\Gamma)^{-\frac{1}{2}} \exp\left(-\frac{1}{2} x^T \Gamma^{-1} x\right)$$

If n is large, it can be too computer intensive to find Γ^{-1} . Additionally, the determinant of Γ may explode or vanish, leading to numerical problems. For these reasons, some more work is needed to derive a easily computable exact likelihood. One clever way to do this is through a **lower unit triangular decomposition** (LUTD) of Γ :

$$\Gamma = ADA^T,$$

where A is a lower triangular matrix with ones on the diagonal and D is a diagonal matrix.

Remark: Since Γ is a covariance matrix, it is also positive semi definite. It can be shown that all real positive semi definite matrices can be decomposed by the Cholesky decomposition $\Gamma = BB^T$, where B is a lower triangular matrix (not necessarily with ones on the diagonal) of non-negative numbers (Higham, 2009). Let $B = AD^{\frac{1}{2}}$, for A and D as defined above. Then it is easy to verify that the Cholesky decomposition is equivalent to the lower unit triangular decomposition. Thus, all covariance matrices can be decomposed in the above fashion.

For an ARMA(p, q) process, the LUTD simplifies the calculation of Γ^{-1} (defined as the inverse of all the elements on its diagonal) and $\det(\Gamma)$, since $\Gamma_{i1} = 0$ for $i > q + 1$. This implies that $A_{i,1} = 0$ for $i > q + 1$. The other columns of A has similar characteristics, resulting in the following matrix:

$$A = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 & 0 \\ a_{21} & 1 & 0 & \cdots & 0 & 0 \\ a_{31} & a_{32} & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \cdots & \vdots & \vdots \\ a_{q+1,1} & a_{q+1,2} & a_{q+1,3} & \cdots & 0 & 0 \\ 0 & a_{q+2,2} & a_{q+2,3} & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \cdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & a_{n,n-1} & 1 \end{bmatrix}. \quad (10)$$

Note that

$$\det(\Gamma) = \det(ADA^T) = \det(ADA^T) = \det(A)\det(D)\det(A^T) = \det(A)^2\det(D) = \det(D)$$

and that $(AA^T) = I_n$, where I_n is the $n \times n$ identity matrix. Thus, the likelihood can be rewritten as

$$L(\theta | x_1, \dots, x_n) = (2\pi)^{-\frac{n}{2}} \det(D)^{-\frac{1}{2}} \exp\left(-\frac{1}{2}(Ax)^T D^{-1} Ax\right)$$

Since D is a diagonal matrix, this is much simpler to maximize than the former expression of the likelihood. There are several ways to construct A and D . It can, for example, be done via the innovations algorithm, described in (Brockwell et al., 2016, p. 87). It is not clear how to do this exactly if $Z_t \sim \text{GARCH}(p, q)$. The series $\{\sigma_t^2\}$ has to start somewhere, but conditioning on an early iteration of σ_t^2 defeats the purpose of an exact likelihood. It might be possible to utilize the weak stationarity property of $\{Z_t\}$ in a similar fashion as what is done with $\{X_t\}$. This is outside the scope of this thesis.

2.5 Prediction with ARMA models

The previous section shows how to obtain good estimates for the parameters in the ARMA model. In this section, it is assumed that the parameters of the ARMA model are known with certainty. Otherwise, the uncertainty of the estimates has to be accounted for when calculating the prediction error.

The goal of time series analysis is often to find the best possible prediction of X_1, \dots, X_{n+h} based on the observations x_1, x_2, \dots, x_n . This requires more concrete notation. Let $P_n Y = f(X_n, X_{n-1}, \dots)$ be the prediction operator. It is the function that minimizes the **mean squared error** (MSE)

$$E[(Y - P_n Y | X_n, X_{n-1}, \dots)^2].$$

It can be shown that $P_n Y = E[X_{n+1} | X_n, X_{n-1}, \dots]$ minimizes the MSE (Brockwell et al., 2016, p. 34). Ideally, the task should be to find this function. In practice, however, this task is generally too difficult, and the objective is instead to find

$$\tilde{P}_n Y = \hat{E}[X_{n+1} | X_n, X_{n-1}, \dots],$$

the best linear predictor of Y based on X_n, X_{n-1}, \dots . Thus, $\tilde{P}_n Y$ is on the form $\sum_{i=1}^n c_i X_{n-i}$. Note that if Z_t is indeed Gaussian, $\tilde{P}_n = P_n$. The innovations algorithm gives a way to find $\tilde{P}_n X_{t+h}$ exactly (Brockwell et al., 2016, p. 87). It is unclear how this algorithm works if the noise is assumed to be $\text{GARCH}(p, q)$. In this thesis, the programming language Matlab was used to generate predictions from the ARMA models. Matlab uses a **large sample prediction procedure** to predict with ARMA models (The MathWorks, 2022b). For a causal and invertible ARMA process, the following identities hold:

$$X_{n+h} = \sum_{i=0}^{\infty} \psi_i Z_{n+h-i}$$

$$Z_{n+h} = X_{n+h} + \sum_{i=1}^{\infty} \pi_i Z_{n+h-i}$$

Applying \tilde{P}_n on both sides of both equations gives

$$\tilde{P}_n X_{n+h} = \sum_{i=h}^{\infty} \psi_i Z_{n+h-i} \quad (11)$$

$$\tilde{P}_n X_{n+h} = - \sum_{i=1}^{\infty} \pi_i \tilde{P}_n X_{n+h-i}, \quad (12)$$

since $\tilde{P}_n Z_{n+h-i} = 0$, for $i \leq h$, and $\tilde{P}_n Z_{n+h-i} = Z_{n+h-i}$, for $i > h$. Note that $\tilde{P}_n X_{n+h-i} = X_{n+h-i}$ for $i > h$. Since X_1, \dots, X_n are observed, Equation (12) defines a recursive way to obtain $\tilde{P}_n X_{t+h}$. Given the infinite past of X_n , Equation (12) gives an exact prediction of X_{n+h} . However, good approximations can be made by truncating the sum after n observations when the coefficients $|\pi_i|$ are small and n is large. The above equations give a neat expression of $X_{n+h} - \tilde{P}_n X_{n+h}$, which is the first step towards calculating the MSE of the forecast:

$$X_{n+h} - \tilde{P}_n X_{n+h} = \sum_{i=0}^{h-1} \psi_i Z_{n+h-i}.$$

For an ARMA process with $Z_t \sim WN(0, \sigma^2)$, the MSE of the forecast after h steps is

$$E[(X_{n+h} - \tilde{P}_n X_{n+h})^2 | X_n, X_{n-1}, \dots] = \sigma^2 \sum_{i=0}^{h-1} \psi_i^2. \quad (13)$$

If the coefficients $|\psi_i|$ are small, this quickly converges when h is large enough. When assuming $Z_t \sim \text{GARCH}(1, 1)$, Equation (12) still holds. However, the MSE changes in the following way:

$$E[(X_{n+h} - \tilde{P}_n X_{n+h})^2 | X_n, X_{n-1}, \dots] = \sigma_t^2 \sum_{i=0}^{h-1} \psi_i^2 = (\alpha_0 + \alpha_1 \epsilon_{t-1} + \beta \sigma_{t-1}^2) \sum_{i=0}^{h-1} \psi_i^2$$

where σ_t^2 can be calculated by the recursions described in Equation (9). Matlab uses Equation (12) to generate predictions and use Equation (13) (or Equation (2.5) with GARCH noise) to generate the associated prediction error (The MathWorks, 2022b).

2.6 Simulation with ARMA models

Simulation can be used to study the properties of a specific ARMA(p, q) process

$$X_t - \phi_1 X_{t-1} - \dots - \phi_p X_{t-p} = Z_t - \theta_1 Z_{t-1} - \dots - \theta_q Z_{t-q}$$

with known parameters $\phi_1, \dots, \phi_p, \theta_1, \dots, \theta_q, \sigma^2$. Assume that the task is to simulate a realization $\tilde{X}_{n+1}, \dots, \tilde{X}_{n+h}$ from an ARMA(p, q) based on the observations x_1, \dots, x_n . The residuals, ϵ_t , for $q < t \leq n$, function as estimators for the unobserved Z_{q+1}, \dots, Z_n , assuming $Z_t = 0$, for $t \leq q$. This is not a wild assumption as it is their unconditional expectation (with or without GARCH noise). Simulating from X_{n+1} reduces to simulating from Z_{n+1} . In principle, no assumption of the particular distribution of Z_t needs to be assumed for an ARMA process. There are at least two ways to deal with this. The first way is by bootstrapping the already observed residuals $\epsilon_{q+1}, \dots, \epsilon_n$. The second method is to assume a distribution on Z_t . The most popular choices are the Gaussian distribution and Student's t-distribution. These distributions can easily be simulated from. The latter simulation method, with an assumed normal distribution, was used to simulate ARMA processes in this thesis, as this is the default approach in Matlab (The MathWorks, 2022d). Once Z_{n+1} is simulated, the simulated value of X_{n+1} easily follows from the ARMA equations given the observations and the residuals. Repeating this procedure for Z_{n+2} , gives a simulated realization of X_{n+2} , and so on. Thus, it is possible to simulate as far as necessary into the future. If the noise is assumed to

be a GARCH(p, q) process, Z_{n+1} is simulated from the GARCH process, where the residuals $\epsilon_{q+1}, \dots, \epsilon_n$ are treated similarly as the observations where above. This time, a particular distribution of e_t has to be assumed, where the default, Gaussian distribution was chosen.

2.7 Model Diagnostics for ARMA models

It is useful to qualify how good a certain ARMA model is fitted to the data. Most notably, this makes it possible to compare different model configurations against each other, to find the best model for the data.

To quantify how well the model fits the observed time series data, the **Akaike information criterion** (AIC) can be used. Denote the maximum likelihood estimates as $\hat{\theta}$. The AIC for a model is defined as

$$AIC(\hat{\theta}) = -2l(\hat{\theta}) + 2k,$$

where k is the number of parameters fitted. As an example, for a SARIMA(p, d, q) \times (P, D, Q)_s process, $k = p + q + P + Q$. The **corrected AIC** (AICc) is a bias-corrected version of the AIC and is defined as

$$AICc(\hat{\theta}) = AIC(\hat{\theta}) + \frac{2k(k+1)}{n-k-1},$$

where n is the sample size used to fit the model. In the case of small sample size, the AICc is usually recommended, as this is when the bias is most prevalent. Note that the AICc converges to the AIC when a large sample size is available. The AICc represents the information lost by the model compared to the actual data. Thus, a lower value indicates a better fit.

The **auto-correlation function** (ACF) for a time series $\{X_t\}$ is defined as

$$\alpha(h) = Cor(X_t, X_{t+h})$$

A plot of the ACF is a nice way to visualize the dependencies of a time series and to choose the order of q for an ARMA model. The values of h are displayed on the x-axis and are called lags. Each lag has the associated value $\alpha(h)$ on the y-axis. 95% confidence bands are also included in the plot. Thus, if for example, the value at lag 3 is significant, it may be wise to fit an ARMA model with $q \geq 3$ to account for this. In this thesis, the order of the models were already decided, and the ACF plot was only used to assess the fit of the SARIMA models. This was done by plotting the ACF of the residuals $\{\epsilon_t\}$ of fitted models.

The residuals of a time series model should ideally be independent, as this means that all the dependencies in the data are explained by the model. By looking at the ACF plot, one can assess whether the residuals are uncorrelated. In this case, the lags of the ACF plot should not be significantly different from zero. However, uncorrelated residuals do not imply independence. If, in particular, the squares of the residuals show significant lags, the residuals still exhibit some serial dependence. Applying a GARCH model to the ARMA model might solve this problem, as it models the dependence between the variances of $\{Z_t\}$.

Engel's ARCH test is another way examine if the residuals exhibit conditional heteroscedasticity. Let $\phi(B)X_t = \theta(B)Z_t$, for $Z_t = \sigma_t^2 e_t$, with $e_t \sim IID(0, 1)$. Then

$$Var(X_t | X_{t-1}, X_{t-2}, \dots) = Var(Z_t | X_{t-1}, X_{t-2}, \dots) = E[Z_t^2 | X_{t-1}, X_{t-2}, \dots] = \sigma_t^2,$$

since $E[Z_t Z_{t+h}] = 0$. Thus, to identify heteroscedasticity, one only need to look at the auto-correlation process. Even if the model accounts for all the correlation of $\{X_t\}$, the variables in $\{\epsilon_t\}$ might still be dependent. By expressing the squared residuals as a AR(m) process

$$\epsilon_t^2 = \alpha_1^2 \epsilon_{t-1}^2 + \dots + \alpha_m^2 \epsilon_{t-m}^2 + w_t, \quad w_t \sim WN(0, 1), \quad (14)$$

it is clear that homoscedasticity is obtained only when $\alpha_0 = \dots, \alpha_m = 0$. This is exactly what is done in Engel's ARCH test, where the null hypothesis is

$$H_0 : \alpha_0 = \dots, \alpha_m = 0$$

for α_j as in Equation (14), while the alternative hypothesis is

$$H_a : \epsilon_t^2 = \alpha_1^2 \epsilon_{t-1}^2 + \dots + \alpha_m^2 \epsilon_{t-m}^2 + w_t,$$

with $\alpha_j \neq 0$ for at least one $j = 1, \dots, m$. Under H_0 , it can be shown that $\epsilon_t^2 \sim \chi_m^2$ (Engle, 1982). In Matlab, the amount of significant lags to test for (given by m in Equation (14)) can be specified (The MathWorks, 2022e). The ARCH test can also be used to test for GARCH(p, q) processes, since it is locally equivalent to an ARCH($p + q$) process.

2.8 The Machine Learning Model

Machine learning models usually have a large number of parameters, which in principle gives them the flexibility to model any function. Machine learning models have recently shown good results in the field of time series forecasting, and have become a contender to the more traditional statistical time series models (Kiranyaz et al., 2021). This section covers two types of machine learning models, that will later be combined and used to predict future daily new cases for the two Covid-19 data sets. It should be noted that the details of each operation described in this section can be modified by the user. This section merely describes how the methods were implemented in this thesis, using an API called TensorFlow.

2.8.1 Convolutional Neural Networks

Convolutional neural networks (CNN) have increased in popularity over the last decades. Their first big achievements were in the field of image classification (Kiranyaz et al., 2021). Image data usually have lots and lots of features (one for each pixel). Using traditional fully connected networks to do classification based on all these features is time consuming, as too many weights must be tuned, and usually results in overfitting. The CNNs solve these problems by having sparse connections between the first layers of neurons.




Figure 2: An illustration of a one-dimensional kernel operation with kernel size is 3 (TowardsDataScience, n.d.).

Recently, CNNs has shown great strides on one-dimensional data, e.g., time series data (Kiranyaz et al., 2021). The key to modeling time series data is to capture the dependence between each time step. CNNs do exactly this, since they create new features by convolving neighbouring data points together, as illustrated by Figure 2. With TensorFlow in the programming language Python, there are several settings specifying the exact nature of the convolutional operation. The following definition only addresses the operation with default settings, as this is what will be used later (Wei Zhang et al., 2017):

Definition 2.13 (Convolution in one-dimension). *Let $x = (x_1, \dots, x_n)$ be the observed time series and let K be a $s \times 1$ convolutional kernel with corresponding bias b . Define $Y_i = (x_i, \dots, x_{i+s})$, for $i = 1, \dots, n - s - 1$. The one-dimensional convolution operation can be defined in the following way:*

$$\begin{aligned} F_1 &= f(Y_1 K + b) \\ F_2 &= f(Y_2 K + b) \\ &\vdots \\ F_{n-s+1} &= f(Y_{n-s+1} K + b), \end{aligned}$$

where $f(\cdot)$ is called an activation function and $F = (F_1, \dots, F_{n-s+1})$ is the convolved time series, often called a feature map.

The activation function is usually non-linear. If not, it would be impossible for the CNN to fit a non-linear function to the data, as the convolution operation is linear. Recently, because of its efficiency, the most popular activation function has been the rectified linear unit (ReLU), i.e., the function $f(x) = \max(x, 0)$. However, other functions have their use cases. The activation function completes a convolutional layer.

A convolutional network consists of one or more convolutional layers, exemplified in Figure 3. Each layer may have multiple filters and each filter is defined by its associated kernel. The convolutional operation from Definition 2.13 happens in the filters, where each filter gives a mapping to the next layer of the network. This

results in sparse connections between the layers of the network, which may reduce overfitting. Additionally, the feature maps are in some sense shift-invariant, since each element is a function all neighbouring points of the input time series.




Figure 3: An illustration of a one-dimensional convolutional neural network (Kiranyaz et al., 2021).

The elements inside a kernel are called weights and are tuned through backpropagation. The goal of backpropagation is to minimize a loss function by tuning the weights in the opposite way of its gradient (Z. Zhang, 2016). The choice of loss function depends on the problem.

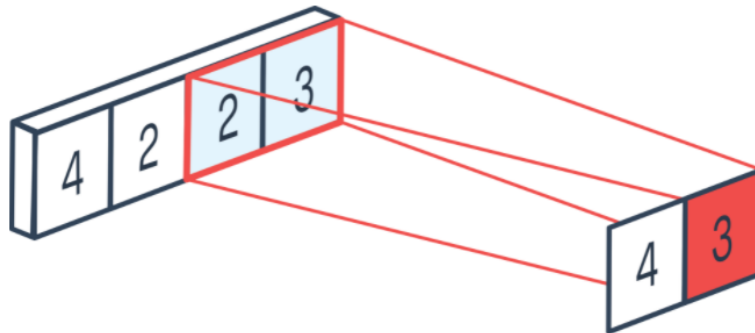



Figure 4: An illustration of a one-dimensional pooling operation, with size 2 and stride 2 (Peltarion, 2022).

The feature maps vary from filter to filter. To reduce the variation between the feature maps, it is common to add a pooling layer after a convolutional layer. They process the incoming feature map in a similar fashion as the convolutional layers. However, instead of applying a kernel section wise, it computes a summarizing statistic over each section, as shown in Figure 4. The dropout layer has a similar purpose. At a specified rate r , the neurons are set to zero, while the remaining neurons are multiplied by $\frac{1}{1-r}$, to keep the sum of the neurons unchanged (TensorFlow, 2022a). This is a regularization technique that forces the model to consider more features when training.



(a) A visualization of how samples are construction for time series (b) The samples are stored in two matrices, the input matrix and the data (MachineLearningMastery, n.d.).
 output matrix (J. Zhang et al., 2021)

Figure 5: Construction of samples from time series data.

To tune the weights of a one-dimensional CNN, the input series has to be reformatted into samples of input-output pairs, as Figure 5 illustrates. The input matrix contains subsequences of the input time series of the same length in each row, and the output matrix contains the ensuing time point(s). For example, if the task is to use the previous n days to predict the next m days, then each row of the input matrix should contain n elements, while the output matrix should contain the next m elements of the input time series. This gives the network input-output samples to learn from and can perform the task at hand once it is trained. Each filter is applied row wise to the input matrix to produce the feature maps. The predictions from the network (one for each row) can then be compared to the respective row in the output matrix, and the weights can be tuned accordingly.

The output of the CNN can then be used as input in a different model, that can interpret the newly processed data. A popular choice in time series modeling is to apply a Long-Short Term Memory (LSTM) model to the convolved data (Luan et al., 2019).

2.8.2 Long Short Term Memory models

The LSTM model is a modified version of a Recurrent Neural Network (RNN) and is a popular option for time series forecasting. A traditional RNN can be seen in Figure 6. The model takes in the input matrix $x = (x_1, \dots, x_n)$ and returns the output matrix $o = (o_1, \dots, o_n)$. Note that x is not the observed time series. For one-dimensional data, x could either be the input matrix created from the transformed data, with each x_i being a sub sequence of observed data (as described in Figure 5), or it could be the output of the previous layer (for example a CNN). The neurons are connected across time steps through the internal state vectors $s = (s_1, \dots, s_n)$. The mathematical description of a RNN is the following (W. Zhang et al., 2019):

$$\begin{aligned} s_t &= g_h (Ux_t + Ws_{t-1} + b_U + b_W) \\ o_t &= g_o (Vs_t + b_V), \end{aligned} \quad (15)$$

where U , V and W are weight matrices and g_h and g_o are activation functions, corresponding to the hidden and output layer, respectively. Excluded from Figure 6 are all the bias vectors, b_U , b_V and b_W corresponding to U , V and W , respectively.




Figure 6: An illustration of a RNN (W. Zhang et al., 2019).

Just as for CNNs, all weight matrices are trained by backpropagation. However, for RNNs, the backpropaga-

tion algorithm also has to traverse all time steps $t = 1, \dots, n$. This is no problem for short sequences of data, however, for longer sequences, the gradient comprises of more and more multiplicative factors. Consider the backpropagation of the weights in W . Let $\mathcal{L}_t := \mathcal{L}(x_t - o_t)$ be the contribution to the loss function at time step t . The loss function is the sum of all its contributions $L := \sum_{t=1}^n \mathcal{L}_t$. Since this is purely a demonstration, the basic gradient decent algorithm was used to calculate the gradient (Ruder, 2016). The error gradient w.r.t. W is (Pascanu et al., 2012):

$$\frac{\partial L}{\partial W} = \sum_{t=1}^n \frac{\partial \mathcal{L}_t}{\partial W}.$$

The gradient of the contribution to the loss at time step t , $\frac{\partial \mathcal{L}_t}{\partial W}$, is propagated back to each time step:

$$\begin{aligned} \frac{\partial \mathcal{L}_t}{\partial W} &= \underbrace{\frac{\partial \mathcal{L}_t}{\partial o_t} \frac{\partial o_t}{\partial s_t} \frac{\partial s_t}{\partial W}}_{\text{contribution to the loss at time } t \text{ from internal state } s_t} + \dots + \underbrace{\frac{\partial \mathcal{L}_t}{\partial o_t} \frac{\partial o_t}{\partial s_t} \frac{\partial s_t}{\partial s_{t-1}} \frac{\partial s_{t-1}}{\partial s_{t-2}} \dots \frac{\partial s_1}{\partial W}}_{\text{contribution to the loss at time } t \text{ from internal state } s_1} \\ &= \frac{\partial \mathcal{L}_t}{\partial o_t} \frac{\partial o_t}{\partial s_t} \frac{\partial s_t}{\partial W} + \sum_{k=2}^t \frac{\partial \mathcal{L}_t}{\partial o_k} \frac{\partial o_k}{\partial s_k} \frac{\partial s_k}{\partial s_{k-1}} \prod_{i=0}^{t-k-1} \frac{\partial s_{t-i}}{\partial s_{t-i-1}} \frac{\partial s_k}{\partial W}. \end{aligned}$$

Thus, the total loss from n observations is

$$\begin{aligned} \frac{\partial L}{\partial W} &= \sum_{t=1}^n \sum_{k=1}^t \frac{\partial \mathcal{L}_k}{\partial o_k} \frac{\partial o_k}{\partial s_k} \frac{\partial s_k}{\partial s_{k-1}} \prod_{i=0}^{t-k-1} \frac{\partial s_{t-i}}{\partial s_{t-i-1}} \frac{\partial s_k}{\partial W} \\ &= \sum_{t=1}^n \sum_{k=1}^t \frac{\partial \mathcal{L}_k}{\partial o_k} \frac{\partial o_k}{\partial s_k} \frac{\partial s_k}{\partial s_{k-1}} \prod_{i=0}^{t-k-1} \left[\frac{\partial}{\partial s_{t-i-1}} g_h(Ux_t + Ws_{t-i-1} + b_U + b_W) \right] \frac{\partial s_k}{\partial W} \\ &= \sum_{t=1}^n \sum_{k=1}^t \frac{\partial \mathcal{L}_k}{\partial o_k} \frac{\partial o_k}{\partial s_k} \frac{\partial s_k}{\partial s_{k-1}} W^{t-k} \prod_{i=0}^{t-k-1} g'_h(Ux_t + Ws_{t-i-1} + b_U + b_W) \frac{\partial s_k}{\partial W}. \end{aligned}$$

The long term contributions to the gradient suffer from the growing exponent in the weight matrix W . If not all weights are equal to unity, this effectively results in a vanishing or an exploding contribution to the full gradient. The same problem occurs when backpropagating U and the biases b_W and b_U . This illustrates why the traditional RNN struggles to learn long term dependencies. This is known as the vanishing or exploding gradient problem.

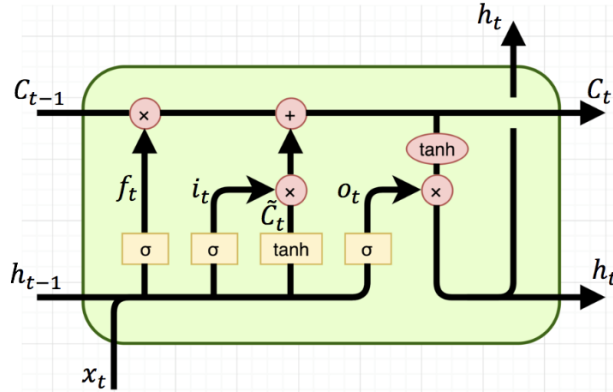


Figure 7: An illustration of a LSTM cell (Varsamopoulos et al., 2018).

A remedy to this problem is the LSTM model, which is illustrated in Figure 7. It uses memory cells to decide what information to forget from the previous cells, and what new information to keep. This information is then sent through to the internal state, just as for traditional RNNs. Let $x = (x_1, \dots, x_n)$ be the same input matrix as described earlier. The following set of equations define the LSTM model (Varsamopoulos et al., 2018):

$$\begin{aligned}
f_t &= g_1(W_f \cdot [h_{t-1}, x_t] + b_f) \\
i_t &= g_1(W_i \cdot [h_{t-1}, x_t] + b_i) \\
o_t &= g_1(W_o \cdot [h_{t-1}, x_t] + b_o) \\
\tilde{C}_t &= g_2(W_C \cdot [h_{t-1}, x_t] + b_C) \\
C_t &= f_t * C_{t-1} + i_t * \tilde{C}_t \\
h_t &= o_t * g_2(C_t),
\end{aligned} \tag{16}$$

where g_1 and g_2 are activation functions.

In the above equation, the forget gate f_t controls what information from C_{t-1} to "forget" based on the new information from h_{t-1} and x_t . The input gate at time t , called i_t , specifies what new information should be passed through to the cell state C_t at the current time step. The output gate o_t controls what information from C_t to output at time step t . The output vector is denoted h_t . It is used as input in the cell corresponding to the next time step, $t + 1$.

The vectors C_t, C_{t-1}, o_t and h_t have the same length, often referred to as the number of *units* u of the LSTM-layer. The number of units has to be specified during implementation. Similar to other machine learning models, the weight matrices W_* and the bias vectors b_* are randomly initialized and are trained through **back propagation in time**.

The construction of the LSTM-cell results in an additive structure of the factors of the gradient. Let $L = \sum_{t=1}^n \mathcal{L}_t$ be the loss function. To show that the LSTM model solve the vanishing and exploding gradient problem, it has to be shown that all partial gradients manage the problem. The derivation of the gradient for W_f, W_i and W_C and the corresponding biases are all similar to that of W in a RNN. For W_f , the gradient is

$$\frac{\partial L}{\partial W_f} = \sum_{t=1}^n \sum_{k=1}^t \frac{\partial \mathcal{L}_k}{\partial h_k} \frac{\partial h_k}{\partial C_k} \frac{\partial C_k}{\partial C_{k-1}} \prod_{i=0}^{t-k-1} \frac{\partial C_{t-i}}{\partial C_{t-i-1}} \frac{\partial C_k}{\partial f_k} \frac{\partial f_k}{\partial W_f}.$$

However, the terms $\frac{\partial C_{t-i}}{\partial C_{t-i-1}}$ from the LSTM model behaves nicer than $\frac{\partial s_{t-i}}{\partial s_{t-i-1}}$ from the RNN:

$$\frac{\partial C_{t-i}}{\partial C_{t-i-1}} = \frac{\partial f_{t-i}}{\partial C_{t-i-1}} C_{t-i-1} + \frac{\partial C_{t-i-1}}{\partial C_{t-i-1}} f_{t-i} + \frac{\partial i_{t-i}}{\partial C_{t-i-1}} \tilde{C}_{t-i} + \frac{\partial \tilde{C}_{t-i}}{\partial C_{t-i-1}} i_{t-i},$$

where

$$\begin{aligned}
\frac{\partial f_{t-i}}{\partial C_{t-i-1}} &= W_f \frac{\partial h_{t-i-1}}{\partial C_{t-i-1}} g'_1(W_f \cdot [h_{t-i-1}, x_{t-i}] + b_f) \\
&= W_f \cdot o_{t-i-1} \cdot C_{t-i-1} \cdot g'_2(C_{t-i-1}) \cdot g'_1(W_f \cdot [h_{t-i-1}, x_{t-i}] + b_f),
\end{aligned}$$

$$\begin{aligned}
\frac{\partial i_{t-i}}{\partial C_{t-i-1}} &= W_i \frac{\partial h_{t-i-1}}{\partial C_{t-i-1}} g'_1(W_i \cdot [h_{t-i-1}, x_{t-i}] + b_i) \\
&= W_i \cdot o_{t-i-1} \cdot C_{t-i-1} \cdot g'_2(C_{t-i-1}) \cdot g'_1(W_i \cdot [h_{t-i-1}, x_{t-i}] + b_i),
\end{aligned}$$

and

$$\begin{aligned}\frac{\partial \tilde{C}_{t-i}}{\partial C_{t-i-1}} &= W_C \frac{\partial h_{t-i-1}}{\partial C_{t-i-1}} g'_1(W_C \cdot [h_{t-i-1}, x_{t-i}] + b_C) \\ &= W_C \cdot o_{t-i-1} \cdot C_{t-i-1} \cdot g'_2(C_{t-i-1}) \cdot g'_1(W_C \cdot [h_{t-i-1}, x_{t-i}] + b_C).\end{aligned}$$

Hence

$$\begin{aligned}\frac{\partial C_{t-i}}{\partial C_{t-i-1}} &= W_f \cdot o_{t-i-1} \cdot C_{t-i-1} \cdot g'_2(C_{t-i-1}) \cdot g'_1(W_f \cdot [h_{t-i-1}, x_{t-i}] + b_f) \cdot C_{t-i-1} \\ &\quad + f_{t-i} \\ &\quad + W_i \cdot o_{t-i-1} \cdot C_{t-i-1} \cdot g'_2(C_{t-i-1}) \cdot g'_1(W_i \cdot [h_{t-i-1}, x_{t-i}] + b_i) \cdot \tilde{C}_{t-i} \\ &\quad + W_C \cdot o_{t-i-1} \cdot C_{t-i-1} \cdot g'_2(C_{t-i-1}) \cdot g'_1(W_C \cdot [h_{t-i-1}, x_{t-i}] + b_C) \cdot i_{t-i}.\end{aligned}$$

Since f_t has direct influence on the value of $\frac{\partial C_{t-i}}{\partial C_{t-i-1}}$ and has a direct update mechanism, the model can better regulate the long term gradients. This enables the model to learn long term dependencies in the direction of W_f .

The weights W_i , W_C and the biases b_f , b_i and b_C can be calculated in a similar way:

$$\frac{\partial L}{\partial W_i} = \sum_{t=1}^n \sum_{k=1}^t \frac{\partial \mathcal{L}_k}{\partial h_k} \frac{\partial h_k}{\partial C_k} \frac{\partial C_k}{\partial C_{k-1}} \prod_{i=0}^{t-k-1} \frac{\partial C_{t-i}}{\partial C_{t-i-1}} \frac{\partial C_k}{\partial i_k} \frac{\partial i_k}{\partial W_i}.$$

$$\frac{\partial L}{\partial W_C} = \sum_{t=1}^n \sum_{k=1}^t \frac{\partial \mathcal{L}_k}{\partial h_k} \frac{\partial h_k}{\partial C_k} \frac{\partial C_k}{\partial C_{k-1}} \prod_{i=0}^{t-k-1} \frac{\partial C_{t-i}}{\partial C_{t-i-1}} \frac{\partial C_k}{\partial \tilde{C}_k} \frac{\partial \tilde{C}_k}{\partial W_C}.$$

$$\frac{\partial L}{\partial b_f} = \sum_{t=1}^n \sum_{k=1}^t \frac{\partial \mathcal{L}_k}{\partial h_k} \frac{\partial h_k}{\partial C_k} \frac{\partial C_k}{\partial C_{k-1}} \prod_{i=0}^{t-k-1} \frac{\partial C_{t-i}}{\partial C_{t-i-1}} \frac{\partial C_k}{\partial f_k} \frac{\partial f_k}{\partial b_f}.$$

$$\frac{\partial L}{\partial b_i} = \sum_{t=1}^n \sum_{k=1}^t \frac{\partial \mathcal{L}_k}{\partial h_k} \frac{\partial h_k}{\partial C_k} \frac{\partial C_k}{\partial C_{k-1}} \prod_{i=0}^{t-k-1} \frac{\partial C_{t-i}}{\partial C_{t-i-1}} \frac{\partial C_k}{\partial i_k} \frac{\partial i_k}{\partial b_i}.$$

$$\frac{\partial L}{\partial b_C} = \sum_{t=1}^n \sum_{k=1}^t \frac{\partial \mathcal{L}_k}{\partial h_k} \frac{\partial h_k}{\partial C_k} \frac{\partial C_k}{\partial C_{k-1}} \prod_{i=0}^{t-k-1} \frac{\partial C_{t-i}}{\partial C_{t-i-1}} \frac{\partial C_k}{\partial \tilde{C}_k} \frac{\partial \tilde{C}_k}{\partial b_C}.$$

In each of these error gradient, the long term memory is ensured by the forget gate. Finally, the gradients for the output gate are the following:

$$\frac{\partial L}{\partial W_o} = \sum_{t=1}^n \sum_{k=1}^t \frac{\partial \mathcal{L}_k}{\partial h_k} \frac{\partial h_k}{\partial o_k} \frac{\partial o_k}{\partial W_o} = \sum_{t=1}^n \sum_{k=1}^t \frac{\partial \mathcal{L}_k}{\partial h_k} g_2(C_k) \frac{\partial o_k}{\partial W_o}$$

$$\frac{\partial L}{\partial b_o} = \sum_{t=1}^n \sum_{k=1}^t \frac{\partial \mathcal{L}_k}{\partial h_k} \frac{\partial h_k}{\partial o_k} \frac{\partial o_k}{\partial b_o} = \sum_{t=1}^n \sum_{k=1}^t \frac{\partial \mathcal{L}_k}{\partial h_k} g_2(C_k) \frac{\partial o_k}{\partial b_o},$$

which is no problem since it does not contain the multiplicative term. This shows that, in principle, the gradient manages to learn long term dependencies in each direction.

The series of vectors $[h_1, \dots, h_n]$ is the output of the model, but the default behaviour in TensorFlow (2022b) is to return only the last vector h_n , since this vector are associated with x_{n+1} . This is regulated

by the input parameter `return_sequences`. It is important to remark that h_n is not supposed to be the prediction of X_{n+1} . Remember that h_t are vectors of length u . A fully connected layer is often used after the LSTM layer(s) to interpret h_n , and to return a single predicted value for each ensuing day. This forecast can then be compared to the observed values, to compute the model's accuracy.

2.9 Measures of accuracy

To compare the performance of each model, a measure of the accuracy of the forecasts from each model is needed. Two accuracy measures are considered in this thesis. The first measure is the **relative root mean squared error** (RRMSE). Given a n -day prediction \hat{y} with corresponding test data y , the RRMSE is calculated by

$$\text{RRMSE}(\hat{y}, y) = \frac{\sqrt{(1/n) \sum_{i=1}^n (\hat{y}_i - y_i)^2}}{(1/n) \sum_{i=1}^n |y_i|} \times 100\%$$

RRMSE has the advantage of being comparable between data sets of different scales. By dividing the root mean squared error by the total scale of the testing data, a data set with millions of new cases of Covid-19, like the global data set, can be compared with models for the Norwegian data set. According to Zain et al. (2021) and Li et al. (2013), the accuracy of a forecast is deemed excellent when the RRMSE is less than 10 percent, good when the RRMSE is between 10 and 20 percent, fair when the RRMSE is between 20 and 30 percent, and poor when the RRMSE is greater than 30 percent.

The **mean absolute percentage error** (MAPE) is calculated in the following way:

$$\text{MAPE}(\hat{y}, y) = \frac{1}{n} \sum_{i=1}^n \frac{|\hat{y}_i - y_i|}{|y_i|} \times 100\%$$

The MAPE is similar to the RRMSE in that it divides the error by the values in the observed test data. It can therefore, in principle, be compared across data sets with differing scales. The difference between the MAPE and the RRMSE is that the division happens inside the summation, and that the absolute value is used to measure the deviation between the prediction and the observed value. This means that extreme outliers do not contribute as much to the MAPE as it would when using the RRMSE, especially when accompanied with higher values in the test data.

3 Methods

This section specifies the three models used in this thesis and gives the necessary details to replicate the upcoming results.

3.1 Defining the three models

The previous section described general results for $\text{ARMA}(p, q)$ models, CNNs and LSTM models. The three specific models used in this thesis are now introduced.

3.1.1 The SARIMA model

In the winter of 2020, Taraldsen (2020b) published a so-called toy model that fit the Covid-19 new-cases data set from Norway nicely. After transforming the new-cases time series $\{Y_t\}$:

$$X_t = \log(\max(Y_t, 0.1)), \quad \text{for } t = 1, \dots, n,$$

he fitted the following model to $\{X_t\}$:

$$(1 - B)(1 - B^7)X_t = (1 - \theta_{MA}B)(1 - \Theta_{SMA}B^7)Z_t, \quad Z_t \sim N(0, \sigma^2), \quad (17)$$

i.e., a $\text{SARIMA}(0, 1, 1) \times (0, 1, 1)_7$ model. The model is completely defined by the three parameters θ_{MA} , Θ_{SMA} and σ^2 . The model will later be referred to as **the SARIMA model**.

Remark: An equivalent formulation of the above SARIMA model is to model the differenced time series $\{\tilde{X}_t\} = (1 - B)(1 - B^7)X_t$ by the following MA(8) process:

$$\tilde{X}_t = (1 - \theta_{MA}B)(1 - \Theta_{SMA}B^7)Z_t. \quad (18)$$

A precondition for fitting the SARIMA model to $\{X_t\}$ is that the time series $\{\tilde{X}_t\}$ is stationary. This is explored in Section 4.2.

3.1.2 The Gandalf model

The SARIMA model fits surprisingly well to the Norwegian data set, but Taraldsen discovered that the squared residuals had a significant correlation at lag 1 (replicated in Figure 15a). To incorporate this into the model, he used the same SARIMA model, but with the assumption that $Z_t \sim \text{GARCH}(1, 1)$. Section 2.3.1 established that $\alpha_1 + \beta_1 < 1$. Additionally, Taraldsen fixed $\alpha_0 = 0.001$, because of numeric problems when α_0 approached zero. The model is described by the following equations:

$$\begin{aligned} (1 - B)(1 - B^7)X_t &= (1 - \theta_{MA}B)(1 - \Theta_{SMA}B^7)Z_t, \\ Z_t &= \sigma_t^2 e_t, \quad \sigma_t^2 = 0.001 + \alpha_1 Z_{t-1}^2 + \beta_1 \sigma_{t-1}^2. \end{aligned}$$

This model is also a SARIMA model. However, to not confuse the reader, this model is called the **Gandalf model**. Sometimes it is useful to refer both the Gandalf model and the SARIMA model together. The plural **SARIMA models** will be used for this purpose.

3.1.3 The CNN-LSTM model

At the time of publishing (July 23rd 2021), a hybrid between a CNN and a LSTM model gave one of the best predictions of the global new-cases data set (Zain et al., 2021). An outline of the model can be seen in Figure 8. First observe that the model is designed to take sequences of new cases of seven consecutive days as input, and outputs the amount of new cases on the ensuing eight day. The input-output pairs are constructed accordingly. The seven-day sequences are passed through two one-dimensional convolutional layers with 64 filters each, with all filters having kernel size 3. Both convolutional layers use ReLU as activation function. After the convolutional layers, a pooling layer of size 1 is added, followed by a dropout layer. However, the dropout rate was not unspecified.

The remaining filters were then flattened into one long vector, which was used as input to the LSTM layer with 200 units. The default activation functions were used, i.e., $g_1(\cdot) = \text{sigmoid}(\cdot)$ and $g_2(\cdot) = \text{tanh}(\cdot)$ (TensorFlow, 2022b). Finally, there are two fully connected layers; the number of units in the first of these layers was not specified and was set to 50, while the second one has 1 unit to be able to output one single value. The fully connected layers used their default activation functions, i.e., ReLU.

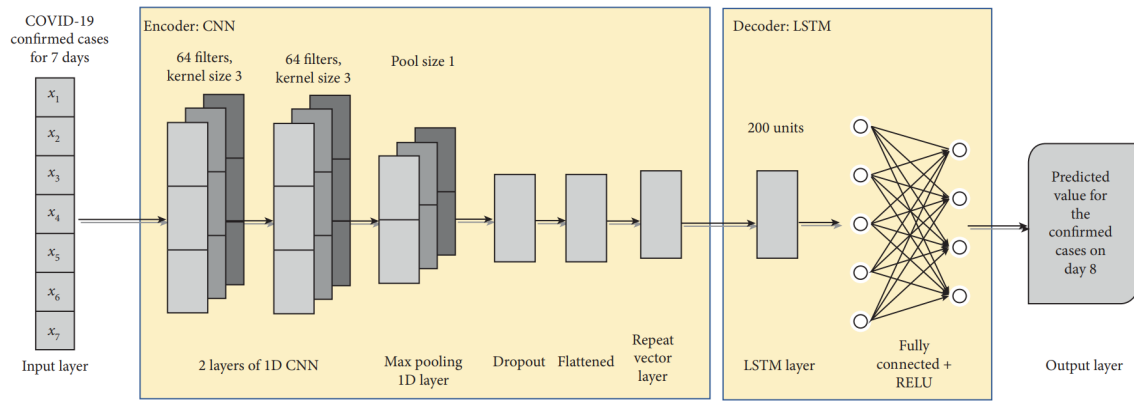


Figure 8: The architecture of the specific CNN-LSTM model from Zain et al. (2021).

It should be noted that the observed new-cases data $x = (x_1, \dots, x_n)$ was min-max-normalized, i.e., the transformation

$$\text{min-max}(x_i) = \frac{x_i - \min(x)}{\max(x) - \min(x)},$$

was applied to x . The inverse transformation was done to the final prediction, using the same min and max value. Some sort of normalization of the input data is commonplace in machine learning models, as it transforms all features to the same scale.

Adaptive Moment Estimation (ADAM), a modified version of gradient decent, was used for training the model. The MSE was used as loss function. The learning rate was not specified, hence the default learning rate for ADAM was chosen, i.e., 0.001 (TensorFlow, 2022c). The batch size of the model specifies the number of input-output pairs to train on before updating the weights, while the number of epochs is defined as the number of times the whole training data set is fed to the network. Zain et. al used a batch size of 22 with 472 epochs. The choice was based on a grid search algorithm using the Optuna framework. The same number of epochs was initially adopted to the model in this thesis, and the (unspecified) dropout rate was set to 0.2 from initial testing.

The model implemented in this thesis had 347109 weights (or parameters) in total. Hence, it is quite a complicated function. On the other hand, at the end of the day, it is simply a function that takes in seven values, and outputs a single value. Mathematically, the general modeling task can be described in the following way: Considering each eight-day sequence (x_i, \dots, x_{i+7}) of observations, find a function $\mathcal{F}(x_i, \dots, x_{i+6} | P) = \hat{X}_{i+7}$ that minimizes the loss function

$$l(\mathcal{F}(x_i, \dots, x_{i+6} | P) - x_{i+7}),$$

w.r.t. the large parameter space P , for all $i = 1, \dots, n - 7$. Both the Covid-19 data sets have clear weekly patterns, with higher prevalence in the weekdays than the weekends. The two SARIMA models were constructed with this in mind. Can a model with such a simple problem description learn this pattern as well? This is explored in Section 4.

Remark: Fitting a CNN-LSTM model will not guarantee the same parameter estimates each time since the fitting procedure involves a lot of randomness: Firstly, the parameters are randomly initialized before the training begins. Secondly, the dropout layer randomly excludes 20% of the data in each batch in each epoch. However, for an already fitted model, the prediction is deterministic, as the dropout rate is excluded when performing this operation (TensorFlow, 2022a).

Remark: Each of these stochastic parts has their own seed. To ensure reproducibility, all these seeds must be fixed. The first function in the provided Python script, called `set_all_seeds`, fixes all the necessary seed values, and is applied at the beginning of each predictive function. Even after doing this, the results may vary depending on software and hardware. The exact specifications used to produce the results in this thesis are soon provided.

3.2 On the different ways of forecasting

While a general method for forecasting with the SARIMA models were established in Section 2.5, a forecasting method for the CNN-LSTM model has yet to be specified. In this thesis, two different methods were used to forecast with each of the models.

Taraldsen and Zain et. al. generated the forecast from their models in fundamentally different ways. Taraldsen used the SARIMA models to predict the next seven days directly, as in Section 2.5, while the CNN-LSTM model published by Zain et. al. predicted 28 days into the future. However, the exact method used to predict these 28 days was not specified. Judging by the accuracy of the model prediction, and the fact that the prediction seems to be as accurate for the final days of the prediction as for the first few days (see Figure 7 in the article by Zain et al. (2021)), the model has probably been given updated information of the actual test data along the way. An alternative is to update the model with the newest prediction, which is more similar to how the SARIMA models generate their h -step prediction directly. This section specifies these two prediction schemes.

3.2.1 Prediction scheme 1

A trained CNN-LSTM model takes seven consecutive amounts of new cases from the Covid-19 data set and returns the predicted amount on the eighth, ensuing day. To predict the new cases on day nine, the amounts for the seven preceding days are needed as input. When it is specified that the CNN-LSTM model uses prediction scheme 1, the most recent input of the model comes from the test data, i.e., the actual value of the day that was previously predicted. In Figure 9, only the grey-colored squares are used as input in for prediction, and the red-colored x es were originally part of the test data. Hence, they were therefore and was therefore not seen by the model under training. Thus, the model gets updated information before generating the next prediction. In this fashion, a series of one-step predictions can be used to generate a h -step prediction. Note that h is limited to the amount of test data.

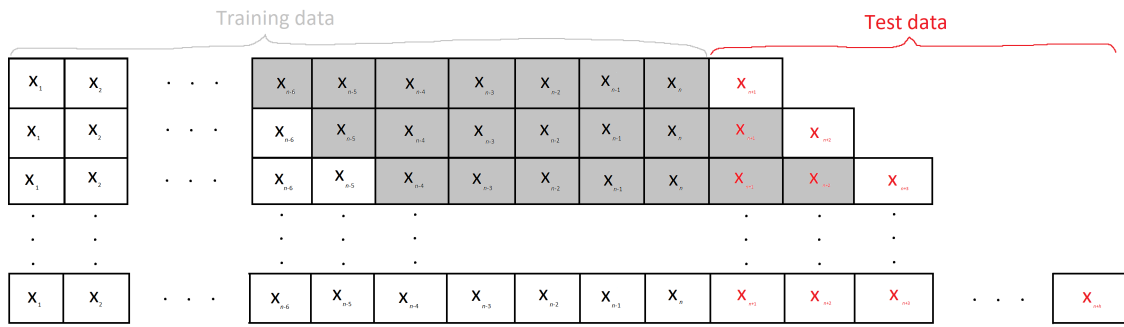


Figure 9: An illustration of prediction scheme 1 for the CNN-LSTM model.

The SARIMA models are much less computer intensive than the CNN-LSTM model. For this reason, it was decided that when the SARIMA models used prediction scheme 1, they make a series of one-step predictions, and fit a new model with all the available data at each time step. This is reflected in Figure 10, where the only difference from Figure 9 is that all data points preceding the next value to be forecasted are highlighted in grey.

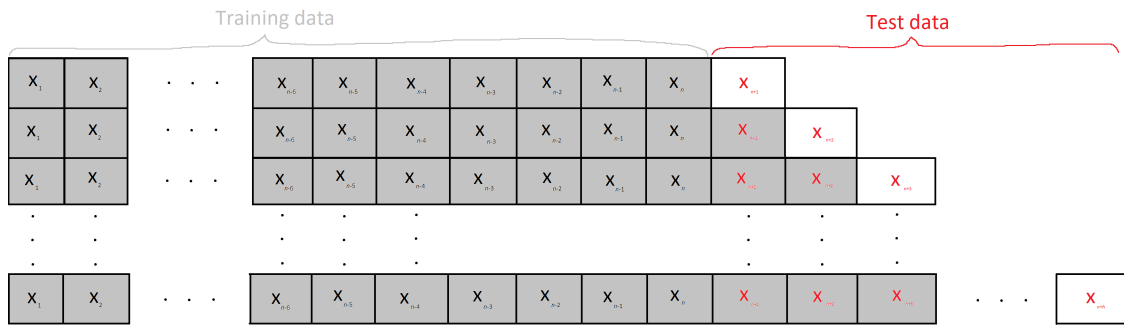


Figure 10: Illustration of prediction scheme 1 for the SARIMA models.

3.2.2 Prediction scheme 2

Using prediction scheme 1 implied that the trained models performed a series of one-step predictions with up-to-date information. When it is specified that prediction scheme 2 is used, the models are never exposed to the test data. For the CNN-LSTM model, a series of one-step predictions are executed to perform a h -step prediction. However, the predictions themselves are used as input to predict the next day, as opposed to the test data. This is illustrated in Figure 11, since the red test data is replaced by green \hat{X}_t s. When prediction scheme 2 is used for the SARIMA models, they simply perform the h -step prediction described in Section 2.5. This is similar to the CNN-LSTM model's prediction scheme 2 since the predictions replace the observations from the training data as the prediction unfolds (see Equation (12)). Eventually, the input of the forecasts using prediction scheme 2 are solely predicted values. By construction, this scheme should result in less precise predictions, and the prediction error should rapidly increase, as explained by Equation (13) for the SARIMA models. This is a more realistic prediction scheme for the task at hand, as one often wants to predict multiple steps ahead, to have time to form a response based on it.

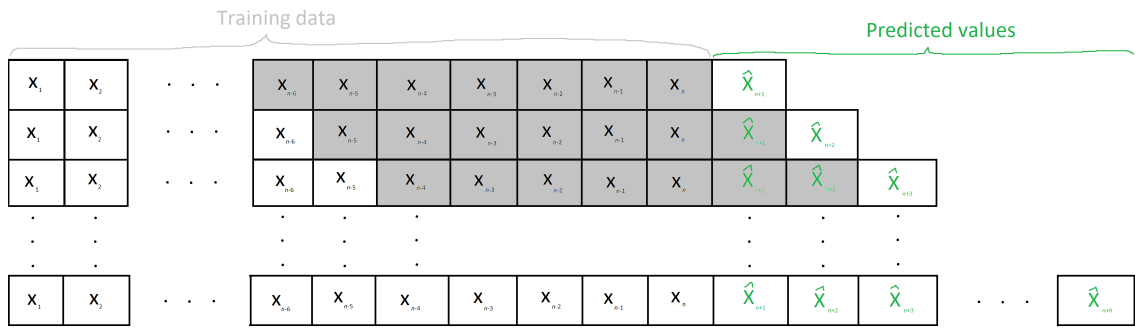


Figure 11: Illustration of prediction scheme 2 for the CNN-LSTM model.

3.3 Programming languages

Matlab was used to implement the SARIMA models and to create all the plots in the thesis. Python was used for the CNN-LSTM model, along with TensorFlow. Finally, R was utilized to pre-process the two data sets. The relevant hardware and software specifications are given in Table 1. All code used to produce the result of this thesis can be found in Appendix B.

Table 1: Relevant computer specifications.

Processor	Intel Core i7-3820
CPU	3.60 GHz
Operating system	Windows 10 Pro
Matlab version	R2022a
Python version	3.7
TensorFlow version	2.8
R version	3.6.1

4 Applying the models on Covid-19 data

In this section, the models were applied to the two Covid-19 data sets. Initially they were applied on the training and test sets used in the two introductory articles published by Zain et. al. and Taraldsen, to hopefully replicate their results. Concrete definitions of these data sets will now be given. **Taraldsen's data set** denotes the training set Taraldsen used to present the two SARIMA models. It consists of new cases in Norway from February 21st 2020 to November 10th 2020, which amounts to 264 days. The CNN-LSTM model was originally evaluated on **Zain's data set**, i.e., the global new cases from January 4th 2020 to July 17th, which is 196 days of data. The models were later compared on other partitions of both the Norwegian and the global data set. Finally, a simulation study was conducted to explore the behaviour of the models in more depth.

4.1 Data

The two data sets used in this thesis were introduced in Section 1. Some additional details about the data gathering process are now given.

The worldwide new cases data set, provided by WHO (2022), contains daily new Covid-19 cases from over 200 countries. A provided R-script was constructed to sum up the new cases on each day for each country, and can be found in Appendix B. This data set is called the global data set. The last day included in both data sets was February 20th 2022. The Norwegian data set started on February 21st 2020, while the global data set begun some weeks earlier, on January 4th. This amounts to 731 and 779 days of data for the two models, respectively. However, the last week of the data sets are especially uncertain and was ignored. Zain et. al. reference the same data set from WHO. According to their article, their data set contained new cases from 216 countries. However, at the time of writing, WHO's new cases data set contains data from 236 countries. This may cause minor deviations in the results compared to the results reported by Zain et. al.

4.2 Exploring the fit of the SARIMA models on the Taraldsen's data set

ARMA models assume stationary input data. To approach homoscedasticity, the input data for the SARIMA model and the Gandalf model was log-transformed, preceded by a $\max(\cdot, 0.1)$ operation to deal with zero valued inputs. The two difference operators $(1 - B)(1 - B^7)$ was applied to the log-transformed data. As remarked in Section 3.1, this transformed time series should be stationary to fulfil the SARIMA model assumptions. Figure 12 shows the result of applying these transformations to the Norwegian and the global data set. The first row shows the two time series. A rapid increase of the daily new cases can be observed in the last few months of both time series. For the log-transformed time series in the next row, the variance is drastically reduced, and the rapid increase in the final months of the data sets is not as notable. On the contrary, a fast upward trend can be observed the earlier parts of both data sets. Both time series in the last row of the figure look quite stationary, which implies that the difference operations did what they are supposed to. The exception is at some of the earlier dates, where both time series have a much larger fluctuations than for later dates. These fluctuations seem to be conditional on previous fluctuations since high fluctuations are followed by high fluctuations and vice versa.

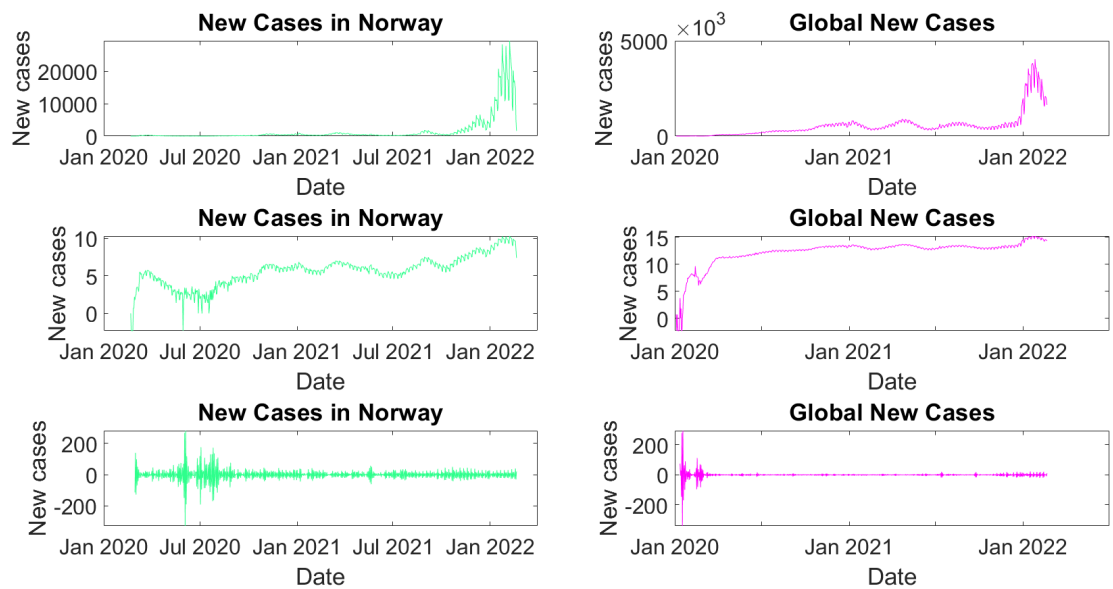




Figure 12: The Norwegian and the global new cases time series with transformations.

Figure 13 shows the seven-day forecasts from the SARIMA models when fitted on Taraldsen’s data set. The plot includes the predictions and the prediction intervals for the SARIMA model and the Gandalf model. Like Taraldsen, prediction scheme 2 was used to generate the seven-day forecasts of new cases from November 11th to November 18th. The accuracy measures and the parameters estimates have been included in the plot. The forecast from the Gandalf model is slightly closer to the observed values than for the SARIMA model, as indicated by both measures of accuracy. The prediction interval corresponding to the Gandalf model is narrower than the SARIMA model’s, almost by a factor of four. The parameter estimates for both models were also included and are close to Taraldsen’s estimates. These minuscule differences can be explained by small adjustments in the data set provided by FHI since Taraldsen extracted the data set. Figure 2 in Taraldsen’s article was used for reference.



(a) seven-day forecast of Norwegian data from November 11 2020



(b) Zoomed version of (a). The upper limit of the prediction interval for the SARIMA model is not in frame

Figure 13: Replication of results from Taraldsen (2020b) with the SARIMA model and the Gandalf model.

The better performance of the Gandalf model on the Taraldsen's data set suggests that the underlying noise during this partition of the Norwegian data set is closer to GARCH(1,1) noise than Gaussian noise. This was explored further by looking at the ACF plot of the residuals in Figure 14.

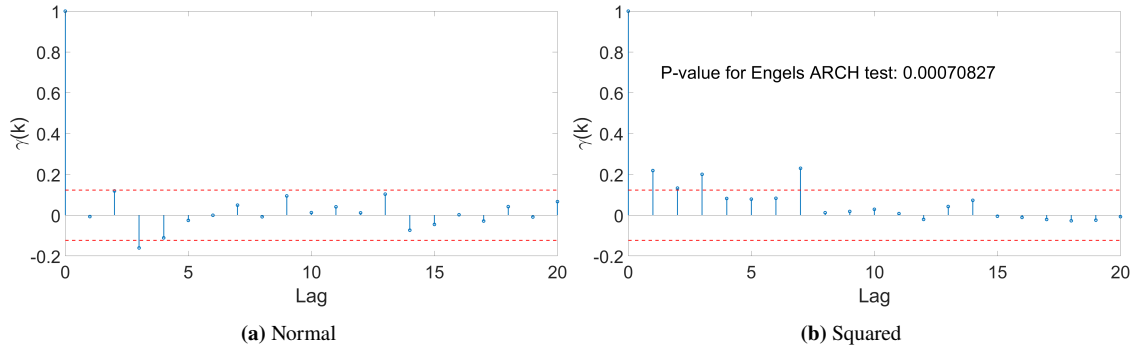


Figure 14: ACF plot of residuals and squared residuals for SARIMA model on Taraldsen’s data set.

The ACF plot of the residuals in Figure 14(a) barley shows one significant lag, indicating that the SARIMA model fits the data well. Some lack of fit can be observed in Figure 14(b). The p-value produced by Engel’s ARCH test with two significant lags (to test for GARCH(1, 1) noise) is lower than 0.05. These results suggest conditional heteroscedasticity in the training data, and motivates the application of the Gandalf model to the original Norwegian data set. The ACF plots of the squared residuals from the Gandalf model can be seen in Figure 15(a). The plot displays three significant lags, specifically at lag 1, 3 and 14. While they are barely significant, this is worse than for the SARIMA model. Additionally, the p-value generated by the ARCH test on the residuals is even lower than its SARIMA equivalent.

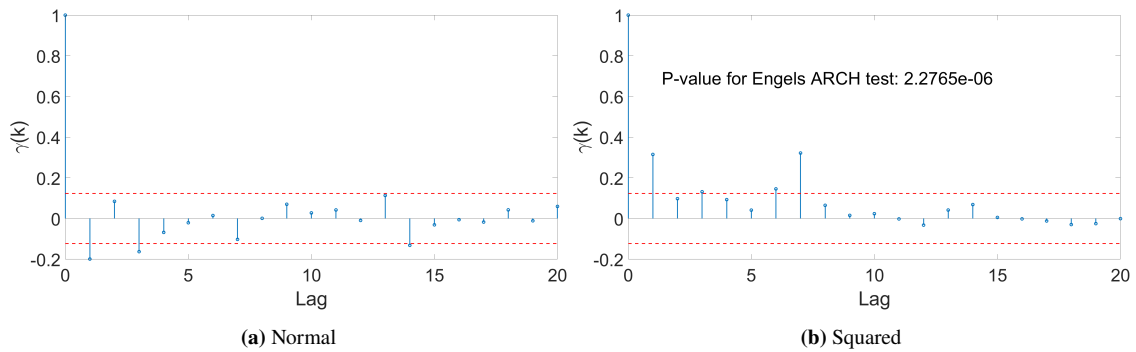


Figure 15: ACF plot of residuals and squared residuals for Gandalf model on Taraldsen’s data set.

4.3 Exploring the fit of the CNN-LSTM model on Zain’s data set

In the article from Zain et al. (2021), the CNN-LSTM model was initially trained in Zain’s data set. As remarked in Section 3.1.3, the process of fitting a CNN-LSTM model involves randomness. Thus, the predictions may vary each time the model was fitted. To account for this, Zain et. al fitted their model ten times. With each of the ten fitted models, a 28-day forecast with prediction scheme 1 was generated. From these forecasts, ten RRMSE scores and MAPE scores can be calculated. Zain et. al. reported the mean of these ten RRMSE scores and of the ten MAPE scores to be 5.30 and 0.19, respectively. By inspection of the prediction and the test data from Figure 7 in their article, the mean MAPE score seems way too low, and was not used for reference in this thesis.




Figure 16: Mean loss function of ten equally specified CNN-LSTM models, trained on Zain’s data set.

Initially, 472 epochs were run to fit the ten CNN-LSTM model. However, this resulted in a mean RRMSE of 8.35, which is worse than Zain et. al. The loss function for the models might explain why. Figure 16 shows the mean of the ten loss functions (one for each model). The red line is the value of the mean loss function at 100 epochs. Observe that the mean loss quickly decreases to around 0.0015. While it continues to decrease slightly afterwards, using all the 472 epochs might cause overfitting. To prevent this, the number of epochs for all CNN-LSTM models used in this thesis was reduced to 100. A more practical benefit of this change was that the time used to fit each model was substantially reduced. Figure 17 shows the predictions of the ten CNN-LSTM models as specified above, together with a RRMSE and MAPE score for each model. The mean RRMSE and MAPE from the ten models is also included. Observe that the mean RRMSE is lower than what Zain et. al reported.




Figure 17: Forecast with ten CNN-LSTM models from July 18th to August 14th 2020 on the global new-cases data with prediction scheme 1. The models were trained on Zain’s data set.

4.4 Comparing the three models

All three models performed well on the training and test sets used in their respective articles. This was to be expected; after all, they were constructed for those particular data sets. This section explores how good the models predict on their opposing original data sets, how they perform across other partitions of the data sets, and the effect on the models' performances with reduced samples sizes.

4.4.1 Comparisons on Taraldsen's data set

Initially, ten CNN-LSTM model were trained on Taraldsen's data set and used to forecast the new cases on the seven ensuing days. The forecasts are shown in Figure 18. Observe that the ten forecasts of the new cases on November 18th are more spread out than the foretasted values on November 11th. This is expected when using prediction 2.




Figure 18: Forecast with ten CNN-LSTM models from November 11th to November 18th 2020 on the Norwegian new-cases data with prediction scheme 1. Taraldsen's data set was used for training.

To compare the SARIMA models with the CNN-LSTM models, it would have been possible to fit ten of each SARIMA models and generate one forecast for each of them. This would give ten forecasts for each model. However, since the fitting and prediction of the SARIMA models are deterministic operations (when a deterministic numerical optimization algorithm is used), this is not meaningful for these models. It is also impractical to have ten forecast as opposed to one. There are many possible ways to display one single prediction from the ten CNN-LSTM models. A tempting choice was to choose the CNN-LSTM model with the best RRMSE and MAPE. However, this can only be determined when the test data is known. Hence, a measure of central tendency of some sort was a better alternative. The mean was the measure of choice, calculated as the mean of each of the ten predicted values of each day. This will later be called **the mean CNN-LSTM prediction**. The accuracy of this single prediction must not be confused with the mean of the ten accuracy measures.




Figure 19: Forecast from the three models with prediction scheme 2, from November 11th to November 18th 2020 on the Norwegian new-cases data. Taraldsen’s data set was used for training.

The three seven-day forecasts generated from the SARIMA model, the Gandalf model, and the mean CNN-LSTM prediction, can be seen in Figure 19. Note that the upper limit of the prediction intervals for the SARIMA model and the Gandalf model are mostly outside the plot. More importantly, note that there is no prediction interval for the machine learning model. The Gandalf model performed best with respect to the two accuracy measures, and the CNN-LSTM model was more accurate than the SARIMA model.

Zain et. al. presented the CNN-LSTM model with a 28-day prediction using prediction scheme 1. One such prediction was generated by all three models, after they were trained on Taraldsen’s data set. These 28-day forecasts are shown in Figure 20. The Gandalf model gave the most accurate forecast, and the SARIMA model trailed closely behind. The CNN-LSTM model was worse than its competitors. The ten individual predictions from the CNN-LSTM model can be found in Appendix A (Figure 41).




Figure 20: Forecast from the three models with prediction scheme 1, from November 11th to December 8th 2020 on the Norwegian new-cases data. Taraldsen’s data set was used for training.

The time to fit and predict with each model may also be an important factor to consider when differentiating the models. To fit a single CNN-LSTM model to the above training data, Python used around 7.3 seconds, while generating a seven-day and 28-day forecast from a pretrained model took around 1.27 and 2.55 seconds, respectively. Each of these processes are repeated ten times. The SARIMA model took 1.3 seconds to fit and used 0.04 and 14.39 seconds to generate seven and 28-day forecasts on the fitted model, respectively. To fit the Gandalf model on Taraldsen’s data set took around 1.30 seconds, while the forecasts took 0.05 and 37.90 seconds, respectively. The large deviation from the plain SARIMA model is a result of the slightly longer time to fit. When using prediction scheme 1, this difference increases with the length of the forecast, since they fit a new model on each consecutive day. The results are summarized in Table 2. Note that these time measurements depend on a lot of factors, like the amount of samples in the data and the hardware used. However, for the sample sizes used in this thesis, the relative differences in time between the models did not change notably.

Table 2: Time used to fit and forecast with each of the models, trained on Taraldsen’s data set.

Model	Time to fit	Seven-day forecast	28-day forecast
SARIMA	0.32 seconds	0.04 seconds	14.39 seconds
Gandalf	1.30 seconds	0.05 seconds	37.90 seconds
CNN-LSTM	64.57 seconds	12.67 seconds	25.45 seconds

When using prediction scheme 1 to predict the 28 days ensuing Taraldsen’s data set (i.e., November 11th to December 8th 2020), the SARIMA models estimate a new set of parameters each time a new value from the test data is available. Figure 21 displays the change in parameter estimates for each of the parameters of the Gandalf model. No significant deviation from the estimates obtained on the original training set could be observed in any of the four parameters. This is not that surprising, since adding 28 days to the original 264 days is a relatively small addition. A possible linear upward trend can be observed for the estimated α_1 , mirrored by a linear decrease in the β_1 estimates. These may eventually deviate significantly from their original estimates.

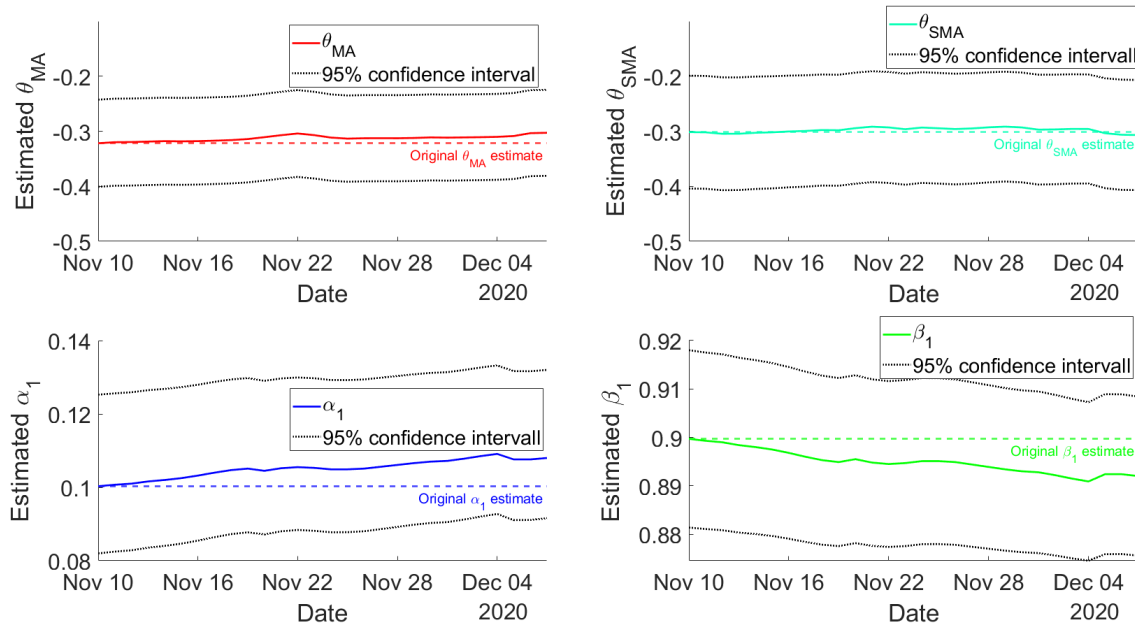


Figure 21: The parameter estimates from the Gandalf models used to predict the 28 days shown in Figure 20. The dates along the x-axis is the last day used to train the model used to predict the day after. The colored dotted lines are the parameter estimates of the Gandalf model fitted on Taraldsen’s data set.

Figure 22 displays the equivalent plot for the SARIMA model. While the estimates for θ_{MA} and θ_{SMA} are stable, the σ^2 estimates follows a clear linear trend, eventually reaching a significant difference from the original σ^2 estimate.

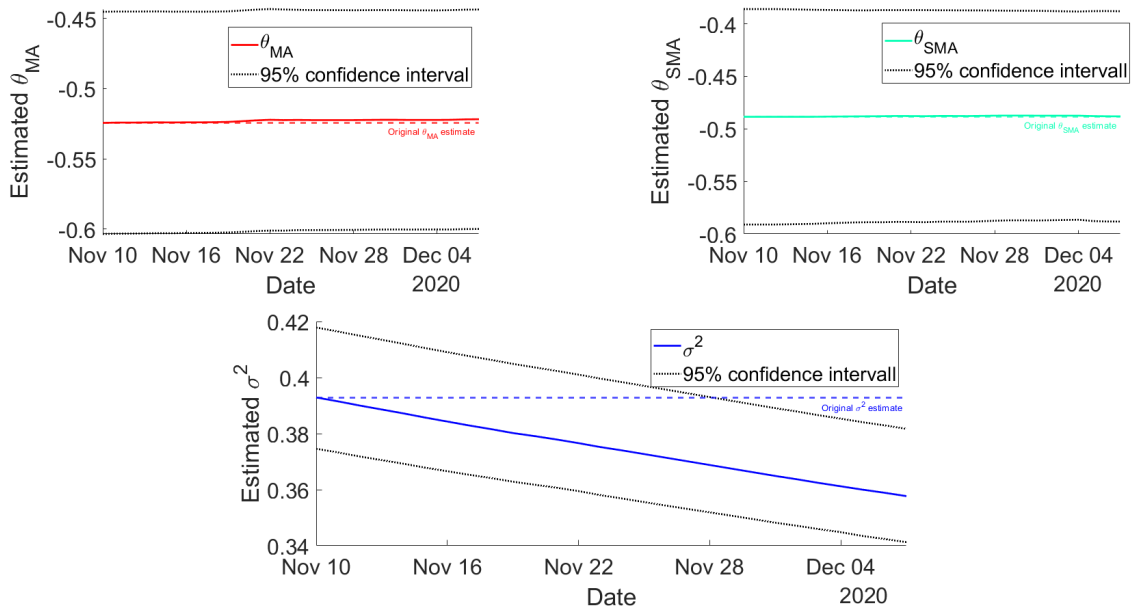


Figure 22: The parameter estimates from the SARIMA models used to predict the 28 days shown in Figure 20. The dates along the x-axis is the last day used to train the model used to predict the day after. The colored dotted lines are the parameter estimates of the SARIMA model fitted on Taraldsen’s data set.

4.4.2 Comparisons on Zain’s data set

The two SARIMA models were tested on Zain’s data set. Figure 23 shows the seven-day forecasts from the SARIMA models, while Figure 24 shows the ACF plots of their residuals and their squared residuals. Both models were more accurate than they were on Taraldsen’s data set, and the SARIMA model scored slightly lower than the Gandalf model on both measures. Some of the parameter estimates are quite different from Taraldsen’s data set. In particular, the estimated variance for the SARIMA model were almost twice as high, and the estimated θ_{MA} and θ_{SMA} are quite different between the Gandalf models. The SARIMA model is more accurate than the Gandalf model, but it also has a much wider prediction interval.




Figure 23: Forecast from the SARIMA models with prediction scheme 2 from July 18th to July 24th on the global data set. Both models were trained on Zain’s data set.

The ACF plot in Figure 24(a) of the residuals from the SARIMA model showed several significant lags, suggesting that some of the correlation in the data have not been captured by the model. The ACF of the Gandalf-residuals (Figure 24(c)) were similar. When the residuals contain first order dependencies, it is not surprising that the squared residuals are serially dependent (see Figure 24(b) and Figure 24(d)). The added GARCH(1, 1) model for the noise component was not enough to account for this. Engel's ARCH test supports these inferences, with p-values close to 0 for both the SARIMA and the Gandalf model's residuals.

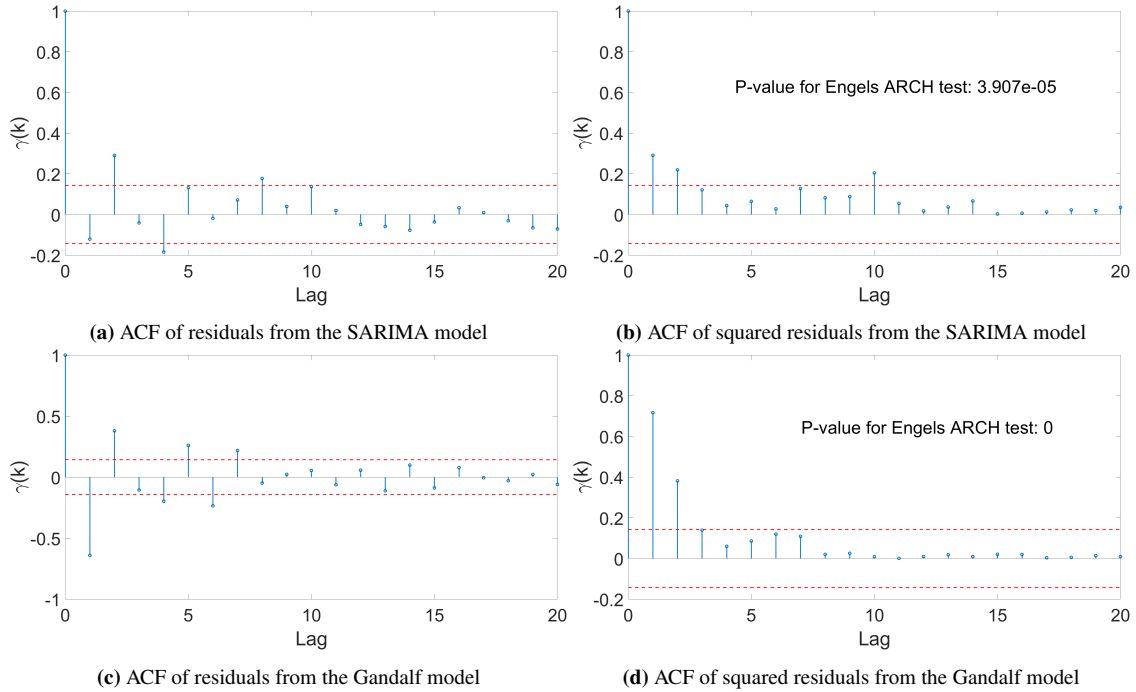


Figure 24: ACF plot of residuals and squared residuals for the SARIMA model and the Gandalf model, trained on Zain's data set.

Figure 25 and Figure 26 shows the seven-day and the 28-day forecasts from the three models trained on Zain's data set, respectively. The seven-day forecast used prediction scheme 2 (like Taraldsen), while the 28-day prediction used prediction scheme 1 (like Zain et al.). On the seven-day forecast, both SARIMA models obtain a slightly lower RRMSE score than the machine learning model. However, all results are excellent according to Zain et al. (2021) and Li et al. (2013). Note that the prediction interval for the SARIMA model is outside the field of interest, but can be observed in Figure 23.

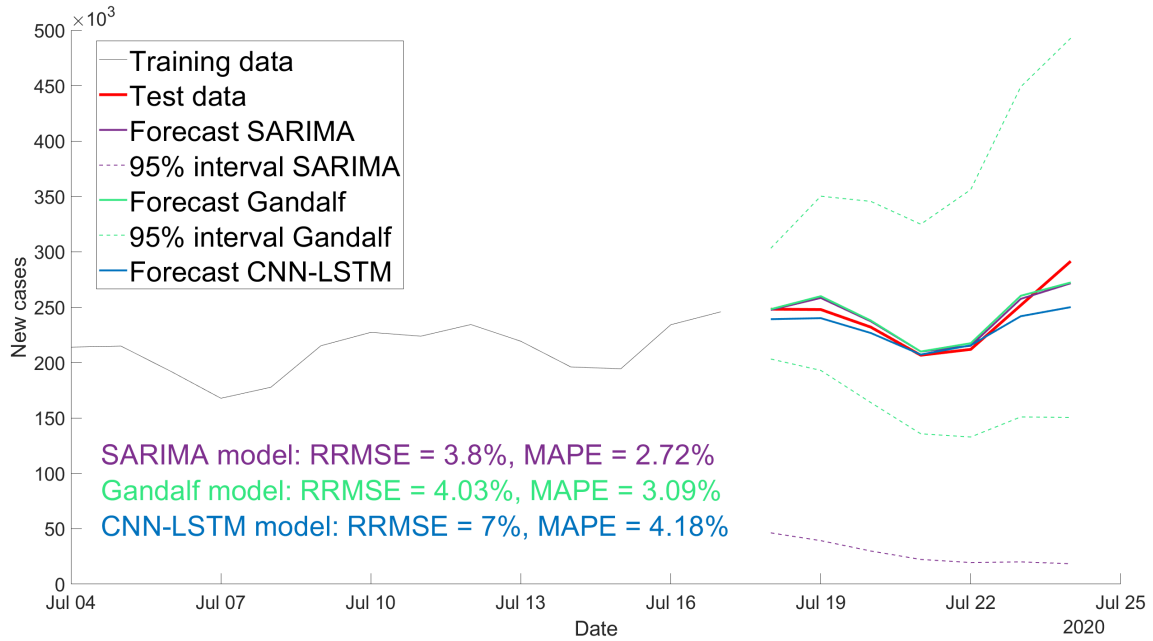


Figure 25: Forecast from the three models with scheme 2, from July 18th to July 25th 2020 on the global new-cases data. All models were trained on Zain’s data set.

On the 28-day forecasts, all three models obtain a slightly lower RRMSE score than what Zain et. al. reported with their CNN-LSTM model. Note that the RRMSE and the MAPE for the mean prediction are slightly lower than the mean of the RRMSE and MAPE for the ten models in Figure 17.

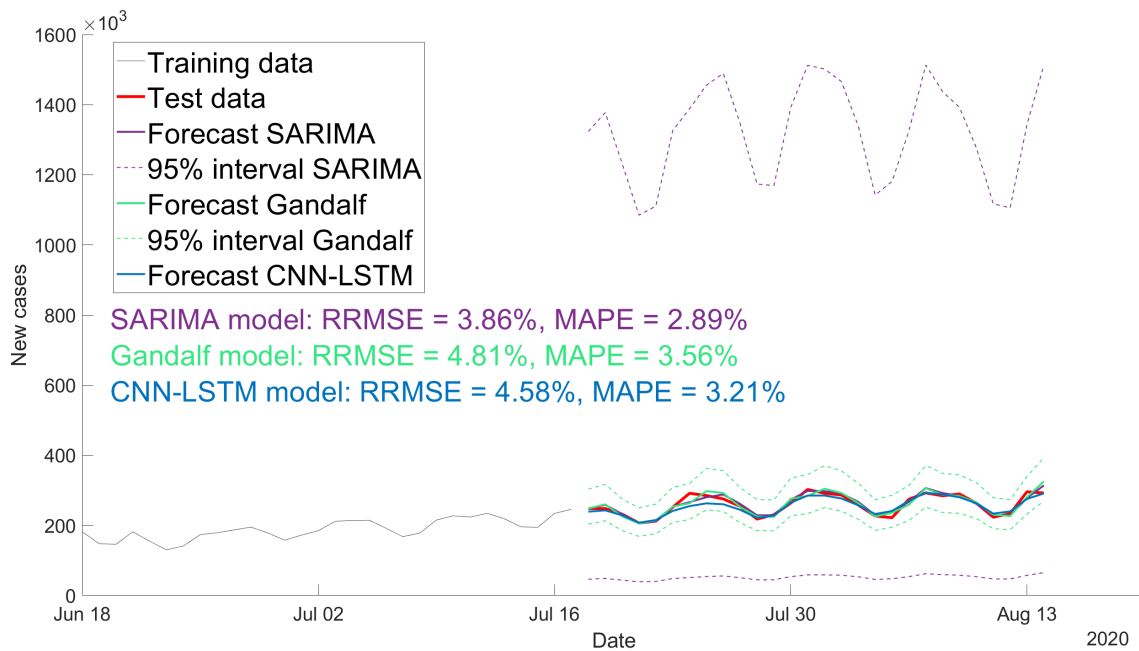



Figure 26: Forecast from the three models with scheme 1, from July 18th to August 14th 2020 on the global new-cases data. All models were trained on Zain’s data set.

The 28 sets of estimated parameters for both models were collected and plotted in Figure 27. Similar phenomena to their equivalent plots based on Taraldsen’s data set can be observed. However, the deviation of the σ^2 estimate is even more drastic.



(a) The parameter estimates from the SARIMA model. (b) The parameter estimates from the Gandalf model.

Figure 27: The parameter estimates for the two SARIMA models used to predict the 28 days shown in Figure 26. The dates along the x-axis is the last day used to train the model used to predict the day after. The colored dotted lines are the parameter estimates of the SARIMA model fitted on Zain’s data set.

As a curiosity, 28-day forecasts with prediction scheme 2 were compared between the three models. The prediction from the ten individual CNN-LSTM models, in Figure 28(a), seems to converge to straight lines. This suggests that the CNN-LSTM model predictions does not maintain the weekly structure of the test data for longer forecasts. For the mean CNN-LSTM prediction in Figure 28(b), the same phenomenon is amplified. On the contrary, the forecasts from the two SARIMA models keep the weekly structure throughout their forecasts. Surprisingly, the machine learning model was more accurate than the SARIMA models. Naturally, the prediction intervals for the SARIMA models expand since prediction scheme 2 was used.




Figure 28: Forecast from the three models with scheme 2, from July 17th 2020 on the global data set. The models were trained on Zain’s data set.

4.4.3 Comparisons with less training data and at different partitions

Until now, all data available before the final day of the training data sets have been used under training. More data is usually preferred in statistics. On the other hand, the handling of a new positive Covid-19 test has changed throughout the pandemic. Say, for example, that in an earlier period, a newly infected individual and all people close to them were obligated to stay at home, whereas in a later period, a new ruling says that neither the infected nor their closest people need to stay home if they did not feel ill. Giving the models information from both of these periods might confuse them, and hence, make them more inaccurate.




Figure 29: Forecasts using the 100 previous Norwegian data training data from November 10th 2020.

Figure 29 displays the seven-day and 28-day forecasts from the three models when reducing the training data set to the 100 days prior to November 10th on the Norwegian data set. The most notable differences are that the Gandalf model is less accurate on the seven-day forecast, and that the machine learning model is less accurate on the 28-day forecast. The CNN-LSTM model is more accurate than the SARIMA models on the seven-day forecast, but less accurate on the ensuing 28 days. Both the forecasts and the associated prediction intervals from the SARIMA models are closer together than in Figure 19 and Figure 20, where all 264 data points were utilized. Overall, the results are worse than with all 264 data points, but not by much.




Figure 30: Forecasts using the 50 previous Norwegian data training data from November 10th 2020.

Figure 30 displays the same forecasts, but trained on the 50 data points prior to November 10th 2020. The measures of accuracy are all higher than in Figure 19 and Figure 20, where all previous data was used. The accuracies are close to what is displayed in Figure 29. The SARIMA models are inseparable, as both their predictions and their prediction intervals align perfectly. The associated parameter estimates across the 28-day forecasts can be seen in Figure 31. As opposed to the earlier plots from Taraldsen’s data set, the estimated σ^2 is stable, and the estimated θ_{MA} and θ_{SMA} between the two models are almost identical. Also note that the estimated α_1 is close to zero, while β_1 stays close to 1.

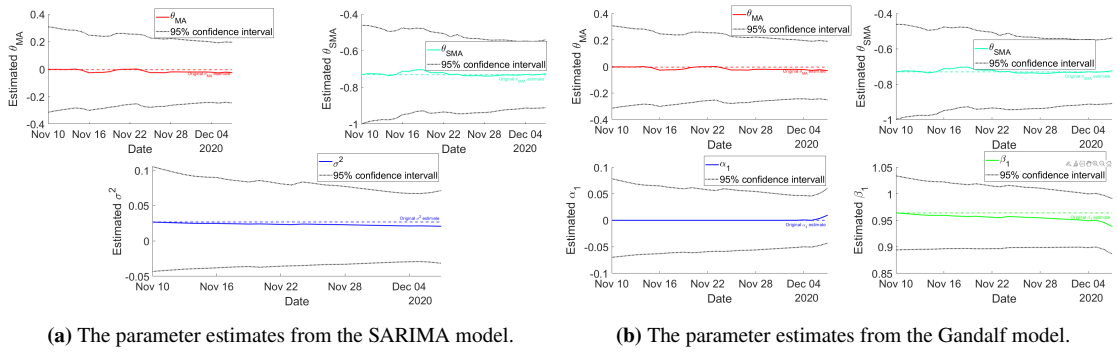


Figure 31: The parameter estimates for the two SARIMA models used to predict the 28 days shown in Figure 30. The dates along the x-axis is the last day used to train the model used to predict the day after. The colored dotted lines are the parameter estimates of the SARIMA model fitted on Zain’s data set.

The models have only been fitted on one particular ending date for each data set, while some experimentation has been done with the sample size. In the remaining paragraphs of this section, they will also be compared on the ending dates corresponding to day 100, 300, 500 and the last possible day, for both data sets. For all these ending dates, the models were fitted with all preceding data points, but also with the preceding 100 and 50 samples. In total, this gives 14 training sets from the Norwegian and 14 training sets the global data set. On each training set, the three models were evaluated on both a seven-day and a 28-day forecast. Note that the ending dates are different for the Norwegian data set than for the global data set, since the latter starts at an earlier date. Since the last week of the data set may have inaccurate measurements, the last possible day to include in test data was February 13th 2022 (for both data sets). Thus, the last possible day to train for a 28-day forecast is 28 days earlier than February 13th, i.e., January 16th 2022. For simplicity, this was also the last training day for seven-day forecasts. This amounts to a total of 56 forecasts for each model. In Table 5 in Appendix A, the mean prediction of the ten CNN-LSTM models universally achieves lower RRMSE and MAPE scores than the mean RRMSE and MAPE of the ten prediction, which was reported by Zain et. al. For this reason, only the RRMSE and MAPE of the mean prediction were compared to the other models in the following table. All the results are shown into Table 3.

Table 3: All forecast results from the SARIMA model and the Gandalf model, and the mean CNN-LSTM prediction. Note that the RRMSE and MAPE from the CNN-LSTM model is based on the mean prediction of ten models. All 28-day predictions used prediction scheme 1, while prediction scheme 2 was used for all seven-day predictions.

Data set	Last training date	Sample size	Days ahead prediction	SARIMA model		Gandalf model		CNN-LSTM model	
				RRMSE	MAPE	RRMSE	MAPE	RRMSE	MAPE
Norway	30.5 2020	100	7	118.52	100.00	118.52	100.00	64.23	88.08
			28	74.12	60.63	75.13	61.46	73.05	87.49
		50	7	87.02	56.39	87.02	56.39	64.52	40.73
			28	60.35	40.00	262.40	118.28	52.72	41.21
	10.11 2020	264	7	17.46	15.63	6.37	5.52	8.03	7.91
			28	11.73	9.21	10.00	7.89	17.91	16.96
		100	7	15.57	12.4	14.64	11.20	10.63	10.73
			28	10.93	8.55	10.87	8.20	25.01	25.98
		50	7	12.18	9.06	12.18	9.06	11.37	11.27
			28	11.29	8.00	11.29	8.00	27.71	28.36
	16.12 2020	300	7	17.22	13.71	16.82	13.25	14.18	11.83
			28	25.26	21.62	26.71	20.25	28.85	22.54
		100	7	19.03	15.16	19.89	16.02	7.41	6.65
			28	24.77	20.03	23.93	19.16	29.22	22.85
		50	7	16.24	12.70	16.24	12.70	18.28	14.00
			28	28.14	22.61	28.41	22.49	25.81	20.82
	04.07 2021	500	7	10.29	7.35	14.15	10.82	7.85	7.72
			28	14.48	11.23	14.67	12.02	17.06	11.72
		100	7	5.55	4.09	7.79	5.05	11.56	11.95
			28	14.65	11.62	14.53	10.95	30.38	20.10
50		7	21.25	17.26	20.87	16.97	10.79	9.23	
		28	15.88	12.51	15.88	12.32	35.02	23.32	
16.01 2022		696	7	8.66	6.84	8.39	6.83	23.05	21.81
			28	10.93	11.06	9.41	9.38	18.76	16.57
	100	7	9.43	7.72	11.18	8.98	24.04	18.24	
		28	8.90	8.78	9.19	9.12	29.56	30.89	
	50	7	12.97	11.03	3.59	3.08	23.76	22.69	
		28	8.86	8.51	8.87	8.28	45.55	46.06	
Global	12.04 2020	100	7	30.39	28.25	35.32	31.34	10.72	9.85
			28	6.76	5.52	6.95	5.45	7.34	6.33
		50	7	15.65	14.44	15.63	14.42	14.38	13.47
			28	6.24	5.19	6.33	5.32	9.86	8.51
	17.07 2020	196	7	6.77	5.67	6.06	4.86	8.01	6.48
			28	3.86	2.89	4.81	3.56	4.58	3.21
		100	7	4.10	3.31	4.11	3.22	5.90	3.59
			28	4.07	3.14	4.06	3.06	4.17	3.20
		50	7	4.27	3.79	4.26	3.71	4.90	3.04
			28	4.29	3.45	4.09	3.21	3.97	3.12
	29.10 2020	300	7	6.79	6.16	4.75	3.52	11.19	9.90
			28	5.10	4.15	6.54	4.79	12.19	11.07
		100	7	6.79	6.16	6.61	6.01	7.97	6.79
			28	5.12	4.23	5.14	4.16	9.03	7.59
		50	7	6.69	5.88	6.67	5.85	9.44	8.41
			28	5.64	4.44	5.74	4.44	10.46	9.34
	17.05 2021	500	7	3.95	3.20	10.15	9.26	13.36	12.99
			28	4.36	3.61	5.82	5.19	9.29	7.89
		100	7	6.44	5.53	7.08	6.18	20.43	20.33
			28	4.61	3.92	4.71	4.00	15.30	13.83
50		7	2.84	2.34	2.86	2.37	18.24	17.88	
		28	4.56	3.84	4.58	3.86	38.96	39.08	
14.01 2022		725	7	14.81	11.74	10.49	9.42	12.17	12.08
			28	8.56	6.98	7.28	5.96	20.59	20.08
	100	7	12.24	10.52	13.64	10.86	10.35	10.30	
		28	7.06	6.02	7.30	5.93	20.64	20.11	
	50	7	12.47	10.66	13.16	10.39	11.38	10.64	
		28	6.62	5.80	6.64	5.49	85.18	82.83	

To get a grasp of the results in the above table, all the MAPE scores for each model were plotted in Figure 32. Some outliers can be observed for all models, but they mostly score below 25. The MAPE scores for the machine learning model seem to be a bit more spread out compared to the SARIMA models.




Figure 32: MAPE scores for each model from Table 3.

The MAPE scores of the models were examined further in Figure 33. To account for the outliers shown in Figure 32, the median MAPE for different subgroups of the above table, rather than the mean. Figure 33(a) shows that the SARIMA model obtains the lowest median MAPE score, with the Gandalf model closely behind. In Figure 33(b), the models were compared on a seven-day and a 28-day forecast. The SARIMA model scores lower on the 28-day forecasts, while the Gandalf model gave the best performance on the seven-day forecasts. The machine learning model struggles substantially more on the 28-day forecasts. Figure 33(c) compared the MAPE between the Norwegian and the global data set. All models did significantly better on the global data set, but the SARIMA models obtained an almost twice as low MAPE score as the machine learning model on this data set. Figure 33(d) demonstrates that reducing the sample size to 100 did not affect the MAPE much. When only 50 data points were available, the SARIMA models received lower MAPE scores than previously, while the CNN-LSTM models increased slightly.

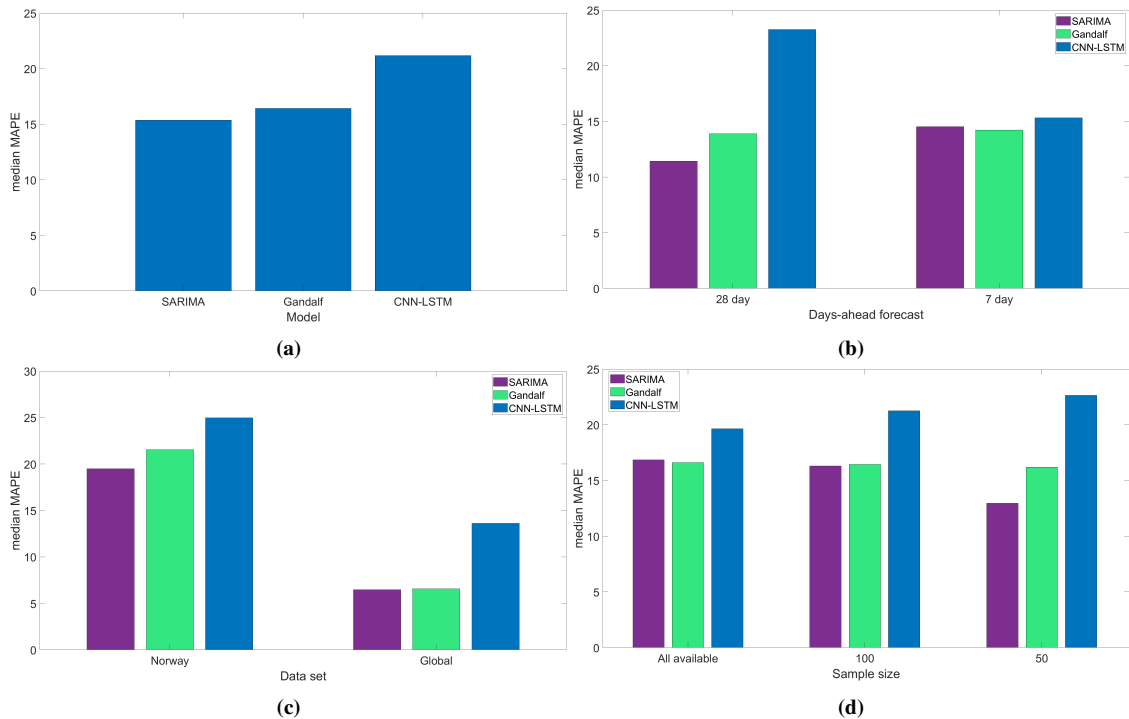


Figure 33: Comparisons of models based on MAPE scores from Table 3

The same analysis was also done on the RRMSE scores from Table 3. These results were placed in

Appendix A (Figure 42 and Figure 43), since they gave more or less the same information as the analysis of the MAPE. The main difference from the MAPE scores is that the Gandalf model performed worse relative to the other two models, and that the CNN-LSTM model had a slightly better performance than the other two on the seven-day forecasts.

4.5 The models on simulated realizations

While no analytical expression for constructing a probability interval of the CNN-LSTM model exists, by using a parametric bootstrap technique, it is possible to estimate a distribution around the RRMSE scores and the MAPE scores for the models. The spread of this distribution is an alternative measure of the spread of the predictions from the CNN-LSTM model. Assuming that the time series are stochastic in nature and that the underlying model is known and can be simulated from, it would be possible to use the models to forecast on each of these simulations. This could be used to control the previous results. Maybe the SARIMA model was lucky in the prediction on the one known realization? Unfortunately, the underlying model for the Covid-19 new cases will forever be unknown. However, as the above analysis suggests, the Gandalf model was a good model for the time series. Using a Gandalf model trained on Zain’s data set, 1000 simulated realizations of the global data from July 18th to August 14th 2020 were generated. The specific parameter estimates of this Gandalf model are equivalent to the estimates in Figure 23, i.e., $\theta_{MA} = 0.09$, $\theta_{SMA} = -0.66$, $\alpha_1 = 0.04$, and $\beta_1 = 0.89$. In Figure 34, three of these 1000 realizations were plotted together with the forecast, the theoretical prediction interval for the Gandalf model, and the 95% quantile range based on the 1000 realizations. Observe that the three realizations are contained inside the 95% prediction interval, and that the quantile range approximates the theoretical prediction interval quite well.

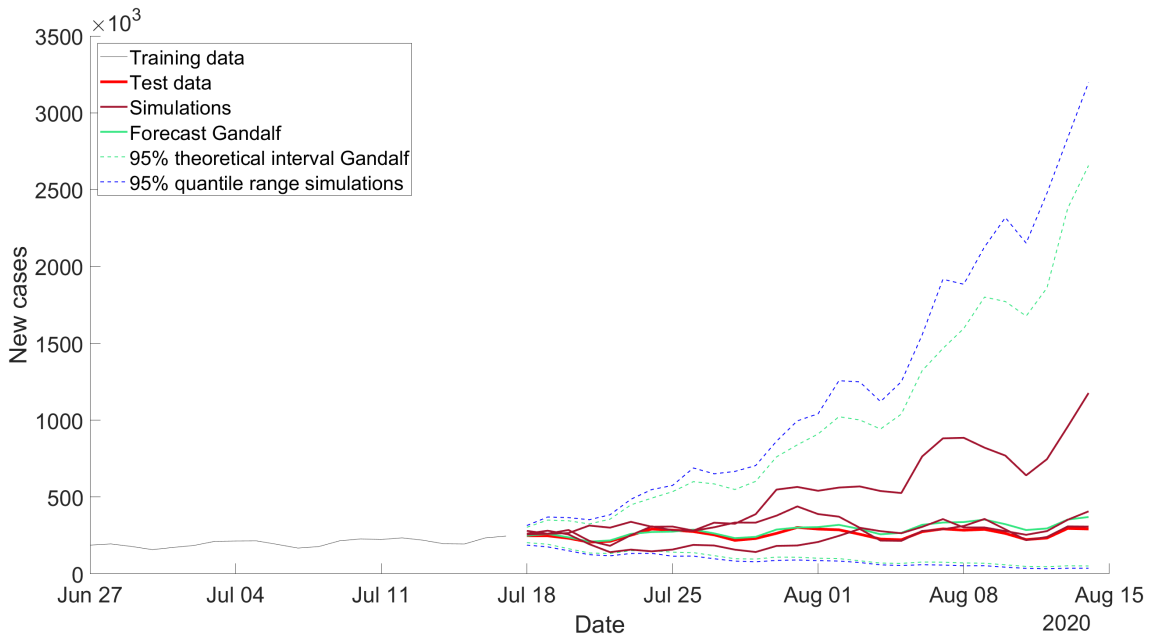


Figure 34: Three simulations from the Gandalf model. The model was trained on the original global training data. The 28 ensuing days were then simulated.

Under the assumption that the underlying process is the Gandalf model, each of these simulated realizations can be treated as a different realization of the global time series. Using prediction scheme 1, Figure 35 shows the prediction from the three models on the three realizations. Of particular interest was Figure 35(b), where the machine learning model struggled to keep up with the fast upward trend of the realization.




Figure 35: Predictions on three simulated realizations from the Gandalf model with prediction scheme 1.

For each of the 1000 simulations, one such forecast was made from each of the models, and the RRMSE and MAPE were evaluated. For 1000 simulated realizations, this gives 1000 RRMSE scores and 1000 MAPE scores for each model. Figure 40 shows the result of this operation for the Gandalf model, displayed as a histogram. The Gandalf model is expected to deliver the lowest RRMSE and MAPE scores, since it generated the "test data". Hence, these distributions serve as a baseline for the performance of the two other models. The measures of accuracy on the actual test data can also be seen in the plot. The Gandalf model was expected to deliver more accurate forecasts on the simulated realizations (from the Gandalf model) than on the test data. It would therefore be reasonable to expect the density of this histogram to center below the original accuracy measures. However, this is not the case. In fact, none of the scores were lower than the scores on the original test set.




Figure 36: Distribution of RRMSE and MAPE based on predictions from the Gandalf model on 1000 simulated realizations from the Gandalf model. Prediction scheme 1 was used to generate the predictions, once for each simulation.

A closer inspection of the parameters for the Gandalf model could explain the above result. Since the predictions were generated using prediction scheme 1, the Gandalf model was fitted 28 times for each of the 1000 simulated realizations. For the 200 first simulations, the parameter estimates were plotted together with a dotted line where the parameter estimates on the training data were. The confidence interval are also based on the original data; they are therefore constant over time. Figure 37 displays the result. The estimates are mostly contained inside the confidence intervals of the original estimates. Some bias can be observed for α_1 and β_1 . This supports the fact that the simulations were generated correctly, since they suggest the same model as the one originally fitted.

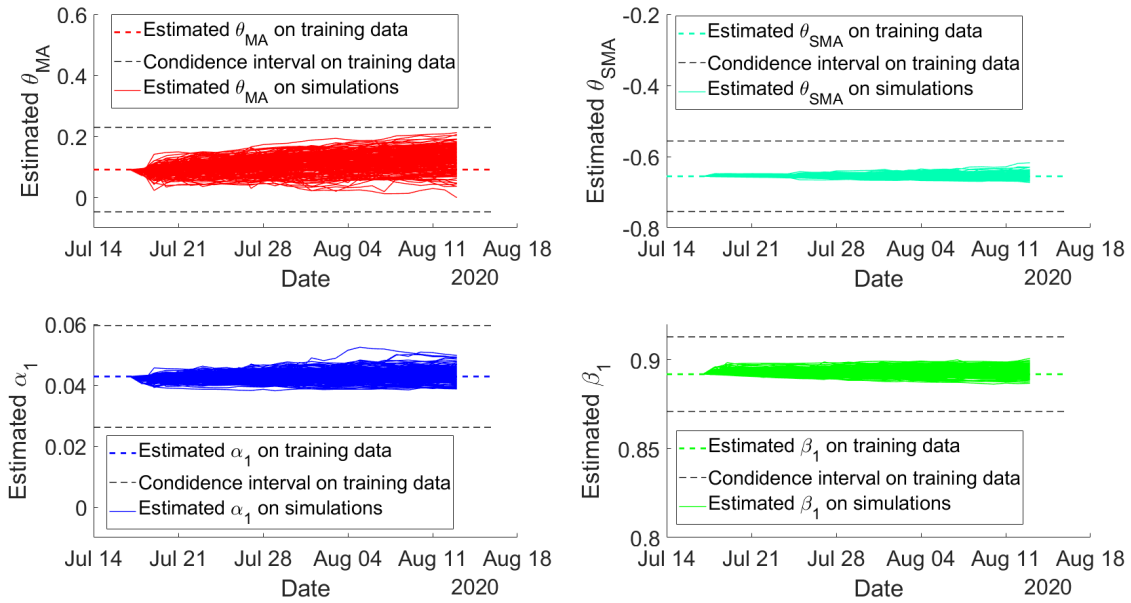


Figure 37: The deviation of the four Gandalf parameters on the 200 simulated realizations. For each simulated realization, prediction scheme 1 was used to fit models and predict one day at a time. This gives 28 fitted models and parameter estimates for each simulation.

The Gandalf model is fitted and predicts the new cases on log-scale. Hence, it was hypothesised that the transformation back to normal scale caused the above result. The log-transformed prediction of the simulated data was evaluated, and the results are shown in Figure 38. Once again, all accuracy measures are above the original score.

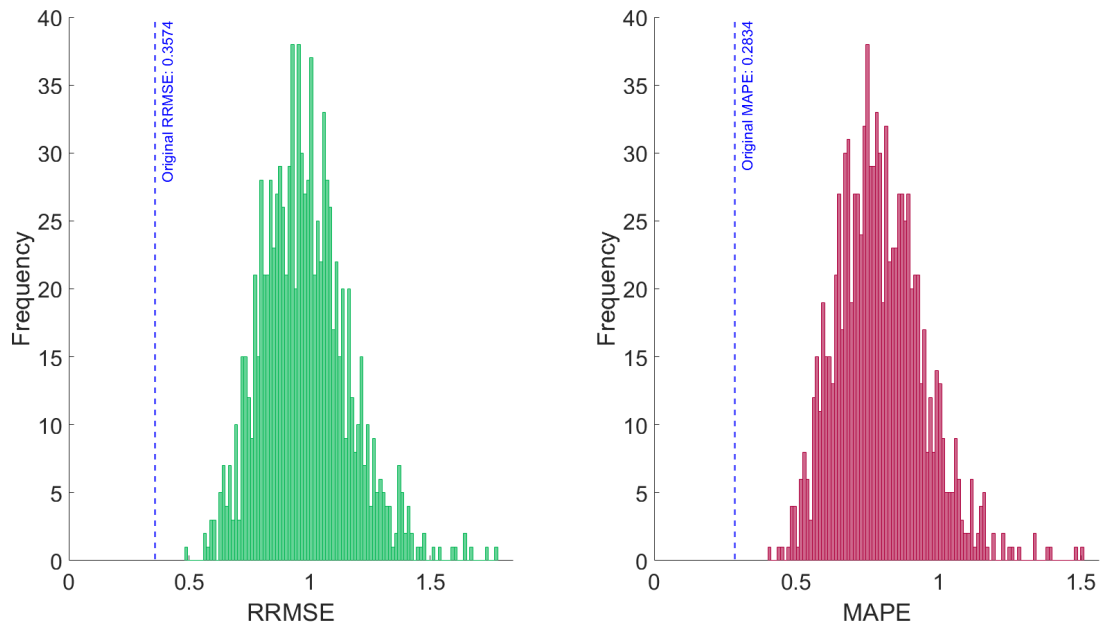


Figure 38: Distribution of RRMSE and MAPE based on log-scaled predictions from the Gandalf model using simulated realizations from the Gandalf model on log scale. Prediction scheme 1 was used to generate the predictions, once for each of the 1000 realizations.

Even though the accuracy of the Gandalf model on the Gandalf simulations were worse than on the actual test data, Figure 36 can serve as a baseline for the two other models. The SARIMA model was expected to predict the results quite well, as the model is closely related to the Gandalf model. By inspection of Figure 39, this was the case.




Figure 39: Distribution of RRMSE and MAPE based on predictions from the SARIMA model on 1000 simulated realizations from the Gandalf model. Prediction scheme 1 was used to generate the predictions, once for each realization.

The main reason for conducting the simulation study in the first place was to see how well the CNN-LSTM model predicted the simulated realizations. Again, there is no analytical expression for the prediction error for the machine learning model. Assuming the Gandalf model is a perfect model for the Covid-19 time series, this was another way to monitor the spread of the model's prediction. Observe that most of the RRMSE scores for the CNN-LSTM model around 40. Once again, none of the forecasts were more accurate on the simulated realizations than the actual test data w.r.t. the two measures.




Figure 40: Distribution of RRMSE and MAPE based on predictions from the CNN-LSTM model on 1000 simulated realizations from the Gandalf model. Prediction scheme 1 was used to generate the predictions, once for each realization. The mean CNN-LSTM prediction was used to generate the forecasts.

To make the comparisons easier, it was helpful to have concrete measures to summarize the three pairs of

histograms. The arithmetic mean gives precise information about how accurate each model was in total, while the standard deviation (Std) measures the spread of the distributions. However, for the purposes of this study, the measure of spread of each distribution should be relative to the magnitude of the errors. **The relative standard deviation (RStd)** satisfies these criteria. It is calculated by dividing the standard deviation of the distributions by their mean. Table 4 displays the RStd for both distributions of measures, for each model. The Gandalf model has the lowest relative spread for both measures, and the lowest mean. The SARIMA model accuracy measures are only slightly more spread out. The RStd of the RRMSE scores for the mean CNN-LSTM predictions are about twice as high as the other two, while it was four times higher for the MAPE scores.

Table 4: Relative standard deviation of accuracy measures of 28-day forecasts on 1000 simulated realizations from the Gandalf model. The models were initially trained on Zain’s data set.

Models	RRMSE			MAPE		
	Mean	Std	RStd	Mean	Std	RStd
Gandalf	13.08	3.10	0.24	9.97	1.89	0.19
SARIMA	15.61	4.42	0.28	11.96	2.76	0.23
CNN-LSTM	47.57	26.24	0.55	39.48	30.12	0.76

5 Discussion

This section discusses the results from the previous section, suggests possible improvements, and addresses some problems with the applicability of the models.

5.1 The performance of the three models

The previous section studied the performance of the three models across the data sets, which are now discussed.

5.1.1 Accuracy

The models were developed for the early parts of the Covid-19 data sets, where they all delivered fairly accurate forecasts overall. Across the two years of the data set, a lot has happened to the way a positive Covid-19 test is treated. Everything from the different variations of the virus to the rapid antigen tests and the vaccines have contributed to this. A worry was that this might alter the behaviour of the time series, and thus hamper the performance of the models. As this thesis have shown, this was not the case; the models continued to deliver good forecast across both time series (in the sense that the mean RRMSE for all models were around 20).

The SARIMA model gave the most accurate forecasts overall. However, the gap between the SARIMA model and the Gandalf model was narrow when comparing MAPE scores. The prediction interval for the Gandalf model were consistently much narrower than the SARIMA model's. Thus, an argument can be made for choosing either of these models. The CNN-LSTM model has the overall highest RRMSE and MAPE score and does not come with a prediction interval. This is major drawback of the machine learning model. It is still quite fascinating that the CNN-LSTM manages to predict with such good accuracies, since it includes more than 300000 parameters: especially for the lower sample sizes.

One of the largest gaps between the models was illuminated when the RRMSE scores were grouped by forecast length (and prediction scheme, since all 28-day forecasts used prediction scheme 1 and vice versa). The SARIMA model performed much better than the other models on 28-day forecasts. On the seven-day predictions, the models performed equally good. Moreover, the SARIMA model was the only model that performed worse on these forecasts. This has a natural explanation, that will be discussed later. On both measures of accuracy, the machine learning model gave the worse performance on the 28-day forecasts, while it was just as good at forecasting seven days with prediction scheme 2. A possible flaw was that the CNN-LSTM models were only fitted initially for the 28-day predictions with predictions scheme 1. On the other hand, the SARIMA models were fitted at each step. The decision was based on the fact that the CNN-LSTM models had a much longer time to fit (see Table 2). This seems to have had a clear effect on the measures of accuracy in Figure 33 and Figure 43. These high inaccuracies of the 28-day forecast translates over to the other comparisons in Figure 33 and Figure 43. In cases where only 50 training data were available, this effect was quite clear, since the machine learning model has to forecast half the length of the training data. If the CNN-LSTM model were forecasted in the same way, the accuracies on the 28-day forecasts might have been more comparable to the SARIMA model, and thus, the other comparisons would also go more in the favour of the CNN-LSTM model.

It is important to note that the above statements are only based on the 28 partitions of each data set in Table 3. The results might have changed if the analysis were done with different ending dates, with a different set of sample sizes, and with other prediction lengths.

Overall, the models were more precise on the global data set. This data set has a much higher volume of new cases, and the relative changes are not as great as in the Norwegian data set. This means that the RRMSE and MAPE often will be lower for the global data set. This was despite the fact that some lack of fit was observed for the SARIMA model. Ideally, these relative measures of accuracy should be comparable across the two data sets. In practice, this remark shows that these comparisons may be flawed.

5.1.2 The changing volatility and the effect of sample reduction

Figure 12 showed that the series $(1 - B)(1 - B^7)X_t$ exhibits conditional heteroscedasticity, with higher variations in the early parts of both data sets. This suggests that a conditional variance model should be used for all SARIMA models when the earliest parts of the data sets were included. This might explain why the Gandalf model performed better than the SARIMA model on Taraldsen's data set. At the later periods, the series $(1 - B)(1 - B^7)X_t$ looks quite stationary for both data sets. Thus, a GARCH(1, 1) model should have little to no advantage here; it might even cause overfitting. For lower sample sizes, there is a greater chance of avoiding the heteroscedastic parts of the data sets. On the other hand, if all previous data is used, these parts are sure to be included. This is reflected in Table 3 and the associated Figure 33 and Figure 43, where the SARIMA model is more accurate for lower sample sizes. The opposite relation is true for the Gandalf model and the CNN-LSTM model.

The estimated θ_{MA} and θ_{SMA} were stable for both the SARIMA models on both original data sets, seen Figure 21, Figure 22 and Figure 27, with the exception of θ_{MA} for the Gandalf model on the global data set. Together with Figure 12, this suggests that the moving average structure of these models fits the data well. Otherwise, they should vary slightly with each additional data point. The major differences between the parameters of the two models were in the estimates associated with the noise component. The same figures show significant declines in the estimated σ^2 parameters. The models were trained on the parts of the data with conditional heteroscedasticity (see Figure 12). Hence, the SARIMA model might have to reduce its estimated variance when fitted on the test data, which is more stable than the average variation of the training data. With prediction scheme 1, the models were fitted for each new day. The variance can therefore be adjusted to the new data before each new value is predicted. This is not the case when using prediction scheme 2, where the initial estimated σ^2 is used for the entire forecast. This may explain why the SARIMA model is the only model with increased RRMSE and MAPE on the seven-day forecasts. In Figure 21 and Figure 27, the noise estimates α_1 and β_1 in the Gandalf model, are more stable for the two original data sets. However, if this were to be a problem, Figure 28 should then be the SARIMA model's downfall, since the parameter estimates were never updated on a 28-day forecast. This was not the case, which suggests that the persistence volatility does not cause any significant drawback to the performance of the SARIMA model for Zain's data set. Additionally, the SARIMA model performs better than the Gandalf model on the seven-day forecast trained on Zain's data set. An explanation might be that the end of this data set is far enough away from the part with higher variation. There is some evidence of the Gandalf model achieving more accurate seven-day forecasts when the elevated variance is closer to the end of the training data, where the SARIMA model has not had time to compensate. A closer investigation of the performances of each model should be made around these problem areas.

The downwards trajectory of the σ^2 estimates for the SARIMA model on the two original data sets explains the narrowing prediction intervals for 28-day forecasts in Figure 20 and Figure 26. The earlier assessment that the Gandalf model have narrower prediction intervals might not be true for later data partitions, where the time series are close to homoscedastic and has less variation overall.

Figure 30 and Figure 29 suggest that the similarity between the forecast of the two models increases as the sample size reduces. A close inspection of Figure 12 gives a more reasonable interpretation. In the figure, the part of the data with increased volatility is excluded from the training set when only considering the past 50 observations, and mostly excluded with the most recent 100 observations. Thus, the added GARCH(1, 1) process should not differ much from the regular white noise process. Figure 31 supports this statement, since the estimated variance for the SARIMA model stays constant across the 28-day forecast. Additionally, the estimated α_1 and β_1 stayed close to 0 and 1, respectively. This gives the following conditional variance process for the Gandalf model:

$$\sigma_t^2 = \alpha_0 + \alpha_1 Z_{t-1}^2 + \beta_1 \sigma_{t-1}^2 \approx \sigma_{t-1}^2,$$

since α_0 was fixed at 0.001. Thus, the GARCH(1, 1) model approximates conditional homoscedasticity. If α_0 had not been fixed, an alternative set of estimates would have been that $\alpha_1 = \beta_1 \approx 0$ and $\alpha_0 \approx \sigma^2$. The similarity between the estimated θ_{MA} and θ_{SMA} in the two models supports the fact that the variance models are similar. This explains why the models' forecasts and prediction intervals align perfectly in Figure 30.

Even if the models performed great overall, there were some partitions of the data sets that were seemingly

harder for the models to model. This might simply be due to chance, but when all models performed badly, it is reasonable to assume that the time series in question changed behaviour on those partitions. Maybe the vaccines started to become available to the masses. Or maybe a new variant increased in prominence. These connections are interesting in themselves and should be investigated.

5.1.3 Results from the simulation study

The main motivation behind the simulation study was to quantify the spread of the predictions from the CNN-LSTM model. None of the RRMSE scores nor the MAPE scores from the predictions on the simulated realizations were lower than they were on original test data, with the Gandalf model. This was strange since all the simulated realizations came from the Gandalf model. Most of the predictions on these were therefore expected to be more accurate than on the original test data. An initial hypothesis was that this was due to the transformation back to normal scale. After all, the SARIMA models were built to perform on the log-scaled data sets. However, the same phenomenon was observed. One cannot exclude the possibility of this being due to a programming error. However, the simulated prediction interval aligns quite well with the Figure 34, and the parameter estimates from the Gandalf model do not go astray. These results provide some assurance that the simulated realizations and the associated forecasts were generated correctly, which gives credibility to later findings.

By inspection of Figure 34, an alternative explanation could be that the actual test data behaves nicer than the simulated realizations, since these are observably much wilder. The high volatility in the beginning of the data set might partially explain this. If this is the correct explanation, comparing the spread of these predictions might still be fruitful. The SARIMA model is closely related to the Gandalf model and was expected to be about as accurate on the simulated realizations. However, the SARIMA model showed a lot wider prediction intervals on the forecasts ensuing Zain's data set. Hence, it was reasonable to assume a wider spread of the RRMSE scores for the SARIMA model than the Gandalf model. For the Gandalf model, the mean RRMSE score was 13.08, and the relative standard deviation was 0.24. On the MAPE scores, the mean and the RStd were slightly lower. The SARIMA model achieved almost as good results but has a slightly higher mean and a wider spread for both measures. However, the relative spread of this forecast is nowhere close the relative width of the prediction intervals on the actual test data, where the prediction interval for the SARIMA model was consistently three to four times wider than the Gandalf model's interval. Whatever spread is measured by the RStd on these measures of accuracy, it should not be compared directly to the prediction intervals. The machine learning models struggled to predict the simulated realizations, with a mean RRMSE around 50 and MAPE around 40. The RStd was two and four times higher than the Gandalf model, respectively, suggesting that the spread of its forecast on this partition of the data is higher.

5.2 Possible improvements

The models used in this thesis were, for the most part, replicated from the articles by Taraldsen and Zain et al. While the models performed fairly well on both data sets, there are several plausible ways to improve the models.

5.2.1 The machine learning model

CNN-LSTM models has a lot of hyper parameters that might have a big impact on the results. Zain et al. did their own hyper parameter tuning with the Optuna framework. Sadly, they did not publish all their optimal hyper parameters. Two of these parameters were essential to know when building the model, namely the dropout rate and the number of neurons in the fully connected layer. As noted earlier, the global data set used in this thesis was not identical to what Zain et al. used. Thus, even the published hyper parameters might not be optimal. The same parameters were used across all partitions of the global data set. The same hyper parameters were also used for the Norwegian data set, where they are even less likely to be optimal. This was intentional, since if the models managed both data sets well using the same architecture, they effectively capture some ground truth about the spread of Covid-19. However, to increase the performance of the CNN-LSTM models, their hyper parameters model should be individually tuned for the Norwegian and the global data set. Ideally, they should be re-tuned for each evaluated partition of the two data sets.

Alterations to the entire model structure should also be included in the hyper parameter searches. This includes the number of layers of each type (CNN layers, LSTM layers and fully connected layers, pooling layers, and dropout layers) and the number of neurons in each of these. This would of course result in a lot longer training times, but if they do not take several days, this should not be a hindrance in practice.

Using the MSE as loss function might not be ideal for this problem. After all, the forecasts were evaluated based on the RRMSE and the MAPE. Some experimentation with the use of MAPE as loss function was made. Ironically, this resulted in a worse MAPE of the forecasts. In hindsight, this loss function might have required a completely different set of hyper parameters to give optimal forecasts. To choose the optimal number of epochs, the models should incorporate early stopping, e.g., with a stopping criterion set to 0.0015. Lastly, the median might be a better choice of measure of central tendency for the model's forecasts, as it deals better with outliers than the mean.

Figure 28 demonstrates a weakness of the CNN-LSTM model in its current form: it does not capture the weekly variation and the models fades out as the prediction length increases. This did not reduce the performance of the model in this one example, but this might not be consistent on other partitions of the data sets. This motivates a variation of the model: The weekday model. This model is really seven sub models, one for each day of the week. Thus, one of these models only considers the new cases on each Monday, another considers all Tuesdays, and so on. Together, they can be used to predict the new cases on each ensuing day. This could potentially be a better alternative, as it would capture the weekly structure by construction. A problem with this proposed model is that this reduces the amount of training data for each of these sub models by a factor of seven. The model was tested (and the code for the model are included in the Python-file in Appendix B), but the initial tests did not achieve great accuracies, which is why it was excluded from the thesis. Looking back, each sub model should be treated individually, with potentially unique model architectures (determined by a unique hyper parameter search).

5.2.2 The SARIMA models

The structure of the SARIMA models was natural for modeling the Covid-19 new cases data, as both data sets have a clear weekly pattern and an ever-changing mean. This is supported by the highly desirable behaviour of the residuals on Taraldsen's data set in Figure 14. Figure 12 showed that the series $(1 - B)(1 - B^7)X_t$ exhibits higher-than-usual variance in the early parts of this data sets. Together with the ACF plot of the squared residuals, this suggests that a conditional variance model would have a better fit to the noise than a white noise model. Even if the residuals and the accuracy of the two forecasts on Taraldsen's data set improved with the added GARCH(1, 1) model, the squared residuals on this data set was not improved by the added GARCH(1, 1) model. It is possible that a different GARCH model might have been a better alternative. This could be determined by a grid search for p and q , where each combination (up to some limit) was compared based on the AICc. This must be done with caution. As should be clear by now, the SARIMA model fits the later partitions of the data set better than the Gandalf model. Even the GARCH(1, 1) model seem to cause overfitting, as the predictions are slightly less accurate overall. The ACF plot of the residuals based on Zain's data set from the SARIMA model showed that many lags were significant. A similar grid search as described above might have improved the results on the global data set even further and should have been attempted.

As described in Section 2.4, Matlab uses a conditional likelihood to estimate the parameters for the SARIMA models. The exact likelihood should give slightly more precise estimates, especially for the reduced sample sizes. Another suggestion for lower sample sizes is to assume that either Z_t (or e_t , when a GARCH model is used) follows Student's t-distribution instead of the Gaussian distribution.

In all plots of the parameters for the Gandalf model (see Figure 21, Figure 27 and Figure 37) $\alpha_1 + \beta_1$ was consistently quite close to 1. The Exponential GARCH(1, 1) model fixates this sum to be one, and should be considered as the noise process for these data sets.

None of the models presented in this thesis takes into account that the new cases are integers. In an earlier article, Taraldsen (2020a) gives several alternatives for modeling new Covid-19 cases. Interer-valued GARCH (INGARCH) models are one such example. The COM-Poisson process has been shown to deal well with both overdispersion and underdispersion, which are commonalities when modeling count data. Zhu (2012) proposes a COM-Poisson INGARCH model for modeling time series of counts. This model

should be tested on both data sets used in this thesis.

5.3 A note on the application of the models in the real world

For the most part, all three models manage to forecast the daily new Covid-19 cases data set fairly well. However, the utility of these forecasts can be questioned. The new cases data only includes reported cases. The regular weekly fluctuations of the time series reflect this, with a drastic decrease in the reported cases at the weekends. These are not correct representations of the spread of the virus. Neither are the forecasts of these series. How useful is this if the goal is to foresee the next outbreak? For this to bring utility, multiple data sources should probably be used in conjunction. Examples of this are mobility data and more detailed models of the person-to-person trends. However, each data source introduces more uncertainty in the final prediction. The models might be used on top of other types of models to account for the weekly fluctuations in the new cases data sets. There might also be more utility in creating models based on the weekly new cases, as these gives a better representation of the actual number of infected people. The hospitalizations due to Covid-19 might also paint a more correct picture of the severity of the situation, especially in the post-vaccine era, where most infected people only have minor symptoms. The problem with both of these alternatives is that they contain less samples than the daily new-cases data set.

The different partitions of the data sets might require different models for optimal results. Thus, the models should continually be tested on the newest data if they were to be used in a real-world setting.

6 Conclusion

This thesis presented promising results for the three models. Overall, the SARIMA model performed better than the other two models. For the CNN-LSTM model, this might be due to sub-optimal hyper parameters. There seems to be persistence of volatility in the beginning of both the Norwegian and the global data set. On these partitions, the Gandalf model is more accurate than the SARIMA model and has a much tighter prediction interval. On most of both data sets, however, the variance is homoscedastic and much lower in magnitude. Here, the added GARCH effect merely introduces more uncertainty, possibly resulting in overfitting. Additionally, the reduced variance estimate will reduce the width of the prediction interval for the SARIMA model on the majority of both data sets. The simulation study showed that the spread of the two SARIMA models was about as low, while the spread of the machine learning model was multiple times higher. However, the simulated realizations were wilder than the test data, so the results are not directly comparable to the results on this data set. While the CNN-LSTM model competes with the other models, there were few scenarios where it achieved better results. It is still quite remarkable that a model with over 300000 parameters gives accurate forecasts for such low sample sizes. If one had to choose one of these models for modeling daily Covid-19 new cases, the SARIMA model with Gaussian white noise would most likely give the best results overall. However, if clear signs of conditional heteroscedasticity are present at during the last few days of the training data, the Gaussian white noise may be replaced by GARCH(1, 1) noise.

References

- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31, 307–327.
- Brockwell, P. J. & Davis, R. A. (2016). *Introduction to time series and forecasting*. Springer.
- Dehning, J., Zierenberg, J., Spitzner, F., Wibral, M., Neto, J., Wilczek, M. & Priesemann, V. (2020). Predictions, role of interventions and effects of a historic national lockdown in india’s response to the covid-19 pandemic: Data science call to arms. *Harvard data science review*. <https://doi.org/https://doi.org/10.1162/99608f92.60e08ed5>
- Ekinci, A. (2021). Modelling and forecasting of growth rate of new covid-19 cases in top nine affected countries: Considering conditional variance and asymmetric effect. *Chaos, Solitons Fractals*, 151, 111227. <https://doi.org/https://doi.org/10.1016/j.chaos.2021.111227>
- Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *The Econometric Society*, 50, 987–1007.
- Hamilton, J. D. (1994). *Time series analysis*. Princeton University Press.
- Higham, N. J. (2009). Cholesky factorization. *WIREs computational statistics*, 1, 251–254. <https://doi.org/https://doi.org/10.1002/wics.18>
- Kiranyaz, S., Avci, O., Abdeljaber, O., Ince, T., Gabbouj, M. & Inman, D. J. (2021). 1d convolutional neural networks and applications: A survey. *Mechanical Systems and Signal Processing*, 151, 107398. <https://doi.org/https://doi.org/10.1016/j.ymssp.2020.107398>
- Kumar, K. A., Kalaga, D. V., Sai Kumar, C. M., Chilkoor, G., Kawaji, M. & Brenza, T. M. (2021). Forecasting the dynamics of cumulative covid-19 cases (confirmed, recovered and deaths) for top-16 countries using statistical machine learning models: Auto-regressive integrated moving average (arima) and seasonal auto-regressive integrated moving average (sarima). *Applied Soft Computing*, 103, 107161. <https://doi.org/https://doi.org/10.1016/j.asoc.2021.107161>
- Li, M.-F., Tang, X.-P., Wu, W. & Liu, H.-B. (2013). General models for estimating daily global solar radiation for different solar radiation zones in mainland china. *Energy Conversion and Management*, 70, 139–148. <https://doi.org/https://doi.org/10.1016/j.enconman.2013.03.004>
- Luan, Y. & Lin, S. (2019). Research on text classification based on cnn and lstm. *2019 IEEE International Conference on Artificial Intelligence and Computer Applications (ICAICA)*, 352–355. <https://doi.org/10.1109/ICAICA.2019.8873454>
- MachineLearningMastery. (n.d.). *Lstm model architecture for rare event time series forecasting*. <https://machinelearningmastery.com/lstm-model-architecture-for-rare-event-time-series-forecasting/>
- Pascanu, R., Mikolov, T. & Bengio, Y. (2012). Understanding the exploding gradient problem. *CoRR*, abs/1211.5063. <http://arxiv.org/abs/1211.5063>
- Peltarion. (2022, May 20). *Max pooling block 1d*. <https://peltarion.com/knowledge-center/documentation/modeling-view/build-an-ai-model/blocks/max-pooling-block-1d>
- Ruder, S. (2016). An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*.
- Taraldsen, G. (2020a). A toy model for covid-19 data. *Researchgate*.
- Taraldsen, G. (2020b). A wizard predicts new covid-19 cases. *Researchgate*.
- TensorFlow. (2022a, May 20). *Tf.keras.layers.dropout*. https://www.tensorflow.org/api_docs/python/tf/keras/layers/Dropout
- TensorFlow. (2022b, May 29). *Tf.keras.layers.lstm*. https://www.tensorflow.org/api_docs/python/tf/keras/layers/LSTM
- TensorFlow. (2022c, May 20). *Tf.keras.optimizers.adam*. https://www.tensorflow.org/api_docs/python/tf/keras/optimizers/Adam
- The MathWorks, I. (2022a, May 13). *Maximum likelihood estimation for conditional mean models*. <https://se.mathworks.com/help/econ/arima-maximum-likelihood-estimation.html>
- The MathWorks, I. (2022b, May 13). *Mmse forecasting of conditional mean models*. <https://se.mathworks.com/help/econ/mmse-forecasting-for-arima-models.html>
- The MathWorks, I. (2022c, May 13). *Model comparison tests*. <https://se.mathworks.com/help/econ/model-comparison-tests.html#btb4y0n>
- The MathWorks, I. (2022d, May 19). *Simulate stationary processes*. <https://se.mathworks.com/help/econ/simulate-stationary-arma-processes.html>
- The MathWorks, I. (2022e, May 19). *Simulate stationary processes*. <https://se.mathworks.com/help/econ/engles-arch-test.html>

-
- TowardsDataScience. (n.d.). *Types of convolution kernels : Simplified*. Retrieved, from <https://towardsdatascience.com/types-of-convolution-kernels-simplified-f040cb307c37>
- Varsamopoulos, S., Bertels, K. & Almudever, C. (2018). *Designing neural network based decoders for surface codes*.
- Wei Zhang, Gaoliang Peng & Chuanhao Li. (2017). Bearings fault diagnosis based on convolutional neural networks with 2-d representation of vibration signals as input. *MATEC Web Conf.*, 95, 13001. <https://doi.org/10.1051/mateconf/20179513001>
- WHO. (2022, February 20). *Who coronavirus (covid-19) dashboard*. <https://covid19.who.int/data>
- Zain, Z. M. & Alturki, N. M. (2021). On the true number of covid-19 infections: Effect of sensitivity, specificity and number of tests on prevalence ratio estimation. *Journal of Control Science and Engineering*, 2021.
- Zhang, J., Peng, Y., Ren, B. & Li, T. (2021). Pm2.5 concentration prediction based on cnn-bilstm and attention mechanism. *Algorithms*, 14, 208. <https://doi.org/10.3390/a14070208>
- Zhang, W., Huang, G., Wang, G. & Wang, Y. (2019). Prediction high frequency parameters based on neural network. *IOP Conference Series: Materials Science and Engineering*, 631, 052035.
- Zhang, Z. (2016). Derivation of backpropagation in convolutional neural network (cnn). *University of Tennessee, Knoxville*.
- Zhu, F. (2012). Modeling time series of counts with com-poisson ingarch models. *Mathematical and Computer Modelling*, 56(9), 191–203. <https://doi.org/https://doi.org/10.1016/j.mcm.2011.11.069>

Appendix

A Figures and tables




Figure 41: Forecast with ten CNN-LSTM models from November 11th to December 8th 2020 on the Norwegian new-cases data. The previous 264 days there used as training data.

Table 5: All forecast results from the mean prediction of ten CNN-LSTM models and the mean RRMSE of the ten (individual) models. All 28 day predictions used prediction scheme 1, while prediction scheme 2 was used for all seven day predictions.

Data set	Last training date	Sample size	Days ahead prediction	CNN-LSTM mean prediction		Ten CNN-LSTM models		
				RRMSE	MAPE	Mean RRMSE	Mean MAPE	
Norway	30.05 2020	100	7	64.23	88.08	70.75	92.84	
			28	73.05	87.49	75.77	89.19	
		50	7	64.52	40.73	65.04	42.56	
			28	52.72	41.21	53.58	41.95	
	10.11 2020	264	7	8.03	7.91	12.68	12.36	
			28	17.91	16.96	18.73	17.81	
		100	7	10.63	10.73	11.16	11.26	
			28	25.01	25.98	25.48	26.07	
	16.12 2020	300	7	14.18	11.83	15.43	13.45	
			28	28.85	22.54	29.26	22.76	
		100	7	7.41	6.65	12.72	12.38	
			28	29.22	22.85	30.24	23.82	
	04.07 2021	500	7	7.85	7.72	11.23	10.38	
			28	17.06	11.72	17.57	12.66	
		100	7	11.56	11.95	13.11	13.03	
			28	30.38	20.10	30.78	21.08	
	16.01 2022	696	7	23.05	21.81	23.17	21.81	
			28	18.76	16.57	20.72	18.27	
		100	7	24.04	18.24	25.53	21.51	
			28	29.56	30.89	29.89	31.11	
	Global	12.05 2020	100	7	10.72	9.85	10.79	9.89
				28	7.34	6.33	7.54	6.44
			50	7	14.38	13.47	14.41	13.47
				28	9.86	8.51	9.93	8.56
17.07 2020		196	7	7.00	4.18	7.16	4.51	
			28	4.58	3.21	4.86	3.59	
		100	7	5.90	3.59	6.17	3.93	
			28	4.17	3.20	4.46	3.41	
29.10 2020		300	7	4.90	3.04	5.03	3.93	
			28	3.79	3.12	4.21	3.31	
		100	7	11.19	9.90	11.27	9.94	
			28	12.19	11.07	12.28	11.14	
17.05 2021		500	7	7.97	6.79	8.05	7.02	
			28	9.03	7.59	9.13	7.71	
		100	7	9.44	8.41	9.49	8.46	
			28	10.46	9.34	10.53	9.39	
14.01 2022		500	7	13.36	12.99	13.42	12.99	
			28	9.29	7.89	9.43	8.19	
		100	7	20.43	20.33	20.46	20.33	
			28	15.30	13.83	15.36	13.87	
14.01 2022		743	7	18.24	17.88	18.26	17.88	
			28	38.96	39.08	39.03	39.08	
		100	7	12.17	12.08	12.80	12.32	
			28	20.59	20.08	21.08	20.42	
14.01 2022	100	7	10.35	10.30	10.58	10.40		
		28	20.64	20.11	20.75	20.22		
	50	7	11.38	10.64	11.76	11.08		
		28	85.18	82.83	85.18	82.83		




Figure 42: RRMSE scores for each model from Table 3.

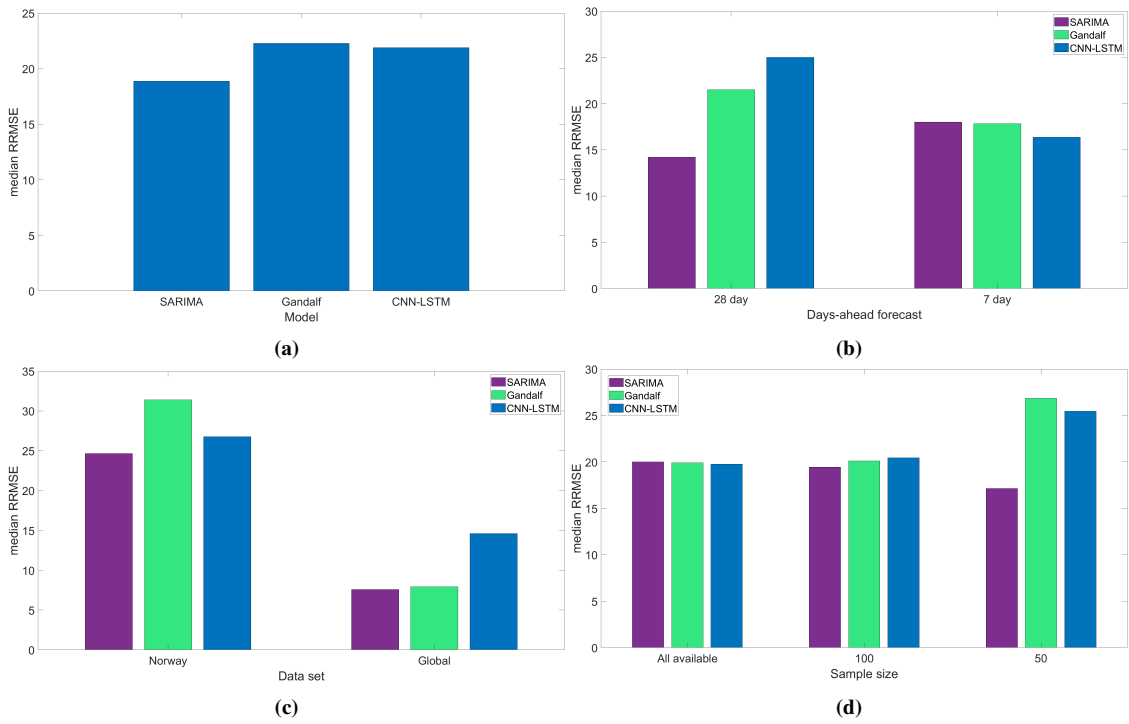


Figure 43: Comparison of models based on RRMSE scores in Table 3

B Code

In this thesis, three programming languages were utilized. Matlab was used for the econometric model and for all the generated plots, while Python was used for the CNN-LSTM model. Finally, R was used for pre-processing of the Norwegian and the global data set. To plot the results from the CNN-LSTM model, the forecasts made in Python was saved and loaded in Matlab. It should be noted that most of the functionality is commented out. To run the code, simply uncomment the desired section of code. Additionally, the file paths should be replaced to fit the user's computer. The code files and the two data sets are available on request.

R-code

The following code was used to pre-process the global and the Norwegian data new cases data sets.

```
1 # Pre-process global
2 df = read.csv("WHO-COVID-19-global-data.csv")
3 df = df[, c(1, 2, 5)]
4 colnames(df) = c("Date", "Country", "New_cases")
5 # Cut-off at February 20th 2022
6 df$Date = as.Date(df$Date, format = "%Y-%m-%d")
7 last_date = as.Date("2022-02-20", format = "%Y-%m-%d")
8 df$Date_diff = last_date - df$Date
9 df_ = df[df$Date_diff >= 0,]
10
11 new_cases_one_country = df_[df$Country == "AF",]
12 worldwide_aggregated = new_cases_one_country[
13   !is.na(new_cases_one_country$Country), 3]
14
15 countries_list = unique(df[df$Country != "AF",]$Country)
16 countries_list = countries_list[!is.na(countries_list)]
17
18 for (country in countries_list){
19   new_cases_one_country = df_[df$Country == country,]
20   worldwide_aggregated = worldwide_aggregated + new_cases_one_country[
21     !is.na(new_cases_one_country$Country), 3]
22 }
23 write.csv(worldwide_aggregated[-1], "worldwide_aggregated_new.csv",
24   row.names = FALSE)
25
26 # Pre-process Norway
27 df = read.csv("antall-meldte-covid-19-t.csv", sep = ";")
28 df = df[, c(1, 3)]
29 colnames(df) = c("Date", "New_cases")
30 # Cut-off at February 20th 2022
31 df$Date = as.Date(df$Date, format = "%d.%m.%Y")
32 last_date = as.Date("2022-02-20", format = "%Y-%m-%d")
33 df$Date_diff = last_date - df$Date
34 df_ = df[df$Date_diff >= 0,]
35 write.csv(df$Nye.tilfeller, "new_cases_Norway.csv", row.names = FALSE)
```

Python-code

The following code was used to derive all predictions from the CNN-LSTM model. The results were then transferred to the Matlab file and was from there plotted.

```
1 import random
2 import time
3 from numpy import array
4 import numpy as np
5
6 import pandas as pd
7 import os
8 import tensorflow as tf
9 from keras.models import Sequential
10 from keras.layers import LSTM
11 from keras.layers import Dense
12 from keras.layers import Dropout
13 from keras.layers import Flatten
14 from keras.layers import RepeatVector
15 from keras.layers import TimeDistributed
16 from keras.layers.convolutional import Conv1D
17 from keras.layers.convolutional import MaxPooling1D
18 from sklearn.preprocessing import MinMaxScaler
19 import matplotlib.pyplot as plt
20
21
22 seed_value = 1234
23
24 # Set seed in all environments to help reproducibility
25 def set_all_seeds(seed_value):
26     # 1. Set `PYTHONHASHSEED` environment variable at a fixed value
27     os.environ['PYTHONHASHSEED']=str(seed_value)
28     # 2. Set `python` built-in pseudo-random generator at a fixed value
29     random.seed(seed_value)
30     # 3. Set `numpy` pseudo-random generator at a fixed value
31     np.random.seed(seed_value)
32     # 4. Set `tensorflow` pseudo-random generator at a fixed value
33     tf.random.set_seed(seed_value)
34
35
36 # import data
37
38 # Set current working directory one layer up
39 path_parent = os.path.dirname(os.getcwd())
40 os.chdir(path_parent)
41
42 ts_global = np.array(pd.read_csv(os.getcwd() + r"\Data\new_cases_global.csv"))
43 ts_norway = np.array(pd.read_csv(os.getcwd() + r"\Data\new_cases_Norway.csv"))
44
45
46 def split_sequence(sequence, n_steps):
47     # Formats the data for supervised learning
48     X, y = list(), list()
49     for i in range(len(sequence)):
50         # find the end of this pattern
51         end_ix = i + n_steps
52         # check if we are beyond the sequence
53         if end_ix > len(sequence) - 1:
54             break
55         # gather input and output parts of the pattern
56         seq_x, seq_y = sequence[i:end_ix], sequence[end_ix]
57         X.append(seq_x)
58         y.append(seq_y)
59     return array(X), array(y)
60
61 def k_step_prediction(k, x_train, model, scaler, prev_days = 7):
62     '''
63     Iteratively performs a one step prediction and use the newly
64     predicted as input to forecast the next day
65     :param k: amount of days ahead to predict
66     :param x_train: the time series used
```

```

67     :param model: pre-trained model
68     :param scaler: the pre-trained min-max scaler,
69     used to revert the prediction back to normal scale
70     :param prev_days: Amount of previous days used to predict the next day.
71     :return: k-step prediction
72     '''
73     predictions = []
74     x_input = x_train[-prev_days:] # last prev_days days
75     x_input = x_input.reshape((1, prev_days, 1))
76     yhat = model.predict(x_input, verbose=0)[0][0].reshape(1,1)
77     predictions.append(yhat)
78     for i in range(1, k):
79         x_input = x_input.reshape(prev_days, 1)
80         # remove first day and add prediction to the end of the input
81         x_input = np.concatenate((x_input[1:], yhat), axis = 0)
82         x_input = x_input.reshape((1, prev_days, 1))
83         yhat = model.predict(x_input, verbose=0)[0][0].reshape(1,1)
84         predictions.append(yhat)
85
86     for i in range(len(predictions)):
87         predictions[i] = scaler.inverse_transform(predictions[i])[0][0]
88     return predictions
89
90
91 def train_model(X, y, days_ahead, prev_days = 7):
92     '''
93     :param X: Data matrix
94     :param y: Response vector
95     :param days_ahead: days ahead to be forecasted after last input
96     :param n_steps: amount of previous days used as input in forecasting
97     :return:
98     '''
99     model = Sequential()
100    model.add(Conv1D(filters=64, kernel_size=3, activation='relu', ...
101                input_shape=(prev_days, 1)))
102    model.add(Conv1D(filters=64, kernel_size=3, activation='relu'))
103    model.add(MaxPooling1D(pool_size=1, strides=1))
104    model.add(Dropout(0.2, seed=seed_value))
105    model.add(Flatten())
106    model.add(RepeatVector(1)) # days_ahead
107    model.add(LSTM(200, return_sequences=True))
108    model.add(TimeDistributed(Dense(50))) # 100
109    model.add(TimeDistributed(Dense(1)))
110    model.compile(optimizer='adam', loss='mse') # 'mean_absolute_percentage_error'
111    # fit model
112    history = model.fit(X, y, epochs=100, verbose=0, batch_size=22) # skal v re ...
113    472, 22
114    loss = history.history['loss']
115    # print(model.summary()) # uncomment to view amount of variables
116    return model, loss
117
118 def prediction_global(ts, endpoint, prev_days, len_train = -1, days_ahead = 7):
119     # The new version uses the repeat vector layer, and uses the
120     # Time Distributed Wrapper at the end insted of at the beginning.
121
122     '''
123     :param ts: Input time series to be forecasted
124     :param endpoint: Last index used for training the model
125     :param prev_days: Amount of days used to predict the next days
126     :param len_train: Amount of days back from endpoint used to train the model
127     :param days_ahead: Amount of days forecasted after endpoint
128     :return: days_ahead-forecast
129     '''
130     set_all_seeds(seed_value)
131     if len_train == -1:
132         len_train = endpoint
133     scaler = MinMaxScaler()
134     # only use train data to fit scaler:
135     x_train = ts[endpoint - len_train:endpoint].reshape(len_train, 1)
136     scaler.fit(x_train)
137     x_train = scaler.transform(x_train)
138     X, y = split_sequence(x_train, prev_days)

```

```

138 # reshape from [samples, timesteps] into
139 # [samples, subsequences, timesteps, features]
140 X = X.reshape((X.shape[0], prev_days, 1))
141
142 preds = np.zeros((10, days_ahead))
143 loss = []
144 for i in range(10):
145     start_time = time.time()
146     model, loss_i = train_model(X, y, days_ahead, prev_days=prev_days)
147     loss.append(loss_i)
148     print("Time used to fit model", i+1, ": ", (time.time() - start_time), " ...
149           sec")
150     start_time = time.time()
151     preds[i] = np.round_(k_step_prediction(
152         days_ahead, x_train, model, scaler, prev_days=prev_days))
153     print("Time used to forecast with model", i + 1, ": ",
154           (time.time() - start_time), " sec")
155
156 np.savetxt(os.getcwd() + r"\Predictions/glob_preds_" + str(days_ahead) +
157            "_days_ahead_from_" + str(endpoint) + "_with_" + str(len_train) +
158            "_train_data.csv", preds)
159
160 # This part can be uncommented to inspect and save the mean loss function
161 # loss = array(loss)
162 # mean_loss = loss.mean(axis = 0)
163 # print(mean_loss.shape)
164 # np.savetxt(os.getcwd() +
165 #            r"\Predictions/mean_loss_on_original_test_data.csv", mean_loss)
166 # print(mean_loss.shape)
167 # plt.plot(mean_loss)
168 # plt.ylabel('loss')
169 # plt.xlabel('epoch')
170 # plt.show()
171 return preds
172
173 def prediction_norway(ts, endpoint, prev_days, len_train = -1, days_ahead = 7):
174     # The new version uses the repeat vector layer, and uses the
175     # Time Distributed Wrapper at the end instead of at the beginning.
176     # Thus, the model trains on a seven day output?
177
178     """
179     :param ts: Input time series to be forecasted
180     :param endpoint: Last index used for training the model
181     :param prev_days: Amount of days used to predict the next days
182     :param len_train: Amount of days back from endpoint used to train the model
183     :param days_ahead: Amount of days forecasted after endpoint
184     :return: days_ahead-forecast
185     """
186     set_all_seeds(seed_value)
187     if len_train == -1:
188         len_train = endpoint
189     scaler = MinMaxScaler()
190     x_train = ts[endpoint - len_train:endpoint].reshape(len_train, 1) # 17th July
191     scaler.fit(x_train)
192     x_train = scaler.transform(x_train)
193     X, y = split_sequence(x_train, prev_days)
194     # reshape from [samples, timesteps] into [samples, subsequences, timesteps, ...
195     # features]
196     X = X.reshape((X.shape[0], prev_days, 1))
197
198 preds = np.zeros((10, days_ahead))
199 for i in range(10):
200     start_time = time.time()
201     model, loss = train_model(X, y, days_ahead, prev_days=prev_days)
202     print("Time used to fit model", i+1, ": ", (time.time() - start_time), " ...
203           sec")
204     start_time = time.time()
205     preds[i] = np.round_(k_step_prediction(
206         days_ahead, x_train, model, scaler, prev_days=prev_days))
207     print("Time used to forecast with model", i + 1, ": ",
208           (time.time() - start_time), " sec")
209
210 np.savetxt(os.getcwd() + r"\Predictions/nor_preds_" + str(days_ahead) +

```

```

208         "_days_ahead_from_" + str(endpoint) + "_with_" + str(len_train) +
209         "_train_data.csv", preds)
210     return preds
211
212
213 def single_pred_with_test_data(days_ahead, x_train_and_test, model, scaler, ...
    prev_days = 7):
214     """
215     Performs a series of one step predictions, one step at a time to generate a ...
        k-step prediction.
216     At each iteration, the newly observed day is incorporated into the input of ...
        the prediction.
217     :return: k-step prediction of covid-19 new cases
218     """
219     predictions = []
220     length_relevant_data = prev_days + days_ahead - 1 # how far back to start the ...
        input
221
222     # the first input to the prediction is the last prev_days elements of the ...
        train data
223     for i in range(0, days_ahead-1):
224         x_input = x_train_and_test[-length_relevant_data + ...
            i:-length_relevant_data + prev_days + i]
225         x_input = x_input.reshape((1, prev_days, 1))
226         yhat = model.predict(x_input, verbose=0)[0][0].reshape(1, 1)
227         predictions.append(yhat)
228
229     x_input = x_train_and_test[-prev_days:]
230     x_input = x_input.reshape((1, prev_days, 1))
231     yhat = model.predict(x_input, verbose=0)[0][0].reshape(1, 1)
232     predictions.append(yhat)
233
234     for i in range(len(predictions)):
235         predictions[i] = scaler.inverse_transform(predictions[i])[0][0]
236     return predictions
237
238 def one_step_preds_global(ts, endpoint, prev_days, len_train = -1, days_ahead = 7):
239     set_all_seeds(seed_value)
240     if len_train == -1:
241         len_train = endpoint
242     scaler = MinMaxScaler()
243     x_train = ts[endpoint - len_train:endpoint].reshape(len_train, 1) # last ...
        len_train days used to predict
244     scaler.fit(x_train)
245     x_train = scaler.transform(x_train)
246     # test data on day d is used to predict new cases on day d+1
247     x_train_and_test = ts[endpoint - ...
        len_train:(endpoint+days_ahead-1)].reshape(len_train+days_ahead-1, 1)
248     x_train_and_test = scaler.transform(x_train_and_test)
249     # split train data into input-output matrix
250     X, y = split_sequence(x_train, prev_days)
251     X = X.reshape((X.shape[0], prev_days, 1))
252
253     # Make 10 predictions and take mean forecast, to reduce variability in forecast
254     preds = np.zeros((10, days_ahead))
255
256     for i in range(10):
257         start_time = time.time()
258         model, loss = train_model(X, y, days_ahead, prev_days=prev_days) # The ...
            model is only trained in the train data!
259         print("Time used to fit model", i+1, ":", " ", (time.time() - start_time), " ...
            sec")
260         start_time = time.time()
261         preds[i] = np.round_(single_pred_with_test_data(days_ahead, ...
            x_train_and_test, model, scaler, prev_days=prev_days))
262         print("Time used to forecast with model", i+1, ":", " ", (time.time() - ...
            start_time), " sec")
263
264     np.savetxt(os.getcwd() + r"\Predictions/glob_test_preds_" + str(days_ahead) +
        "_days_ahead_from_" + str(endpoint) + "_with_" + str(len_train) + ...
        "_train_data.csv", preds)
265
266     return preds
267

```

```

268 def one_step_preds_norway(ts, endpoint, prev_days, len_train = -1, days_ahead = 7):
269     set_all_seeds(seed_value)
270     if len_train == -1:
271         len_train = endpoint
272     scaler = MinMaxScaler()
273     x_train = ts[endpoint - len_train:endpoint].reshape(len_train, 1) # last ...
274         len_train days used to predict
275     scaler.fit(x_train)
276     x_train = scaler.transform(x_train)
277     # test data on day d is used to predict new cases on day d+1
278     x_train_and_test = ts[endpoint - ...
279         len_train:(endpoint+days_ahead-1)].reshape(len_train+days_ahead-1, 1)
280     x_train_and_test = scaler.transform(x_train_and_test)
281     # split train data into input-output matrix
282     X, y = split_sequence(x_train, prev_days)
283     X = X.reshape((X.shape[0], prev_days, 1))
284
285     # Make 10 predictions and take mean forecast, to reduce variability in forecast
286     preds = np.zeros((10, days_ahead))
287     for i in range(10):
288         start_time = time.time()
289         model, loss = train_model(X, y, days_ahead, prev_days=prev_days) # The ...
290             model is only trained in the train data!
291         print("Time used to fit model", i+1, ":", (time.time() - start_time), " ...
292             sec")
293         start_time = time.time()
294         preds[i] = np.round_(single_pred_with_test_data(days_ahead, ...
295             x_train_and_test, model, scaler, prev_days=prev_days))
296         print("Time used to forecast with model", i+1, ":", (time.time() - ...
297             start_time), " sec")
298
299     np.savetxt(os.getcwd() + r"\Predictions/nor_test_preds_" + str(days_ahead) +
300         "_days_ahead_from_" + str(endpoint) + "_with_" + str(len_train) + ...
301         "_train_data.csv", preds)
302
303     return preds
304
305 def one_step_preds_on_train_global(ts, endpoint, prev_days, len_train = -1, ...
306     days_ahead = 7):
307     set_all_seeds(seed_value)
308     if len_train == -1:
309         len_train = endpoint
310     scaler = MinMaxScaler()
311     x_train = ts[endpoint - len_train:endpoint].reshape(len_train, 1) # last ...
312         len_train days used to predict
313     scaler.fit(x_train)
314     x_train = scaler.transform(x_train)
315     # test data on day d is used to predict new cases on day d+1
316     x_train_and_test = ts[endpoint - ...
317         len_train:(endpoint+days_ahead-1)].reshape(len_train+days_ahead-1, 1)
318     x_train_and_test = scaler.transform(x_train_and_test)
319     # split train data into input-output matrix
320     X, y = split_sequence(x_train, prev_days)
321     X = X.reshape((X.shape[0], prev_days, 1))
322     # Make 10 predictions and take mean forecast, to reduce variability in forecast
323     preds = np.zeros((10, days_ahead))
324     for i in range(10):
325         start_time = time.time()
326         model, loss = train_model(X, y, days_ahead, prev_days=prev_days) # The ...
327             model is only trained in the train data!
328         print("Time used to fit model", i+1, ":", (time.time() - start_time), " ...
329             sec")
330         start_time = time.time()
331         preds[i] = np.round_(single_pred_with_test_data(days_ahead, ...
332             x_train_and_test[:-days_ahead], model, scaler, prev_days=prev_days))
333         print("Time used to forecast with model", i+1, ":", (time.time() - ...
334             start_time), " sec")
335
336     np.savetxt(os.getcwd() + r"\Predictions/glob_test_preds_train_" + ...
337         str(days_ahead) +
338         "_from_" + str(endpoint) + "_with_" + str(len_train) + ...
339         "_train_data.csv", preds)
340
341     return preds
342
343

```

```

325 def one_step_preds_on_train_norway(ts, endpoint, prev_days, len_train = -1, ...
326     days_ahead = 7):
327     set_all_seeds(seed_value)
328     if len_train == -1:
329         len_train = endpoint
330     scaler = MinMaxScaler()
331     x_train = ts[endpoint - len_train:endpoint].reshape(len_train, 1) # last ...
332         len_train days used to predict
333     scaler.fit(x_train)
334     x_train = scaler.transform(x_train)
335     # test data on day d is used to predict new cases on day d+1
336     x_train_and_test = ts[endpoint - ...
337         len_train:(endpoint+days_ahead-1)].reshape(len_train+days_ahead-1, 1)
338     x_train_and_test = scaler.transform(x_train_and_test)
339     # split train data into input-output matrix
340     X, y = split_sequence(x_train, prev_days)
341     X = X.reshape((X.shape[0], prev_days, 1))
342     # Make 10 predictions and take mean forecast, to reduce variability in forecast
343     preds = np.zeros((10, days_ahead))
344     for i in range(10):
345         start_time = time.time()
346         model, loss = train_model(X, y, days_ahead, prev_days=prev_days) # The ...
347             model is only trained in the train data!
348         print("Time used to fit model", i+1, ":", (time.time() - start_time), " ...
349             sec")
350         start_time = time.time()
351         preds[i] = np.round(single_pred_with_test_data(days_ahead, ...
352             x_train_and_test[:-days_ahead], model, scaler, prev_days=prev_days))
353         print("Time used to forecast with model", i+1, ":", (time.time() - ...
354             start_time), " sec")
355
356     np.savetxt(os.getcwd() + r"\Predictions/nor_test_preds_on_train_" + ...
357         str(days_ahead) +
358         "_from_" + str(endpoint) + "_with_" + str(len_train) + ...
359         "_train_data.csv", preds)
360
361     return preds
362
363 def all_in_one_global(endpoint, days_ahead, len_train):
364     # Wrapper function to systemize predictions results
365     # The function will output both the one-step prediction and the days_ahead ...
366     prediction
367     prediction_global(ts=ts_global, endpoint=endpoint, prev_days=7, ...
368         days_ahead=days_ahead, len_train=len_train)
369     one_step_preds_global(ts=ts_global, endpoint=endpoint, prev_days=7, ...
370         days_ahead = days_ahead, len_train = len_train)
371
372
373 def all_in_one_norway(endpoint, days_ahead, len_train):
374     # Wrapper function to systemize predictions,
375     # The function will output both the one-step prediction and the days_ahead ...
376     prediction
377     one_step_preds_norway(ts=ts_norway, endpoint=endpoint, prev_days=7, ...
378         days_ahead = days_ahead, len_train = len_train)
379     prediction_norway(ts=ts_norway, endpoint=endpoint, prev_days=7, days_ahead = ...
380         days_ahead, len_train = len_train)
381
382
383 def accuracy_with_simulations_global(ts, simulations, endpoint, days_ahead, ...
384     prev_days = 7, len_train = -1):
385     '''
386     Forecasts with scheme 1, but use simulations instead of observations
387     :param ts: The time series of interest
388     :param simulations: Simulated observations from Gandalf model, imported from ...
389         Matlab
390     :param endpoint: Last index of dataset used in training
391     :param prev_days: how many days put in to the model (7 is default)
392     :param days_ahead: days ahead forecast
393     :return: the RRMSEs and MAPEs associated with each simulation
394     '''
395     set_all_seeds(seed_value)
396
397     simulations = np.array(simulations)
398     if len_train == -1:
399         len_train = endpoint

```

```

381     scaler = MinMaxScaler()
382     x_train = ts[endpoint - len_train:endpoint].reshape(len_train, 1) # last ...
383         len_train days used to predict
384     scaler.fit(x_train)
385     x_train = scaler.transform(x_train)
386     # split train data into input-output matrix
387     X, y = split_sequence(x_train, prev_days)
388     X = X.reshape((X.shape[0], prev_days, 1))
389
390     # Make 10 predictions and take mean forecast, to reduce variability in forecast
391     RRMSEs = np.zeros(simulations.shape[0])
392     MAPEs = np.zeros(simulations.shape[0])
393
394     models = [] # This saves a lot of time!
395     for j in range(10):
396         model, loss = train_model(X, y, days_ahead,
397                                 prev_days=prev_days) # The model is only ...
398                                     trained in the train data!
399
400     models.append(model)
401
402     for i in range(simulations.shape[0]):
403         start_time = time.time()
404         simulation = scaler.transform(simulations[i].reshape(-1, 1))
405         input_in_prediction = np.concatenate((x_train, simulation[0:-1]), axis=0)
406         ten_preds = np.zeros((10, days_ahead))
407         for j in range(10):
408             ten_preds[j] = np.round_(
409                 single_pred_with_test_data(days_ahead, input_in_prediction, ...
410                                             models[j], scaler, prev_days=prev_days))
411         mean_pred = np.mean(ten_preds, axis = 0)
412
413         simulation = scaler.inverse_transform(simulation.reshape(-1, 1))
414         mean_sim = simulation.mean()
415         MSE = np.square(np.subtract(simulation, mean_pred)).mean()
416         RRMSEs[i] = np.round_(np.sqrt(MSE) / mean_sim * 100, decimals=3)
417         mape = tf.keras.losses.MeanAbsolutePercentageError()
418         MAPEs[i] = np.round_(mape(simulation, mean_pred).numpy(), decimals=3)
419         print("Time used to predict with simulation", i, ":", (time.time() - ...
420                 start_time), " sec")
421
422     np.savetxt(os.getcwd() + r"\Predictions/RRMSEs_from_simulation_global_" + ...
423               str(days_ahead) +
424               "_days_ahead_with_" + str(len_train) + "_train_data_from" + ...
425               str(endpoint) + ".csv", RRMSEs)
426     np.savetxt(os.getcwd() + r"\Predictions/MAPEs_from_simulation_global_" + ...
427               str(days_ahead) +
428               "_days_ahead_with_" + str(len_train) + "_train_data_from" + ...
429               str(endpoint) + ".csv", MAPEs)
430
431 def predictions_with_simulations_global(ts, simulations, endpoint, days_ahead, ...
432     prev_days = 7, len_train = -1):
433     """
434     Forecasts with scheme 1, but use simulations instead of observations, the ...
435     accuracy measures are returned
436     :param ts: The time series of interest
437     :param simulations: Simulated observations from Gandalf model, imported from ...
438         Matlab
439     :param endpoint: Last index of dataset used in training
440     :param prev_days: how many days put in to the model (7 is default)
441     :param days_ahead: days ahead forecast
442     :return: the RRMSEs and MAPEs associated with each simulation
443     """
444     set_all_seeds(seed_value)
445
446     simulations = np.array(simulations)
447     if len_train == -1:
448         len_train = endpoint
449     scaler = MinMaxScaler()
450     x_train = ts[endpoint - len_train:endpoint].reshape(len_train, 1) # last ...
451         len_train days used to predict
452     scaler.fit(x_train)
453     x_train = scaler.transform(x_train)

```



```

442 # split train data into input-output matrix
443 X, y = split_sequence(x_train, prev_days)
444 X = X.reshape((X.shape[0], prev_days, 1))
445
446 # Make 10 predictions and take mean forecast, to reduce variability in forecast
447 preds_matlab = np.zeros((3, 28))
448
449 start_time = time.time()
450
451 models = [] # This saves a lot of time!
452 for j in range(10):
453     model, loss = train_model(X, y, days_ahead,
454                             prev_days=prev_days) # The model is only ...
455                                                     trained in the train data!
456
457     models.append(model)
458 print("time to fit ten models: ", time.time() - start_time)
459
460 for i in range(3): # for the first three predictions
461     preds = np.zeros((10, 28))
462     print("simulations[i]", simulations[i])
463     simulation = scaler.transform(simulations[i].reshape(-1, 1))
464     input_in_prediction = np.concatenate((x_train, simulation[0:-1]), axis=0)
465     for j in range(10):
466         pred = np.round_(
467             single_pred_with_test_data(days_ahead, input_in_prediction, ...
468                                       models[j], scaler, prev_days=prev_days))
469         preds[j] = pred
470     preds_matlab[i] = np.mean(preds, axis = 0)
471 print("Time used to predict all simulations: ", (time.time() - start_time), " ...
472       sec")
473
474 print(preds_matlab, preds_matlab.shape)
475 np.savetxt(os.getcwd() + r"\Predictions/pred_three_preds_on_sims.csv", ...
476           preds_matlab)
477
478 def predict_weekday_no_test(Xs, Ys, x_trains, days_ahead, scaler, prev_days = 7):
479     '''
480     The function predicts a days_ahead forecast for seven models, one for each ...
481     weekday.
482     :param datas: list of seven scaled train data, one containing all Mondays etc...
483     :param days_ahead: how many days ahead to predict
484     :param scaler: pre-trained scaler function
485     :param prev_days: how many days used to predict preceding day for each submodel
486     :return: days_ahead forecast
487     '''
488     prediction = np.zeros(days_ahead)
489     prediction_days = [i for i in range(days_ahead)] # sequence of numbers
490     for i in range(len(Xs)):
491         days_to_predict = prediction_days[i::7] # divides what days to predict by ...
492             each submodel, the length is the amount of days to predict
493         X = Xs[i]
494         y = Ys[i]
495         x_train = x_trains[i]
496         model, loss = train_model(X, y, len(days_to_predict), prev_days)
497         prediction[days_to_predict] = k_step_prediction(len(days_to_predict), ...
498               x_train, model, scaler, prev_days) # prediction of all of a single ...
499             weekday
500     return prediction
501
502 def weekday_models_no_test_data_global(ts, endpoint, days_ahead, prev_days = 7, ...
503 len_train = -1):
504     set_all_seeds(seed_value)
505     if len_train == -1:
506         len_train = endpoint
507     scaler = MinMaxScaler()
508     x_train = ts[endpoint - len_train:endpoint].reshape(len_train, 1) # use ...
509         train data to fit scaler
510     scaler.fit(x_train)

```

```

505     x_train = scaler.transform(x_train)
506     x_train_1 = x_train[::7] # Every 7th day
507     x_train_2 = x_train[1::7]
508     x_train_3 = x_train[2::7]
509     x_train_4 = x_train[3::7]
510     x_train_5 = x_train[4::7]
511     x_train_6 = x_train[5::7]
512     x_train_7 = x_train[6::7]
513     print(x_train_1.shape, x_train_2.shape, x_train_3.shape, x_train_4.shape, ...
           x_train_5.shape, x_train_6.shape, x_train_7.shape)
514
515     x_trains = [x_train_1, x_train_2, x_train_3, x_train_4, x_train_5, x_train_6, ...
                x_train_7]
516     X_1, y_1 = split_sequence(x_train_1, prev_days)
517     X_2, y_2 = split_sequence(x_train_2, prev_days)
518     X_3, y_3 = split_sequence(x_train_3, prev_days)
519     X_4, y_4 = split_sequence(x_train_4, prev_days)
520     X_5, y_5 = split_sequence(x_train_5, prev_days)
521     X_6, y_6 = split_sequence(x_train_6, prev_days)
522     X_7, y_7 = split_sequence(x_train_7, prev_days)
523
524     Ys = [y_1, y_2, y_3, y_4, y_5, y_6, y_7]
525     # reshape from [samples, timesteps] into [samples, subsequences, timesteps, ...
           features]
526     X_1 = X_1.reshape((X_1.shape[0], prev_days, 1))
527     X_2 = X_2.reshape((X_2.shape[0], prev_days, 1))
528     X_3 = X_3.reshape((X_3.shape[0], prev_days, 1))
529     X_4 = X_4.reshape((X_4.shape[0], prev_days, 1))
530     X_5 = X_5.reshape((X_5.shape[0], prev_days, 1))
531     X_6 = X_6.reshape((X_6.shape[0], prev_days, 1))
532     X_7 = X_7.reshape((X_7.shape[0], prev_days, 1))
533     Xs = [X_1, X_2, X_3, X_4, X_5, X_6, X_7]
534     preds = np.zeros((10, days_ahead))
535
536     for i in range(10):
537         start_time = time.time()
538         preds[i] = np.round_(predict_seven_no_test(Xs, Ys, x_trains, days_ahead, ...
           scaler, prev_days))
539         print("Time used to fit and forecast with model", i + 1, ":", " ", ...
              (time.time() - start_time), " sec")
540
541     np.savetxt(os.getcwd() + r"\Predictions/glob_weekday_preds_" + ...
               str(days_ahead) +
542               "_days_ahead_from_" + str(endpoint) + "_with_" + str(len_train) + ...
               "_train_data.csv", preds)
543
544     def predict_weekday_with_test(Xs, Ys, x_train_and_tests, days_ahead, scaler, ...
           prev_days = 7):
545         """
546         The function predicts a days_ahead forecast for seven models, one for each ...
           weekday.
547         :param datas: list of seven scaled train data, one containing all Mondays etc...
548         :param days_ahead: how many days ahead to predict
549         :param scaler: pre-trained scaler function
550         :param prev_days: how many days used to predict preceding day for each submodel
551         :return: days_ahead forecast
552         """
553
554         prediction = np.zeros(days_ahead)
555         prediction_days = [i for i in range(days_ahead)] # sequence of numbers
556         for i in range(len(Xs)):
557             days_to_predict = prediction_days[i::7] # divides what days to predict by ...
                each submodel, the length is the amount of days to predict
558             X = Xs[i]
559             y = Ys[i]
560             x_train_and_test = x_train_and_tests[i]
561             model, loss = train_model(X, y, len(days_to_predict), prev_days)
562             # prediction of all of a single weekday:
563             prediction[days_to_predict] = \
564                 single_pred_with_test_data(len(days_to_predict), x_train_and_test, ...
                model, scaler, prev_days)
565         return prediction
566

```

```

567 def weekday_models_with_test_data_global(ts, endpoint, days_ahead, prev_days = 7, ...
    len_train = -1):
568     set_all_seeds(seed_value)
569     if len_train == -1:
570         len_train = endpoint
571     scaler = MinMaxScaler()
572     x_train = ts[endpoint - len_train:endpoint].reshape(len_train, 1) # use ...
        train data to fit scaler
573     scaler.fit(x_train)
574     x_train = scaler.transform(x_train)
575     x_train_1 = x_train[::7] # Every 7th day
576     x_train_2 = x_train[1::7]
577     x_train_3 = x_train[2::7]
578     x_train_4 = x_train[3::7]
579     x_train_5 = x_train[4::7]
580     x_train_6 = x_train[5::7]
581     x_train_7 = x_train[6::7]
582
583     x_train_and_test = ts[endpoint - ...
        len_train:(endpoint+days_ahead-1)].reshape(len_train+days_ahead-1, 1)
584     x_train_and_test = scaler.transform(x_train_and_test)
585     x_train_and_test_1 = x_train_and_test[::7] # Every 7th day
586     x_train_and_test_2 = x_train_and_test[1::7]
587     x_train_and_test_3 = x_train_and_test[2::7]
588     x_train_and_test_4 = x_train_and_test[3::7]
589     x_train_and_test_5 = x_train_and_test[4::7]
590     x_train_and_test_6 = x_train_and_test[5::7]
591     x_train_and_test_7 = x_train_and_test[6::7]
592
593     x_train_and_tests = [x_train_and_test_1, x_train_and_test_2, ...
        x_train_and_test_3, x_train_and_test_4,
        x_train_and_test_5, x_train_and_test_6, x_train_and_test_7]
594
595
596     X_1, y_1 = split_sequence(x_train_1, prev_days)
597     X_2, y_2 = split_sequence(x_train_2, prev_days)
598     X_3, y_3 = split_sequence(x_train_3, prev_days)
599     X_4, y_4 = split_sequence(x_train_4, prev_days)
600     X_5, y_5 = split_sequence(x_train_5, prev_days)
601     X_6, y_6 = split_sequence(x_train_6, prev_days)
602     X_7, y_7 = split_sequence(x_train_7, prev_days)
603
604     Ys = [y_1, y_2, y_3, y_4, y_5, y_6, y_7]
605     # reshape from [samples, timesteps] into [samples, subsequences, timesteps, ...
        features]
606     X_1 = X_1.reshape((X_1.shape[0], prev_days, 1))
607     X_2 = X_2.reshape((X_2.shape[0], prev_days, 1))
608     X_3 = X_3.reshape((X_3.shape[0], prev_days, 1))
609     X_4 = X_4.reshape((X_4.shape[0], prev_days, 1))
610     X_5 = X_5.reshape((X_5.shape[0], prev_days, 1))
611     X_6 = X_6.reshape((X_6.shape[0], prev_days, 1))
612     X_7 = X_7.reshape((X_7.shape[0], prev_days, 1))
613     Xs = [X_1, X_2, X_3, X_4, X_5, X_6, X_7]
614     preds = np.zeros((10, days_ahead))
615
616     for i in range(10):
617         start_time = time.time()
618         preds[i] = np.round_(predict_weekday_with_test(Xs, Ys, x_train_and_tests, ...
            days_ahead, scaler, prev_days))
619         print("Time used to fit and forecast with model", i + 1, ":", ...,
            (time.time() - start_time), " sec")
620
621     np.savetxt(os.getcwd() + r"\Predictions/glob_weekday_preds_test_" + ...
        str(days_ahead) +
        "_days_ahead_from_" + str(endpoint) + "_with_" + str(len_train) + ...
        "_train_data.csv", preds)
622
623
624
625
626
627 ##### Global #####
628
629 # 28 day predictions global
630

```

```

631
632 #####
633 all_in_one_global( endpoint = 40, days_ahead = 28, len_train = 40)
634 # all_in_one_global( endpoint = 196, days_ahead = 28, len_train = 196)
635 # all_in_one_global( endpoint = 196, days_ahead = 28, len_train = 50)
636 # all_in_one_global( endpoint = 196, days_ahead = 28, len_train = 100)
637 #
638 # #####
639 # all_in_one_global(endpoint = 100, days_ahead = 28, len_train = 100)
640 # all_in_one_global(endpoint = 300, days_ahead = 28, len_train = 300)
641 # all_in_one_global(endpoint = 500, days_ahead = 28, len_train = 500)
642 # all_in_one_global(endpoint = 743, days_ahead = 28, len_train = 743) # newest
643 #
644
645 # # Fewer observations: How robust is the model?
646 # all_in_one_global(endpoint = 100, days_ahead = 28, len_train = 50)
647 # all_in_one_global(endpoint = 300, days_ahead = 28, len_train = 50)
648 # all_in_one_global(endpoint = 300, days_ahead = 28, len_train = 100)
649 # all_in_one_global(endpoint = 500, days_ahead = 28, len_train = 50)
650 # all_in_one_global(endpoint = 500, days_ahead = 28, len_train = 100)
651 # all_in_one_global(endpoint = 743, days_ahead = 28, len_train = 50)
652 # all_in_one_global(endpoint = 743, days_ahead = 28, len_train = 100)
653
654 ##### Norway #####
655
656 # An equivalent set of predictions can be made on the norwegian data set.
657 # Start at 264 to compare with original Gandalf model.
658
659 # all_in_one_norway(endpoint = 264, days_ahead = 28, len_train = 264)
660 # all_in_one_norway(endpoint = 264, days_ahead = 28, len_train = 50)
661 # all_in_one_norway(endpoint = 264, days_ahead = 28, len_train = 100)
662 #
663 # all_in_one_norway(endpoint = 100, days_ahead = 28, len_train = 100)
664 # all_in_one_norway(endpoint = 300, days_ahead = 28, len_train = 300)
665 # all_in_one_norway(endpoint = 500, days_ahead = 28, len_train = 500)
666 # all_in_one_norway(endpoint = 696, days_ahead = 28, len_train = 696)
667 #
668 # # Fewer observations: How robust is the model?
669 # all_in_one_norway(endpoint = 100, days_ahead = 28, len_train = 50)
670 # all_in_one_norway(endpoint = 300, days_ahead = 28, len_train = 50)
671 # all_in_one_norway(endpoint = 300, days_ahead = 28, len_train = 100)
672 # all_in_one_norway(endpoint = 500, days_ahead = 28, len_train = 50)
673 # all_in_one_norway(endpoint = 500, days_ahead = 28, len_train = 100)
674 # all_in_one_norway(endpoint = 696, days_ahead = 28, len_train = 50)
675 # all_in_one_norway(endpoint = 696, days_ahead = 28, len_train = 100)
676
677
678 ##### Simulation study #####
679
680 simulations = pd.read_csv(
681     os.getcwd() + r"\Matlab\Saved variables\simulations_from_gandalf.csv", ...
682     header=None)
683
684 # print("simulation.shape = ", simulations.shape)
685 #print("simulations.iloc[0:2, :] = ", simulations.iloc[0:3, :])
686 # accuracy_with_simulations_global(ts_global, simulations, endpoint=196, ...
687     days_ahead=28)
687 # predictions_with_simulations_global(ts_global, simulations.iloc[0:3, ...
688     :], endpoint=196, days_ahead=28)
688 # predictions_with_simulations_global(ts_norway, simulations.iloc[0:3, ...
689     :], endpoint=196, days_ahead=28)
689
690 ##### Predict with weekday model #####
691
692 # weekday_models_no_test_data_global(ts = ts_global, endpoint = 196, days_ahead = 7)
693 # weekday_models_with_test_data_global(ts = ts_global, endpoint = 600, days_ahead ...
694     = 7)

```

Matlab-code

The following code was used to derive and present all the results of the thesis, except for the predictions from the CNN-LSTM model.

```
1 warning('off','all')
2 %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% Import data sets %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
3
4 Norwegian_data_set = table2array(readtable("Data\new_cases_Norway.csv"));
5 global_data_set = table2array(readtable("Data\new_cases_global.csv"));
6
7 % Transform to log scale immediately
8 ts = log(max(0.1, Norwegian_data_set));
9 ts_global = log(max(0.1, global_data_set));
10
11 %%%% General dictionary: %%%%
12
13 % endpoint: last index included as training data
14
15 % len_train: amount of training data
16
17 % days_ahead: length of forecast ensuing endpoint
18
19 % With_observed: prediction scheme 1
20
21 % no_observed: prediction scheme 2
22
23 %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
24
25
26
27
28 % Collecting predictions from CNN-LSTM model from Python
29
30 %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% Gather Global Forecasts %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
31
32 loss_func = table2array(readtable(...
33     "Predictions\mean_loss_on_original_test_data.csv"));
34
35 % global 28 day forecast with prediction scheme 1
36 glob_test_preds_28_days_ahead_from_100_with_50_train_data = ...
37     table2array(readtable(...
38         'Predictions\glob_test_preds_28_days_ahead_from_100_with_50_train_data.csv'));
39 glob_test_preds_28_days_ahead_from_100_with_100_train_data = ...
40     table2array(readtable(...
41         'Predictions\glob_test_preds_28_days_ahead_from_100_with_100_train_data.csv'));
42
43 glob_test_preds_28_days_ahead_from_196_with_50_train_data = ...
44     table2array(readtable(...
45         'Predictions\glob_test_preds_28_days_ahead_from_196_with_50_train_data.csv'));
46 glob_test_preds_28_days_ahead_from_196_with_100_train_data = ...
47     table2array(readtable(...
48         'Predictions\glob_test_preds_28_days_ahead_from_196_with_100_train_data.csv'));
49 glob_test_preds_28_days_ahead_from_196_with_196_train_data = ...
50     table2array(readtable(...
51         'Predictions\glob_test_preds_28_days_ahead_from_196_with_196_train_data.csv'));
52
53 glob_test_preds_28_days_ahead_from_300_with_50_train_data = ...
54     table2array(readtable(...
55         'Predictions\glob_test_preds_28_days_ahead_from_300_with_50_train_data.csv'));
56 glob_test_preds_28_days_ahead_from_300_with_100_train_data = ...
57     table2array(readtable(...
58         'Predictions\glob_test_preds_28_days_ahead_from_300_with_100_train_data.csv'));
59 glob_test_preds_28_days_ahead_from_300_with_300_train_data = ...
60     table2array(readtable(...
61         'Predictions\glob_test_preds_28_days_ahead_from_300_with_300_train_data.csv'));
62
63 glob_test_preds_28_days_ahead_from_500_with_50_train_data = ...
64     table2array(readtable(...
65         'Predictions\glob_test_preds_28_days_ahead_from_500_with_50_train_data.csv'));
66 glob_test_preds_28_days_ahead_from_500_with_100_train_data = ...
```

```

67     table2array(readtable(...
68     'Predictions\glob_test_preds_28_days_ahead_from_500_with_100_train_data.csv'));
69 glob_test_preds_28_days_ahead_from_500_with_500_train_data = ...
70     table2array(readtable(...
71     'Predictions\glob_test_preds_28_days_ahead_from_500_with_500_train_data.csv'));
72
73 glob_test_preds_28_days_ahead_from_743_with_50_train_data = ...
74     table2array(readtable(...
75     'Predictions\glob_test_preds_28_days_ahead_from_500_with_500_train_data.csv'));
76 glob_test_preds_28_days_ahead_from_743_with_100_train_data = ...
77     table2array(readtable(...
78     'Predictions\glob_test_preds_28_days_ahead_from_743_with_100_train_data.csv'));
79 glob_test_preds_28_days_ahead_from_743_with_743_train_data = ...
80     table2array(readtable(...
81     'Predictions\glob_test_preds_28_days_ahead_from_743_with_743_train_data.csv'));
82
83
84 % global seven and 28 day forecasts with prediction scheme 2
85 glob_preds_28_days_ahead_from_100_with_50_train_data = ...
86     table2array(readtable(...
87     'Predictions\glob_preds_28_days_ahead_from_100_with_50_train_data.csv'));
88 glob_preds_7_days_ahead_from_100_with_50_train_data = ...
89     glob_preds_28_days_ahead_from_100_with_50_train_data(:, 1:7);
90 glob_preds_28_days_ahead_from_100_with_100_train_data = ...
91     table2array(readtable(...
92     'Predictions\glob_preds_28_days_ahead_from_100_with_100_train_data.csv'));
93 glob_preds_7_days_ahead_from_100_with_100_train_data = ...
94     glob_preds_28_days_ahead_from_100_with_100_train_data(:, 1:7);
95
96 glob_preds_28_days_ahead_from_196_with_50_train_data = ...
97     table2array(readtable(...
98     'Predictions\glob_preds_28_days_ahead_from_196_with_50_train_data.csv'));
99 glob_preds_7_days_ahead_from_196_with_50_train_data = ...
100     glob_preds_28_days_ahead_from_196_with_50_train_data(:, 1:7);
101 glob_preds_28_days_ahead_from_196_with_100_train_data = ...
102     table2array(readtable(...
103     'Predictions\glob_preds_28_days_ahead_from_196_with_100_train_data.csv'));
104 glob_preds_7_days_ahead_from_196_with_100_train_data = ...
105     glob_preds_28_days_ahead_from_196_with_100_train_data(:, 1:7);
106 glob_preds_28_days_ahead_from_196_with_196_train_data = ...
107     table2array(readtable(...
108     'Predictions\glob_preds_28_days_ahead_from_196_with_196_train_data.csv'));
109 glob_preds_7_days_ahead_from_196_with_196_train_data = ...
110     glob_preds_28_days_ahead_from_196_with_196_train_data(:, 1:7);
111
112 glob_preds_28_days_ahead_from_300_with_50_train_data = ...
113     table2array(readtable(...
114     'Predictions\glob_preds_28_days_ahead_from_300_with_50_train_data.csv'));
115 glob_preds_7_days_ahead_from_300_with_50_train_data = ...
116     glob_preds_28_days_ahead_from_300_with_50_train_data(:, 1:7);
117 glob_preds_28_days_ahead_from_300_with_100_train_data = ...
118     table2array(readtable(...
119     'Predictions\glob_preds_28_days_ahead_from_300_with_100_train_data.csv'));
120 glob_preds_7_days_ahead_from_300_with_100_train_data = ...
121     glob_preds_28_days_ahead_from_300_with_100_train_data(:, 1:7);
122 glob_preds_28_days_ahead_from_300_with_300_train_data = ...
123     table2array(readtable(...
124     'Predictions\glob_preds_28_days_ahead_from_300_with_300_train_data.csv'));
125 glob_preds_7_days_ahead_from_300_with_300_train_data = ...
126     glob_preds_28_days_ahead_from_300_with_300_train_data(:, 1:7);
127
128 glob_preds_28_days_ahead_from_500_with_50_train_data = ...
129     table2array(readtable(...
130     'Predictions\glob_preds_28_days_ahead_from_500_with_50_train_data.csv'));
131 glob_preds_7_days_ahead_from_500_with_50_train_data = ...
132     glob_preds_28_days_ahead_from_500_with_50_train_data(:, 1:7);
133 glob_preds_28_days_ahead_from_500_with_100_train_data = ...
134     table2array(readtable(...
135     'Predictions\glob_preds_28_days_ahead_from_500_with_100_train_data.csv'));
136 glob_preds_7_days_ahead_from_500_with_100_train_data = ...
137     glob_preds_28_days_ahead_from_500_with_100_train_data(:, 1:7);
138 glob_preds_28_days_ahead_from_500_with_500_train_data = ...
139     table2array(readtable(...

```

```

140     'Predictions\glob_preds_28_days_ahead_from_500_with_500_train_data.csv'));
141 glob_preds_7_days_ahead_from_500_with_500_train_data = ...
142     glob_preds_28_days_ahead_from_500_with_500_train_data(:, 1:7);
143
144 glob_preds_28_days_ahead_from_743_with_50_train_data = ...
145     table2array(readtable(...
146     'Predictions\glob_preds_28_days_ahead_from_743_with_50_train_data.csv'));
147 glob_preds_7_days_ahead_from_743_with_50_train_data = ...
148     glob_preds_28_days_ahead_from_743_with_50_train_data(:, 1:7);
149 glob_preds_28_days_ahead_from_743_with_100_train_data = ...
150     table2array(readtable(...
151     'Predictions\glob_preds_28_days_ahead_from_743_with_100_train_data.csv'));
152 glob_preds_7_days_ahead_from_743_with_100_train_data = ...
153     glob_preds_28_days_ahead_from_743_with_100_train_data(:, 1:7);
154 glob_preds_28_days_ahead_from_743_with_743_train_data = ...
155     table2array(readtable(...
156     'Predictions\glob_preds_28_days_ahead_from_743_with_743_train_data.csv'));
157 glob_preds_7_days_ahead_from_743_with_743_train_data = ...
158     glob_preds_28_days_ahead_from_743_with_743_train_data(:, 1:7);
159
160
161 %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% Gather Norwegian Forecasts %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
162
163 % Norway 28 day forecast with prediction scheme 1
164 nor_test_preds_28_days_ahead_from_100_with_50_train_data = ...
165     table2array(readtable(...
166     'Predictions\nor_test_preds_28_days_ahead_from_100_with_50_train_data.csv'));
167 nor_test_preds_28_days_ahead_from_100_with_100_train_data = ...
168     table2array(readtable(...
169     'Predictions\nor_test_preds_28_days_ahead_from_100_with_100_train_data.csv'));
170
171 nor_test_preds_28_days_ahead_from_264_with_50_train_data = ...
172     table2array(readtable(...
173     'Predictions\nor_test_preds_28_days_ahead_from_264_with_50_train_data.csv'));
174 nor_test_preds_28_days_ahead_from_264_with_100_train_data = ...
175     table2array(readtable(...
176     'Predictions\nor_test_preds_28_days_ahead_from_264_with_100_train_data.csv'));
177 nor_test_preds_28_days_ahead_from_264_with_264_train_data = ...
178     table2array(readtable(...
179     'Predictions\nor_test_preds_28_days_ahead_from_264_with_264_train_data.csv'));
180
181 nor_test_preds_28_days_ahead_from_300_with_50_train_data = ...
182     table2array(readtable(...
183     'Predictions\nor_test_preds_28_days_ahead_from_300_with_50_train_data.csv'));
184 nor_test_preds_28_days_ahead_from_300_with_100_train_data = ...
185     table2array(readtable(...
186     'Predictions\nor_test_preds_28_days_ahead_from_300_with_100_train_data.csv'));
187 nor_test_preds_28_days_ahead_from_300_with_300_train_data = ...
188     table2array(readtable(...
189     'Predictions\nor_test_preds_28_days_ahead_from_300_with_300_train_data.csv'));
190
191 nor_test_preds_28_days_ahead_from_500_with_50_train_data = ...
192     table2array(readtable(...
193     'Predictions\nor_test_preds_28_days_ahead_from_500_with_50_train_data.csv'));
194 nor_test_preds_28_days_ahead_from_500_with_100_train_data = ...
195     table2array(readtable(...
196     'Predictions\nor_test_preds_28_days_ahead_from_500_with_100_train_data.csv'));
197 nor_test_preds_28_days_ahead_from_500_with_500_train_data = ...
198     table2array(readtable(...
199     'Predictions\nor_test_preds_28_days_ahead_from_500_with_500_train_data.csv'));
200
201 nor_test_preds_28_days_ahead_from_696_with_50_train_data = ...
202     table2array(readtable(...
203     'Predictions\nor_test_preds_28_days_ahead_from_696_with_50_train_data.csv'));
204 nor_test_preds_28_days_ahead_from_696_with_100_train_data = ...
205     table2array(readtable(...
206     'Predictions\nor_test_preds_28_days_ahead_from_696_with_100_train_data.csv'));
207 nor_test_preds_28_days_ahead_from_696_with_696_train_data = ...
208     table2array(readtable(...
209     'Predictions\nor_test_preds_28_days_ahead_from_696_with_696_train_data.csv'));
210
211
212 % Norway seven (and 28) day forecasts with prediction scheme 2

```

```

213 nor_preds_28_days_ahead_from_100_with_50_train_data = ...
214     table2array(readtable(...
215         'Predictions\nor_preds_28_days_ahead_from_100_with_50_train_data.csv'));
216 nor_preds_7_days_ahead_from_100_with_50_train_data = ...
217     nor_preds_28_days_ahead_from_100_with_50_train_data(:, 1:7);
218 nor_preds_28_days_ahead_from_100_with_100_train_data = ...
219     table2array(readtable(...
220         'Predictions\nor_preds_28_days_ahead_from_100_with_100_train_data.csv'));
221 nor_preds_7_days_ahead_from_100_with_100_train_data = ...
222     nor_preds_28_days_ahead_from_100_with_100_train_data(:, 1:7);
223
224 nor_preds_28_days_ahead_from_264_with_50_train_data = ...
225     table2array(readtable(...
226         'Predictions\nor_preds_28_days_ahead_from_264_with_50_train_data.csv'));
227 nor_preds_7_days_ahead_from_264_with_50_train_data = ...
228     nor_preds_28_days_ahead_from_264_with_50_train_data(:, 1:7);
229 nor_preds_28_days_ahead_from_264_with_100_train_data = ...
230     table2array(readtable(...
231         'Predictions\nor_preds_28_days_ahead_from_264_with_100_train_data.csv'));
232 nor_preds_7_days_ahead_from_264_with_100_train_data = ...
233     nor_preds_28_days_ahead_from_264_with_100_train_data(:, 1:7);
234 nor_preds_28_days_ahead_from_264_with_264_train_data = ...
235     table2array(readtable(...
236         'Predictions\nor_preds_28_days_ahead_from_264_with_264_train_data.csv'));
237 nor_preds_7_days_ahead_from_264_with_264_train_data = ...
238     nor_preds_28_days_ahead_from_264_with_264_train_data(:, 1:7);
239
240 nor_preds_28_days_ahead_from_300_with_50_train_data = ...
241     table2array(readtable(...
242         'Predictions\nor_preds_28_days_ahead_from_300_with_50_train_data.csv'));
243 nor_preds_7_days_ahead_from_300_with_50_train_data = ...
244     nor_preds_28_days_ahead_from_300_with_50_train_data(:, 1:7);
245 nor_preds_28_days_ahead_from_300_with_100_train_data = ...
246     table2array(readtable(...
247         'Predictions\nor_preds_28_days_ahead_from_300_with_100_train_data.csv'));
248 nor_preds_7_days_ahead_from_300_with_100_train_data = ...
249     nor_preds_28_days_ahead_from_300_with_100_train_data(:, 1:7);
250 nor_preds_28_days_ahead_from_300_with_300_train_data = ...
251     table2array(readtable(...
252         'Predictions\nor_preds_28_days_ahead_from_300_with_300_train_data.csv'));
253 nor_preds_7_days_ahead_from_300_with_300_train_data = ...
254     nor_preds_28_days_ahead_from_300_with_300_train_data(:, 1:7);
255
256 nor_preds_28_days_ahead_from_500_with_50_train_data = ...
257     table2array(readtable(...
258         'Predictions\nor_preds_28_days_ahead_from_500_with_50_train_data.csv'));
259 nor_preds_7_days_ahead_from_500_with_50_train_data = ...
260     nor_preds_28_days_ahead_from_500_with_50_train_data(:, 1:7);
261 nor_preds_28_days_ahead_from_500_with_100_train_data = ...
262     table2array(readtable(...
263         'Predictions\nor_preds_28_days_ahead_from_500_with_100_train_data.csv'));
264 nor_preds_7_days_ahead_from_500_with_100_train_data = ...
265     nor_preds_28_days_ahead_from_500_with_100_train_data(:, 1:7);
266 nor_preds_28_days_ahead_from_500_with_500_train_data = ...
267     table2array(readtable(...
268         'Predictions\nor_preds_28_days_ahead_from_500_with_500_train_data.csv'));
269 nor_preds_7_days_ahead_from_500_with_500_train_data = ...
270     nor_preds_28_days_ahead_from_500_with_500_train_data(:, 1:7);
271
272 nor_preds_28_days_ahead_from_696_with_50_train_data = ...
273     table2array(readtable(...
274         'Predictions\nor_preds_28_days_ahead_from_696_with_50_train_data.csv'));
275 nor_preds_7_days_ahead_from_696_with_50_train_data = ...
276     nor_preds_28_days_ahead_from_696_with_50_train_data(:, 1:7);
277 nor_preds_28_days_ahead_from_696_with_100_train_data = ...
278     table2array(readtable(...
279         'Predictions\nor_preds_28_days_ahead_from_696_with_100_train_data.csv'));
280 nor_preds_7_days_ahead_from_696_with_100_train_data = ...
281     nor_preds_28_days_ahead_from_696_with_100_train_data(:, 1:7);
282 nor_preds_28_days_ahead_from_696_with_696_train_data = ...
283     table2array(readtable(...
284         'Predictions\nor_preds_28_days_ahead_from_696_with_696_train_data.csv'));
285 nor_preds_7_days_ahead_from_696_with_696_train_data = ...

```



```

286     nor_preds_28_days_ahead_from_696_with_696_train_data(:, 1:7);
287
288
289 %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% Plots %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
290
291 % plot_norway_and_global_time_series_and_transform(ts, ts_global, false)
292 % %
293 % plot_norway_and_global_time_series_and_transform(ts, ts_global, true)
294
295
296 %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% Norwegian Forecasts %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
297
298 % plot_preds_sarima_gandalf(ts, 264, 264, 7, false)
299 % plot_acf_of_res(ts, 264, 264, false)
300 % plot_acf_of_res(ts, 264, 264, true)
301
302 % plot_10_norway_CNN_LSTM_Preds(floor(exp(ts)), ...
303 %     nor_preds_7_days_ahead_from_264_with_264_train_data, 264, 264, 7, false)
304 % plot_preds_without_observed(ts, ...
305 %     mean(nor_preds_7_days_ahead_from_264_with_264_train_data, 1), ...
306 %     264, 264, 7, false, 14)
307 % plot_10_norway_CNN_LSTM_Preds(floor(exp(ts)), ...
308 %     nor_test_preds_28_days_ahead_from_264_with_264_train_data, ...
309 %     264, 264, 28, false)
310 % plot_preds_with_observed(ts, ...
311 %     mean(nor_test_preds_28_days_ahead_from_264_with_264_train_data, 1), ...
312 %     264, 264, 28, false, 30)
313 %
314 %
315 % plot_parameters_from_prediction_scheme_1(ts, 264, 264, 28, false)
316 %
317 % plot_preds_without_observed(ts, ...
318 %     mean(nor_preds_7_days_ahead_from_264_with_50_train_data, 1), ...
319 %     264, 50, 7, false, 14)
320 % plot_preds_with_observed(ts, ...
321 %     mean(nor_test_preds_28_days_ahead_from_264_with_50_train_data, 1), ...
322 %     264, 50, 28, false, 21)
323 % plot_parameters_from_prediction_scheme_1(ts, 264, 50, 28, false)
324 % plot_preds_without_observed(ts, ...
325 %     mean(nor_preds_7_days_ahead_from_264_with_100_train_data, 1), ...
326 %     264, 100, 7, false, 14)
327 % plot_preds_with_observed(ts, ...
328 %     mean(nor_test_preds_28_days_ahead_from_264_with_100_train_data, 1), ...
329 %     264, 100, 28, false, 30)
330 % plot_parameters_from_prediction_scheme_1(ts, 264, 100, 28, false)
331
332
333 %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% Global Forecasts %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
334
335 % plot_loss_function(loss_func)
336
337 % plot_10_global_CNN_LSTM_Preds(floor(exp(ts_global)), ...
338 %     glob_test_preds_28_days_ahead_from_196_with_196_train_data, ...
339 %     196, 196, 28, true)
340 % plot_preds_sarima_gandalf(ts_global, 196, 196, 7, true)
341 % plot_acf_of_res(ts_global, 196, 196, false)
342 % plot_acf_of_res(ts_global, 196, 196, true)
343 %
344 % plot_preds_without_observed(ts_global, ...
345 %     mean(glob_preds_7_days_ahead_from_196_with_196_train_data, 1), ...
346 %     196, 196, 7, true, 14)
347 % plot_preds_with_observed(ts_global, ...
348 %     mean(glob_test_preds_28_days_ahead_from_196_with_196_train_data, 1), ...
349 %     196, 196, 28, true, 30)
350 % plot_parameters_from_prediction_scheme_1(ts_global, 196, 196, 28, true)
351 %
352 % plot_10_global_CNN_LSTM_Preds(floor(exp(ts_global)), ...
353 %     glob_preds_28_days_ahead_from_196_with_196_train_data, ...
354 %     196, 196, 28, true)
355 % plot_preds_without_observed(ts_global, ...
356 %     mean(glob_preds_28_days_ahead_from_196_with_196_train_data, 1), ...
357 %     196, 196, 28, true, 30)
358

```

```

359
360 %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% Results for table 2 %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
361
362
363 % Generate all numbers to be inserted in table
364 % Norway:
365 % extract_percision_results_without_observed(ts, ...
      nor_preds_7_days_ahead_from_100_with_50_train_data, 100, 50, 7)
366 % extract_percision_results_without_observed(ts, ...
      nor_preds_7_days_ahead_from_100_with_100_train_data, 100, 100, 7)
367 %
368 % extract_percision_results_without_observed(ts, ...
      nor_preds_7_days_ahead_from_264_with_50_train_data, 264, 50, 7)
369 % extract_percision_results_without_observed(ts, ...
      nor_preds_7_days_ahead_from_264_with_100_train_data, 264, 100, 7)
370 % extract_percision_results_without_observed(ts, ...
      nor_preds_7_days_ahead_from_264_with_264_train_data, 264, 264, 7)
371 %
372 % extract_percision_results_without_observed(ts, ...
      nor_preds_7_days_ahead_from_300_with_50_train_data, 300, 50, 7)
373 % extract_percision_results_without_observed(ts, ...
      nor_preds_7_days_ahead_from_300_with_100_train_data, 300, 100, 7)
374 % extract_percision_results_without_observed(ts, ...
      nor_preds_7_days_ahead_from_300_with_300_train_data, 300, 300, 7)
375 %
376 % extract_percision_results_without_observed(ts, ...
      nor_preds_7_days_ahead_from_500_with_50_train_data, 500, 50, 7)
377 % extract_percision_results_without_observed(ts, ...
      nor_preds_7_days_ahead_from_500_with_100_train_data, 500, 100, 7)
378 % extract_percision_results_without_observed(ts, ...
      nor_preds_7_days_ahead_from_500_with_500_train_data, 500, 500, 7)
379 %
380 % extract_percision_results_without_observed(ts, ...
      nor_preds_7_days_ahead_from_696_with_50_train_data, 696, 50, 7)
381 % extract_percision_results_without_observed(ts, ...
      nor_preds_7_days_ahead_from_696_with_100_train_data, 696, 100, 7)
382 % extract_percision_results_without_observed(ts, ...
      nor_preds_7_days_ahead_from_696_with_696_train_data, 696, 696, 7)
383 %
384 %
385 %
386 % extract_percision_results_with_observed(ts, ...
      nor_test_preds_28_days_ahead_from_100_with_50_train_data, 100, 50, 28)
387 % extract_percision_results_with_observed(ts, ...
      nor_test_preds_28_days_ahead_from_100_with_100_train_data, 100, 100, 28)
388 %
389 % extract_percision_results_with_observed(ts, ...
      nor_test_preds_28_days_ahead_from_264_with_50_train_data, 264, 50, 28)
390 % extract_percision_results_with_observed(ts, ...
      nor_test_preds_28_days_ahead_from_264_with_100_train_data, 264, 100, 28)
391 % extract_percision_results_with_observed(ts, ...
      nor_test_preds_28_days_ahead_from_264_with_264_train_data, 264, 264, 28)
392 %
393 % extract_percision_results_with_observed(ts, ...
      nor_test_preds_28_days_ahead_from_300_with_50_train_data, 300, 50, 28)
394 % extract_percision_results_with_observed(ts, ...
      nor_test_preds_28_days_ahead_from_300_with_100_train_data, 300, 100, 28)
395 % extract_percision_results_with_observed(ts, ...
      nor_test_preds_28_days_ahead_from_300_with_300_train_data, 300, 300, 28)
396 %
397 % extract_percision_results_with_observed(ts, ...
      nor_test_preds_28_days_ahead_from_500_with_50_train_data, 500, 50, 28)
398 % extract_percision_results_with_observed(ts, ...
      nor_test_preds_28_days_ahead_from_500_with_100_train_data, 500, 100, 28)
399 % extract_percision_results_with_observed(ts, ...
      nor_test_preds_28_days_ahead_from_500_with_500_train_data, 500, 500, 28)
400 %
401 % extract_percision_results_with_observed(ts, ...
      nor_test_preds_28_days_ahead_from_696_with_50_train_data, 696, 50, 28)
402 % extract_percision_results_with_observed(ts, ...
      nor_test_preds_28_days_ahead_from_696_with_100_train_data, 696, 100, 28)
403 % extract_percision_results_with_observed(ts, ...
      nor_test_preds_28_days_ahead_from_696_with_696_train_data, 696, 696, 28)

```

```

404 %
405 % Global:
406 % extract_percision_results_without_observed(ts_global, ...
      glob_preds_7_days_ahead_from_100_with_50_train_data, 100, 50, 7)
407 % extract_percision_results_without_observed(ts_global, ...
      glob_preds_7_days_ahead_from_100_with_100_train_data, 100, 100, 7)
408 %
409 % extract_percision_results_without_observed(ts_global, ...
      glob_preds_7_days_ahead_from_196_with_50_train_data, 196, 50, 7)
410 % extract_percision_results_without_observed(ts_global, ...
      glob_preds_7_days_ahead_from_196_with_100_train_data, 196, 100, 7)
411 % extract_percision_results_without_observed(ts_global, ...
      glob_preds_7_days_ahead_from_196_with_196_train_data, 196, 196, 7)
412 %
413 % extract_percision_results_without_observed(ts_global, ...
      glob_preds_7_days_ahead_from_300_with_50_train_data, 300, 50, 7)
414 % extract_percision_results_without_observed(ts_global, ...
      glob_preds_7_days_ahead_from_300_with_100_train_data, 300, 100, 7)
415 % extract_percision_results_without_observed(ts_global, ...
      glob_preds_7_days_ahead_from_300_with_300_train_data, 300, 300, 7)
416 %
417 % extract_percision_results_without_observed(ts_global, ...
      glob_preds_7_days_ahead_from_500_with_50_train_data, 500, 50, 7)
418 % extract_percision_results_without_observed(ts_global, ...
      glob_preds_7_days_ahead_from_500_with_100_train_data, 500, 100, 7)
419 % extract_percision_results_without_observed(ts_global, ...
      glob_preds_7_days_ahead_from_500_with_500_train_data, 500, 500, 7)
420 %
421 % extract_percision_results_without_observed(ts_global, ...
      glob_preds_7_days_ahead_from_743_with_50_train_data, 743, 50, 7)
422 % extract_percision_results_without_observed(ts_global, ...
      glob_preds_7_days_ahead_from_743_with_100_train_data, 743, 100, 7)
423 % extract_percision_results_without_observed(ts_global, ...
      glob_preds_7_days_ahead_from_743_with_743_train_data, 743, 743, 7)
424 %
425 %
426 %
427 % extract_percision_results_with_observed(ts_global, ...
      glob_test_preds_28_days_ahead_from_100_with_50_train_data, 100, 50, 28)
428 % extract_percision_results_with_observed(ts_global, ...
      glob_test_preds_28_days_ahead_from_100_with_100_train_data, 100, 100, 28)
429 %
430 % extract_percision_results_with_observed(ts_global, ...
      glob_test_preds_28_days_ahead_from_196_with_50_train_data, 196, 50, 28)
431 % extract_percision_results_with_observed(ts_global, ...
      glob_test_preds_28_days_ahead_from_196_with_100_train_data, 196, 100, 28)
432 % extract_percision_results_with_observed(ts_global, ...
      glob_test_preds_28_days_ahead_from_196_with_196_train_data, 196, 196, 28)
433 %
434 % extract_percision_results_with_observed(ts_global, ...
      glob_test_preds_28_days_ahead_from_300_with_50_train_data, 300, 50, 28)
435 % extract_percision_results_with_observed(ts_global, ...
      glob_test_preds_28_days_ahead_from_300_with_100_train_data, 300, 100, 28)
436 % extract_percision_results_with_observed(ts_global, ...
      glob_test_preds_28_days_ahead_from_300_with_300_train_data, 300, 300, 28)
437 %
438 % extract_percision_results_with_observed(ts_global, ...
      glob_test_preds_28_days_ahead_from_500_with_50_train_data, 500, 50, 28)
439 % extract_percision_results_with_observed(ts_global, ...
      glob_test_preds_28_days_ahead_from_500_with_100_train_data, 500, 100, 28)
440 % extract_percision_results_with_observed(ts_global, ...
      glob_test_preds_28_days_ahead_from_500_with_500_train_data, 500, 500, 28)
441 % %
442 % extract_percision_results_with_observed(ts_global, ...
      glob_test_preds_28_days_ahead_from_743_with_50_train_data, 743, 50, 28)
443 % extract_percision_results_with_observed(ts_global, ...
      glob_test_preds_28_days_ahead_from_743_with_100_train_data, 743, 100, 28)
444 % extract_percision_results_with_observed(ts_global, ...
      glob_test_preds_28_days_ahead_from_743_with_743_train_data, 743, 743, 28)
445
446
447
448 %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% Study table with bar plots %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

```

```

449
450 % Gather measures from Table 2 to plot summarized results:
451 % RRMSE:
452 results_RRMSE_28_day_preds = [74 75 73; 60 262 53; 12 10 18; 11 11 25;
453     11 11 28; 25 27 29; 25 24 17; 28 28 26; 14 15 30; 15 15 35; 16 16 19;
454     11 9 19; 9 9 30; 9 9 46; 7 7 7; 6 6 10; 4 5 5; 4 4 4; 4 4 4; 5 7 12;
455     5 5 9; 6 6 11; 4 6 9; 5 5 15; 5 5 39; 9 7 21; 7 7 21; 7 7 85];
456
457 results_RRMSE_7_day_preds = [119 119 64; 87 87 65; 17 6 8; 16 15 11;
458     12 12 12; 17 17 14; 19 20 7; 16 16 18; 10 14 8; 5 8 12;
459     21 21 11; 9 8 23; 9 11 24; 13 4 24;
460
461     30 35 11; 16 16 14; 7 6 8; 4 4 6; 4 4 5; 7 5 11; 7 7 8; 7 7 9; 4 10 13;
462     6 7 20; 3 3 18; 15 10 12; 12 14 10; 12 13 12];
463
464 input_bar_plot_RRMSE_days_ahead = [mean(results_RRMSE_28_day_preds, 1); ...
465     mean(results_RRMSE_7_day_preds, 1)];
466
467 results_RRMSE_Norway = [results_RRMSE_7_day_preds(1:14, :); ...
468     results_RRMSE_28_day_preds(1:14, :)];
469 results_RRMSE_Global = [results_RRMSE_7_day_preds(15:28, :); ...
470     results_RRMSE_28_day_preds(15:28, :)];
471
472 input_bar_plot_RRMSE_data_set = [mean(results_RRMSE_Norway, 1); ...
473     mean(results_RRMSE_Global, 1)];
474
475 results_RRMSE_full_data = ...
476     [results_RRMSE_7_day_preds([1, 3, 6, 9, 12, 15, 17, 20, 23, 26], :);
477     results_RRMSE_28_day_preds([1, 3, 6, 9, 12, 15, 17, 20, 23, 26], :)];
478
479 results_RRMSE_100_data = ...
480     [results_RRMSE_7_day_preds([1, 4, 7, 10, 13, 15, 18, 21, 24, 27], :);
481     results_RRMSE_28_day_preds([1, 4, 7, 10, 13, 15, 18, 21, 24, 27], :)];
482
483 results_RRMSE_50_data = ...
484     [results_RRMSE_7_day_preds([2, 5, 8, 11, 14, 16, 19, 22, 25, 28], :);
485     results_RRMSE_28_day_preds([2, 5, 8, 11, 14, 16, 19, 22, 25, 28], :)];
486
487 input_bar_plot_RRMSE_sample_size = [mean(results_RRMSE_full_data, 1); ...
488     mean(results_RRMSE_100_data, 1); mean(results_RRMSE_50_data, 1)];
489
490 input_bar_plot_RRMSE_models = mean([results_RRMSE_full_data; ...
491     results_RRMSE_100_data; results_RRMSE_50_data], 1);
492
493 % MAPE:
494 results_MAPE_7_day_preds = [100 100 88; 57 57 41; 16 6 8; 12 11 11;
495     9 9 11; 14 13 12; 15 16 7; 13 13 14; 7 11 8; 4 5 12; 17 17 9; 7 7 22;
496     7 9 18; 11 3 23; 28 31 10; 14 14 13; 6 6 6; 3 3 4; 4 4 3; 6 4 10;
497     6 6 7; 6 6 8; 3 9 13; 6 6 20; 2 2 18; 12 9 12; 11 11 10; 11 10 11];
498
499 results_MAPE_28_day_preds = [61 62 87; 40 118 41; 9 8 17; 9 8 26; 8 8 28;
500     22 20 23; 20 19 23; 23 22 21; 11 12 12; 12 11 20; 13 12 23; 11 9 17;
501     9 9 31; 9 8 46;
502
503     6 5 6; 5 5 9; 3 4 3; 3 3 3; 3 3 3; 4 5 11; 4 4 8;
504     4 4 9; 4 5 8; 4 4 14; 4 4 39; 7 6 20; 6 6 20; 6 5 83];
505
506 input_bar_plot_MAPE_days_ahead = [mean(results_MAPE_28_day_preds, 1); ...
507     mean(results_MAPE_7_day_preds, 1)];
508
509 results_MAPE_Norway = ...
510     [results_MAPE_7_day_preds(1:14, :); results_MAPE_28_day_preds(1:14, :)];
511 results_MAPE_Global = ...
512     [results_MAPE_7_day_preds(15:28, :); results_MAPE_28_day_preds(15:28, :)];
513
514 input_bar_plot_MAPE_data_set = [mean(results_MAPE_Norway, 1); ...
515     mean(results_MAPE_Global, 1)];
516
517 results_MAPE_full_data = ...
518     [results_MAPE_7_day_preds([1, 3, 6, 9, 12, 15, 17, 20, 23, 26], :);
519     results_MAPE_28_day_preds([1, 3, 6, 9, 12, 15, 17, 20, 23, 26], :)];
520
521 results_MAPE_100_data = ...

```

```

522     [results_MAPE_7_day_preds([1, 4, 7, 10, 13, 15, 18, 21, 24, 27], :);
523     results_MAPE_28_day_preds([1, 4, 7, 10, 13, 15, 18, 21, 24, 27], :)];
524
525 results_MAPE_50_data = ...
526     [results_MAPE_7_day_preds([2, 5, 8, 11, 14, 16, 19, 22, 25, 28], :);
527     results_MAPE_28_day_preds([2, 5, 8, 11, 14, 16, 19, 22, 25, 28], :)];
528
529 input_bar_plot_MAPE_sample_size = [mean(results_MAPE_full_data, 1); ...
530     mean(results_MAPE_100_data, 1); mean(results_MAPE_50_data, 1)];
531
532
533 input_bar_plot_MAPE_models = mean([results_MAPE_full_data; ...
534     results_MAPE_100_data; results_MAPE_50_data],1).';
535
536 % % Creating histogram for all RRMSEs:
537 % RRMSEs_from_table = ...
538 %     [results_RRMSE_28_day_preds; results_RRMSE_7_day_preds];
539 % RRMSEs_sarima_table = RRMSEs_from_table(:, 1);
540 % RRMSEs_gandalf_table = RRMSEs_from_table(:, 2);
541 % RRMSEs_cnn_table = RRMSEs_from_table(:, 3);
542 % figure
543 % sar = histogram(RRMSEs_sarima_table, 40, 'FaceColor', [0.4940 0.1840 0.5560]);
544 % mean(RRMSEs_sarima_table)
545 % mean_sar = xline(mean(RRMSEs_sarima_table),'--',...
546 %     'Color', [1 0 0], 'LineWidth',2);
547 % legend([sar mean_sar], 'SARIMA model','Mean RRMSE')
548 % xlabel('RRMSE')
549 % ylabel('Occurrences')
550 % set(gcf,'color','w')
551 % set(gca,'FontSize',24)
552 % % xticks(min(RRMSEs_sarima_table):5:max(RRMSEs_sarima_table))
553 % figure
554 % gand = histogram(RRMSEs_gandalf_table, 40, 'FaceColor', [.2, .9, .5]);
555 % mean(RRMSEs_gandalf_table)
556 % mean_gand = xline(mean(RRMSEs_gandalf_table),'--',...
557 %     'Color', [1 0 0], 'LineWidth',2);
558 % legend([gand mean_gand], 'Gandalf model','Mean RRMSE')
559 % xlabel('RRMSE')
560 % ylabel('Occurrences')
561 % set(gcf,'color','w')
562 % set(gca,'FontSize',24)
563 %
564 % figure
565 % cnn = histogram(RRMSEs_cnn_table, 40, 'FaceColor', [0 0.4470 0.7410]);
566 % mean(RRMSEs_cnn_table)
567 % mean_cnn = xline(mean(RRMSEs_cnn_table),'--',...
568 %     'Color', [1 0 0], 'LineWidth',2);
569 % legend([cnn mean_cnn], 'CNN-LSTM model','Mean RRMSE')
570 % xlabel('RRMSE')
571 % ylabel('Occurrences')
572 % set(gcf,'color','w')
573 % set(gca,'FontSize',24)
574
575
576 % Creating histogram for all MAPEs:
577 % MAPEs_from_table = [results_MAPE_28_day_preds; results_MAPE_7_day_preds];
578 % MAPEs_sarima_table = MAPEs_from_table(:, 1);
579 % MAPEs_gandalf_table = MAPEs_from_table(:, 2);
580 % MAPEs_cnn_table = MAPEs_from_table(:, 3);
581 % figure
582 % sar = histogram(MAPEs_sarima_table, 40, 'FaceColor', [0.4940 0.1840 0.5560]);
583 % mean(MAPEs_sarima_table)
584 % mean_sar = xline(mean(MAPEs_sarima_table),'--',...
585 %     'Color', [1 0 0], 'LineWidth',2);
586 % legend([sar mean_sar], 'SARIMA model','Mean MAPE')
587 % xlabel('MAPE')
588 % ylabel('Occurrences')
589 % set(gcf,'color','w')
590 % set(gca,'FontSize',24)
591 % % xticks(min(MAPEs_sarima_table):5:max(MAPEs_sarima_table))
592 % figure
593 % gand = histogram(MAPEs_gandalf_table, 40, 'FaceColor', [.2, .9, .5]);
594 % mean(MAPEs_gandalf_table)

```

```

595 % mean_gand = xline(mean(MAPEs_gandalf_table),'--',...
596 %     'Color', [1 0 0], 'LineWidth',2);
597 % legend([gand mean_gand], 'Gandalf model','Mean MAPE')
598 % xlabel('MAPE')
599 % ylabel('Occurrences')
600 % set(gcf,'color','w')
601 % set(gca,'FontSize',24)
602 %
603 % figure
604 % cnn = histogram(MAPEs_cnn_table, 40, 'FaceColor', [0 0.4470 0.7410]);
605 % mean(MAPEs_cnn_table)
606 % mean_cnn = xline(mean(MAPEs_cnn_table),'--',...
607 %     'Color', [1 0 0], 'LineWidth',2);
608 % legend([cnn mean_cnn], 'CNN-LSTM model','Mean MAPE')
609 % xlabel('MAPE')
610 % ylabel('Occurrences')
611 % set(gcf,'color','w')
612 % set(gca,'FontSize',24)
613
614
615
616 %
617 % bin_range = min(min(RRMSEs_from_table)):max(max(RRMSEs_from_table));
618 % counts_sarima = histc(RRMSEs_sarima, bin_range);
619 % counts_gandalf = histc(RRMSEs_gandalf, bin_range);
620 % counts_cnn = histc(RRMSEs_cnn, bin_range);
621 % counts_sarima_gandalf = counts_sarima + counts_gandalf;
622 % counts_total = counts_sarima + counts_gandalf + counts_cnn;
623 %
624 % clr = [0 0.4470 0.7410; .2, .9, .5; 0.4940 0.1840 0.5560];
625 % colormap(clr);
626 % figure
627 % cnn = bar(bin_range, counts_total);
628 % hold on
629 % gand = bar(bin_range, counts_sarima_gandalf);
630 % sar = bar(bin_range, counts_sarima);
631 % hold off
632 % xlabel('RRMSE')
633 % ylabel('Occurrences')
634 % legend([sar, gand, cnn], 'SARIMA model', 'Gandalf model', 'CNN-LSTM model')
635 % set(gcf,'color','w')
636 % set(gca,'FontSize',24)
637 % set(h, {'DisplayName'}, {'SARIMA', 'Gandalf', 'CNN-LSTM'})
638
639
640 % % Creating histogram for all MAPEs:
641 % MAPEs_from_table = [results_MAPE_28_day_preds; results_MAPE_7_day_preds];
642 % MAPEs_sarima_table = MAPEs_from_table(:, 1);
643 % MAPEs_gandalf_table = MAPEs_from_table(:, 2);
644 % MAPEs_cnn_table = MAPEs_from_table(:, 3);
645 %
646 % bin_range = min(min(MAPEs_from_table)):max(max(MAPEs_from_table));
647 % counts_sarima = histc(MAPEs_sarima, bin_range);
648 % counts_gandalf = histc(MAPEs_gandalf, bin_range);
649 % counts_cnn = histc(MAPEs_cnn, bin_range);
650 % counts_sarima_gandalf = counts_sarima + counts_gandalf;
651 % counts_total = counts_sarima + counts_gandalf + counts_cnn;
652 %
653 % clr = [0 0.4470 0.7410; .2, .9, .5; 0.4940 0.1840 0.5560];
654 % colormap(clr);
655 % figure
656 % cnn = bar(bin_range, counts_total);
657 % hold on
658 % gand = bar(bin_range, counts_sarima_gandalf);
659 % sar = bar(bin_range, counts_sarima);
660 % hold off
661 % xlabel('MAPE')
662 % ylabel('Occurrences')
663 % legend([sar, gand, cnn], 'SARIMA model', 'Gandalf model', 'CNN-LSTM model')
664 % set(gcf,'color','w')
665 % set(gca,'FontSize',24)
666 % set(h, {'DisplayName'}, {'SARIMA', 'Gandalf', 'CNN-LSTM'})
667

```

```

668
669
670
671 % Figure to create bar plots, switch out text and so on depending on the
672 % specific plot
673 % % figure
674 % h = bar(input_bar_plot_RRMSE_sample_size); % , 'stacked')
675 % % ylabel('median MAPE')
676 % ylabel('median RRMSE')
677 % % xlabel('Model')
678 % % xticklabels({'SARIMA', 'Gandalf', 'CNN-LSTM'})
679 % % xlabel('Data set')
680 % % xticklabels({'Norway', 'Global'})
681 % % xlabel('Days-ahead forecast')
682 % % xticklabels({'28 day', '7 day'})
683 % xlabel('Sample size')
684 % xticklabels({'All available', '100', '50'})
685 % set(gcf, 'color', 'w')
686 % set(gca, 'FontSize', 24)
687 % set(h, {'DisplayName'}, {'SARIMA', 'Gandalf', 'CNN-LSTM'})
688 % legend()
689
690
691
692
693 %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% Simulation study %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
694
695 % simulations = simulate_from_Gandalf(ts_global, 28, 196, 1000).';
696 % plot_simulations(ts_global, simulations.', 28, 196);
697 % predictions_on_simulations(ts_global, simulations(1:3,:), ...
698 %     pred_three_preds_on_sims, 28, 196, true);
699
700
701 RRMSEs_cnn = table2array(readtable(...
702     'Predictions\RRMSEs_from_simulation_global_28_days_ahead_with_196_train_data_from196.csv'));
703 MAPEs_cnn = table2array(readtable(...
704     'Predictions\MAPEs_from_simulation_global_28_days_ahead_with_196_train_data_from196.csv'));
705 pred_three_preds_on_sims = table2array(readtable(...
706     'Predictions\pred_three_preds_on_sims.csv'));
707
708
709 %The below functions take approximately 20 hours to run with 1000 sims
710 % [RRMSEs_sarima, MAPEs_sarima, RRMSEs_gandalf, MAPEs_gandalf, ...
711 %     ma_sim, sma_sim, arch_sim, garch_sim] = ...
712 %     RRMSE_MAPE_from_simulation_sarima_and_gandalf...
713 %     (ts_global, simulations, 196, false);
714 %
715 % [log_RRMSEs_sarima, log_MAPEs_sarima, ...
716 %     log_RRMSEs_gandalf, log_MAPEs_gandalf] = ...
717 %     RRMSE_MAPE_from_simulation_sarima_and_gandalf...
718 %     (ts_global, simulations, 196, true);
719 %
720 %
721 % % save('RRMSEs_sarima.mat', 'RRMSEs_sarima');
722 % % save('MAPEs_sarima.mat', 'MAPEs_sarima');
723 % % save('RRMSEs_gandalf.mat', 'RRMSEs_gandalf');
724 % % save('MAPEs_gandalf.mat', 'MAPEs_gandalf');
725 % %
726 % % save('log_RRMSEs_gandalf.mat', 'log_RRMSEs_gandalf');
727 % % save('log_MAPEs_gandalf.mat', 'log_MAPEs_gandalf');
728 %
729 %
730 % load('RRMSEs_sarima.mat', 'RRMSEs_sarima');
731 % load('MAPEs_sarima.mat', 'MAPEs_sarima');
732 % load('RRMSEs_gandalf.mat', 'RRMSEs_gandalf');
733 % load('MAPEs_gandalf.mat', 'MAPEs_gandalf');
734 % load('log_RRMSEs_gandalf.mat', 'log_RRMSEs_gandalf');
735 % load('log_MAPEs_gandalf.mat', 'log_MAPEs_gandalf');
736 %
737 % Plot the results:
738 % plot_simulated_RRMSE_MAPE(RRMSEs_cnn, MAPEs_cnn)
739 % plot_simulated_RRMSE_MAPE(RRMSEs_sarima, MAPEs_sarima)
740 % plot_simulated_RRMSE_MAPE(RRMSEs_gandalf, MAPEs_gandalf)

```

```

741 % plot_simulated_RRMSE_MAPE(log_RRMSEs_gandalf, log_MAPEs_gandalf)
742 % plot_simulated_parameters(ts_global, 196, ma_sim(1:200,:), ...
743 %     sma_sim(1:200,:), arch_sim(1:200,:), garch_sim(1:200,:))
744
745 function extract_percision_results_without_observed...
746     (ts, preds_cnn, endpoint, len_train, days_ahead)
747     % Does not show plot of predictions, only results. This was used to
748     % fill out the summarized table.
749     sarima = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
750     gandalf = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
751     garchmod = garch('Constant',0.001,'GARCHLags',1,'ARCHLags',1);
752     gandalf.Variance = garchmod;
753
754     train_data = ts(endpoint-len_train+1:endpoint);
755     fitted_sarima = estimate(sarima, train_data, 'Display','off');
756     fitted_gandalf= estimate(gandalf, train_data, 'Display','off');
757     preds_sarima= forecast(fitted_sarima, days_ahead, train_data);
758     preds_gandalf = forecast(fitted_gandalf, days_ahead, train_data);
759
760     mean_preds_cnn = mean(preds_cnn, 1);
761     %mse_cnn = mse_cnn_lstm(ts, mean_preds_cnn, endpoint, len_train, true);
762
763     % convert to normal scale where needed
764     ts = floor(exp(ts));
765     test_data = ts(endpoint+1: endpoint + days_ahead);
766     preds_sarima = floor(exp(preds_sarima));
767     preds_gandalf = floor(exp(preds_gandalf));
768
769     % calculate RRMSE and MAPE for single predictions
770     RRMSE_sarima = sqrt(mean((test_data - preds_sarima).^2)/mean(test_data)*100;
771     MAPE_sarima = mean(abs(test_data - preds_sarima)./abs(test_data))*100;
772     RRMSE_gandalf = sqrt(mean((test_data - preds_gandalf).^2)/mean(test_data)*100;
773     MAPE_gandalf = mean(abs(test_data - preds_gandalf)./abs(test_data))*100;
774     RRMSE_cnn = sqrt(mean((test_data - mean_preds_cnn).^2)/mean(test_data)*100;
775     MAPE_cnn = mean(abs(test_data - mean_preds_cnn)./abs(test_data))*100;
776
777     % calculate RRMSE and MAPE for the 10 CNN-LSTM models
778     RRMSE = zeros(10, 1);
779     MAPE = zeros(10, 1);
780     for i = 1:10
781         RRMSE(i) = sqrt(mean((test_data.' - preds_cnn(i, ...
782             :)).^2)/mean(test_data)*100;
783         MAPE(i) = mean(abs(test_data.' - preds_cnn(i, :))./abs(test_data.'))*100;
784     end
785
786     mean_RRMSE_cnn = round(mean(RRMSE), 2);
787     mean_MAPE_cnn = round(mean(MAPE), 2);
788
789     disp(".....")
790     disp(["Displaying the results from the ", days_ahead,"-day predictions ...
791         starting at ", endpoint, ...
792         " using ", len_train, " data without test data:"])
793     disp(["RRMSE SARIMA: ", round(RRMSE_sarima, 2), ", MAPE SARIMA: ", ...
794         round(MAPE_sarima, 2)])
795     disp(["RRMSE Gandalf: ", round(RRMSE_gandalf, 2), ", MAPE Gandalf: ", ...
796         round(MAPE_gandalf, 2)])
797     disp(["RRMSE mean CNN-LSTM: ", round(RRMSE_cnn, 2), ", MAPE mean CNN-LSTM: ", ...
798         round(MAPE_cnn, 2)])
799     disp(["Mean RRMSE of ten CNN-LSTM models: ", round(mean_RRMSE_cnn, 2), ", ...
800         Mean MAPE of ten CNN-LSTM models: ", round(mean_MAPE_cnn, 2)])
801
802 end
803
804 function extract_percision_results_with_observed...
805     (ts, preds_cnn, endpoint, len_train, days_ahead)
806     % Does not show plot of predictions, only results. This was used to
807     % fill out the summarizinf table.
808     sarima = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
809     gandalf = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
810     garchmod = garch('Constant',0.001,'GARCHLags',1,'ARCHLags',1);
811     gandalf.Variance = garchmod;
812

```



```

808     preds_sarima = [];
809     preds_gandalf = [];
810
811     train_data = ts(endpoint-len_train+1:endpoint);
812     for i = 1:days_ahead
813         fitted_sarima = estimate(sarima, train_data, 'Display','off');
814         fitted_gandalf= estimate(gandalf, train_data, 'Display','off');
815         preds_sarima(end+1) = forecast(fitted_sarima, 1, train_data);
816         preds_gandalf(end+1) = forecast(fitted_gandalf, 1, train_data);
817         train_data = [train_data; ts(endpoint+i)]; % add next observation to ...
            training data
818     end
819
820     mean_preds_cnn = mean(preds_cnn, 1);
821     %mse_cnn = mse_cnn_lstm(ts, mean_preds_cnn, endpoint, len_train, true);
822
823     % convert to normal scale where needed
824     ts = floor(exp(ts));
825     test_data = ts(endpoint+1: endpoint + days_ahead);
826     preds_sarima = floor(exp(preds_sarima));
827     preds_gandalf = floor(exp(preds_gandalf));
828
829     % calculate RRMSE and MAPE for single predictions
830     RRMSE_sarima = sqrt(mean((test_data - preds_sarima).^2))/mean(test_data)*100;
831     MAPE_sarima = mean(abs(test_data - preds_sarima)./abs(test_data))*100;
832     RRMSE_gandalf = sqrt(mean((test_data - preds_gandalf).^2))/mean(test_data)*100;
833     MAPE_gandalf = mean(abs(test_data - preds_gandalf)./abs(test_data))*100;
834     RRMSE_cnn = sqrt(mean((test_data - mean_preds_cnn).^2))/mean(test_data)*100;
835     MAPE_cnn = mean(abs(test_data - mean_preds_cnn)./abs(test_data))*100;
836
837     % calculate RRMSE and MAPE for the 10 CNN-LSTM models
838     RRMSE = zeros(10, 1);
839     MAPE = zeros(10, 1);
840     for i = 1:10
841         RRMSE(i) = sqrt(mean((test_data.' - preds_cnn(i, ...
            :).^2))/mean(test_data)*100;
842         MAPE(i) = mean(abs(test_data.' - preds_cnn(i, :))./abs(test_data.'))*100;
843     end
844
845     mean_RRMSE_cnn = round(mean(RRMSE), 2);
846     mean_MAPE_cnn = round(mean(MAPE), 2);
847     disp(".....")
848     disp(["Displaying the results from the ", days_ahead,"-day predictions ...
            starting at ", ...
            endpoint, " using ", len_train, " data with test data:"])
849     disp(["RRMSE SARIMA: ", round(RRMSE_sarima, 2), ", MAPE SARIMA: ", ...
            round(MAPE_sarima, 2)])
850     disp(["RRMSE Gandalf: ", round(RRMSE_gandalf, 2), ", MAPE Gandalf: ", ...
            round(MAPE_gandalf, 2)])
851     disp(["RRMSE mean CNN-LSTM: ", round(RRMSE_cnn, 2), ", MAPE mean CNN-LSTM: ", ...
            round(MAPE_cnn, 2)])
852     disp(["Mean RRMSE of ten CNN-LSTM models: ", round(mean_RRMSE_cnn, 2), ", ...
            Mean MAPE of ten CNN-LSTM models: ", round(mean_MAPE_cnn, 2)])
853
854
855
856
857     end
858
859     function predictions_on_simulations...
860         (ts, simulations, preds_sim_cnn, days_ahead, endpoint, is_global)
861         % Displays three forecasts on simulated realizations
862         sarima = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
863         gandalf = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
864         garchmod = garch('Constant',0.001,'GARCLags',1,'ARCLags',1);
865         gandalf.Variance = garchmod;
866
867         % simulation was on normal scale
868         simulations = log(max(simulations, 0.1));
869         tic
870         for i = 1:size(simulations,1) % iterate through each simulation
871
872             preds_sim_cnn(i,:)
873             train_data = ts(1:endpoint);

```

```

874
875 % Extract params from simulation model
876 preds_sarima = [];
877 preds_gandalf = [];
878 mse_sarima = [];
879 mse_gandalf = [];
880
881 % make predictions, iterating through the current simulation
882 for j = 1:size(simulations,2)
883     fitted_sarima = estimate(sarima, train_data, 'Display','off');
884     fitted_gandalf= estimate(gandalf, train_data, 'Display','off');
885     [preds_sarima(end+1), mse_sarima(end+1)] = forecast(fitted_sarima, 1, ...
886         train_data);
887     [preds_gandalf(end+1), mse_gandalf(end+1)] = forecast(fitted_gandalf, 1, ...
888         train_data);
889     train_data = [train_data; simulations(i,j)]; % add next simulation to ...
890         training data
891 end
892 % revert predictions and simulations back to normal scale
893 upper_sarima = preds_sarima + 1.96*sqrt(mse_sarima);
894 lower_sarima = preds_sarima - 1.96*sqrt(mse_sarima);
895 upper_gandalf = preds_gandalf + 1.96*sqrt(mse_gandalf);
896 lower_gandalf = preds_gandalf - 1.96*sqrt(mse_gandalf);
897
898 % convert to normal scale where needed
899 simulations(i,:) = floor(exp(simulations(i,:)));
900 preds_sarima = floor(exp(preds_sarima));preds_gandalf = ...
901     floor(exp(preds_gandalf));
902 upper_sarima = floor(exp(upper_sarima));lower_sarima = floor(exp(lower_sarima));
903 upper_gandalf = floor(exp(upper_gandalf));lower_gandalf = ...
904     floor(exp(lower_gandalf));
905
906 plot_length = 21;
907
908 if is_global
909     dates_train = index_to_date_global(endpoint-plot_length+1:endpoint);
910     dates_test = index_to_date_global(endpoint:endpoint+days_ahead-1);
911 else
912     dates_train = index_to_date_norway(endpoint-plot_length+1:endpoint);
913     dates_test = index_to_date_norway(endpoint:endpoint+days_ahead-1);
914 end
915
916 RRMSE_sarima = sqrt(mean((simulations(i,:) - ...
917     preds_sarima).^2)/mean(simulations(i,:))*100;
918 MAPE_sarima = mean(abs(simulations(i,:) - ...
919     preds_sarima)./abs(simulations(i,:))*100;
920 RRMSE_gandalf = sqrt(mean((simulations(i,:) - ...
921     preds_gandalf).^2)/mean(simulations(i,:))*100;
922 MAPE_gandalf = mean(abs(simulations(i,:) - ...
923     preds_gandalf)./abs(simulations(i,:))*100;
924 RRMSE_cnn = sqrt(mean((simulations(i,:) - ...
925     preds_sim_cnn(i,:).^2)/mean(simulations(i,:))*100;
926 MAPE_cnn = mean(abs(simulations(i,:) - ...
927     preds_sim_cnn(i,:))./abs(simulations(i,:))*100;
928
929 figure
930 hold on
931 data = plot(dates_train, ...
932     floor(exp(ts(endpoint-plot_length+1:endpoint))), 'Color', [0.25, 0.25, 0.25]);
933 obs = plot(dates_test, simulations(i,:), 'Color', '#A2142F', 'LineWidth', 3);
934 predictions_sarima = plot(dates_test, preds_sarima, 'Color', [0.4940 0.1840 ...
935     0.5560], 'LineWidth', 2);
936 u_sarima = plot(dates_test, upper_sarima, '--', 'Color', [0.4940 0.1840 ...
937     0.5560], 'LineWidth', 1);
938 l_sarima = plot(dates_test, lower_sarima, '--', 'Color', [0.4940 0.1840 ...
939     0.5560], 'LineWidth', 1);
940 predictions_gandalf = plot(dates_test, preds_gandalf, 'Color', [.2, .9, .5], ...
941     'LineWidth', 2);
942 u_gandalf = plot(dates_test, upper_gandalf, '--', 'Color', [.2, .9, .5], ...
943     'LineWidth', 1);
944 l_gandalf = plot(dates_test, lower_gandalf, '--', 'Color', [.2, .9, .5], ...
945     'LineWidth', 1);

```

```

929     cnn = plot(dates_test,preds_sim_cnn(i,:), 'Color', [0 0.4470 0.7410], ...
930             'LineWidth',2);
931
932     % gtext(['RRMSE sarima: ', num2str(RRMSE_sarima), newline,'RRMSE gandalf: ', ...
933           num2str(RRMSE_gandalf)],'FontSize', 24)
934
935     gtext([ ...
936           '\color[rgb]{' sprintf('%f,%f,%f', [0.4940 0.1840 0.5560] ) '}' SARIMA ...
937           model: RRMSE = ', num2str(round(RRMSE_sarima, 2)), '%', ...
938           ', MAPE = ', num2str(round(MAPE_sarima, 2)), '%',newline, ...
939           '\color[rgb]{' sprintf('%f,%f,%f', [.2, .9, .5]) '}' Gandalf model: RRMSE ...
940           = ', num2str(round(RRMSE_gandalf, 2)), '%', ...
941           ', MAPE = ', num2str(round(MAPE_gandalf, 2)), '%', newline, ...
942           '\color[rgb]{' sprintf('%f,%f,%f', [0 0.4470 0.7410]) '}' CNN-LSTM model: ...
943           RRMSE = ', num2str(round(RRMSE_cnn, 2)), '%', ...
944           ', MAPE = ', num2str(round(MAPE_cnn, 2)), '%',
945           ], 'Interpreter', 'tex','FontSize', 30);
946
947     legend([data, obs, predictions_sarima, u_sarima, predictions_gandalf, ...
948           u_gandalf cnn],...
949           'Training data', 'Simulated realization', 'Forecast SARIMA','95% interval ...
950           SARIMA', ...
951           'Forecast Gandalf','95% interval Gandalf', 'Forecast CNN-LSTM', ...
952           'Location', 'NorthWest','FontSize', 30)
953
954     ylabel('New cases')
955     xlabel('Date')
956     set(gcf,'Color','w')
957     set(gca,'FontSize',24)
958     hold off
959     ax = gca;
960     if is_global
961         ax.YAxis.Exponent = 3;
962     else
963         ax.YAxis.Exponent = 0;
964     end
965     hold off
966 end
967 toc
968 end
969
970
971
972 function plot_simulated_parameters...
973 (ts, endpoint, ma_sim, sma_sim, arch_sim, garch_sim)
974 %ma_sim = 28x1000 array
975 gandalf = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
976 garchmod = garch('Constant',0.001,'GARCHLags',1, 'ARCHLags',1);
977 gandalf.Variance = garchmod;
978
979 [~, variances]= estimate(gandalf, ts(1:196), 'Display','off');
980 SDs = diag(variances).^0.5;
981 SDs = SDs([2 3 6 5]).';
982
983
984 % first element contain parametes fitted on training data only:
985 original_ma = repelem(ma_sim(1,1),34);
986 upper_ma = original_ma +2*SDs(1);
987 lower_ma = original_ma -2*SDs(1);
988 original_sma = repelem(sma_sim(1,1),34);
989 upper_sma = original_sma +2*SDs(2);
990 lower_sma = original_sma -2*SDs(2);
991 original_arch = repelem(arch_sim(1,1),34);
992 upper_arch = original_arch +2*SDs(3);
993 lower_arch = original_arch -2*SDs(3);
994 original_garch = repelem(garch_sim(1,1),34);
995 upper_garch = original_garch +2*SDs(4);
996 lower_garch = original_garch -2*SDs(4);
997
998 % Plotting
999 dates_before = index_to_date_global(endpoint-3:endpoint-1);
1000 dates_sim = index_to_date_global(endpoint:endpoint+27);
1001 dates_after = index_to_date_global(endpoint+28:endpoint + 30);
1002
1003 figure
1004 subplot(2,2,1)
1005 hold on

```

```

994 orig_ma = plot([dates_before dates_sim dates_after], original_ma, '--', ...
995               'Color', 'red', 'LineWidth', 2);
996 up_ma = plot([dates_before dates_sim dates_after], upper_ma, '--', 'Color', ...
997             'black', 'LineWidth', 1);
998 low_ma = plot([dates_before dates_sim dates_after], lower_ma, '--', 'Color', ...
999             'black', 'LineWidth', 1);
1000 ma = plot(dates_sim, ma_sim, 'LineWidth', 1, 'Color', 'red');
1001 ylabel('Estimated \theta_{MA}')
1002 xlabel('Date')
1003 set(gcf, 'Color', 'w')
1004 set(gca, 'FontSize', 24)
1005 legend([orig_ma up_ma ma(1)], 'Estimated \theta_{MA} on training data', ...
1006        'Confidence interval on training data', ...
1007        'Estimated \theta_{MA} on simulations', 'Location', 'northwest');
1008 hold off
1009
1010 subplot(2,2,2)
1011 hold on
1012 orig_sma = plot([dates_before dates_sim dates_after], original_sma, '--', ...
1013               'Color', [0 1 0.7], 'LineWidth', 2);
1014 up_sma = plot([dates_before dates_sim dates_after], upper_sma, '--', 'Color', ...
1015             'black', 'LineWidth', 1);
1016 low_sma = plot([dates_before dates_sim dates_after], lower_sma, '--', ...
1017             'Color', 'black', 'LineWidth', 1);
1018 sma = plot(dates_sim, sma_sim, 'LineWidth', 1, 'Color', [0 1 0.7]);
1019 ylabel('Estimated \theta_{SMA}')
1020 xlabel('Date')
1021 set(gcf, 'Color', 'w')
1022 set(gca, 'FontSize', 24)
1023 legend([orig_sma up_sma sma(1)], 'Estimated \theta_{SMA} on training data', ...
1024        'Confidence interval on training data', ...
1025        'Estimated \theta_{SMA} on simulations', 'Location', 'northwest');
1026 hold off
1027
1028 subplot(2,2,3)
1029 hold on
1030 orig_arch = plot([dates_before dates_sim dates_after], original_arch, '--', ...
1031                'Color', 'blue', 'LineWidth', 2);
1032 arch = plot(dates_sim, arch_sim, 'LineWidth', 1, 'Color', 'blue');
1033 up_arch = plot([dates_before dates_sim dates_after], upper_arch, '--', ...
1034              'Color', 'black', 'LineWidth', 1);
1035 low_arch = plot([dates_before dates_sim dates_after], lower_arch, '--', ...
1036              'Color', 'black', 'LineWidth', 1);
1037 ylabel('Estimated \alpha_1')
1038 xlabel('Date')
1039 set(gcf, 'Color', 'w')
1040 set(gca, 'FontSize', 24)
1041 legend([orig_arch up_arch arch(1)], 'Estimated \alpha_1 on training data', ...
1042        'Confidence interval on training data', ...
1043        'Estimated \alpha_1 on simulations', 'Location', 'northwest');
1044 hold off
1045
1046 subplot(2,2,4)
1047 hold on
1048 orig_garch = plot([dates_before dates_sim dates_after], original_garch, '--', ...
1049                 'Color', 'green', 'LineWidth', 2);
1050 garch_ = plot(dates_sim, garch_sim, 'LineWidth', 1, 'Color', 'green');
1051 up_garch = plot([dates_before dates_sim dates_after], upper_garch, '--', ...
1052               'Color', 'black', 'LineWidth', 1);
1053 low_garch = plot([dates_before dates_sim dates_after], lower_garch, '--', ...
1054               'Color', 'black', 'LineWidth', 1);
1055 ylabel('Estimated \beta_1')
1056 xlabel('Date')
1057 set(gcf, 'Color', 'w')
1058 set(gca, 'FontSize', 24)
1059 legend([orig_garch up_garch garch_(1)], 'Estimated \beta_1 on training data', ...
1060        'Confidence interval on training data', ...
1061        'Estimated \beta_1 on simulations', 'Location', 'northwest');
1062 hold off
1063 end
1064
1065 function plot_simulated_RRMSE_MAPE(RRMSEs, MAPEs)
1066

```

```

1055 % Uncomment the corresponding accuracy measures on the original data set
1056 %     mean_RRMSE_orig = 3.86; % RRMSE on original SARIMA
1057 %     mean_MAPE_orig = 2.89; % MAPE on original SARIMA
1058 %     mean_RRMSE_orig = 4.81; % RRMSE on original Gandalf
1059 %     mean_MAPE_orig = 3.56; % MAPE on original Gandalf
1060 mean_RRMSE_orig = 0.3574; % RRMSE on original Gandalf log scale
1061 mean_MAPE_orig = 0.2834; % MAPE on original Gandalf log scale
1062
1063 %     mean_RRMSE_orig = 4.58; % RRMSE on original CNN-LSTM
1064 %     mean_MAPE_orig = 3.21; % MAPE on original CNN-LSTM
1065
1066 figure
1067 set(gcf,'color','w')
1068 subplot(1, 2, 1)
1069 hold on
1070 h = histogram(RRMSEs, 100, 'FaceColor', '#00b551', 'EdgeColor','#00b551');
1071 orig_RRMSE = xline(mean_RRMSE_orig,'--', ['Original RRMSE: ', ...
1072     num2str(mean_RRMSE_orig)], ...
1073     'Color', [0 0 1], 'LineWidth',2, 'FontSize',16);
1074 ylabel("Frequency")
1075 xlabel("RRMSE")
1076 hold off
1077 set(gca,'FontSize',24)
1078 subplot(1, 2, 2)
1079 hold on
1080 h = histogram(MAPEs, 100, 'FaceColor', '#ad003d', 'EdgeColor','#ad003d');
1081 orig_MAPE = xline(mean_MAPE_orig,'--', ['Original MAPE: ', ...
1082     num2str(mean_MAPE_orig)], ...
1083     'Color', [0 0 1], 'LineWidth',2, 'FontSize',16);
1084 ylabel("Frequency")
1085 xlabel("MAPE")
1086 hold off
1087 set(gca,'FontSize',24)
1088 end
1089
1090 function [RRMSEs_sarima, MAPEs_sarima, RRMSEs_gandalf, MAPEs_gandalf, ...
1091     ma_sim, sma_sim, arch_sim, garch_sim] = ...
1092     RRMSE_MAPE_from_simulation_sarima_and_gandalf(ts, simulations, endpoint, ...
1093         log_scale)
1094 % This function predicts one step ahead with simulated data from
1095 % Gandalf model using the SARIMA model and the Gandalf model
1096 sarima = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
1097 gandalf = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
1098 garchmod = garch('Constant',0.001,'GARCHLags',1, 'ARCHLags',1);
1099 gandalf.Variance = garchmod;
1100
1101 RRMSEs_sarima = [];
1102 MAPEs_sarima = [];
1103 RRMSEs_gandalf = [];
1104 MAPEs_gandalf = [];
1105
1106 ma_sim = [];
1107 sma_sim = [];
1108 arch_sim = [];
1109 garch_sim = [];
1110
1111 simulations = log(max(simulations, 0.1));
1112
1113 % RRMSE and MAPE on the real test data, for reference later:
1114 train_data = ts(1:endpoint);
1115 preds_sarima = [];
1116 preds_gandalf = [];
1117
1118 for i = 1:size(simulations,2)
1119     fitted_sarima = estimate(sarima, train_data, 'Display','off');
1120     fitted_gandalf= estimate(gandalf, train_data, 'Display','off');
1121     preds_sarima(end+1) = forecast(fitted_sarima, 1, train_data);
1122     preds_gandalf(end+1) = forecast(fitted_gandalf, 1, train_data);
1123     train_data = [train_data; ts(endpoint+i)]; % add next simulation to training data
1124 end

```

```

1125
1126 test_data = ts(endpoint+1:endpoint+size(simulations,2));
1127
1128 % Id : Ta ut parameterverdiane i loopen over
1129
1130 RRMSE_sarima_orig = sqrt(mean((test_data - preds_sarima.').^2))/mean(test_data)*100
1131 MAPE_sarima_orig = mean(abs(test_data - preds_sarima.')./abs(test_data))*100
1132 RRMSE_gandalf_orig = sqrt(mean((test_data - preds_gandalf.').^2))/mean(test_data)*100
1133 MAPE_gandalf_orig = mean(abs(test_data - preds_gandalf.')./abs(test_data))*100
1134
1135
1136 % simulation was on normal scale, need to be on log-scale to fulfill the
1137 % SARIMA models assumptions.
1138
1139 tic
1140 for i = 1:size(simulations,1) % iterate through each simulation
1141     train_data = ts(1:endpoint);
1142     preds_sarima = [];
1143     preds_gandalf = [];
1144
1145     ma = [];
1146     sma = [];
1147     arch = [];
1148     garch_ = [];
1149     % make predictions, iterating through the current simulation
1150     for j = 1:size(simulations,2)
1151         fitted_sarima = estimate(sarima, train_data, 'Display','off');
1152         fitted_gandalf = estimate(gandalf, train_data, 'Display','off');
1153
1154         ma(end+1) = round(cell2mat(fitted_gandalf.MA), 4);
1155         sma(end+1) = round(cell2mat(fitted_gandalf.SMA(7)), 4);
1156         arch(end+1) = round(cell2mat(fitted_gandalf.Variance.ARCH), 4);
1157         garch_(end+1) = round(cell2mat(fitted_gandalf.Variance.GARCH), 4);
1158
1159         preds_sarima(end+1) = forecast(fitted_sarima, 1, train_data);
1160         preds_gandalf(end+1) = forecast(fitted_gandalf, 1, train_data);
1161         train_data = [train_data; simulations(i,j)]; % add next simulation to ...
            training data
1162     end
1163
1164     if log_scale == false
1165         % revert predictions and simulations back to normal scale
1166         preds_sarima = floor(exp(preds_sarima));
1167         preds_gandalf = floor(exp(preds_gandalf));
1168         simulations(i,:) = floor(exp(simulations(i,:)));
1169     end
1170
1171     % add parameters
1172     ma_sim = [ma_sim; ma];
1173     sma_sim = [sma_sim; sma];
1174     arch_sim = [arch_sim; arch];
1175     garch_sim = [garch_sim; garch_];
1176
1177     % calculate RRMSE and MAPE
1178     RRMSEs_sarima(end+1) = sqrt(mean((simulations(i,:) - ...
            preds_sarima).^2))/mean(simulations(i,:))*100;
1179     MAPEs_sarima(end+1) = mean(abs(simulations(i,:) - ...
            preds_sarima)./abs(simulations(i,:)))*100;
1180     RRMSEs_gandalf(end+1) = sqrt(mean((simulations(i,:) - ...
            preds_gandalf).^2))/mean(simulations(i,:))*100;
1181     MAPEs_gandalf(end+1) = mean(abs(simulations(i,:) - ...
            preds_gandalf)./abs(simulations(i,:)))*100;
1182     length(MAPEs_gandalf) % To monitor progress
1183 end
1184 toc
1185 end
1186
1187 function plot_simulations(ts, simulations, days_ahead, endpoint)
1188     gandalf = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
1189     garchmod = garch('Constant',0.001,'GARCHLags',1,'ARCHLags',1);
1190     gandalf.Variance = garchmod;
1191
1192     train_data = ts(1:endpoint);

```

```

1193 %     test_data = floor(exp(ts(endpoint+1:endpoint+28)));
1194 rng(123) % Set seed for reproducibility
1195 fitted_model = estimate(gandalf, train_data, 'Display', 'off');
1196 [prediction, YMSE]= forecast(fitted_model, days_ahead, 'Y0', train_data);
1197 upper = prediction + 1.96*sqrt(YMSE);
1198 lower = prediction - 1.96*sqrt(YMSE);
1199 prediction = floor(exp(prediction)); upper = floor(exp(upper)); lower = ...
        floor(exp(lower));

1200
1201 % Plot
1202 plot_length = 21;
1203 dates_train = index_to_date_global(endpoint-plot_length+1:endpoint);
1204 dates_test = index_to_date_global(endpoint+1:endpoint+days_ahead);
1205 figure
1206 hold on
1207 data = plot(dates_train, ...
        floor(exp(ts(endpoint-plot_length+1:endpoint))), 'Color', [0.25, 0.25, 0.25]);
1208 obs = plot(dates_test, floor(exp(ts(endpoint+1:endpoint+days_ahead))), ...
        'Color', [1, 0, 0], 'LineWidth', 3);
1209 pred = plot(dates_test, prediction, 'Color', [.2, .9, .5], 'LineWidth', 2);
1210 up = plot(dates_test, upper, '--', 'Color', [.2, .9, .5], 'LineWidth', 1);
1211 low = plot(dates_test, lower, '--', 'Color', [.2, .9, .5], 'LineWidth', 1);
1212 sims = plot(dates_test, simulations(:,1:3), 'Color', '#A2142F', 'LineWidth', 2);
1213 up_sims = plot(dates_test, prctile(simulations, 97.5, 2), '--', 'Color', 'b', ...
        'LineWidth', 1);
1214 low_sims = plot(dates_test, prctile(simulations, 2.5, 2), '--', 'Color', 'b', ...
        'LineWidth', 1);
1215 %med_sims = plot(dates_test, median(simulations, 2), 'Color', 'b', 'LineWidth', 2);
1216 % Hvis du vil plotte alle med gjennomsiktighet
1217 %     for i=1:length(sims)
1218 %         sims(i).Color = [sims(i).Color 0.05]; % alpha=0.1
1219 %     end
1220 legend([data obs sims(1) pred up up_sims], 'Training data', 'Test data', ...
        'Simulations', 'Forecast Gandalf', ...
        '95% theoretical interval Gandalf', '95% quantile range simulations', ...
        'Location', 'northwest')
1221 ylabel('New cases')
1222 xlabel('Date')
1223 set(gcf, 'color', 'w')
1224 ax = gca;
1225 ax.YAxis.Exponent = 3;
1226 set(gca, 'FontSize', 24)
1227 hold off
1228 end
1229
1230
1231 function simulations = simulate_from_Gandalf...
1232     (ts, days_ahead, endpoint, amount_simulations)
1233     gandalf = arima('Constant', 0, 'D', 1, 'Seasonality', 7, 'MALags', 1, 'SMALags', 7);
1234     garchmod = garch('Constant', 0.001, 'GARCHLags', 1, 'ARCHLags', 1);
1235     gandalf.Variance = garchmod;
1236
1237     train_data = ts(1:endpoint);
1238     rng(123) % Set seed for reproducibility
1239     fitted_model = estimate(gandalf, train_data, 'Display', 'off');
1240     [prediction, YMSE]= forecast(fitted_model, days_ahead, 'Y0', train_data);
1241     upper = prediction + 1.96*sqrt(YMSE);
1242     lower = prediction - 1.96*sqrt(YMSE);
1243     prediction = floor(exp(prediction));
1244     upper = floor(exp(upper));
1245     lower = floor(exp(lower));
1246     simulations = floor(exp(simulate(fitted_model, ...
1247         days_ahead, 'NumPaths', amount_simulations, 'Y0', train_data)));
1248     % Save the results. These will be used in Python later
1249     writematrix(simulations.', ...
1250         'Matlab\Saved variables\simulations_from_gandalf.csv')
1251 end
1252
1253 function plot_loss_function(loss_func)
1254     [%minimum, minimum_ind] = min(loss_func); % find minimal loss and its ...
        corresponding epoc
1255     figure
1256     hold on
1257     loss = plot(loss_func, 'LineWidth', 3);

```

```

1258     yline(loss_func(100), '--', 'Mean loss after 100 epochs', ...
1259           'Color', 'red', 'LineWidth',2, 'LabelVerticalAlignment', 'top', ...
1260           'FontSize', 24);
1261     ylim([-0.001 0.01])
1262     ylabel('Mean loss')
1263     xlabel('Epoocs')
1264     set(gcf, 'color', 'w')
1265     set(gca, 'FontSize', 24)
1266     hold off
1267 end
1268 function plot_acf_of_res(ts, endpoint, len_train, garch_noise)
1269     % plot ACF of residuals and squared residuals
1270
1271     train_data = ts(endpoint-len_train+1:endpoint);
1272
1273     mod = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
1274     if garch_noise
1275         noise = garch('Constant',0.001,'GARCHLags',1, 'ARCHLags',1);
1276         mod.Variance = noise
1277     end
1278     % fit model
1279     fitted_model = estimate(mod, train_data, 'Display','off');
1280     residuals = infer(fitted_model, train_data);
1281
1282     % res = plot(residuals)
1283     [r, p_arch] = archtest(residuals, 'Lags', 2);
1284     % [r, p_lbq] = lbqtest(residuals);
1285     figure
1286
1287     [acf,lags,bounds] = autocorr(residuals);
1288
1289     stem(lags,acf, 'LineWidth',2); xlabel('Lag'); ylabel('\gamma(k)');
1290
1291     hold on;
1292     h = plot(lags,bounds(1)*ones(length(acf),1), '--','LineWidth',2, 'Color',[1 0 ...
1293           0]);
1294     h1 = plot(lags,bounds(2)*ones(length(acf),1), '--', 'LineWidth',2, 'Color',[1 ...
1295           0 0]);
1296     title([''])
1297     % title(['Sample ACF of residuals of model from ',...
1298           % datestr(index_to_date_norway(endpoint-len_train+1)), ' to ' ...
1299           % datestr(index_to_date_norway(endpoint))])
1300     set(gcf, 'color', 'w')
1301     set(gca, 'FontSize', 40)
1302     figure
1303     [acf,lags,bounds] = autocorr(residuals.^2);
1304     set(gca, 'FontSize', 40)
1305     stem(lags,acf, 'LineWidth',2); xlabel('Lag'); ylabel('\gamma(k)');
1306     hold on;
1307     h = plot(lags,bounds(1)*ones(length(acf),1), '--','LineWidth',2, 'Color',[1 0 ...
1308           0]);
1309     h1 = plot(lags,bounds(2)*ones(length(acf),1), '--', 'LineWidth',2, 'Color',[1 ...
1310           0 0]);
1311     gtext(['P-value for Engels ARCH test: ', num2str(p_arch)], 'FontSize', 40) % , ...
1312           newline, 'P-value for Ljung-Box Q-test: ', num2str(p_lbq)]
1313     title([''])
1314     % title(['Sample ACF of squared residuals of model from ', ...
1315           % datestr(index_to_date_norway(endpoint-len_train+1)), ' to ' ...
1316           % datestr(index_to_date_norway(endpoint))])
1317     set(gcf, 'color', 'w')
1318     set(gca, 'FontSize', 40)
1319 end
1320 function plot_preds_sarima_gandalf...
1321     (ts, endpoint, len_train, days_ahead, is_global)
1322     % Directly predicts the next days_ahead days after endpoint
1323     % These are then plotted and compared to the CNN-LSTM model with the
1324     % same prediction sceeme
1325
1326     % ts should be on log-scale
1327
1328     plot_length = 14;

```



```

1323
1324 sarima = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
1325 gandalf = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
1326 garchmod = garch('Constant',0.001,'GARCHLags',1,'ARCHLags',1);
1327 gandalf.Variance = garchmod;
1328
1329 preds_sarima = [];
1330 mse_sarima = [];
1331 preds_gandalf = [];
1332 mse_gandalf = [];
1333
1334 train_data = ts(endpoint-len_train+1:endpoint);
1335
1336 [fitted_sarima, ~, logL_sarima] = estimate(sarima, train_data, 'Display','off');
1337 [fitted_gandalf, ~, logL_gandalf] = estimate(gandalf, train_data, ...
1338     'Display','off');
1339 [preds_sarima, mse_sarima] = forecast(fitted_sarima, days_ahead, train_data);
1340 [preds_gandalf, mse_gandalf] = forecast(fitted_gandalf, days_ahead, train_data);
1341
1342 % Display the AICc values for both models
1343 [~, ~, AICc_sarima] = aicbic(logL_sarima, 3, len_train);
1344 [~, ~, AICc_gandalf] = aicbic(logL_gandalf, 4, len_train);
1345 AICc_sarima = AICc_sarima.aicc
1346 AICc_gandalf = AICc_gandalf.aicc
1347
1348 % get parameter estimates
1349 MA_sarima = cell2mat(fitted_sarima.MA); SMA_sarima = ...
1350     cell2mat(fitted_sarima.SMA(7));
1351 sigma2 = fitted_sarima.Variance;
1352 round(SMA_sarima, 2)
1353 MA_gandalf = cell2mat(fitted_gandalf.MA); SMA_gandalf = ...
1354     cell2mat(fitted_gandalf.SMA(7));
1355 GARCH = cell2mat(fitted_gandalf.Variance.GARCH); ARCH = ...
1356     cell2mat(fitted_gandalf.Variance.ARCH);
1357
1358 % create 95% intervals
1359 upper_sarima = preds_sarima + 1.96*sqrt(mse_sarima);
1360 lower_sarima = preds_sarima - 1.96*sqrt(mse_sarima);
1361 upper_gandalf = preds_gandalf + 1.96*sqrt(mse_gandalf);
1362 lower_gandalf = preds_gandalf - 1.96*sqrt(mse_gandalf);
1363
1364 % convert to normal scale where needed
1365 ts = floor(exp(ts));
1366 test_data = ts(endpoint+1: endpoint + days_ahead);
1367 preds_sarima = floor(exp(preds_sarima));preds_gandalf = ...
1368     floor(exp(preds_gandalf));
1369 upper_sarima = floor(exp(upper_sarima));lower_sarima = floor(exp(lower_sarima));
1370 upper_gandalf = floor(exp(upper_gandalf));lower_gandalf = ...
1371     floor(exp(lower_gandalf));
1372
1373 % calculate RRMSE and MAPE
1374 RRMSE_sarima = sqrt(mean((test_data - preds_sarima).^2)/mean(test_data)*100;
1375 MAPE_sarima = mean(abs(test_data - preds_sarima)./abs(test_data))*100;
1376 RRMSE_gandalf = sqrt(mean((test_data - preds_gandalf).^2)/mean(test_data)*100;
1377 MAPE_gandalf = mean(abs(test_data - preds_gandalf)./abs(test_data))*100;
1378
1379 % set dates
1380 if is_global
1381     dates_train = index_to_date_global(endpoint-plot_length+1:endpoint);
1382     dates_test = index_to_date_global(endpoint+1:endpoint+days_ahead);
1383 else
1384     dates_train = index_to_date_norway(endpoint-plot_length+1:endpoint);
1385     dates_test = index_to_date_norway(endpoint+1:endpoint+days_ahead);
1386 end
1387
1388 % plotting
1389 figure
1390 hold on
1391
1392 data = plot(dates_train, ts(endpoint-plot_length+1:endpoint), 'Color', [0.25, ...
1393     0.25, 0.25]);
1394 obs = plot(dates_test, test_data, 'Color', [1, 0, 0], 'LineWidth', 3);

```

```

1389     predictions_sarima = plot(dates_test,preds_sarima, 'Color', [0.4940 0.1840 ...
1390         0.5560], 'LineWidth',2);
1391     u_sarima = plot(dates_test,upper_sarima, '--','Color', [0.4940 0.1840 ...
1392         0.5560], 'LineWidth',1);
1393     l_sarima = plot(dates_test,lower_sarima, '--','Color', [0.4940 0.1840 ...
1394         0.5560], 'LineWidth',1);
1395     predictions_gandalf = plot(dates_test,preds_gandalf, 'Color', [.2, .9, .5], ...
1396         'LineWidth',2);
1397     u_gandalf = plot(dates_test,upper_gandalf, '--', 'Color', [.2, .9, .5], ...
1398         'LineWidth',1);
1399     l_gandalf = plot(dates_test,lower_gandalf, '--', 'Color', [.2, .9, .5], ...
1400         'LineWidth',1);
1401
1402     %ylim([0 1700])
1403
1404     % insert percision results in plot
1405     set(gca,'FontSize',24)
1406     gtext([ ...
1407         '\color[rgb]{', sprintf('%f,%f,%f', [0.4940 0.1840 0.5560] ) ,'} SARIMA ...
1408         model: RRMSE = ', num2str(round(RRMSE_sarima, 2)), '%', ...
1409         ', MAPE = ', num2str(round(MAPE_sarima, 2)), '%', newline, ...
1410         '\theta_{MA} = ', num2str(round(MA_sarima, 2)), ', \theta_{SMA} = ', ...
1411         num2str(round(SMA_sarima, 2)), ', \sigma^{2} = ', num2str(round(sigma2, ...
1412         2)), newline, ...
1413         '\color[rgb]{', sprintf('%f,%f,%f', [.2, .9, .5]) ,'} Gandalf model: RRMSE ...
1414         = ', num2str(round(RRMSE_gandalf, 2)), '%', ...
1415         ', MAPE = ', num2str(round(MAPE_gandalf, 2)), '%', newline, ...
1416         '\theta_{MA} = ', num2str(round(MA_gandalf, 2)), ', \theta_{SMA} = ', ...
1417         num2str(round(SMA_gandalf, 2)), ', \alpha_{1} = ', num2str(round(ARCH, ...
1418         2)), ...
1419         ', \beta_{1} = ', num2str(round(GARCH, 2))], 'Interpreter', ...
1420         'tex','FontSize', 30);
1421
1422     legend([data, obs, predictions_sarima, u_sarima, predictions_gandalf, ...
1423         u_gandalf],...
1424         'Training data', 'Test data', 'Forecast SARIMA','95% interval SARIMA', ...
1425         'Forecast Gandalf','95% interval Gandalf', 'NorthWest','FontSize', 30)
1426     ylabel('New cases','FontSize', 30)
1427     xlabel('Date','FontSize', 30)
1428
1429     if is_global
1430         title(['Forecasts of Global data from ', ...
1431             datestr(index_to_date_global(endpoint+1),'FontSize', 20)])
1432     else
1433         title(['Forecasts of Norwegian data from ', ...
1434             datestr(index_to_date_norway(endpoint+1)),'FontSize', 20)])
1435     end
1436     set(gcf,'color','w')
1437
1438     ax = gca;
1439
1440     if is_global
1441         ax.YAxis.Exponent = 3;
1442     else
1443         ax.YAxis.Exponent = 0;
1444     end
1445     hold off
1446 end
1447
1448 function plot_parameters_from_prediction_scheme_1...
1449     (ts, endpoint, len_train, days_ahead, is_global)
1450     sarima = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
1451     gandalf = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
1452     garchmod = garch('Constant',0.001,'GARCHLags',1, 'ARCHLags',1);
1453     gandalf.Variance = garchmod;
1454
1455     preds_sarima = [];
1456     preds_gandalf = [];
1457
1458     train_data = ts(endpoint-len_train+1:endpoint);
1459
1460     % Gather initial SARIMA parameters

```

```

1448 fitted_sarima = estimate(sarima, train_data, 'Display','off');
1449 inintial_MA_sarima = round(cell2mat(fitted_sarima.MA), 4);
1450 inintial_SMA_sarima = round(cell2mat(fitted_sarima.SMA(7)), 4);
1451 inintial_SIGMA_sarima = round(fitted_sarima.Variance, 4);
1452
1453 % Gather initial gandalf parameters
1454 fitted_gandalf = estimate(gandalf, train_data, 'Display','off');
1455 inintial_MA_gandalf = round(cell2mat(fitted_gandalf.MA), 4);
1456 inintial_SMA_gandalf = round(cell2mat(fitted_gandalf.SMA(7)), 4);
1457 inintial_ARCH_gandalf = round(cell2mat(fitted_gandalf.Variance.ARCH), 4);
1458 inintial_GARCH_gandalf = round(cell2mat(fitted_gandalf.Variance.GARCH), 4);
1459
1460 preds_sarima = [];
1461 preds_gandalf = [];
1462
1463 MA_sarima = [];
1464 SMA_sarima = [];
1465 SIGMA_sarima = [];
1466 MA_upper_sarima = [];
1467 SMA_upper_sarima = [];
1468 SIGMA_upper_sarima = [];
1469 MA_lower_sarima = [];
1470 SMA_lower_sarima = [];
1471 SIGMA_lower_sarima = [];
1472 MA_lower_gandalf = [];
1473 SMA_lower_gandalf = [];
1474 ARCH_lower_gandalf = [];
1475 GARCH_lower_gandalf = [];
1476 MA_gandalf = [];
1477 SMA_gandalf = [];
1478 ARCH_gandalf = [];
1479 GARCH_gandalf = [];
1480 MA_upper_gandalf = [];
1481 SMA_upper_gandalf = [];
1482 ARCH_upper_gandalf = [];
1483 GARCH_upper_gandalf = [];
1484 MA_lower_gandalf = [];
1485 SMA_lower_gandalf = [];
1486 ARCH_lower_gandalf = [];
1487 GARCH_lower_gandalf = [];
1488
1489 for i = 1:days_ahead
1490     [fitted_sarima, varmat_sarima] = estimate(sarima, train_data, ...
1491         'Display','off');
1492     [fitted_gandalf, varmat_gandalf] = estimate(gandalf, train_data, ...
1493         'Display','off');
1494
1495     SDs = diag(varmat_sarima).^0.5;
1496     SDs = SDs([2 3 4]).';
1497     SDs = diag(varmat_gandalf).^0.5;
1498     SDs = SDs([2 3 6 5]).';
1499
1500     % Collect parameter estimates from SARIMA model
1501     MA_new_sarima = round(cell2mat(fitted_sarima.MA), 4);
1502     MA_sarima = [MA_sarima MA_new_sarima];
1503     MA_upper_sarima = [MA_upper_sarima MA_new_sarima+2*SDs(1)];
1504     MA_lower_sarima = [MA_lower_sarima MA_new_sarima-2*SDs(1)];
1505     SMA_new_sarima = round(cell2mat(fitted_sarima.SMA(7)), 4);
1506     SMA_sarima = [SMA_sarima SMA_new_sarima];
1507     SMA_upper_sarima = [SMA_upper_sarima SMA_new_sarima+2*SDs(2)];
1508     SMA_lower_sarima = [SMA_lower_sarima SMA_new_sarima-2*SDs(2)];
1509     SIGMA_new_sarima = round(fitted_sarima.Variance, 4);
1510     SIGMA_sarima = [SIGMA_sarima SIGMA_new_sarima];
1511     SIGMA_upper_sarima = [SIGMA_upper_sarima SIGMA_new_sarima+2*SDs(3)];
1512     SIGMA_lower_sarima = [SIGMA_lower_sarima SIGMA_new_sarima-2*SDs(4)];
1513
1514     % Collect parameter estimates from Gandalf model
1515     MA_new_gandalf = round(cell2mat(fitted_gandalf.MA), 4);
1516     MA_gandalf = [MA_gandalf MA_new_gandalf];
1517     MA_upper_gandalf = [MA_upper_gandalf MA_new_gandalf+2*SDs(1)];
1518     MA_lower_gandalf = [MA_lower_gandalf MA_new_gandalf-2*SDs(1)];
1519     SMA_new_gandalf = round(cell2mat(fitted_gandalf.SMA(7)), 4);
1520     SMA_gandalf = [SMA_gandalf SMA_new_gandalf];

```

```

1519     SMA_upper_gandalf = [SMA_upper_gandalf SMA_new_gandalf+2*SDs(2)];
1520     SMA_lower_gandalf = [SMA_lower_gandalf SMA_new_gandalf-2*SDs(2)];
1521     ARCH_new_gandalf = round(cell2mat(fitted_gandalf.Variance.ARCH), 4);
1522     ARCH_gandalf = [ARCH_gandalf ARCH_new_gandalf];
1523     ARCH_upper_gandalf = [ARCH_upper_gandalf ARCH_new_gandalf+2*SDs(3)];
1524     ARCH_lower_gandalf = [ARCH_lower_gandalf ARCH_new_gandalf-2*SDs(4)];
1525     GARCH_new_gandalf = round(cell2mat(fitted_gandalf.Variance.GARCH), 4);
1526     GARCH_gandalf = [GARCH_gandalf GARCH_new_gandalf];
1527     GARCH_upper_gandalf = [GARCH_upper_gandalf GARCH_new_gandalf+2*SDs(4)];
1528     GARCH_lower_gandalf = [GARCH_lower_gandalf GARCH_new_gandalf-2*SDs(4)];
1529
1530     train_data = [train_data; ts(endpoint+i)]; % add next observation to ...
1531     training data
1532
1533     end
1534
1535     if is_global
1536         dates_test = index_to_date_global(endpoint:endpoint+days_ahead-1);
1537     else
1538         dates_test = index_to_date_norway(endpoint:endpoint+days_ahead-1);
1539     end
1540
1541     % Plot SARIMA parameters
1542     figure
1543     subplot(2,5,[1 2])
1544     hold on
1545     ma = plot(dates_test, MA_sarima, 'Color', 'red','LineWidth',2);
1546     ma_upper = plot(dates_test, MA_upper_sarima,'black:','LineWidth',2);
1547     ma_lower = plot(dates_test, MA_lower_sarima,'black:','LineWidth',2);
1548     yline(inintial_MA_sarima, '--', 'Original \theta_{MA} estimate', ...
1549         'Color', 'red', 'LineWidth',2, 'LabelVerticalAlignment', 'bottom');
1550     legend([ma, ma_upper, ma_lower], '\theta_{MA}','95% confidence interval')
1551     ylabel('Estimated \theta_{MA}')
1552     xlabel('Date')
1553     set(gcf,'Color','w')
1554     set(gca,'FontSize',24)
1555     hold off
1556
1557     subplot(2,5, [4 5])
1558     hold on
1559     sma = plot(dates_test, SMA_sarima, 'Color', [0 1 0.7],'LineWidth',2);
1560     sma_upper = plot(dates_test, SMA_upper_sarima,'black:','LineWidth',2);
1561     sma_lower = plot(dates_test, SMA_lower_sarima,'black:','LineWidth',2);
1562     yline(inintial_SMA_sarima, '--', 'Original \theta_{SMA} estimate', ...
1563         'Color', [0 1 0.7], 'LineWidth',2, 'LabelVerticalAlignment', 'bottom');
1564     legend([sma sma_upper], '\theta_{SMA}','95% confidence intervall')
1565     ylabel('Estimated \theta_{SMA}')
1566     xlabel('Date')
1567     set(gcf,'Color','w')
1568     set(gca,'FontSize',24)
1569     hold off
1570
1571     subplot(2,5, [7 8 9])
1572     hold on
1573     sigma = plot(dates_test, SIGMA_sarima, 'Color', 'blue','LineWidth',2);
1574     sigma_upper = plot(dates_test, SIGMA_upper_sarima,'black:','LineWidth',2);
1575     sigma_lower = plot(dates_test, SIGMA_lower_sarima,'black:','LineWidth',2);
1576     yline(inintial_SIGMA_sarima, '--', 'Original \sigma^2 estimate', ...
1577         'Color', 'blue', 'LineWidth',2, 'LabelVerticalAlignment', 'bottom');
1578     legend([sigma sigma_upper], '\sigma^2','95% confidence intervall')
1579     ylabel('Estimated \sigma^2')
1580     xlabel('Date')
1581     set(gcf,'Color','w')
1582     set(gca,'FontSize',24)
1583     hold off
1584
1585     % Plot Gandalf parameters
1586     figure
1587     subplot(2,2,1)
1588     hold on
1589     ma = plot(dates_test, MA_gandalf, 'Color', 'red','LineWidth',2);
1590     ma_upper = plot(dates_test, MA_upper_gandalf,'black:','LineWidth',2);
1591     ma_lower = plot(dates_test, MA_lower_gandalf,'black:','LineWidth',2);

```

```

1591 yline(inintial_MA_gandalf, '--', 'Original \theta_{MA} estimate', ...
1592       'Color', 'red', 'LineWidth',2, 'LabelVerticalAlignment', 'bottom');
1593 legend([ma, ma_upper, ma_lower], '\theta_{MA}', '95% confidence interval')
1594 ylabel('Estimated \theta_{MA}')
1595 xlabel('Date')
1596 set(gcf, 'Color', 'w')
1597 set(gca, 'FontSize', 24)
1598 hold off
1599
1600
1601 subplot(2,2,2)
1602 hold on
1603 sma = plot(dates_test, SMA_gandalf, 'Color', [0 1 0.7], 'LineWidth', 2);
1604 sma_upper = plot(dates_test, SMA_upper_gandalf, 'black:', 'LineWidth', 2);
1605 sma_lower = plot(dates_test, SMA_lower_gandalf, 'black:', 'LineWidth', 2);
1606 yline(inintial_SMA_gandalf, '--', 'Original \theta_{SMA} estimate', ...
1607       'Color', [0 1 0.7], 'LineWidth', 2, 'LabelVerticalAlignment', 'bottom');
1608 legend([sma sma_upper], '\theta_{SMA}', '95% confidence intervall')
1609 ylabel('Estimated \theta_{SMA}')
1610 xlabel('Date')
1611 set(gcf, 'Color', 'w')
1612 set(gca, 'FontSize', 24)
1613 hold off
1614
1615 subplot(2,2,3)
1616 hold on
1617 arch = plot(dates_test, ARCH_gandalf, 'Color', 'blue', 'LineWidth', 2);
1618 arch_upper = plot(dates_test, ARCH_upper_gandalf, 'black:', 'LineWidth', 2);
1619 arch_lower = plot(dates_test, ARCH_lower_gandalf, 'black:', 'LineWidth', 2);
1620 yline(inintial_ARCH_gandalf, '--', 'Original \alpha_{1} estimate', ...
1621       'Color', 'blue', 'LineWidth', 2, 'LabelVerticalAlignment', 'bottom');
1622 legend([arch arch_upper], '\alpha_{1}', '95% confidence intervall')
1623 ylabel('Estimated \alpha_{1}')
1624 xlabel('Date')
1625 set(gcf, 'Color', 'w')
1626 set(gca, 'FontSize', 24)
1627 hold off
1628
1629 subplot(2,2,4)
1630 hold on
1631 garch_plot = plot(dates_test, GARCH_gandalf, 'Color', 'green', 'LineWidth', 2);
1632 garch_upper = plot(dates_test, GARCH_upper_gandalf, 'black:', 'LineWidth', 2);
1633 garch_lower = plot(dates_test, GARCH_lower_gandalf, 'black:', 'LineWidth', 2);
1634 yline(inintial_GARCH_gandalf, '--', 'Original \beta_{1} estimate', ...
1635       'Color', 'green', 'LineWidth', 2, 'LabelVerticalAlignment', 'bottom');
1636 legend([garch_plot garch_upper], '\beta_{1}', '95% confidence intervall')
1637 ylabel('Estimated \beta_{1}')
1638 xlabel('Date')
1639 set(gcf, 'Color', 'w')
1640 set(gca, 'FontSize', 24)
1641 hold off
1642 end
1643
1644 function plot_all_predictions_with_observed...
1645 (ts, preds_cnn, endpoint, len_train, days_ahead, is_global)
1646 % Plots the two previous functions in the same figure
1647 % ts is on log-scale
1648 plot_length = 14;
1649 sarima = arima('Constant', 0, 'D', 1, 'Seasonality', 7, 'MALags', 1, 'SMALags', 7);
1650 gandalf = arima('Constant', 0, 'D', 1, 'Seasonality', 7, 'MALags', 1, 'SMALags', 7);
1651 garchmod = garch('Constant', 0.001, 'GARCHLags', 1, 'ARCHLags', 1);
1652 gandalf.Variance = garchmod;
1653
1654 preds_sarima = [];
1655 mse_sarima = [];
1656 preds_gandalf = [];
1657 mse_gandalf = [];
1658
1659 train_data = ts(endpoint-len_train+1:endpoint);
1660 for i = 1:days_ahead
1661     fitted_sarima = estimate(sarima, train_data, 'Display', 'off');
1662     fitted_gandalf = estimate(gandalf, train_data, 'Display', 'off');
1663     [preds_sarima(end+1), mse_sarima(end+1)] = ...

```

```

1664         forecast(fitted_sarima, 1, train_data);
1665         [preds_gandalf(end+1), mse_gandalf(end+1)] = ...
1666         forecast(fitted_gandalf, 1, train_data);
1667         % add next observation to training data:
1668         train_data = [train_data; ts(endpoint+i)];
1669     end
1670
1671     mean_preds_cnn = mean(preds_cnn, 1);
1672     %mse_cnn = mse_cnn_lstm(ts, mean_preds_cnn, endpoint, len_train, true);
1673
1674     % create 95% intervals
1675     upper_sarima = preds_sarima + 1.96*sqrt(mse_sarima);
1676     lower_sarima = preds_sarima - 1.96*sqrt(mse_sarima);
1677     upper_gandalf = preds_gandalf + 1.96*sqrt(mse_gandalf);
1678     lower_gandalf = preds_gandalf - 1.96*sqrt(mse_gandalf);
1679     %upper_cnn = mean_preds_cnn + 1.96*round(sqrt(mse_cnn));
1680     %lower_cnn = mean_preds_cnn - 1.96*round(sqrt(mse_cnn));
1681
1682     % convert to normal scale where needed
1683     ts = floor(exp(ts));
1684     test_data = ts(endpoint+1: endpoint + days_ahead);
1685     preds_sarima = floor(exp(preds_sarima));preds_gandalf = ...
1686         floor(exp(preds_gandalf));
1687     upper_sarima = floor(exp(upper_sarima));lower_sarima = ...
1688         floor(exp(lower_sarima));
1689     upper_gandalf = floor(exp(upper_gandalf));lower_gandalf = ...
1690         floor(exp(lower_gandalf));
1691
1692     % calculate RRMSE and MAPE for single predictions
1693     RRMSE_sarima = ...
1694         sqrt(mean((test_data - preds_sarima).^2)/mean(test_data)*100;
1695     MAPE_sarima = ...
1696         mean(abs(test_data - preds_sarima)./abs(test_data))*100;
1697     RRMSE_gandalf = ...
1698         sqrt(mean((test_data - preds_gandalf).^2)/mean(test_data)*100;
1699     MAPE_gandalf = ...
1700         mean(abs(test_data - preds_gandalf)./abs(test_data))*100;
1701     RRMSE_cnn = ...
1702         sqrt(mean((test_data - mean_preds_cnn).^2)/mean(test_data)*100;
1703     MAPE_cnn = ...
1704         mean(abs(test_data - mean_preds_cnn)./abs(test_data))*100;
1705
1706     % calculate RRMSE and MAPE for the 10 CNN-LSTM models
1707     RRMSE = zeros(10, 1);
1708     MAPE = zeros(10, 1);
1709     for i = 1:10
1710         RRMSE(i) = sqrt(mean((test_data.' - preds_cnn(i, :)).^2)/...
1711             mean(test_data)*100;
1712         MAPE(i) = mean(abs(test_data.' - preds_cnn(i, :))./...
1713             abs(test_data.'))*100;
1714     end
1715
1716     % set dates
1717     if is_global
1718         dates_train = index_to_date_global(endpoint-plot_length+1:endpoint);
1719         dates_test = index_to_date_global(endpoint+1:endpoint+days_ahead);
1720     else
1721         dates_train = index_to_date_norway(endpoint-plot_length+1:endpoint);
1722         dates_test = index_to_date_norway(endpoint+1:endpoint+days_ahead);
1723     end
1724     % generate RGB colors for plot for the 10 CNN-LSTM models
1725     rng(1234)
1726     cols = [];
1727     for i = 1:10
1728         c = [rand, rand, rand];
1729         cols = [cols; c];
1730     end
1731
1732     % plotting
1733     figure
1734     set(gca, 'FontSize', 24)
1735     subplot(1,2,1)
1736     hold on

```

```

1737 data = plot(dates_train, ts(endpoint-plot_length+1:endpoint),...
1738         'Color',[0.25, 0.25, 0.25]);
1739 obs = plot(dates_test, test_data, 'Color', [1 0 0], 'LineWidth',3);
1740 pred1 = plot(dates_test,preds_cnn(1, :), 'Color', cols(1, :), 'LineWidth',1);
1741 pred2 = plot(dates_test,preds_cnn(2, :), 'Color', cols(2, :), 'LineWidth',1);
1742 pred3 = plot(dates_test,preds_cnn(3, :), 'Color', cols(3, :), 'LineWidth',1);
1743 pred4 = plot(dates_test,preds_cnn(4, :), 'Color', cols(4, :), 'LineWidth',1);
1744 pred5 = plot(dates_test,preds_cnn(5, :), 'Color', cols(5, :), 'LineWidth',1);
1745 pred6 = plot(dates_test,preds_cnn(6, :), 'Color', cols(6, :), 'LineWidth',1);
1746 pred7 = plot(dates_test,preds_cnn(7, :), 'Color', cols(7, :), 'LineWidth',1);
1747 pred8 = plot(dates_test,preds_cnn(8, :), 'Color', cols(8, :), 'LineWidth',1);
1748 pred9 = plot(dates_test,preds_cnn(9, :), 'Color', cols(9, :), 'LineWidth',1);
1749 pred10 = plot(dates_test,preds_cnn(10, :), 'Color', cols(10, :), 'LineWidth',1);
1750
1751 gtext([ ...
1752 '\color{rgb}{', sprintf('%f,%f,%f', cols(1,:) ) '} Model 1: RRMSE = ', ...
1753         num2str(round(RRMSE(1), 2)), '%', ...
1754         ', MAPE = ', num2str(round(MAPE(1), 2)), '%',newline, ...
1755 '\color{rgb}{', sprintf('%f,%f,%f', cols(2,:) ) '} Model 2: RRMSE = ', ...
1756         num2str(round(RRMSE(2), 2)), '%', ...
1757         ', MAPE = ', num2str(round(MAPE(2), 2)), '%',newline, ...
1758 '\color{rgb}{', sprintf('%f,%f,%f', cols(3,:) ) '} Model 3: RRMSE = ', ...
1759         num2str(round(RRMSE(3), 2)), '%', ...
1760         ', MAPE = ', num2str(round(MAPE(3), 2)), '%',newline ...
1761 '\color{rgb}{', sprintf('%f,%f,%f', cols(4,:) ) '} Model 4: RRMSE = ', ...
1762         num2str(round(RRMSE(4), 2)), '%', ...
1763         ', MAPE = ', num2str(round(MAPE(4), 2)), '%',newline ...
1764 '\color{rgb}{', sprintf('%f,%f,%f', cols(5,:) ) '} Model 5: RRMSE = ', ...
1765         num2str(round(RRMSE(5), 2)), '%', ...
1766         ', MAPE = ', num2str(round(MAPE(5), 2)), '%',newline ...
1767 '\color{rgb}{', sprintf('%f,%f,%f', cols(6,:) ) '} Model 6: RRMSE = ', ...
1768         num2str(round(RRMSE(6), 2)), '%', ...
1769         ', MAPE = ', num2str(round(MAPE(6), 2)), '%',newline ...
1770 '\color{rgb}{', sprintf('%f,%f,%f', cols(7,:) ) '} Model 7: RRMSE = ', ...
1771         num2str(round(RRMSE(7), 2)), '%', ...
1772         ', MAPE = ', num2str(round(MAPE(7), 2)), '%',newline ...
1773 '\color{rgb}{', sprintf('%f,%f,%f', cols(8,:) ) '} Model 8: RRMSE = ', ...
1774         num2str(round(RRMSE(8), 2)), '%', ...
1775         ', MAPE = ', num2str(round(MAPE(8), 2)), '%',newline ...
1776 '\color{rgb}{', sprintf('%f,%f,%f', cols(9,:) ) '} Model 9: RRMSE = ', ...
1777         num2str(round(RRMSE(9), 2)), '%', ...
1778         ', MAPE = ', num2str(round(MAPE(9), 2)), '%',newline ...
1779 '\color{rgb}{', sprintf('%f,%f,%f', cols(10,:) ) '} Model 10: RRMSE = ', ...
1780         num2str(round(RRMSE(10), 2)), '%', ...
1781         ', MAPE = ', num2str(round(MAPE(10), 2))], 'Interpreter', 'tex','FontSize', 30);
1782
1783 gtext(['Mean RRMSE = ', num2str(round(mean(RRMSE), 2)), '%',', mean MAPE = ', ...
1784         num2str(round(mean(MAPE), 2)), '%'],'FontSize', 30)
1785
1786 legend([data, obs], 'Training data', 'Test data', 'NorthWest','FontSize', 30)
1787 ylabel('New cases')
1788 xlabel('Date')
1789 if is_global
1790     title(['Forecast with 10 CNN-LSTM models on Global data from ', ...
1791           datestr(index_to_date_global(endpoint+1)),...
1792           ' with sample size ', num2str(len_train), ' and necessary observed ...
1793           data'],'FontSize', 60)
1794 else
1795     title(['Forecast with 10 CNN-LSTM models on Norwegian data from ', ...
1796           datestr(index_to_date_norway(endpoint+1)),...
1797           ' with sample size ', num2str(len_train), ' and necessary observed ...
1798           data'],'FontSize', 60)
1799 end
1800 ax = gca;
1801 ax.YAxis.Exponent = 3;
1802 set(gca,'FontSize',20)
1803 hold off
1804
1805 subplot(1,2,2)
1806 hold on
1807 data = plot(dates_train, ts(endpoint-plot_length+1:endpoint), 'Color',[0.25, ...
1808         0.25, 0.25]);
1809 obs = plot(dates_test,test_data, 'Color', [1, 0, 0], 'LineWidth',3);

```

```

1794 predictions_sarima = plot(dates_test,preds_sarima, 'Color', [0.4940 0.1840 ...
    0.5560], 'LineWidth',2);
1795 u_sarima = plot(dates_test,upper_sarima, '--','Color', [0.4940 0.1840 ...
    0.5560], 'LineWidth',1);
1796 l_sarima = plot(dates_test,lower_sarima, '--','Color', [0.4940 0.1840 ...
    0.5560], 'LineWidth',1);
1797 predictions_gandalf = plot(dates_test,preds_gandalf, 'Color', [.2, .9, .5], ...
    'LineWidth',2);
1798 u_gandalf = plot(dates_test,upper_gandalf, '--', 'Color', [.2, .9, .5], ...
    'LineWidth',1);
1799 l_gandalf = plot(dates_test,lower_gandalf, '--', 'Color', [.2, .9, .5], ...
    'LineWidth',1);
1800 predictions_cnn = plot(dates_test,mean_preds_cnn, 'Color', [0 0.4470 0.7410], ...
    'LineWidth',2);
1801 %u_cnn = plot(dates_test,upper_cnn, '--', 'Color', [0 0.4470 0.7410], ...
    'LineWidth',1);
1802 %l_cnn = plot(dates_test,lower_cnn, '--', 'Color', [0 0.4470 0.7410], ...
    'LineWidth',1);
1803
1804 ylim([0 800000])
1805
1806 % insert percision results in plot
1807 gtext([ ...
1808     '\color{rgb}' sprintf('%f,%f,%f', [0.4940 0.1840 0.5560] ) ') SARIMA ...
        model: RRMSE = ', num2str(round(RRMSE_sarima, 2)), '%', ...
1809     ', MAPE = ', num2str(round(MAPE_sarima, 2)), '%',newline, ...
1810     '\color{rgb}' sprintf('%f,%f,%f', [.2, .9, .5]) ') Gandalf model: RRMSE ...
        = ', num2str(round(RRMSE_gandalf, 2)), '%', ...
1811     ', MAPE = ', num2str(round(MAPE_gandalf, 2)), '%', newline, ...
1812     '\color{rgb}' sprintf('%f,%f,%f', [0 0.4470 0.7410]) ') CNN-LSTM model: ...
        RRMSE = ', num2str(round(RRMSE_cnn, 2)), '%', ...
1813     ', MAPE = ', num2str(round(MAPE_cnn, 2)), '%',
1814     ], 'Interpreter', 'tex','FontSize', 30);
1815
1816 legend([data, obs, predictions_sarima, u_sarima, predictions_gandalf, ...
    u_gandalf, predictions_cnn, u_cnn],...
1817     'Training data', 'Test data', 'Forecast SARIMA','95% interval SARIMA', ...
1818     'Forecast Gandalf','95% interval Gandalf', 'Forecast CNN-LSTM', ...
        'NorthWest','FontSize', 30) % '95% interval CNN-LSTM',
1819 ylabel('New cases','FontSize', 30)
1820 xlabel('Date','FontSize', 30)
1821 if is_global
1822     title(['Forecasts of Global data from ', ...
        datestr(index_to_date_global(endpoint+1)), ...
1823         ' using only one-step predictions with observed values'],'FontSize', 30)
1824 else
1825     title(['Forecasts of Norwegian data from ', ...
        datestr(index_to_date_norway(endpoint+1)), ...
1826         ' using only one-step predictions with observed values'],'FontSize', 30)
1827 end
1828 set(gcf,'color','w')
1829 set(gca,'FontSize',20)
1830 ax = gca;
1831 ax.YAxis.Exponent = 3;
1832 if is_global
1833     ax.YAxis.Exponent = 3;
1834 else
1835     ax.YAxis.Exponent = 0;
1836 end
1837 hold off
1838 end
1839
1840 function plot_all_predictions_without_observed(ts, preds_cnn, endpoint, ...
    len_train, days_ahead, is_global)
1841 % Plots the two previous functions in the same figure
1842 % ts is on log-scale
1843 plot_length = 14;
1844 sarima = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
1845 gandalf = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
1846 garchmod = garch('Constant',0.001,'GARCHLags',1, 'ARCHLags',1);
1847 gandalf.Variance = garchmod;
1848
1849 %[sarima_1, ~, logL_without] = estimate(mod, ts_train, 'Display','off');

```



```

1850
1851 preds_sarima = [];
1852 mse_sarima = [];
1853 preds_gandalf = [];
1854 mse_gandalf = [];
1855
1856 train_data = ts(endpoint-len_train+1:endpoint);
1857
1858 tic
1859 fitted_sarima = estimate(sarima, train_data, 'Display','off');
1860 fitted_gandalf= estimate(gandalf, train_data, 'Display','off');
1861 [preds_sarima, mse_sarima] = forecast(fitted_sarima, days_ahead, train_data);
1862 [preds_gandalf, mse_gandalf] = forecast(fitted_gandalf, days_ahead, train_data);
1863 toc
1864
1865 mean_preds_cnn = mean(preds_cnn, 1);
1866 %mse_cnn = mse_cnn_lstm(ts, mean_preds_cnn, endpoint, len_train, false);
1867
1868 % create 95% intervals
1869 upper_sarima = preds_sarima + 1.96*sqrt(mse_sarima);
1870 lower_sarima = preds_sarima - 1.96*sqrt(mse_sarima);
1871 upper_gandalf = preds_gandalf + 1.96*sqrt(mse_gandalf);
1872 lower_gandalf = preds_gandalf - 1.96*sqrt(mse_gandalf);
1873 %upper_cnn = mean_preds_cnn + 1.96*round(sqrt(mse_cnn));
1874 %lower_cnn = mean_preds_cnn - 1.96*round(sqrt(mse_cnn));
1875
1876 % convert to normal scale where needed
1877 ts = floor(exp(ts));
1878 test_data = ts(endpoint+1: endpoint + days_ahead);
1879 preds_sarima = floor(exp(preds_sarima));preds_gandalf = ...
    floor(exp(preds_gandalf));
1880 upper_sarima = floor(exp(upper_sarima));lower_sarima = floor(exp(lower_sarima));
1881 upper_gandalf = floor(exp(upper_gandalf));lower_gandalf = ...
    floor(exp(lower_gandalf));
1882
1883 % calculate RRMSE and MAPE for single predictions
1884 RRMSE_sarima = sqrt(mean((test_data - preds_sarima).^2)/mean(test_data)*100;
1885 MAPE_sarima = mean(abs(test_data - preds_sarima)./abs(test_data))*100;
1886 RRMSE_gandalf = sqrt(mean((test_data - preds_gandalf).^2)/mean(test_data)*100;
1887 MAPE_gandalf = mean(abs(test_data - preds_gandalf)./abs(test_data))*100;
1888 RRMSE_cnn = sqrt(mean((test_data - mean_preds_cnn).^2)/mean(test_data)*100;
1889 MAPE_cnn = mean(abs(test_data - mean_preds_cnn)./abs(test_data))*100;
1890
1891 % calculate RRMSE and MAPE for the 10 CNN-LSTM models
1892 RRMSE = zeros(10, 1);
1893 MAPE = zeros(10, 1);
1894 for i = 1:10
1895     RRMSE(i) = sqrt(mean((test_data.' - preds_cnn(i, ...
1896         :).^2)/mean(test_data)*100;
1897     MAPE(i) = mean(abs(test_data.' - preds_cnn(i, :))./abs(test_data.))*100;
1898 end
1899
1900 % set dates
1901 if is_global
1902     dates_train = index_to_date_global(endpoint-plot_length+1:endpoint);
1903     dates_test = index_to_date_global(endpoint+1:endpoint+days_ahead);
1904 else
1905     dates_train = index_to_date_norway(endpoint-plot_length+1:endpoint);
1906     dates_test = index_to_date_norway(endpoint+1:endpoint+days_ahead);
1907 end
1908
1909 % generate RGB colors for plot for the 10 CNN-LSTM models
1910 rng(1234)
1911 cols = [];
1912 for i = 1:10
1913     c = [rand, rand, rand];
1914     cols = [cols; c];
1915 end
1916
1917 % plotting
1918 figure
1919 set(gca, 'FontSize', 24)

```

```

1920 subplot(1,2,1)
1921 hold on
1922 data = plot(dates_train, ts(endpoint-plot_length+1:endpoint), 'Color', [0.25, ...
    0.25, 0.25]);
1923 obs = plot(dates_test, test_data, 'Color', [1 0 0], 'LineWidth', 3);
1924 pred1 = plot(dates_test, preds_cnn(1, :), 'Color', cols(1, :), 'LineWidth', 1);
1925 pred2 = plot(dates_test, preds_cnn(2, :), 'Color', cols(2, :), 'LineWidth', 1);
1926 pred3 = plot(dates_test, preds_cnn(3, :), 'Color', cols(3, :), 'LineWidth', 1);
1927 pred4 = plot(dates_test, preds_cnn(4, :), 'Color', cols(4, :), 'LineWidth', 1);
1928 pred5 = plot(dates_test, preds_cnn(5, :), 'Color', cols(5, :), 'LineWidth', 1);
1929 pred6 = plot(dates_test, preds_cnn(6, :), 'Color', cols(6, :), 'LineWidth', 1);
1930 pred7 = plot(dates_test, preds_cnn(7, :), 'Color', cols(7, :), 'LineWidth', 1);
1931 pred8 = plot(dates_test, preds_cnn(8, :), 'Color', cols(8, :), 'LineWidth', 1);
1932 pred9 = plot(dates_test, preds_cnn(9, :), 'Color', cols(9, :), 'LineWidth', 1);
1933 pred10 = plot(dates_test, preds_cnn(10, :), 'Color', cols(10, :), 'LineWidth', 1);
1934
1935 gtext([ ...
1936 '\color{rgb}{ ' sprintf('%f,%f,%f', cols(1,:) ) ' } Model 1: RRMSE = ', ...
    num2str(round(RRMSE(1), 2)), '%', ...
1937 ', MAPE = ', num2str(round(MAPE(1), 2)), '%', newline, ...
1938 '\color{rgb}{ ' sprintf('%f,%f,%f', cols(2,:) ) ' } Model 2: RRMSE = ', ...
    num2str(round(RRMSE(2), 2)), '%', ...
1939 ', MAPE = ', num2str(round(MAPE(2), 2)), '%', newline, ...
1940 '\color{rgb}{ ' sprintf('%f,%f,%f', cols(3,:) ) ' } Model 3: RRMSE = ', ...
    num2str(round(RRMSE(3), 2)), '%', ...
1941 ', MAPE = ', num2str(round(MAPE(3), 2)), '%', newline ...
1942 '\color{rgb}{ ' sprintf('%f,%f,%f', cols(4,:) ) ' } Model 4: RRMSE = ', ...
    num2str(round(RRMSE(4), 2)), '%', ...
1943 ', MAPE = ', num2str(round(MAPE(4), 2)), '%', newline ...
1944 '\color{rgb}{ ' sprintf('%f,%f,%f', cols(5,:) ) ' } Model 5: RRMSE = ', ...
    num2str(round(RRMSE(5), 2)), '%', ...
1945 ', MAPE = ', num2str(round(MAPE(5), 2)), '%', newline ...
1946 '\color{rgb}{ ' sprintf('%f,%f,%f', cols(6,:) ) ' } Model 6: RRMSE = ', ...
    num2str(round(RRMSE(6), 2)), '%', ...
1947 ', MAPE = ', num2str(round(MAPE(6), 2)), '%', newline ...
1948 '\color{rgb}{ ' sprintf('%f,%f,%f', cols(7,:) ) ' } Model 7: RRMSE = ', ...
    num2str(round(RRMSE(7), 2)), '%', ...
1949 ', MAPE = ', num2str(round(MAPE(7), 2)), '%', newline ...
1950 '\color{rgb}{ ' sprintf('%f,%f,%f', cols(8,:) ) ' } Model 8: RRMSE = ', ...
    num2str(round(RRMSE(8), 2)), '%', ...
1951 ', MAPE = ', num2str(round(MAPE(8), 2)), '%', newline ...
1952 '\color{rgb}{ ' sprintf('%f,%f,%f', cols(9,:) ) ' } Model 9: RRMSE = ', ...
    num2str(round(RRMSE(9), 2)), '%', ...
1953 ', MAPE = ', num2str(round(MAPE(9), 2)), '%', newline ...
1954 '\color{rgb}{ ' sprintf('%f,%f,%f', cols(10,:) ) ' } Model 10: RRMSE = ', ...
    num2str(round(RRMSE(10), 2)), '%', ...
1955 ', MAPE = ', num2str(round(MAPE(10), 2))], 'Interpreter', 'tex', 'FontSize', 30);
1956
1957 gtext(['Mean RRMSE = ', num2str(round(mean(RRMSE), 2)), '%', ' mean MAPE = ', ...
    num2str(round(mean(MAPE), 2)), '%'], 'FontSize', 30)
1958
1959 legend([data, obs], 'Training data', 'Test data', 'NorthWest', 'FontSize', 30)
1960 ylabel('New cases', 'FontSize', 30)
1961 xlabel('Date', 'FontSize', 30)
1962 if is_global
1963 title(['Forecast with 10 CNN-LSTM models on Global data from ', ...
    datestr(index_to_date_global(endpoint+1)), ...
1964 ' with sample size ', num2str(len_train)], 'FontSize', 30)
1965 else
1966 title(['Forecast with 10 CNN-LSTM models on Norwegian data from ', ...
    datestr(index_to_date_norway(endpoint+1)), ...
1967 ' with sample size ', num2str(len_train)], 'FontSize', 30)
1968 end
1969 ax = gca;
1970 ax.YAxis.Exponent = 3;
1971 set(gca, 'FontSize', 20)
1972 hold off
1973
1974 subplot(1,2,2)
1975 hold on
1976 data = plot(dates_train, ts(endpoint-plot_length+1:endpoint), 'Color', [0.25, ...
    0.25, 0.25]);
1977 obs = plot(dates_test, test_data, 'Color', [1, 0, 0], 'LineWidth', 3);

```

```

1978 predictions_sarima = plot(dates_test,preds_sarima, 'Color', [0.4940 0.1840 ...
1979 0.5560], 'LineWidth',2);
1980 u_sarima = plot(dates_test,upper_sarima, '--','Color', [0.4940 0.1840 ...
1981 0.5560], 'LineWidth',1);
1982 l_sarima = plot(dates_test,lower_sarima, '--','Color', [0.4940 0.1840 ...
1983 0.5560], 'LineWidth',1);
1984 predictions_gandalf = plot(dates_test,preds_gandalf, 'Color', [.2, .9, .5], ...
1985 'LineWidth',2);
1986 u_gandalf = plot(dates_test,upper_gandalf, '--', 'Color', [.2, .9, .5], ...
1987 'LineWidth',1);
1988 l_gandalf = plot(dates_test,lower_gandalf, '--', 'Color', [.2, .9, .5], ...
1989 'LineWidth',1);
1990 predictions_cnn = plot(dates_test,mean_preds_cnn, 'Color', [0 0.4470 0.7410], ...
1991 'LineWidth',2);
1992 %u_cnn = plot(dates_test,upper_cnn, '--', 'Color', [0 0.4470 0.7410], ...
1993 'LineWidth',1);
1994 %l_cnn = plot(dates_test,lower_cnn, '--', 'Color', [0 0.4470 0.7410], ...
1995 'LineWidth',1);
1996
1997 ylim([0 800000])
1998
1999 % insert percision results in plot
2000 gtext([ ...
2001 '\color{rgb}' sprintf('%f,%f,%f', [0.4940 0.1840 0.5560] ) ') SARIMA ...
2002 model: RRMSE = ', num2str(round(RRMSE_sarima, 2)), '%', ...
2003 ', MAPE = ', num2str(round(MAPE_sarima, 2)), '%',newline, ...
2004 '\color{rgb}' sprintf('%f,%f,%f', [.2, .9, .5]) ') Gandalf model: RRMSE ...
2005 = ', num2str(round(RRMSE_gandalf, 2)), '%', ...
2006 ', MAPE = ', num2str(round(MAPE_gandalf, 2)), '%', newline, ...
2007 '\color{rgb}' sprintf('%f,%f,%f', [0 0.4470 0.7410]) ') CNN-LSTM model: ...
2008 RRMSE = ', num2str(round(RRMSE_cnn, 2)), '%', ...
2009 ', MAPE = ', num2str(round(MAPE_cnn, 2)), '%',
2010 ], 'Interpreter', 'tex','FontSize', 30);
2011
2012 legend([data, obs, predictions_sarima, u_sarima, predictions_gandalf, ...
2013 u_gandalf, predictions_cnn],...
2014 'Training data', 'Test data', 'Forecast SARIMA','95% interval SARIMA', ...
2015 'Forecast Gandalf','95% interval Gandalf', 'Forecast CNN-LSTM', ...
2016 'NorthWest','FontSize', 30)
2017 ylabel('New cases','FontSize', 30)
2018 xlabel('Date','FontSize', 30)
2019 if is_global
2020 title(['Forecasts of Global data from ', ...
2021 datestr(index_to_date_global(endpoint+1),' with sample size ', ...
2022 num2str(len_train)],'FontSize', 30)
2023 else
2024 title(['Forecasts of Norwegian data from ', ...
2025 datestr(index_to_date_norway(endpoint+1)), ' with sample size ', ...
2026 num2str(len_train)],'FontSize', 30)
2027 end
2028 set(gcf,'color','w')
2029 ax = gca;
2030 if is_global
2031 ax.YAxis.Exponent = 3;
2032 else
2033 ax.YAxis.Exponent = 0;
2034 end
2035 hold off
2036 end
2037
2038 function plot_10_norway_CNN_LSTM_Preds(ts, preds, endpoint, len_train, ...
2039 days_ahead, use_observed)
2040 % Plot the results from all 10 CNN-LSTM models
2041 % preds = 10 days_ahead forecasts
2042 % use_observed should be true if observations are used to predict the
2043 % ensuing day
2044 plot_length = 14;
2045 test_data = ts(endpoint+1: endpoint + days_ahead);
2046 % calculate RRMSE, MAE and OSRE
2047 RRMSE = zeros(10, 1);
2048 MAPE = zeros(10, 1);
2049 for i = 1:10
2050 RRMSE(i) = sqrt(mean((test_data.' - preds(i, :)).^2))/mean(test_data)*100;

```

```

2032     MAPE(i) = mean(abs(test_data.' - preds(i, :))./abs(test_data.'))*100;
2033 end
2034
2035 dates_train = index_to_date_norway(endpoint-plot_length+1:endpoint);
2036 dates_test = index_to_date_norway(endpoint+1:endpoint+days_ahead);
2037
2038
2039 % generate RGB colors for plot
2040 rng(1234)
2041 cols = [];
2042 for i = 1:10
2043     c = [rand, rand, rand];
2044     cols = [cols; c];
2045 end
2046 figure
2047 hold on
2048
2049 data = plot(dates_train, ts(endpoint-plot_length+1:endpoint),'Color',[0.25, ...
2050     0.25, 0.25]);
2051 obs = plot(dates_test, test_data, 'Color', [1 0 0],'LineWidth',3);
2052 pred1 = plot(dates_test,preds(1, :), 'Color', cols(1, :), 'LineWidth',1);
2053 pred2 = plot(dates_test,preds(2, :), 'Color', cols(2, :), 'LineWidth',1);
2054 pred3 = plot(dates_test,preds(3, :), 'Color', cols(3, :), 'LineWidth',1);
2055 pred4 = plot(dates_test,preds(4, :), 'Color', cols(4, :), 'LineWidth',1);
2056 pred5 = plot(dates_test,preds(5, :), 'Color', cols(5, :), 'LineWidth',1);
2057 pred6 = plot(dates_test,preds(6, :), 'Color', cols(6, :), 'LineWidth',1);
2058 pred7 = plot(dates_test,preds(7, :), 'Color', cols(7, :), 'LineWidth',1);
2059 pred8 = plot(dates_test,preds(8, :), 'Color', cols(8, :), 'LineWidth',1);
2060 pred9 = plot(dates_test,preds(9, :), 'Color', cols(9, :), 'LineWidth',1);
2061 pred10 = plot(dates_test,preds(10, :), 'Color', cols(10, :), 'LineWidth',1);
2062
2063 set(gca,'FontSize',24)
2064 gtext([ ...
2065     '\color{rgb}{%f,%f,%f', cols(1,:) ) '} Model 1: RRMSE = ', ...
2066     num2str(round(RRMSE(1), 2)), '%', ...
2067     ', MAPE = ', num2str(round(MAPE(1), 2)), '%',newline, ...
2068     '\color{rgb}{%f,%f,%f', cols(2,:) ) '} Model 2: RRMSE = ', ...
2069     num2str(round(RRMSE(2), 2)), '%', ...
2070     ', MAPE = ', num2str(round(MAPE(2), 2)), '%',newline, ...
2071     '\color{rgb}{%f,%f,%f', cols(3,:) ) '} Model 3: RRMSE = ', ...
2072     num2str(round(RRMSE(3), 2)), '%', ...
2073     ', MAPE = ', num2str(round(MAPE(3), 2)), '%',newline, ...
2074     '\color{rgb}{%f,%f,%f', cols(4,:) ) '} Model 4: RRMSE = ', ...
2075     num2str(round(RRMSE(4), 2)), '%', ...
2076     ', MAPE = ', num2str(round(MAPE(4), 2)), '%',newline, ...
2077     '\color{rgb}{%f,%f,%f', cols(5,:) ) '} Model 5: RRMSE = ', ...
2078     num2str(round(RRMSE(5), 2)), '%', ...
2079     ', MAPE = ', num2str(round(MAPE(5), 2)), '%',newline, ...
2080     '\color{rgb}{%f,%f,%f', cols(6,:) ) '} Model 6: RRMSE = ', ...
2081     num2str(round(RRMSE(6), 2)), '%', ...
2082     ', MAPE = ', num2str(round(MAPE(6), 2)), '%',newline, ...
2083     '\color{rgb}{%f,%f,%f', cols(7,:) ) '} Model 7: RRMSE = ', ...
2084     num2str(round(RRMSE(7), 2)), '%', ...
2085     ', MAPE = ', num2str(round(MAPE(7), 2)), '%',newline, ...
2086     '\color{rgb}{%f,%f,%f', cols(8,:) ) '} Model 8: RRMSE = ', ...
2087     num2str(round(RRMSE(8), 2)), '%', ...
2088     ', MAPE = ', num2str(round(MAPE(8), 2)), '%',newline, ...
2089     '\color{rgb}{%f,%f,%f', cols(9,:) ) '} Model 9: RRMSE = ', ...
2090     num2str(round(RRMSE(9), 2)), '%', ...
2091     ', MAPE = ', num2str(round(MAPE(9), 2)), '%',newline, ...
2092     '\color{rgb}{%f,%f,%f', cols(10,:) ) '} Model 10: RRMSE = ', ...
2093     num2str(round(RRMSE(10), 2)), '%', ...
2094     ', MAPE = ', num2str(round(MAPE(10), 2))], 'Interpreter','tex','FontSize', 30);
2095
2096 gtext(['Mean RRMSE = ', num2str(round(mean(RRMSE), 2)), '%',', mean MAPE = ', ...
2097     num2str(round(mean(MAPE), 2)), '%'],'FontSize', 30)
2098
2099 legend([data, obs], 'Training data', 'Test data', 'NorthWest','FontSize', 30)
2100
2101 ylabel('New cases','FontSize', 30)
2102 xlabel('Date','FontSize', 30)
2103 if use_observed
2104     title(['Predictions\Forecast with 10 CNN-LSTM models on Norwegian data ...

```

```

2093         from ', datestr(index_to_date_norway(endpoint+1)),...
2094         ' using the ', num2str(len_train), ' previous days and ...
2095         observations'],'FontSize', 30)
2096     else
2097         title(['Predictions\Forecast with 10 CNN-LSTM models on Norwegian data ...
2098         from ', datestr(index_to_date_norway(endpoint+1)),...
2099         ' using the ', num2str(len_train), ' previous days'],'FontSize', 30)
2100     end
2101     set(gcf,'color','w')
2102     set(gca,'FontSize',20)
2103     ax = gca;
2104     ax.YAxis.Exponent = 0;
2105     hold off
2106 end
2107
2108 function plot_10_global_CNN_LSTM_Preds(ts, preds, endpoint, len_train, ...
2109     days_ahead, use_observed)
2110     % Plot the results from all 10 CNN-LSTM models
2111     % preds = 10 days_ahead forecasts
2112     % use_observed should be true if observations are used to predict the
2113     % ensuing day
2114     plot_length = 30;
2115     test_data = ts(endpoint+1: endpoint + days_ahead);
2116     % calculate RRMSE, MAE and OSRE
2117     RRMSE = zeros(10, 1);
2118     MAPE = zeros(10, 1);
2119     for i = 1:10
2120         RRMSE(i) = sqrt(mean((test_data.' - preds(i, :)).^2))/mean(test_data)*100;
2121         MAPE(i) = mean(abs(test_data.' - preds(i, :))./abs(test_data.'))*100;
2122     end
2123
2124     dates_train = index_to_date_global(endpoint-plot_length+1:endpoint);
2125     dates_test = index_to_date_global(endpoint+1:endpoint+days_ahead);
2126
2127     % generate RGB colors for plot
2128     rng(1234)
2129     cols = [];
2130     for i = 1:10
2131         c = [rand, rand, rand];
2132         cols = [cols; c];
2133     end
2134     figure
2135     hold on
2136
2137     data = plot(dates_train, ts(endpoint-plot_length+1:endpoint),'Color',[0.25, ...
2138         0.25, 0.25]);
2139
2140     obs = plot(dates_test, test_data, 'Color', [1 0 0],'LineWidth',3);
2141     pred1 = plot(dates_test,preds(1, :), 'Color', cols(1, :), 'LineWidth',1);
2142     pred2 = plot(dates_test,preds(2, :), 'Color', cols(2, :), 'LineWidth',1);
2143     pred3 = plot(dates_test,preds(3, :), 'Color', cols(3, :), 'LineWidth',1);
2144     pred4 = plot(dates_test,preds(4, :), 'Color', cols(4, :), 'LineWidth',1);
2145     pred5 = plot(dates_test,preds(5, :), 'Color', cols(5, :), 'LineWidth',1);
2146     pred6 = plot(dates_test,preds(6, :), 'Color', cols(6, :), 'LineWidth',1);
2147     pred7 = plot(dates_test,preds(7, :), 'Color', cols(7, :), 'LineWidth',1);
2148     pred8 = plot(dates_test,preds(8, :), 'Color', cols(8, :), 'LineWidth',1);
2149     pred9 = plot(dates_test,preds(9, :), 'Color', cols(9, :), 'LineWidth',1);
2150     pred10 = plot(dates_test,preds(10, :), 'Color', cols(10, :), 'LineWidth',1);
2151
2152     set(gca,'FontSize',24)
2153     gtext([ ...
2154         '\color{rgb}{', sprintf('%f,%f,%f', cols(1,:) ) '} Model 1: RRMSE = ', ...
2155         num2str(round(RRMSE(1), 2)), '%', ...
2156         ', MAPE = ', num2str(round(MAPE(1), 2)), '%',newline, ...
2157         '\color{rgb}{', sprintf('%f,%f,%f', cols(2,:) ) '} Model 2: RRMSE = ', ...
2158         num2str(round(RRMSE(2), 2)), '%', ...
2159         ', MAPE = ', num2str(round(MAPE(2), 2)), '%',newline, ...
2160         '\color{rgb}{', sprintf('%f,%f,%f', cols(3,:) ) '} Model 3: RRMSE = ', ...
2161         num2str(round(RRMSE(3), 2)), '%', ...
2162         ', MAPE = ', num2str(round(MAPE(3), 2)), '%',newline, ...
2163         '\color{rgb}{', sprintf('%f,%f,%f', cols(4,:) ) '} Model 4: RRMSE = ', ...
2164         num2str(round(RRMSE(4), 2)), '%', ...
2165         ', MAPE = ', num2str(round(MAPE(4), 2)), '%',newline, ...

```

```

2157     '\color{rgb}{%f,%f,%f', cols(5,:) ) '} Model 5: RRMSE = ', ...
        num2str(round(RRMSE(5), 2)), '%', ...
2158     ', MAPE = ', num2str(round(MAPE(5), 2)), '%',newline ...
2159     '\color{rgb}{%f,%f,%f', cols(6,:) ) '} Model 6: RRMSE = ', ...
        num2str(round(RRMSE(6), 2)), '%', ...
2160     ', MAPE = ', num2str(round(MAPE(6), 2)), '%',newline ...
2161     '\color{rgb}{%f,%f,%f', cols(7,:) ) '} Model 7: RRMSE = ', ...
        num2str(round(RRMSE(7), 2)), '%', ...
2162     ', MAPE = ', num2str(round(MAPE(7), 2)), '%',newline ...
2163     '\color{rgb}{%f,%f,%f', cols(8,:) ) '} Model 8: RRMSE = ', ...
        num2str(round(RRMSE(8), 2)), '%', ...
2164     ', MAPE = ', num2str(round(MAPE(8), 2)), '%',newline ...
2165     '\color{rgb}{%f,%f,%f', cols(9,:) ) '} Model 9: RRMSE = ', ...
        num2str(round(RRMSE(9), 2)), '%', ...
2166     ', MAPE = ', num2str(round(MAPE(9), 2)), '%',newline ...
2167     '\color{rgb}{%f,%f,%f', cols(10,:) ) '} Model 10: RRMSE = ', ...
        num2str(round(RRMSE(10), 2)), '%', ...
2168     ', MAPE = ', num2str(round(MAPE(10), 2))], 'Interpreter', 'tex','FontSize', 30);
2169
2170 gtext(['Mean RRMSE = ', num2str(round(mean(RRMSE), 2)), '%',', mean MAPE = ', ...
        num2str(round(mean(MAPE), 2)), '%'],'FontSize', 30)
2171
2172 legend([data, obs], 'Training data', 'Test data', 'NorthWest','FontSize', 30)
2173
2174 ylabel('New cases','FontSize', 30)
2175 xlabel('Date','FontSize', 30)
2176 if use_observed
2177     title(['Predictions\Forecast with 10 CNN-LSTM models on Global data from ...
2178           ', datestr(index_to_date_global(endpoint+1)),...
           ' using the ', num2str(len_train), ' previous days and ...
           observations'],'FontSize', 30)
2179 else
2180     title(['Predictions\Forecast with 10 CNN-LSTM models on Global data from ...
2181           ', datestr(index_to_date_global(endpoint+1)),...
           ' using the ', num2str(len_train), ' previous days'],'FontSize', 30)
2182 end
2183 set(gcf,'color','w')
2184 set(gca,'FontSize',20)
2185 ax = gca;
2186 ax.YAxis.Exponent = 3;
2187 hold off
2188 end
2189
2190 function plot_preds_without_observed(ts, preds_cnn, endpoint, len_train, ...
        days_ahead, is_global, plot_length)
2191     % Directly predicts the next days_ahead days after endpoint
2192     % These are then plotted and compared to the CNN-LSTM model with the
2193     % same prediction sceeme
2194
2195     % ts should be on log-scale
2196     % preds_cnn should not use observed values after training to forecast
2197
2198
2199     sarima = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
2200     gandalf = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
2201     garchmod = garch('Constant',0.001,'GARCHLags',1, 'ARCHLags',1);
2202     gandalf.Variance = garchmod;
2203
2204     %[sarima_1, ~, logL_without] = estimate(mod, ts_train, 'Display','off');
2205
2206     preds_sarima = [];
2207     mse_sarima = [];
2208     preds_gandalf = [];
2209     mse_gandalf = [];
2210
2211     train_data = ts(endpoint-len_train+1:endpoint);
2212
2213     tic
2214     fitted_sarima = estimate(sarima, train_data, 'Display','off');
2215     fitted_gandalf= estimate(gandalf, train_data, 'Display','off');
2216     [preds_sarima, mse_sarima] = forecast(fitted_sarima, days_ahead, train_data);
2217     [preds_gandalf, mse_gandalf] = forecast(fitted_gandalf, days_ahead, train_data);
2218     toc

```

```

2219
2220 %mse_cnn = mse_cnn_lstm(ts, preds_cnn, endpoint, len_train, false);
2221
2222 % create 95% intervals
2223 upper_sarima = preds_sarima + 1.96*sqrt(mse_sarima);
2224 lower_sarima = preds_sarima - 1.96*sqrt(mse_sarima);
2225 upper_gandalf = preds_gandalf + 1.96*sqrt(mse_gandalf);
2226 lower_gandalf = preds_gandalf - 1.96*sqrt(mse_gandalf);
2227 %upper_cnn = preds_cnn + 1.96*round(sqrt(mse_cnn));
2228 %lower_cnn = preds_cnn - 1.96*round(sqrt(mse_cnn));
2229
2230 % convert to normal scale where needed
2231 ts = floor(exp(ts));
2232 test_data = ts(endpoint+1: endpoint + days_ahead);
2233 preds_sarima = floor(exp(preds_sarima));preds_gandalf = ...
    floor(exp(preds_gandalf));
2234 upper_sarima = floor(exp(upper_sarima));lower_sarima = floor(exp(lower_sarima));
2235 upper_gandalf = floor(exp(upper_gandalf));lower_gandalf = ...
    floor(exp(lower_gandalf));
2236
2237 % calculate RRMSE and MAPE
2238 RRMSE_sarima = sqrt(mean((test_data - preds_sarima).^2))/mean(test_data)*100;
2239 MAPE_sarima = mean(abs(test_data - preds_sarima)./abs(test_data))*100;
2240 RRMSE_gandalf = sqrt(mean((test_data - preds_gandalf).^2))/mean(test_data)*100;
2241 MAPE_gandalf = mean(abs(test_data - preds_gandalf)./abs(test_data))*100;
2242 RRMSE_cnn = sqrt(mean((test_data - preds_cnn).^2))/mean(test_data)*100;
2243 MAPE_cnn = mean(abs(test_data - preds_cnn)./abs(test_data))*100;
2244
2245 % set dates
2246 if is_global
2247     dates_train = index_to_date_global(endpoint-plot_length+1:endpoint);
2248     dates_test = index_to_date_global(endpoint+1:endpoint+days_ahead);
2249 else
2250     dates_train = index_to_date_norway(endpoint-plot_length+1:endpoint);
2251     dates_test = index_to_date_norway(endpoint+1:endpoint+days_ahead);
2252 end
2253
2254 % plotting
2255 figure
2256 set(gca,'FontSize',24)
2257 hold on
2258
2259 data = plot(dates_train, ts(endpoint-plot_length+1:endpoint),'Color',[0.25, ...
    0.25, 0.25]);
2260 obs = plot(dates_test,test_data, 'Color', [1, 0, 0],'LineWidth',3);
2261 predictions_sarima = plot(dates_test,preds_sarima, 'Color', [0.4940 0.1840 ...
    0.5560], 'LineWidth',2);
2262 u_sarima = plot(dates_test,upper_sarima, '--','Color', [0.4940 0.1840 ...
    0.5560], 'LineWidth',1);
2263 l_sarima = plot(dates_test,lower_sarima, '--','Color', [0.4940 0.1840 ...
    0.5560], 'LineWidth',1);
2264 predictions_gandalf = plot(dates_test,preds_gandalf, 'Color', [.2, .9, .5], ...
    'LineWidth',2);
2265 u_gandalf = plot(dates_test,upper_gandalf, '--', 'Color', [.2, .9, .5], ...
    'LineWidth',1);
2266 l_gandalf = plot(dates_test,lower_gandalf, '--', 'Color', [.2, .9, .5], ...
    'LineWidth',1);
2267 predictions_cnn = plot(dates_test,preds_cnn, 'Color', [0 0.4470 0.7410], ...
    'LineWidth',2);
2268 %u_cnn = plot(dates_test,upper_cnn, '--', 'Color', [0 0.4470 0.7410], ...
    'LineWidth',1);
2269 %l_cnn = plot(dates_test,lower_cnn, '--', 'Color', [0 0.4470 0.7410], ...
    'LineWidth',1);
2270
2271 %ylim([0 700000])
2272
2273 % insert percision results in plot
2274 gtext([ ...
2275     '\color{rgb}{%f,%f,%f', [0.4940 0.1840 0.5560] ) ' SARIMA ...
        model: RRMSE = ', num2str(round(RRMSE_sarima, 2)), '%', ...
2276     ', MAPE = ', num2str(round(MAPE_sarima, 2)), '%',newline, ...
2277     '\color{rgb}{%f,%f,%f', [.2, .9, .5]) ' Gandalf model: RRMSE ...
        = ', num2str(round(RRMSE_gandalf, 2)), '%', ...

```

```

2278     ', MAPE = ', num2str(round(MAPE_gandalf, 2)), '%', newline, ...
2279     '\color[rgb]{', sprintf('%f,%f,%f', [0 0.4470 0.7410]) ') CNN-LSTM model: ...
        RRMSE = ', num2str(round(RRMSE_cnn, 2)), '%', ...
2280     ', MAPE = ', num2str(round(MAPE_cnn, 2)), '%',
2281     ], 'Interpreter', 'tex', 'FontSize', 30);
2282
2283 legend([data, obs, predictions_sarima, u_sarima, predictions_gandalf, ...
        u_gandalf, predictions_cnn],...
2284     'Training data', 'Test data', 'Forecast SARIMA', '95% interval SARIMA', ...
2285     'Forecast Gandalf', '95% interval Gandalf', 'Forecast CNN-LSTM', ...
        'NorthWest', 'FontSize', 30) % '95% interval CNN-LSTM',
2286 ylabel('New cases', 'FontSize', 30)
2287 xlabel('Date', 'FontSize', 30)
2288 if is_global
2289     title(['Forecasts of Global data from ', ...
        datestr(index_to_date_global(endpoint+1))], 'FontSize', 30)
2290 else
2291     title(['Forecasts of Norwegian data from ', ...
        datestr(index_to_date_norway(endpoint+1))], 'FontSize', 30)
2292 end
2293 set(gcf, 'color', 'w')
2294 set(gca, 'FontSize', 20)
2295 ax = gca;
2296
2297 if is_global
2298     ax.YAxis.Exponent = 3;
2299 else
2300     ax.YAxis.Exponent = 0;
2301
2302 end
2303 hold off
2304 end
2305
2306 function plot_preds_with_observed(ts, preds_cnn, endpoint, len_train, days_ahead, ...
        is_global, plot_length)
2307     % performs a series of one step predictions for the SARIMA model and
2308     % the Gandalf model, where the next observation is included as train
2309     % data.
2310     % These are then plotted and compared to the CNN-LSTM model with the
2311     % same prediction sceeme
2312
2313     % ts should be on log-scale
2314     % preds_cnn should be a series of one-steps from CNN-LSTM model using
2315     % observations to predict the nex day.
2316
2317
2318     sarima = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
2319     gandalf = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
2320     garchmod = garch('Constant',0.001,'GARCHLags',1, 'ARCHLags',1);
2321     gandalf.Variance = garchmod;
2322
2323     %[sarima_1, ~, logL_without] = estimate(mod, ts_train, 'Display','off');
2324
2325     preds_sarima = [];
2326     mse_sarima = [];
2327     preds_gandalf = [];
2328     mse_gandalf = [];
2329
2330     train_data = ts(endpoint-len_train+1:endpoint);
2331     for i = 1:days_ahead
2332         fitted_sarima = estimate(sarima, train_data, 'Display','off');
2333         fitted_gandalf= estimate(gandalf, train_data, 'Display','off');
2334         [preds_sarima(end+1), mse_sarima(end+1)] = forecast(fitted_sarima, 1, ...
                train_data);
2335         [preds_gandalf(end+1), mse_gandalf(end+1)] = forecast(fitted_gandalf, 1, ...
                train_data);
2336         train_data = [train_data; ts(endpoint+i)]; % add next observation to ...
                training data
2337     end
2338     %mse_cnn = mse_cnn_lstm(ts, preds_cnn, endpoint, len_train, true);
2339
2340     % create 95% intervals
2341     upper_sarima = preds_sarima + 1.96*sqrt(mse_sarima);

```



```

2342 lower_sarima = preds_sarima - 1.96*sqrt(mse_sarima);
2343 upper_gandalf = preds_gandalf + 1.96*sqrt(mse_gandalf);
2344 lower_gandalf = preds_gandalf - 1.96*sqrt(mse_gandalf);
2345 %upper_cnn = preds_cnn + 1.96*round(sqrt(mse_cnn));
2346 %lower_cnn = preds_cnn - 1.96*round(sqrt(mse_cnn));
2347
2348 % convert to normal scale where needed
2349 ts = floor(exp(ts));
2350 test_data = ts(endpoint+1: endpoint + days_ahead);
2351 preds_sarima = floor(exp(preds_sarima));preds_gandalf = ...
    floor(exp(preds_gandalf));
2352 upper_sarima = floor(exp(upper_sarima));lower_sarima = floor(exp(lower_sarima));
2353 upper_gandalf = floor(exp(upper_gandalf));lower_gandalf = ...
    floor(exp(lower_gandalf));
2354
2355 % calculate RRMSE and MAPE
2356 RRMSE_sarima = sqrt(mean((test_data - preds_sarima).^2)/mean(test_data)*100;
2357 MAPE_sarima = mean(abs(test_data - preds_sarima)./abs(test_data))*100;
2358 RRMSE_gandalf = sqrt(mean((test_data - preds_gandalf).^2)/mean(test_data)*100;
2359 MAPE_gandalf = mean(abs(test_data - preds_gandalf)./abs(test_data))*100;
2360 RRMSE_cnn = sqrt(mean((test_data - preds_cnn).^2)/mean(test_data)*100;
2361 MAPE_cnn = mean(abs(test_data - preds_cnn)./abs(test_data))*100;
2362
2363 % set dates
2364 if is_global
2365     dates_train = index_to_date_global(endpoint-plot_length+1:endpoint);
2366     dates_test = index_to_date_global(endpoint+1:endpoint+days_ahead);
2367 else
2368     dates_train = index_to_date_norway(endpoint-plot_length+1:endpoint);
2369     dates_test = index_to_date_norway(endpoint+1:endpoint+days_ahead);
2370 end
2371
2372
2373 % plotting
2374 figure
2375 set(gca,'FontSize',24)
2376 hold on
2377
2378 data = plot(dates_train, ts(endpoint-plot_length+1:endpoint),'Color',[0.25, ...
    0.25, 0.25]);
2379 obs = plot(dates_test,test_data, 'Color', [1, 0, 0],'LineWidth',3);
2380 predictions_sarima = plot(dates_test,preds_sarima, 'Color', [0.4940 0.1840 ...
    0.5560], 'LineWidth',2);
2381 u_sarima = plot(dates_test,upper_sarima, '--','Color', [0.4940 0.1840 ...
    0.5560], 'LineWidth',1);
2382 l_sarima = plot(dates_test,lower_sarima, '--','Color', [0.4940 0.1840 ...
    0.5560], 'LineWidth',1);
2383 predictions_gandalf = plot(dates_test,preds_gandalf, 'Color', [.2, .9, .5], ...
    'LineWidth',2);
2384 u_gandalf = plot(dates_test,upper_gandalf, '--', 'Color', [.2, .9, .5], ...
    'LineWidth',1);
2385 l_gandalf = plot(dates_test,lower_gandalf, '--', 'Color', [.2, .9, .5], ...
    'LineWidth',1);
2386 predictions_cnn = plot(dates_test,preds_cnn, 'Color', [0 0.4470 0.7410], ...
    'LineWidth',2);
2387 %u_cnn = plot(dates_test,upper_cnn, '--', 'Color', [0 0.4470 0.7410], ...
    'LineWidth',1);
2388 %l_cnn = plot(dates_test,lower_cnn, '--', 'Color', [0 0.4470 0.7410], ...
    'LineWidth',1);
2389
2390 gtext([ ...
2391     '\color[rgb]{', sprintf('%f,%f,%f', [0.4940 0.1840 0.5560] ) ,'} SARIMA ...
    model: RRMSE = ', num2str(round(RRMSE_sarima, 2)), '%', ...
2392     ', MAPE = ', num2str(round(MAPE_sarima, 2)), '%',newline, ...
2393     '\color[rgb]{', sprintf('%f,%f,%f', [.2, .9, .5]) ,'} Gandalf model: RRMSE ...
    = ', num2str(round(RRMSE_gandalf, 2)), '%', ...
2394     ', MAPE = ', num2str(round(MAPE_gandalf, 2)), '%', newline, ...
2395     '\color[rgb]{', sprintf('%f,%f,%f', [0 0.4470 0.7410]) ,'} CNN-LSTM model: ...
    RRMSE = ', num2str(round(RRMSE_cnn, 2)), '%', ...
2396     ', MAPE = ', num2str(round(MAPE_cnn, 2)), '%',
2397     ], 'Interpreter', 'tex','FontSize', 30);
2398
2399 legend([data, obs, predictions_sarima, u_sarima, predictions_gandalf, ...

```

```

    u_gandalf, predictions_cnn],...
2400     'Training data', 'Test data', 'Forecast SARIMA','95% interval SARIMA', ...
2401     'Forecast Gandalf','95% interval Gandalf', 'Forecast CNN-LSTM', ...
        'NorthWest','FontSize', 30) % '95% interval CNN-LSTM',
2402 ylabel('New cases','FontSize', 30)
2403 xlabel('Date','FontSize', 30)
2404 if is_global
2405     title(['Forecasts of Global data from ', ...
        datestr(index_to_date_global(endpoint+1)), ...
2406         ' using only observed values to predict each ensuing ...
            day'],'FontSize', 30)
2407 else
2408     title(['Forecasts of Norwegian data from ', ...
        datestr(index_to_date_norway(endpoint+1)), ...
2409         ' using only observed values to predict each ensuing ...
            day'],'FontSize', 30)
2410 end
2411 set(gcf,'color','w')
2412 set(gca,'FontSize',20)
2413 ax = gca;
2414
2415 if is_global
2416     ax.YAxis.Exponent = 3;
2417 else
2418     ax.YAxis.Exponent = 0;
2419 end
2420 hold off
2421 end
2422
2423 function mse = mse_cnn_lstm(ts, preds_cnn, endpoint, len_train, use_observed)
2424 % simulates one step at a time from gandalf model to see how good the
2425 % prediction of the cnn-lstm model is if the underlying time series was
2426 % the gandalf model
2427
2428 % ts should be on log scale
2429 days_ahead = length(preds_cnn);
2430 mod = arima('Constant',0,'D',1,'Seasonality',7,'MALags',1,'SMALags',7);
2431 garchmod = garch('Constant',0.001,'GARCHLags',1,'ARCHLags',1);
2432 mod.Variance = garchmod;
2433 mse = [];
2434 if use_observed
2435     log_preds = log(max(preds_cnn, 0.1)); % transform to log-scale to use ...
        with gandalf model
2436     train_data = ts(endpoint-len_train+1:endpoint); % have to use predictions ...
        as initial observations for simulation??
2437     for i = 1:days_ahead
2438         gandalf_mod = estimate(mod, train_data, 'Display','off');
2439         sim = floor(exp(simulate(gandalf_mod, 1, 'NumPaths', 1000000, 'Y0', ...
            train_data)));
2440         mse(end+1) = mean((sim - preds_cnn(i)).^2);
2441         train_data = [train_data;log_preds(i)]; % add newest prediction to ...
            the reain data
2442     end
2443 else % do everything all at once???
2444     gandalf_mod = estimate(mod, ts(endpoint-len_train+1:endpoint), ...
        'Display','off');
2445     sim = floor(exp(simulate(gandalf_mod, days_ahead, 'NumPaths', 1000000, ...
        'Y0', ts(endpoint-len_train+1:endpoint))));
2446     for i = 1:days_ahead
2447         mse(end+1) = mean((sim(i,:) - preds_cnn(i)).^2);
2448     end
2449 end
2450 end
2451
2452 function plot_norway_and_global_time_series_and_transform(ts, ts_global, ...
    include_trans)
2453 % This function was used to plot both time series and their respective
2454 % log and difference transofrmations.
2455 ts = floor(exp(ts));
2456 ts_global = floor(exp(ts_global));
2457
2458 dates_norway = index_to_date_norway(1:length(ts));
2459 dates_global = index_to_date_global(1:length(ts_global));

```

```

2460
2461     if include_trans
2462         figure
2463         subplot(3,2,1)
2464         plot_norway = plot(dates_norway, ts, 'Col', '#2eff8c');
2465         title('New Cases in Norway')
2466         ylabel('New cases')
2467         xlabel('Date')
2468         ax = gca;
2469         ax.YAxis.Exponent = 0;
2470         set(gca, 'FontSize', 24)
2471         subplot(3,2,2)
2472         plot_global = plot(dates_global, ts_global, 'Col', 'Magenta');
2473         title('Global New Cases')
2474         ylabel('New cases')
2475         xlabel('Date')
2476         set(gcf, 'color', 'w')
2477         ax = gca;
2478         ax.YAxis.Exponent = 3;
2479         set(gca, 'FontSize', 24)
2480
2481         subplot(3,2,3)
2482         plot_norway = plot(dates_norway, log(max(ts, 0.1)), 'Col', '#2eff8c');
2483         title('New Cases in Norway')
2484         ylabel('New cases')
2485         xlabel('Date')
2486         ax = gca;
2487         ax.YAxis.Exponent = 0;
2488         set(gca, 'FontSize', 24)
2489
2490         subplot(3,2,4)
2491         plot_global = plot(dates_global, log(max(ts_global, 0.1)), 'Col', 'Magenta');
2492         title('Global New Cases')
2493         ylabel('New cases')
2494         xlabel('Date')
2495         set(gcf, 'color', 'w')
2496         ax = gca;
2497         ax.YAxis.Exponent = 0;
2498         set(gca, 'FontSize', 24)
2499
2500
2501         % apply differencing on top of log transform:
2502         ts_trans = diff(diff(log(max(ts, 0.1)), 1), 7);
2503         ts_trans_global = diff(diff(log(max(ts_global, 0.1)), 1), 7);
2504         subplot(3,2,5)
2505         plot_norway = plot(dates_norway(9:end), ts_trans, 'Col', '#2eff8c'); % ...
2506             Note that values of the first eight days are removed!
2507         title('New Cases in Norway')
2508         ylabel('New cases')
2509         xlabel('Date')
2510         ax = gca;
2511         ax.YAxis.Exponent = 0;
2512         set(gca, 'FontSize', 24)
2513         subplot(3,2,6)
2514         plot_global = plot(dates_global(9:end), ts_trans_global, 'Col', 'Magenta');
2515         title('Global New Cases')
2516         ylabel('New cases')
2517         xlabel('Date')
2518         ax = gca;
2519         ax.YAxis.Exponent = 0;
2520         set(gcf, 'color', 'w')
2521         set(gca, 'FontSize', 24)
2522
2523     else
2524         figure
2525         subplot(2,1,1)
2526         plot_norway = plot(dates_norway, ts, 'Col', '#2eff8c');
2527         title('New Cases in Norway')
2528         ylabel('New cases')
2529         xlabel('Date')
2530         ax = gca;
2531         ax.YAxis.Exponent = 0;
2532         set(gca, 'FontSize', 24)

```

```
2532     subplot(2,1,2)
2533     plot_global = plot(dates_global, ts_global, 'Col', 'Magenta');
2534     title('Global New Cases')
2535     ylabel('New cases')
2536     xlabel('Date')
2537     set(gcf, 'color', 'w')
2538     ax = gca;
2539     ax.YAxis.Exponent = 3;
2540     set(gca, 'FontSize', 24)
2541
2542     end
2543
2544
2545
2546 end
2547
2548 function date = index_to_date_norway(index)
2549     % get value in date format from index
2550     % index one is February 21th 2020 for Norwegian data.
2551     t0 = datetime(2020,2,21);
2552     date = t0 + days(index-1);
2553 end
2554
2555 function date = index_to_date_global(index)
2556     % get value in date format from index
2557     % index one is January 4th 2020 for Global data.
2558     t0 = datetime(2020,1,4);
2559     date = t0 + days(index-1);
2560 end
```

