

Article

A Statistics and Deep Learning Hybrid Method for Multivariate Time Series Forecasting and Mortality Modeling

Thabang Mathonsi ^{1,*}  and Terence L. van Zyl ^{2,†} 

¹ School of Computer Science and Applied Mathematics, University of the Witwatersrand, Johannesburg 2000, South Africa

² Institute for Intelligent Systems, University of Johannesburg, Johannesburg 2092, South Africa; tvanzyl@uj.ac.za

* Correspondence: thabang.mathonsi@wits.ac.za

† These authors contributed equally to the conceptualisation and presentation of this work.

Abstract: Hybrid methods have been shown to outperform pure statistical and pure deep learning methods at forecasting tasks and quantifying the associated uncertainty with those forecasts (prediction intervals). One example is Exponential Smoothing Recurrent Neural Network (ES-RNN), a hybrid between a statistical forecasting model and a recurrent neural network variant. ES-RNN achieves a 9.4% improvement in absolute error in the Makridakis-4 Forecasting Competition. This improvement and similar outperformance from other hybrid models have primarily been demonstrated only on univariate datasets. Difficulties with applying hybrid forecast methods to multivariate data include (i) the high computational cost involved in hyperparameter tuning for models that are not parsimonious, (ii) challenges associated with auto-correlation inherent in the data, as well as (iii) complex dependency (cross-correlation) between the covariates that may be hard to capture. This paper presents Multivariate Exponential Smoothing Long Short Term Memory (MES-LSTM), a generalized multivariate extension to ES-RNN, that overcomes these challenges. MES-LSTM utilizes a vectorized implementation. We test MES-LSTM on several aggregated coronavirus disease of 2019 (COVID-19) morbidity datasets and find our hybrid approach shows consistent, significant improvement over pure statistical and deep learning methods at forecast accuracy and prediction interval construction.

Keywords: deep learning; multivariate time series forecasting; prediction intervals; mortality modeling



Citation: Mathonsi, T.; van Zyl, T.L. A Statistics and Deep Learning Hybrid Method for Multivariate Time Series Forecasting and Mortality Modeling. *Forecasting* **2022**, *4*, 1–25. <https://doi.org/10.3390/forecast4010001>

Academic Editors: Sonia Leva, Francesca Perla and Salvatore Scognamiglio

Received: 24 October 2021

Accepted: 15 December 2021

Published: 22 December 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Morbidity and mortality modeling is crucial for planning in global economies, national healthcare systems, and other industries such as insurance. Practitioners from statistics, machine learning, and actuarial backgrounds have invested into improving the accuracy of morbidity and mortality forecasting. Some recent advances have emerged from the fields of hybrid models, and interpretable models such as Temporal Fusion Transformers [1]. Despite these advances, the recent devastating global impact of the novel coronavirus disease of 2019 (COVID-19) virus has highlighted the importance of effective planning from government agencies and healthcare bodies across the globe. This kind of planning requires reliable projections into the future, and as a result there exists a need for improved forecasting techniques and methods.

Smyl [2] developed a hybrid method for generating point forecasts and quantifying the uncertainty associated with those point forecasts. The model quantifies uncertainty by producing prediction intervals at the 1%, 5% and 10% levels of significance. The hybrid method combines Exponential Smoothing (ES, [3]) with Recurrent Neural Networks (RNN, [4]) and the resulting scheme is referred to as ES-RNN. This name is a bit of a misnomer as Smyl's [2] methodology actually involves combining ES, with a variant of RNN called Long Short-Term Memory (LSTM, [5,6]).

Smyl [2] purports that their method produces “. . .forecasts that are more accurate than those generated by either pure statistical or pure [machine learning] approaches, thus exploiting their advantages while avoiding their drawbacks”. The hybrid method, ES-RNN, outperformed pure statistical and pure deep learning methods submitted alongside it to the Makridakis-4 (M4) Forecasting Competition [7]. The M4 Competition provides to Competition participants 100,000 datasets from a cross-section of industries and applications for the entrants to apply their techniques to.

Redd et al. [8] extend the Smyl ES-RNN [2] hybrid technique by executing a graphical processing unit (GPU) implementation (GPU-ES-RNN). Using a vectorized and GPU-based implementation of the original ES-RNN, Redd et al. [8] state they achieve up to 322× increase in training speed, largely attributed to batching and parallelization. The authors report that their network produces performance results similar to those reported in the original Smyl submission. Furthermore, the GPU implementation migrates the original C++ code to Python and PyTorch.

In both cases (ES-RNN and GPU-ES-RNN), the authors focus on univariate datasets. However, a significant number of real world forecasting involves multivariate datasets [9]. In some cases, researchers have found it useful to augment their univariate models with multivariate exogenous factors and/or multivariate analogues of their models to improve predictive accuracy [10]. In addition, the complexities associated with modeling morbidity and mortality often require multiple data sources in order to fit a model that better describes the real world situation [11].

This paper extends both the ES-RNN and GPU-ES-RNN research by adapting the ES and LSTM hybrid method to make it applicable to the multivariate case. By incorporating exogenous factors throughout, our extension departs from the classical univariate case and, as such, poses more complications associated with, for instance, auto-correlation inherent in the data and cross-correlation dependencies between covariates.

Our model incorporates exogenous covariates through the use of global and local variables, i.e., those derived from all the data or large segments of it, and those derived from separate covariates. This combination allows the model to cross-learn at the same time leverage information presented at a granular level.

Essentially, we have three workstreams i.e., multivariate exponential smoothing, recurrent neural networks in general and LSTM in particular, and hybrid methods. We review all three workstreams next.

1.1. Literature Review

There are several notable works concerning multivariate extensions of exponential smoothing. Jones [12] applies recursion for estimating the smoothing matrix. Enns et al. [13] consider a class of various exponential smoothing models and use them as a proxy for the univariate exponential smoothing counterpart. Trigg and Leach [14] adapt the smoothing matrix of their so-called adaptive models periodically, with the aid of maximum likelihood estimation. Harvey [15] further simplifies the class of models proposed by Enns et al. [13] and the simplifications enable the use of univariate smoothing models in forecasting tasks for correlated time series. Moreover, Harvey’s [15] results are also valid in the case where the smoothing models exhibit polynomial trend and seasonal components. Pfeffermann and Allon [16] focus on structural models aimed at producing optimal forecasts. By building upon previous multivariate time series exponential smoothing research, Pfeffermann and Allon [16] also offer detailed instructions for parameter initialization and model-fitting.

Harvey’s approach is most suitable for our interests for two reasons. One, because of reduced computational complexity; and two, we find that adapting Harvey’s approach yields similar results to using an adapted version of, for example, the more complex Pfeffermann and Allon [16] approach. The second reason is due to the LSTM’s ability to model complex interdependency between covariates [17,18], which negates the need to model the cross-correlation using the preprocessing layer statistical methods.

The recent COVID-19 pandemic has provided a unique opportunity due to (i) a multitude of open datasets and (ii) extensive related research exploring the performance of multivariate forecasting models. In the domain of COVID-19-related research, LSTM networks have been applied by, for instance, Kirbaş et al. [19], who focus their attention on the spread of the pandemic in countries including Denmark, Belgium, Germany, France, United Kingdom, Finland, Switzerland and Turkey. Kirbaş et al. [19] model using Auto-Regressive Integrated Moving Average (ARIMA) and Nonlinear Autoregression Neural Network (NARNN, [20]) as benchmarks and report that LSTM has superior accuracy in terms of forecasting the cumulative cases of infected individuals.

Chandra et al. [21] use variations of LSTM including Bidirectional LSTM (Bi-LSTM), and encoder-decoder LSTM (ed-LSTM) models for multi-step intra-country forecasting in the short-term in India. They cite challenges in the modeling process caused by difficulties in capturing, in any available COVID-19 dataset, potentially important multivariate factors such as population density, travel logistics, and other societal issues (e.g., the general standard of living). If they were available, these exogenous factors could be useful in improving predictive skill in the modeling process.

Chimmula and Zhang [22] base their forecasts on training data that they acquired from the John Hopkins University and the Canadian Health Authority. They use a standard implementation of LSTM. They forecast the pandemic ending in June 2020 (which we now know to be incorrect).

Shahid et al. [23] model the rises and declines of confirmed cases, deaths and recoveries in ten countries, including China. They rate their model performances from best to poorest as Bi-LSTM, LSTM, Gated Recurrent Units (GRU, [24]), Support Vector Regressor (SVR) and ARIMA. They report Bi-LSTM has superior performance over the other models with the lowest MAE and RMSE values of 0.0070 and 0.0077, respectively.

As evidenced by previous related research, there are successes and failures with using LSTM to forecast COVID-19. Various factors contribute to the at times the poor performance of LSTM and its variations in this regard. One such factor is that the pandemic is relatively recent, and any available dataset is a small fraction of the requisite volume of training data the data-hungry artificial neural networks require. LSTM has been shown to perform well in a variety of applications with datasets that are sufficiently long [25]. In our research, we circumvent the problem of limited data by integrating a small, parsimonious LSTM network in our model. We offer more details about the structure of our model in Section 2.3.

Univariate forecasting with the aid of pure machine learning has been conducted since as far back as the works of Hu and Root [26]. Techniques merging machine learning with statistical methods have since gained popularity.

Recently, the original ES-RNN hybrid technique won over hundreds of other submissions both in terms of (i) point forecasts as well as (ii) their associated prediction intervals. Both components are difficult to model, but perhaps the latter even more so. Deep learning models often do not have a mechanism for quantifying uncertainty [27]. In such cases, prediction intervals must be constructed using computational mechanisms [25].

Oreshkin et al. [28], however, produce forecast machinery comprised of deep, fully connected layers and report that their model outperforms ES-RNN on the M4 Competition dataset. Based on Neural Basis Expansion Analysis (NBEATS), Oreshkin et al. [28] conclude that pure deep learning is better than hybrid methods as their method outperforms the best statistical and the best hybrid method by 11% and 3%, respectively. The technique has since been extended to include exogenous factors (NBEATS-x, [29]).

Further evidence has emerged since the end of the Makridakis-5 (M5) Accuracy and Uncertainty Competitions [30,31] that pure deep learning models may be superior for hierarchical forecasting. However, hybrid methods are still worth investigating in the multivariate with exogenous variables forecasting setting. Starting here as a point of departure, we hypothesise that a hybrid technique such as ES-RNN may work just as well in this setting (exploiting a simple LSTM's ability to model cross-correlation).

Furthermore, extending the current ES-RNN hybrid forecasting research to the multivariate case is crucial as multivariate datasets are usually more representative of realistic forecasting scenarios likely to be encountered in other applications.

1.2. Contribution

Our contribution can be summarized as follows:

- We extend the current research, which focuses on ES and RNN hybrid methods for univariate forecasting, to a multivariate framework. We thus test assertions on multivariate mortality data with exogenous variables previously only tested empirically on univariate data, i.e., are hybrid methods better than pure statistical or pure deep learning methods at (i) forecasting tasks and (ii) quantifying forecast uncertainty? In particular, we present a natural extension of Smyl's ES-RNN to higher dimensions;
- we present our forecast engine MES-LSTM, which is an efficient generalization, and as such, may be applied not only to the multivariate case but also to the univariate setting with ease; and
- whereas previous (univariate) research on forecasting hybrid methods primarily focuses on the multiplicative seasonality case, we consider both multiplicative and additive, with automatic adaptation to the case most applicable to the particular dataset.

The remainder of this paper is organized as follows. In Section 2 we describe our architecture in detail, and how we merge the classical state-space forecast model with the advanced artificial neural network. We also describe how we evaluate our model's performance. In Section 3 we describe the datasets used. The data section is followed by Section 4, where we discuss our results in detail. We conclude in Section 5 with a few key points and possibilities for extending this research to future work.

2. Methods

We employ GPU computation by utilising Tensorflow's eager execution to transform the original Smyl [2] ES-RNN to a generalised multivariate implementation. In their GPU implementation, Redd et al. [8] initialise per-series attributes according to the guidelines of the M4 Competition. We initialise per-covariate parameters, and have a vectorized ES-LSTM, which we term Multivariate ES-LSTM (MES-LSTM).

The other difference between ES-RNN and MES-LSTM is that we use the most suitable between additive or multiplicative seasonality, whereas previous authors have only considered multiplicative seasonality. The most suitable seasonality structure in each case is ascertained by tracking the exponential smoothing fit on the training data.

In the remainder of this chapter, we explicitly define the quintessential aspects of MES-LSTM. The description is analogous to that of Smyl [2], with relevant extensions made to suit our multivariate presentation. Our model comprises two distinct layers: an exponential smoothing layer inside our preprocessing module, and an LSTM layer used for learning the dependency between the covariates. This methodology is consistent with Section 1.1, where we outline our choice to expand upon Harvey's approach [15] over the Pfeiffermann and Allon [16] approach for multivariate exponential smoothing.

Concisely, our model learns the parameters associated with each covariate in the preprocessing step, then learns the correlation between the covariates in the deep learning layer. Parameters optimized in both steps are then used at the inference stage, where the prediction of multivariate point forecasts are produced, and the uncertainty associated with these forecasts is quantified.

2.1. Preprocessing Layer

For each covariate in our multivariate sample space, let $\{y\}_t$ represent, for example, a weekly time series (weekly series for ease of exposition only) which we assume can be decomposed into the additive form

$$y_t = l_t + b_t + s_t + \varepsilon_t, \tag{1}$$

where l_t is the level in week t , s_t the seasonal effect and ε_t the noise term centered at zero with constant variance. Let \hat{l}_t denote the estimated level in week t and \hat{b}_t the corresponding trend estimate. Let the estimate of the seasonal effect at time t corresponding to week $t + i$, $i = 1, 2, \dots, 52$ be denoted by $\hat{s}_{t,t+i}$. The estimates $\hat{s}_{t,t+i}$ satisfy the condition $\sum_{i=1}^{52} \hat{s}_{t,t+i} = 0$. When a new observation y_{t+1} becomes available, all estimated components, i.e., seasonality, trend, and level, are updated with the aid of three smoothing constants $0 \leq \alpha, \gamma, \delta \leq 1$ as follows:

$$\hat{l}_{t+1} = \alpha(y_{t+1} - \hat{s}_{t,t+1}) + (1 - \alpha)(\hat{l}_t + \hat{b}_t) \tag{2}$$

$$\hat{b}_{t+1} = \gamma(\hat{l}_{t+1} - \hat{l}_t) + (1 - \gamma)\hat{b}_t \tag{3}$$

$$\hat{s}_{t+1,t+1}^* = \delta(y_{t+1} - \hat{l}_{t+1}) + (1 - \delta)\hat{s}_{t,t+1} \tag{4}$$

$$\sum_{i=0}^{51} \hat{s}_{t+1,t+1+i}^* = 0. \tag{5}$$

Equations (4) and (5) define a two-step computational procedure of the seasonal effects $\hat{s}_{t+1,t+1+i}$. The second step standardizes the initial values computed in the first (with the asterisk indicating “intermediary value”) so that they sum to zero.

The forecast at time t of a future out-of-sample realization ($m > i$) y_{t+m} is given by

$$\hat{y}_{t,t+m} = \hat{l}_t + m\hat{b}_t + \hat{s}_{t,t+m}. \tag{6}$$

For the multiplicative seasonality case, make the substitution $\{y\}_t = \ln(\{y\}_t)$. This substitution amounts to assuming the original series level changes in approximately constant rates instead of constant increments. The substitution also requires changing the condition imposed on the estimates of the multiplicative seasonal factors over the 52 weeks from summing to zero to their geometric mean equalling one. Equations (2)–(5) then become

$$\hat{l}_{t+1} = \alpha\left(\frac{y_{t+1}}{\hat{s}_{t,t+1}}\right) + (1 - \alpha)\hat{l}_t\hat{b}_t \tag{7}$$

$$\hat{b}_{t+1} = \gamma\left(\frac{\hat{l}_{t+1}}{\hat{l}_t}\right) + (1 - \gamma)\hat{b}_t \tag{8}$$

$$\hat{s}_{t+1,t+1}^* = \delta\frac{y_{t+1}}{\hat{l}_t\hat{b}_t} + (1 - \delta)\hat{s}_{t,t+1} \tag{9}$$

$$\left(\prod_{i=0}^{51} \hat{s}_{t+1,t+1+i}^*\right)^{\frac{1}{51}} = 1, \tag{10}$$

and the predictive Equation (6) becomes

$$\hat{y}_{t,t+m} = \hat{l}_t(\hat{b}_t)^m \hat{s}_{t,t+m}. \tag{11}$$

In both the additive and multiplicative seasonality cases expressed above, the initial smoothing parameters are estimated as follows. The level is taken to be the overall in-sample average. The trend is initialized as the difference between the first and last in-sample observations, divided by the total number of increments. Seasonality is initialized by the

deviations between the first week's observations and the level plus trend fit. More details about the dimensionality of the parameter space are given in Section 2.3.

Now that we have the formulations for both the additive and multiplicative seasonality cases, we can merge the preprocessing layer with our deep learning component as follows. For simplicity, we keep the input size and the size of the output predictions equal.

For the univariate additive case we have

$$\hat{y}_{t, t+1, \dots, t+m} = \hat{l}_t + [t+1, \dots, t+m]^\top \odot \text{LSTM}(X_t) + \hat{s}_{t, t+1, \dots, t+m} \quad (12)$$

where \odot is the Hadamard product, and X_t denotes a vector of de-trended and de-seasonalized observations of which a scalar element x_i is given by:

$$x_i = y_i - \hat{l}_i - \hat{s}_i, \text{ for } i = t+1, \dots, t+m. \quad (13)$$

For the multivariate additive case we have

$$\hat{y}_{t, t+1, \dots, t+m} = \hat{l}_t + [t+1, \dots, t+m]^\top \odot \text{LSTM}(X_t) + \hat{s}_{t, t+1, \dots, t+m}, \quad (14)$$

where (if we have, say, k covariates in total) X_t is a $k \times m$ matrix of de-trended, de-seasonalized features of which a vector component (for each covariate) x_i is calculated via the equation:

$$x_i = y_i - \hat{l}_i - \hat{s}_i, \text{ for } i = t+1, \dots, t+m. \quad (15)$$

Here, the vectors \hat{l}_i and \hat{s}_i are the same dimension as y_i .

The univariate multiplicative seasonality case is given by

$$\hat{y}_{t, t+1, \dots, t+m} = \text{LSTM}(X_t) \odot \hat{l}_t \odot \hat{s}_{t, t+1, \dots, t+m} \quad (16)$$

$$x_i = \frac{y_i}{\hat{l}_i \hat{s}_i} \text{ for } i = t+1, \dots, t+m, \quad (17)$$

where the LSTM takes as input a vector of de-trended and de-seasonalized observations of which a scalar element x_i is computed using Equation (17).

Finally, the multivariate multiplicative case is thus expressed as

$$\hat{y}_{t, t+1 \dots t+m} = \text{LSTM}(X_t) \odot \hat{l}_t \odot \hat{s}_{t, t+1 \dots t+m} \quad (18)$$

$$x_i = y_i \odot (\hat{l}_i \odot \hat{s}_i)^{-1} \text{ for } i = t+1, \dots, t+m \quad (19)$$

where the LSTM takes as input a size $k \times m$ matrix X_t of de-trended, de-seasonalized observations, where each vector component x_i is computed via Equation (19).

2.2. Deep Learning Layer

Our neural network data flow is presented in Figure 1. We input a matrix composed of k vectors each of size m . After preprocessing, we input the matrix to our LSTM layer. The LSTM layer output (composed of j predictands each sized n) feeds into a dense layer, then postprocessing is conducted to produce the model's final output. LSTMs employ gated connections, are good at modeling latent representations with temporal dependency, and as a result, are superior to vanilla RNNs [5]. The deep learning and exponential smoothing layers are optimized consecutively, as depicted in Figure 1.

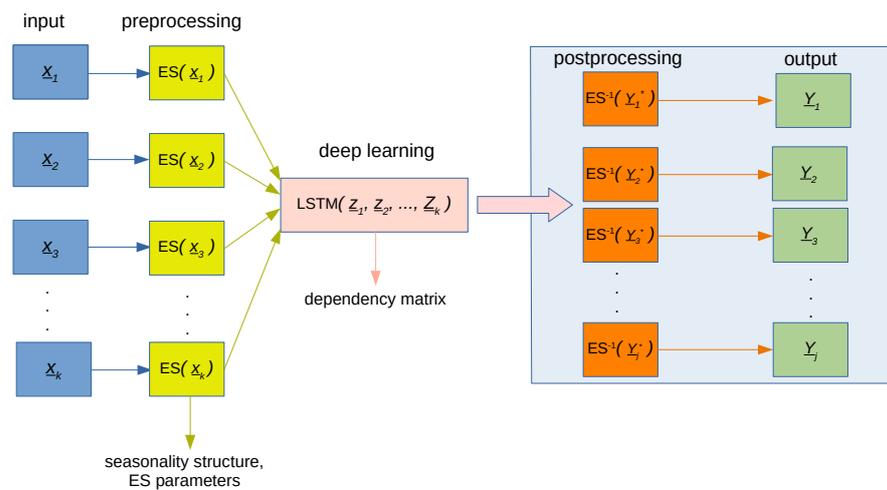


Figure 1. MES-LSTM Data Flowchart.

In order to quantify the uncertainty associated with our point forecasts, we modify the above architecture as follows. Instead of the dense layer that back-propagates and learns scalar sets of weights and biases, we employ a densely-connected layer class with Flipout estimator [32] that learns a *distribution* of weights and biases. In this variational, probabilistic setting, traditional back-propagation is replaced by Bayes-by-backprop [33].

Flipout [32], through Monte Carlo simulation, approximates a distribution of weights by integrating over the kernel and bias. By assuming the kernel and biases are drawn from distributions, the dense-flipout is able to implement the Bayesian Variational inference dense-flipout layer, analogous to the above (traditional) dense layer. Using samples from the kernel and bias posteriors, this dense-flipout layer is able to run stochastic forward passes. Another difference between this stochastic process and the analogous densely connected layer, is we use variational inferencing by minimizing the KL-divergence [34]. The uncertainty quantification is implemented using Tensorflow Probability [35].

In effect, we can now forecast using a sample from the distribution of weights and biases. Applying Monte-Carlo simulation, each time we forecast, we iteratively compile a distribution of forecasts. We then use the appropriate percentiles to extract the desired prediction intervals from the forecast distribution. So, to compute the $(1 - \alpha) \times 100\%$ prediction interval, we use the percentiles at

$$\left(\frac{\alpha}{2}, (1 - \alpha) + \left(\frac{\alpha}{2} \right) \right). \tag{20}$$

The data flow process for quantifying uncertainty is not dissimilar to Figure 1. The only difference is that in the deep learning layer, we now produce probabilistic forecasts instead of deterministic ones.

The variational inference approach described above is similar to two other notable prediction interval construction techniques, i.e., Monte Carlo dropout [36] and the quantile bootstrap [37] also known as the reverse percentile bootstrap [38]. The technique used here, Flipout, is similar to dropout and bootstrapping as they all use quantiles extracted from synthetic distributions to produce the final intervals. Mathonsi and van Zyl [25] offer examples where dropout and bootstrapping have been applied and compared [25].

The key differences in all three methods is how their respective distributions are constructed, and their relative uncertainty quantification skill. See for instance the work of Wen et al. [32], where Flipout has been shown to outperform dropout. In addition, a fortuitous side-effect of the variational inference (and dropout methods) is a reduction in epistemic uncertainty, which in turn regularizes the network and addresses any concerns that may arise from the possibility of a lack of sufficient training data, and issues associated with overfitting.

To summarize, we input a matrix (sized number of training observations by number of predictors) to the model. The preprocessing layer computes all the exponential smoothing parameters as outlined in Section 2.1 above. We then feed the output directly to the deep learning layer, which discerns the inherent dependency within the covariates. The first step in our two-step methodology is inspired by the works of Harvey [15], who, through simplifications, enables forecasting related time series through the use of univariate smoothing models. We find this two-step methodology for MES yields similar results to using the one-step formulation given by Pfeffermann [16], with our method enjoying the added benefit of reduction in computational cost. The output from the LSTM is then fed directly to the dense layer or dense-flipout layer in the case of point forecasts and prediction interval construction, respectively. Finally, seasonality and trend estimates are added back to these intermediate outputs. We now have the final model forecasts and prediction intervals, respectively.

Figure 2 illustrates our model architecture zoomed in to the deep learning layer. The input shape is (k, m) where we have k covariates each with m observations in the training data. The network's intermediate output (before postprocessing) is generated by feeding the output of the LSTM through to a dense layer with a ReLU activation function. The size of the LSTM, S , is deduced empirically using the technique described in Section 2.3 below. The size of the dense layer (or dense-flipout layer for uncertainty quantification) and correspondingly the output layer is determined by the number of predictands j . The output in the schematic then goes through postprocessing to meet our required format. It is re-trended and re-seasonalized (using the parameter estimates from the exponential smoothing equations) to arrive at the format presented by the ground truth data.

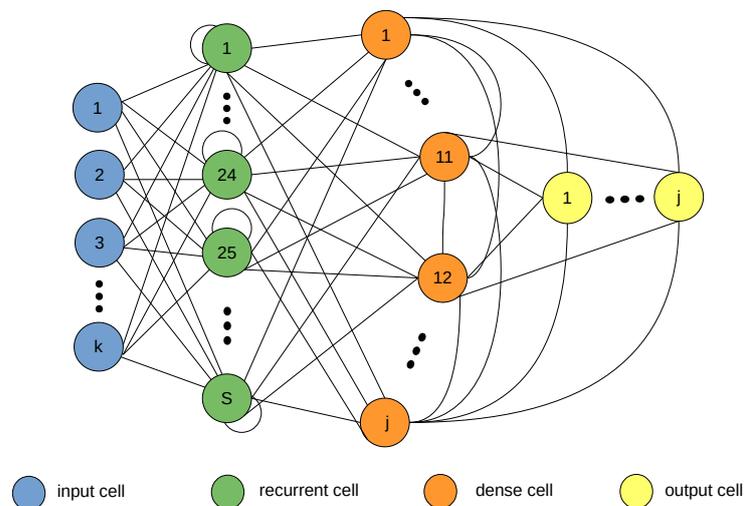


Figure 2. MES-LSTM Architecture Diagram.

2.3. Hyperparameter Tuning

Note from our formulation given in Equations (14) and (15) (multivariate model with additive seasonality), and Equations (18) and (19) (multivariate model with multiplicative seasonality), we do not need to compute the estimates for the trend component \hat{b}_t and this significantly reduces the hyper-parameter search space of the preprocessing layer. The initial estimates of the coefficients for level and seasonality are deduced by calculating primer estimates following the classical exponential smoothing Equations (2) and (4), as well as Equations (7) and (9) for the respective additive and multiplicative seasonality cases. The level is initialized as the in-sample average, and the initial seasonality is the deviations between the level plus trend fit and the first week's observations. If we have k attributes in our dataset, the model computes and stores $(k + 1) \times (2 + P)$ exponential smoothing parameters, where P indicates seasonality length.

For the deep learning layer, we use a small model with relatively few parameters. The reasons for this model configuration are two-fold. First, we have a relatively small amount of data for training, and secondly, a large over-parameterized network might overfit and not generalize well to the test data [39].

We iterate over a subset of LSTM sizes and training epochs and select the model configuration with the best forecast accuracy for our validation data and fewest parameters to fit. We optimize the batch size and number of samples the rolling window looks into the past to forecast the next example. We summarize our hyperparameter search space for the deep learning layer in Table 1. We inspect the best five configurations (w.r.t. the error metrics detailed in the section that follows) in terms of forecasting all of the predictands. We then select the model configuration that is most parsimonious in order to ensure our model best mitigates overfitting over multiple runs. According to this methodology, the best configuration is given by LSTM of size 50, trained over 25 epochs, using batches sized 16, with an input window of 14 days.

Table 1. Configuration Grid for Deep Learning Layer Hyperparameter Search Space.

Hyperparameter	Search Space
LSTM size	50, 55, 60, . . . , 150
epochs	15, 20, 25, . . . , 75
batch size	8, 16, 24, . . . , 64
input window	7, 14, 21

Our train-validation-test data split is 75-15-10. The reason for this is we forecast in the short-to-medium term, so 10% test data is sufficient to evaluate our model's performance.

The order of our VARMAX model, (p, q) , is deduced by conducting a grid search for optimization in much the same way as above for MES-LSTM. We set the maximum iterations for maximum likelihood search to 200 (four times the default in Python) to ensure convergence. The trend component is deduced through a grid search varying the trend polynomial from constant, linear, quadratic, to cubic. We auto-regress the predictors and the predictands are declared as exogenous.

The model hyperparameter optimization process for SARIMAX is exactly the same as for VARMAX. In both cases we perform the grid search and choose the model with the smallest Akaike information criteria (AIC, [40]). For simplicity, we do not enforce invertibility on the moving average polynomials. This also ensures more of the models are estimable. In cases where there are several suitable model configurations (equal AIC values) we choose the model that is most parsimonious (smallest product of p and q).

The MLR is fit using Ordinary Least Squares (OLS, [41]). We plug in the optimal hyperparameters for LSTM into the DeepAR, with additional searches for the dropout rate (0.1, 0.15, 0.2), and likelihood (noise model for probabilistic forecasts, so we can discern uncertainty) between Gaussian and student-T. The LGB model was tuned considering tree maximum depth (capped at 5), maximum leaves (capped at 30), number of estimators (capped at 125), and the learning rate (same search space as DeepAR).

All the models are trained on a cluster of Intel Core i7-9750H processors each with an Nvidia GeForce GTX 1650 GPU and 4GB RAM.

2.4. Metrics

For determining the most suitable seasonality structure in the preprocessing layer, we use the Sum of Squared Errors (SSE). For training in the deep learning layer, we use Mean Absolute Error loss (MAE) with Adam efficient optimizer [42]. For performance

evaluation of the point forecasts, we employ the symmetric Mean Absolute Percent Error (sMAPE, [43]) and Root Mean Squared Error (RMSE),

$$\text{sMAPE} = \frac{2}{m} \sum_{t=1}^m \frac{|y_t - \hat{y}_t|}{|y_t| + |\hat{y}_t|} \quad (21)$$

$$\text{RMSE} = \sqrt{\frac{1}{m} \sum_{t=1}^m (y_t - \hat{y}_t)^2}, \quad (22)$$

where y_t is the post-sample value of the predictand at point t , \hat{y}_t the estimated forecast and m is again the forecast horizon. The former is in line with the performance metrics from the M4 Competition, and is used for continuity as results published by Smyl [2] also used this metric. Further motivation for our choice is sMAPE being a median-based error criterion, it is also useful in instances where there are large amounts of outliers in the data. Lastly, sMAPE is symmetric on an absolute scale [44].

The latter is useful as it gives error in the same units as the forecast variable itself, thus easily interpretable. The other reason we employ RMSE as a complement is although symmetric on an absolute scale, sMAPE has been shown to penalize large positive errors more than negative ones on a percentage scale [45]. Both metrics are mean absolute differences between forecast and actual values. The key difference is how sMAPE is normalized.

For evaluating the prediction intervals we use the Mean Interval Score (MIS, [46]), which is averaged over all out-of-sample observations,

$$\text{MIS}_\alpha = \frac{1}{m} \sum_{t=1}^m (U_t - L_t) + \frac{2}{\alpha} (L_t - y_t) \mathbb{1}_{\{y_t < L_t\}} + \frac{2}{\alpha} (y_t - U_t) \mathbb{1}_{\{y_t > U_t\}}, \quad (23)$$

where U_t (L_t) is the upper (lower) bound of the prediction interval at time t , α is the significance level and $\mathbb{1}$ is the indicator function.

The MIS_α adds a penalty at the points where future observations are outside the specified bounds $[L_t, U_t]$. The width of the interval at t is added to said penalty, if any. As such, the MIS_α also penalizes wide intervals. Finally, this sum at the individual out-of-sample points is averaged.

As a supplementary metric for evaluating the performance of the prediction intervals, we employ the Coverage Score (CS_α) which indicates the percentage of observations that fall within the prediction interval,

$$\text{CS}_\alpha = \frac{1}{m} \sum_{t=1}^m \mathbb{1}_{\{U_t \leq y_t \leq L_t\}}. \quad (24)$$

With MIS_α the subscript α denotes explicit dependence on the significance level, whereas in the case of CS_α the dependence is implicit.

2.5. Benchmarks

For benchmarking MES-LSTM, the statistical methods used are Multiple Linear Regression (MLR, [41]), Vector Autoregression Moving-Average with Exogenous Regressors (VARMAX, [47]) and Seasonal Autoregressive Integrated Moving-Average with Exogenous Regressors (SARIMAX, [48]). For deep learning benchmarks we use a vanilla LSTM (without direct incorporation of a preprocessing layer as described in Section 2.1) Deep Autoregressive Recurrent neural network (DeepAR, [49]), and Light Gradient Boosting machine (LGB, [50]).

The choice of the pure deep learning benchmarks/baselines models is based on, for one, the architecture of our proposed model. Because we have a statistical-and-LSTM hybrid model, it makes sense to focus on LSTM w.r.t. pure deep learning techniques. Secondly, the choice of LGB (and LSTM) is motivated by the findings of the M5 Competitions. Almost

all of the top 50 submissions for the M5 Accuracy [51] and Uncertainty [52] competitions use a variation of LGB. In particular, considering the top five: submissions ranked 1, 2, 4, and 5 from the Accuracy competition [51] incorporate LGB into their model, and the submission ranked third uses DeepAR [49], which is built on multiple LSTMs. From the top five submissions in the M5 Uncertainty competition [52], 1, 2, 3, and 5 incorporate LGB, while the submission ranked fourth incorporates LSTM.

3. Datasets

We use the Our World in Data (OWID) COVID-19 dataset [53,54]. This dataset is aggregated from various sources, and includes historical data on the pandemic up to the date of publication, updated daily. A summary of the data and the aggregated data sources is shown in Table 2. It is important to note that even though the actual databases are updated at daily, weekly, and other time frequencies, the aggregated datasets used in our study are all presented at a daily temporal resolution. This means no frequency harmonization is required.

Table 2. OWID COVID-19 Dataset Summary.

Metrics	Source	Updated	Countries
Vaccinations	Official data collated by the Our World in Data team	Daily	217
Tests & positivity	Official data collated by the Our World in Data team	Weekly	136
Hospital & ICU	Official data collated by the Our World in Data team	Weekly	34
Confirmed cases	JHU CSSE COVID-19 Data	Daily	194
Confirmed deaths	JHU CSSE COVID-19 Data	Daily	194
Reproduction rate	Arroyo-Marioli F, Bullano F, Kucinskas S, Rondón-Moreno C	Daily	184
Policy responses	Oxford COVID-19 Government Response Tracker	Daily	186
Other variables of interest	International organizations (UN, World Bank, OECD, IHME...)	Fixed	240

The variables represent data related to confirmed cases, deaths, hospitalizations, and testing, as well as 56 other variables. We only use a subset of the available attributes, as detailed in the feature list shown in Table 3.

Table 3. OWID COVID-19 Dataset Feature List.

Variable	Description
total cases	Total confirmed cases of COVID-19
new cases	New confirmed cases of COVID-19
total cases per million	Total confirmed cases of COVID-19 per 1,000,000 people
new cases per million	New confirmed cases of COVID-19 per 1,000,000 people
total deaths	Total deaths attributed to COVID-19
new deaths	New deaths attributed to COVID-19
total deaths per million	Total deaths attributed to COVID-19 per 1,000,000 people
new deaths per million	New deaths attributed to COVID-19 per 1,000,000 people
icu patients	Number of COVID-19 patients in intensive care units (ICUs) on a given day
icu patients per million	Number of COVID-19 patients in ICUs on a given day per 1,000,000 people
hosp patients	Number of COVID-19 patients in hospital on a given day
weekly icu admissions	Number of COVID-19 patients newly admitted to ICUs in a given week
weekly icu admissions per million	Number of COVID-19 patients newly admitted to ICUs in a given week per 1,000,000 people
weekly hosp admissions	Number of COVID-19 patients newly admitted to hospitals in a given week
weekly hosp admissions per million	Number of COVID-19 patients newly admitted to hospitals in a given week per 1,000,000 people
stringency index	Government Response Stringency Index: composite measure based on 9 response indicators
reproduction rate	Real-time estimate of the effective reproduction rate (R) of COVID-19
total tests	Total tests for COVID-19
new tests	New tests for COVID-19 (only calculated for consecutive days)
positive rate	Share of COVID-19 tests that are positive, rolling 7-day average (inverse of tests per case)

Table 3. Cont.

Variable	Description
tests per case	Tests conducted per new confirmed case of COVID-19, rolling 7-day average (inverse of positive rate)
total vaccinations	Total number of COVID-19 vaccination doses administered
people vaccinated	Total number of people who received at least one vaccine dose
people fully vaccinated	Total number of people who received all doses
new vaccinations	New COVID-19 vaccination doses administered (only calculated for consecutive days)
total vaccinations per hundred	Total number of COVID-19 vaccination doses administered per 100 people
people vaccinated per hundred	Total number of people who received at least one vaccine dose per 100 people
people fully vaccinated per hundred	Total number of people who received all doses prescribed by the vaccination protocol per 100 people
location	Geographical location
date	Date of observation
population	Population in 2020
population density	Number of people divided by land area, measured in square kilometers
median age	Median age of the population, UN projection for 2020
aged 65 older	Share of the population that is 65 years and older, most recent year available
aged 70 older	Share of the population that is 70 years and older in 2015
gdp per capita	Gross domestic product at purchasing power parity
extreme poverty	Share of the population living in extreme poverty
cardiovasc death rate	Death rate from cardiovascular disease in 2017
diabetes prevalence	Diabetes prevalence (% of population aged 20 to 79) in 2017
female smokers	Share of women who smoke, most recent year available
male smokers	Share of men who smoke, most recent year available
handwashing facilities	Share of the population with basic handwashing facilities on premises
hospital beds per thousand	Hospital beds per 1000 people, most recent year available since 2010
life expectancy	Life expectancy at birth in 2019
human development index	Composite average achievement in (i) a long, healthy life (ii) knowledge (iii) standard of living
excess mortality	Excess mortality P-scores for all ages

We choose this aggregated dataset over other available COVID-19 datasets because it does to some extent mitigate the concerns raised by Chandra et al. [21], for example. As highlighted in Section 1.1, some exogenous factors not directly related to COVID-19 may be useful in the modeling process. For instance, the OWID COVID-19 dataset includes covariates such as *human development index*, *extreme poverty*, and *handwashing facilities*.

The variables that we have omitted from our study are duplicates where numerical data has been smoothed, e.g., *new cases smoothed*. The smoothed attributes are removed for two reasons: OWID offers no insights on how the data is smoothed, and we avoid any duplication as our model also conducts preprocessing.

In particular, we use our model to forecast and quantify the associated prediction uncertainty for two attributes of interest: *total cases* and *total deaths*. The predictions and uncertainty quantification can be conducted for a single country, or easily across a multitude of countries or even averaged over a specific region. Regional or multi-country inferencing can assist policy-makers to make tough decisions like restricting inter-provincial/state travel or closing national borders completely.

Below we present results for the Southern African Development Community (SADC). SADC is a regional economic community comprising 16 member states: Angola, Botswana, Comoros, Democratic Republic of Congo (DRC), Eswatini, Lesotho, Madagascar, Malawi, Mauritius, Mozambique, Namibia, Seychelles, South Africa, Tanzania, Zambia and Zimbabwe.

4. Results and Discussion

In this section, we detail the results for multivariate point forecasts as well as the prediction intervals for our predictands *total cases* and *total deaths*. The results are presented for the analysis conducted for the SADC region. In order to mitigate the stochastic nature of the probabilistic forecasts from MES-LSTM and some of the benchmark models, the experiments are repeated 35 times. The results presented are the aggregates of all the repeated independent trials. The code repository for the empirical experiments is available online at the link provided directly in Supplementary Material after Section 5 below.

4.1. Forecast Performance for SADC

The results for point forecasts for SADC are tabulated in Tables 4 and 5, for each of our predictands. We note MES-LSTM outperforms the benchmarks for all the nations in SADC, except sMAPE for *total cases* in South Africa. The best performer in each instance is highlighted for ease of interpretation. Interestingly, overall, the worst forecast performance from MES-LSTM is in South Africa and this could be a result of the nation presenting the most accurate data. The other nations possibly don't update their data as frequently, and the model learns easier as there isn't a lot of variability. Furthermore, the (less accurate) data from the other nations is closer to the linear assumptions of the statistical benchmark models, i.e., VARMAX, SARIMAX and MLR.

In Figure 3 we average the forecast results across the entirety of the SADC region. We note MES-LSTM is the best aggregate performer. This figure also illustrates the importance of choosing multiple error metrics. In Figure 3a SARIMAX for instance, seems to perform fairly poorly for both predictands, but this as a stand-alone interpretation would be inaccurate. From Figure 3b we note that in terms of regional skill for *total deaths*, SARIMAX is actually competitive. Recall that RMSE gives the error in the same units as the original predictand. South Africa has a total population of 65 million, so any model that over- or under-forecasts *total cases* by a few thousand cases could still be useful for planning and management of the pandemic outbreak in question.

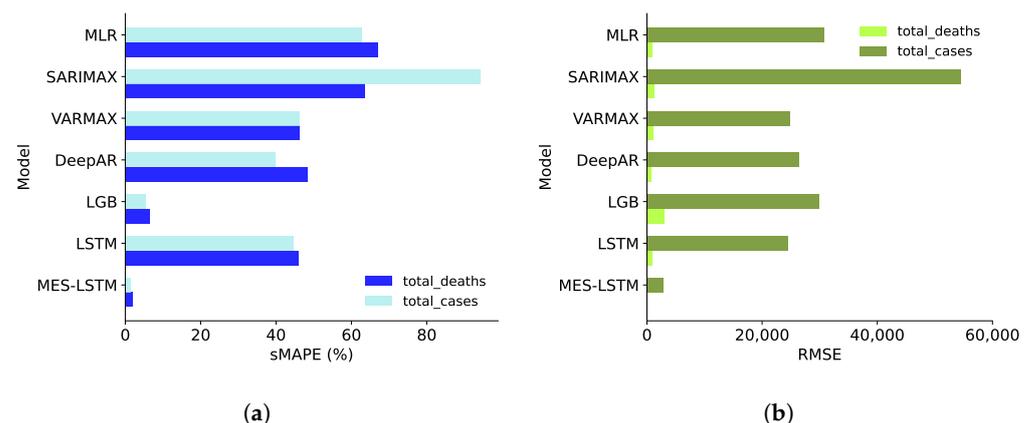


Figure 3. Forecast Accuracy in the SADC Region. (a) sMAPE. (b) RMSE.

4.2. Prediction Interval Performance for SADC

The results for the prediction interval are documented in Tables 6 and 7. The best performer or best tied performers are highlighted. We note that MES-LSTM MIS is superior to the other models for both predictands for almost the entire SADC region.

Recall from Section 2.4, the MIS penalizes large intervals, so our model generally has the narrowest prediction intervals. In Figure 4 we average the results for prediction intervals across the entirety of the SADC region. MES-LSTM consistently has the narrowest aggregate intervals for both predictands at all prediction intervals, as evidenced in Figure 4a–c.

In terms of Coverage, our model outperforms the benchmarks for the individual countries except in a few instances. In most of the exceptions, the difference is minuscule, whereas in other instances (e.g., South Africa) it is admittedly not negligible. What is more important however, is the aggregate performance over the entire region, where MES-LSTM scores better (as can be seen in Figure 4d–f).

Table 4. Forecast Accuracy Averaged Over All Trials for *total cases* in the SADC Region.

Country	MES-LSTM		LSTM		LGB		DeepAR		VARMAX		SARIMAX		MLR	
	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE
Angola	0.7	563.1	68.9	28,023.4	5.6	28,524.5	59.0	23,811.6	76.4	35,442.3	107.1	59,732.6	77.6	35,830.6
Botswana	1.6	5817.7	85.5	94,414.1	5.8	32,635.2	62.9	72,983.0	99.5	125,387.6	71.0	123,392.6	114.6	137,334.7
Comoros	1.1	52.5	16.5	623.3	5.3	32,005.1	50.1	1443.4	6.2	313.8	23.6	1058.6	17.6	704.2
DRC	0.8	468.3	72.6	25,987.3	5.3	29,933.9	41.9	17,302.5	12.8	7047.8	27.6	15,527.6	83.9	34,102.8
Eswatini	1.3	617.7	59.1	18,348.9	5.7	31,572.0	61.2	17,619.4	63.6	23,417.0	37.9	16,687.6	110.6	33,055.4
Lesotho	0.7	167.2	47.2	7324.4	5.5	28,662.4	33.9	5484.3	35.2	6478.7	137.8	24,942.5	42.4	7603.3
Madagascar	1.3	636.6	35.1	12,019.8	5.5	29,387.4	45.6	13,719.5	11.1	5495.1	8.8	3889.1	22.3	9014.0
Malawi	1.3	817.5	22.0	11,967.0	5.9	30,163.2	14.1	8183.4	10.4	7629.0	9.7	5956.9	42.2	23,504.7
Mauritius	3.8	2056.8	99.8	9891.6	5.8	29,677.4	91.1	8643.8	179.5	17,135.7	242.9	32,875.6	171.3	16,742.4
Mozambique	1.1	1685.0	5.0	7457.6	5.2	29,234.1	3.0	4438.9	2.5	3946.2	90.5	125,768.7	10.3	15,089.2
Namibia	1.1	1508.8	5.2	7516.3	5.6	26,149.9	2.3	2969.0	12.0	17,579.1	25.3	32,551.7	7.7	10,010.1
Seychelles	0.9	266.7	59.9	8908.6	5.4	27,596.0	62.1	8596.8	11.4	2428.1	47.4	10,115.2	99.0	14,797.0
South Africa	5.0	26,979.2	4.8	150,416.4	5.6	30,420.9	7.8	212,612.0	2.8	87,376.0	8.2	265,998.0	4.0	118,139.5
Tanzania	0.3	134.6	116.3	15,445.1	6.0	30,900.7	96.5	12,839.8	184.1	25,060.9	636.9	87,188.8	194.9	25,818.8
Zambia	1.2	2499.0	6.4	15,567.7	5.8	28,686.9	2.0	4333.4	7.3	16,766.9	72.1	143,490.3	2.1	6120.9
Zimbabwe	1.1	1518.5	8.6	10,652.2	5.4	32,284.2	5.7	7251.2	8.4	12,296.7	6.4	9733.6	2.6	4164.0

Table 5. Forecast Accuracy Averaged Over All Trials for *total deaths* in the SADC Region.

Country	MES-LSTM		LSTM		LGB		DeepAR		VARMAX		SARIMAX		MLR	
	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE
Angola	0.9	19.8	74.6	783.5	6.6	2886.7	60.8	687.0	89.9	1055.7	71.6	1038.1	89.4	1052.1
Botswana	1.0	26.1	86.1	1203.9	6.7	2909.9	77.2	1128.8	97.6	1576.5	15.7	363.9	120.8	1806.0
Comoros	1.5	2.4	12.5	16.7	6.2	3069.4	72.4	66.8	15.6	27.3	13.5	20.6	10.9	16.0
DRC	1.0	12.0	67.2	470.1	6.1	3059.6	38.2	326.4	2.4	34.2	11.6	128.9	74.0	594.3
Eswatini	1.1	14.2	49.6	435.1	7.0	2594.2	58.8	490.0	46.3	525.6	35.3	403.1	95.1	801.0
Lesotho	0.8	5.4	47.2	222.9	6.9	2790.2	42.4	207.1	46.8	249.7	97.4	516.2	44.5	240.8
Madagascar	1.2	12.2	48.1	330.3	6.6	2820.6	71.5	431.7	4.2	44.6	2.0	20.4	41.7	335.9
Malawi	1.2	28.7	27.1	518.9	6.6	3072.6	20.4	416.1	14.5	349.6	14.2	314.9	46.2	958.7
Mauritius	1.7	39.9	103.0	109.6	6.9	3014.0	104.4	110.3	173.1	183.1	300.2	341.7	172.7	183.1
Mozambique	1.1	21.9	5.1	96.7	6.4	3269.7	4.5	85.6	4.9	98.7	9.2	180.1	15.6	281.4
Namibia	1.1	41.9	7.9	333.8	6.1	3200.6	4.9	174.9	19.6	848.9	17.2	601.2	29.3	913.9
Seychelles	1.7	7.0	71.9	54.0	7.1	3288.2	82.4	58.9	5.9	8.3	43.5	47.7	116.9	88.8
South Africa	4.9	446.2	7.8	7902.4	5.9	2773.4	10.4	8515.8	10.7	10,806.6	16.1	15,438.9	6.5	6276.2
Tanzania	0.3	3.8	114.7	425.6	6.7	3143.2	113.7	423.5	179.3	685.9	362.3	1379.1	193.0	712.9
Zambia	1.3	47.7	6.6	266.8	6.6	2806.5	4.3	174.2	7.5	301.4	1.9	87.8	11.7	420.1
Zimbabwe	1.1	53.4	6.5	298.2	6.3	3067.7	8.1	361.0	14.8	791.5	5.8	289.7	6.7	345.0

Table 6. Prediction Interval Accuracy Averaged Over All Trials for *total cases* in SADC ($\alpha = 0.05$).

Country	MES-LSTM		LSTM		LGB		DeepAR		VARMAX		SARIMAX		MLR	
	MIS	Coverage	MIS	Coverage	MIS	Coverage	MIS	Coverage	MIS	Coverage	MIS	Coverage	MIS	Coverage
Angola	1910.8	95.2	116,739.1	0.0	55,663.7	81.6	54,875.9	0.0	902,714.3	0.0	814,594.0	0.0	660,826.0	0.0
Botswana	11,735.6	85.2	378,936.2	0.0	164,794.4	77.6	171,350.6	0.0	4,572,270.6	0.0	1,203,716.2	0.0	4259,387.7	0.0
Comoros	172.9	97.4	360.0	87.6	3908.4	84.4	5064.1	0.0	3085.7	40.9	1790.1	59.1	5461.2	100.0
DRC	2351.8	97.6	102,410.9	0.0	45,538.0	65.2	68,391.5	0.0	143,127.3	0.0	137,288.5	0.0	413,486.5	0.0
Eswatini	1837.3	91.0	86,314.2	0.0	44,987.2	79.6	44,252.9	0.0	518,283.5	14.3	140,160.1	0.0	683,626.1	0.0
Lesotho	750.4	97.4	24,269.2	0.4	20,065.3	89.0	18,693.3	0.0	218,810.3	0.0	310,359.2	0.0	24,172.3	55.8
Madagascar	2310.7	83.6	25,720.5	16.9	30,953.7	83.1	19,934.0	56.3	139,138.3	0.0	5270.5	100.0	48,987.3	100.0
Malawi	2503.3	88.0	27,408.2	38.8	40,635.9	78.9	31,982.9	0.0	60,588.0	70.2	33,628.7	100.0	303,117.6	55.3
Mauritius	2460.2	90.1	44,971.4	0.0	21,957.8	77.6	43,120.5	0.0	625,931.1	0.0	441,038.4	0.0	592,480.4	0.0
Mozambique	5249.2	94.2	17,924.2	92.1	68,002.5	86.0	3197.7	46.2	23,040.1	100.0	1,364,340.7	0.0	83,378.3	22.9
Namibia	4487.2	93.1	18,930.5	93.9	53,947.4	79.7	1726.9	59.2	48,9847.0	0.0	28,333.2	38.8	45,068.1	67.3
Seychelles	728.2	93.1	32,422.8	0.0	20,364.2	73.7	17,835.9	14.3	68,568.1	0.0	103,106.7	0.0	232,264.3	0.0
South Africa	193,075.6	80.7	348,860.7	97.7	1,418,713.2	86.6	239,009.3	100.0	740,131.9	100.0	1,061,162.0	100.0	828,447.2	100.0
Tanzania	943.6	96.2	67,427.5	0.0	32,459.5	77.3	67,350.0	0.0	990,720.2	0.0	1,279,781.1	0.0	1,018,128.2	0.0
Zambia	8972.0	89.2	24,151.6	96.4	89,095.3	73.5	8316.6	16.6	489,501.1	0.0	1,509,742.0	0.0	36,167.6	95.8
Zimbabwe	4770.4	87.7	23,667.6	85.8	63,299.2	77.9	2987.9	100.0	367,165.9	0.0	14,904.8	95.8	16,359.3	100.0

Table 7. Prediction Interval Accuracy Averaged Over All Trials for *total deaths* in SADC ($\alpha = 0.05$).

Country	MES-LSTM		LSTM		LGB		DeepAR		VARMAX		SARIMAX		MLR	
	MIS	Coverage	MIS	Coverage	MIS	Coverage	MIS	Coverage	MIS	Coverage	MIS	Coverage	MIS	Coverage
Angola	53.2	94.0	3206.0	0.0	3348.1	77.1	2089.5	0.0	29,346.8	0.0	24,022.4	0.0	23,844.1	0.0
Botswana	106.4	93.3	4803.4	0.0	4357.0	75.1	1706.3	10.6	53,087.1	0.0	1079.0	48.9	51,432.7	0.0
Comoros	5.9	95.3	12.5	86.3	318.6	78.3	189.9	0.0	575.5	0.0	34.5	100.0	225.3	100.0
DRC	46.1	95.7	1725.4	0.0	1636.2	84.0	1013.0	20.4	163.4	100.0	336.5	100.0	4297.0	20.4
Eswatini	49.0	90.7	2,236.0	0.0	2350.4	90.2	671.6	0.0	8795.5	42.9	7693.3	0.0	10,342.2	0.0
Lesotho	22.8	97.6	795.3	0.3	1302.8	85.2	677.5	0.0	8941.9	0.0	11,845.1	0.0	1408.6	23.3
Madagascar	46.5	85.4	733.9	11.8	1787.3	77.7	894.2	2.1	244.4	56.3	108.1	100.0	1108.4	100.0
Malawi	89.4	88.8	1223.5	29.5	3526.6	85.7	1202.7	0.0	3309.8	61.7	1153.1	100.0	15,170.2	31.9
Mauritius	75.8	74.0	468.6	0.0	484.5	79.9	454.8	0.0	6978.1	0.0	8379.1	0.0	6861.6	0.0
Mozambique	66.8	94.2	244.7	91.2	1928.7	75.0	70.9	42.1	341.8	100.0	602.0	22.9	2707.0	22.9
Namibia	134.0	91.3	587.2	90.8	3411.9	81.7	101.4	89.7	21,239.2	0.0	23,742.4	0.0	10,490.6	0.0
Seychelles	5.8	91.0	188.7	0.0	258.6	76.1	143.5	0.0	33.4	77.6	694.2	0.0	1984.9	0.0
South Africa	5785.9	79.8	10,713.1	97.4	97,311.6	84.7	9589.2	17.6	63,626.1	48.1	33,910.2	100.0	26,489.6	100.0
Tanzania	28.4	94.7	1853.4	0.0	1911.1	85.9	1844.1	0.0	27,048.4	0.0	35,267.8	0.0	28,062.0	0.0
Zambia	150.0	89.7	429.3	94.9	3701.5	77.6	429.3	19.1	6940.6	6.3	458.3	100.0	1731.8	77.1
Zimbabwe	169.9	87.3	810.4	85.8	5004.7	84.6	119.7	100.0	26,898.9	0.0	746.0	95.8	4778.9	20.8

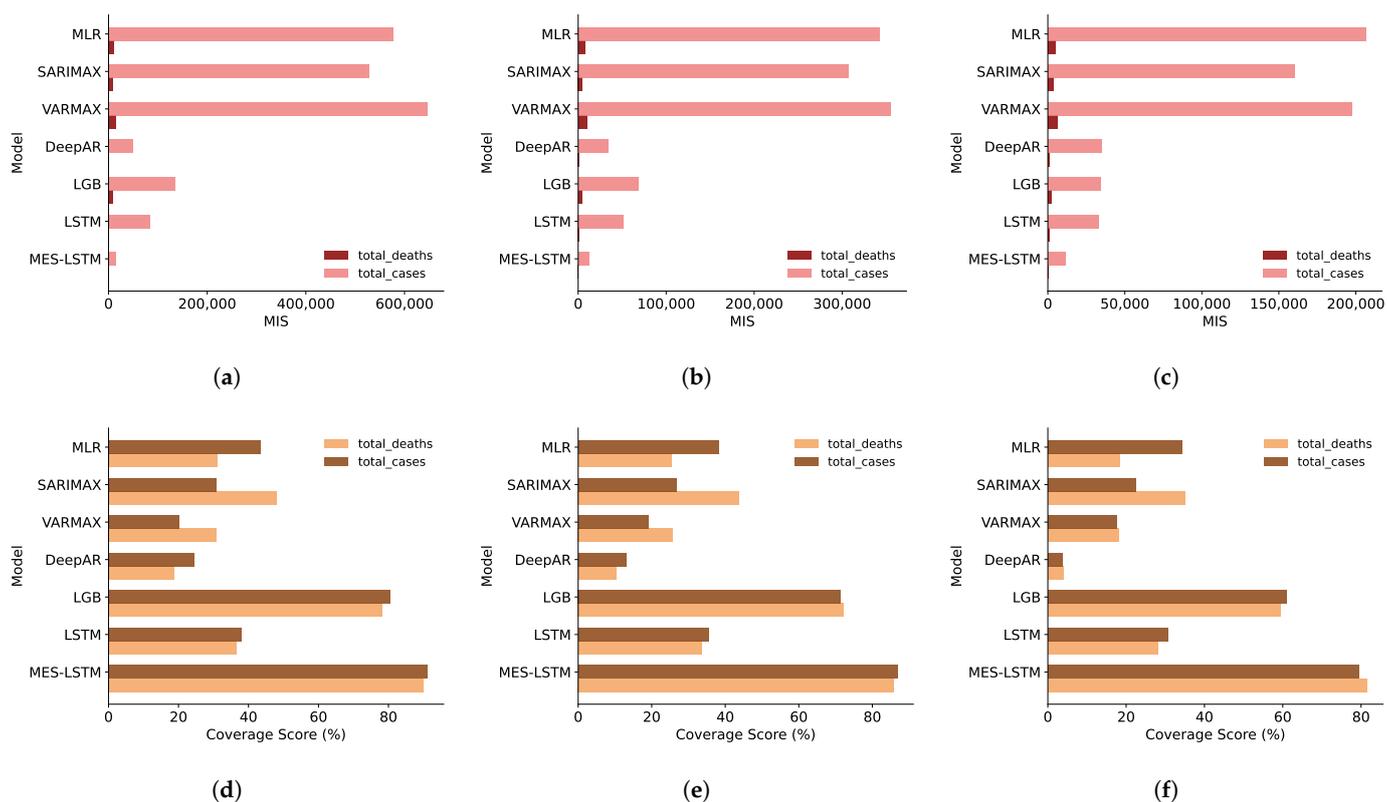


Figure 4. Prediction Interval Accuracy in the SADC Region. (a) MIS ($\alpha = 0.05$). (b) MIS ($\alpha = 0.1$). (c) MIS ($\alpha = 0.2$). (d) Coverage Score ($\alpha = 0.05$). (e) Coverage Score ($\alpha = 0.1$). (f) Coverage Score ($\alpha = 0.2$).

4.3. Forecast Performance for South Africa

After noting the relatively poor performance for South Africa, we believe this requires further probing.

We first present our observed error metrics distribution over the independent trials for all the models for both predictands in South Africa. Tables 8 and 9 show the forecasting error for *total cases* and *total deaths* respectively. The models VARMAX, SARIMAX and MLR are all deterministic, so the output for different trials is identical. As a result, there is no variation in the accuracy of these models and the standard deviation is zero. For *total cases*, VARMAX (DeepAR) has the lowest (highest) average sMAPE, and MES-LSTM (SARIMAX) has the lowest (highest) average RMSE. For *total deaths*, MES-LSTM has the lowest average sMAPE and RMSE, while and SARIMAX has the highest figures for both.

Table 8. Forecast Accuracy Distribution for All Trials for *total cases* in South Africa.

	MES-LSTM		LSTM		LGB		DeepAR		VARMAX		SARIMAX		MLR	
	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE
mean	5.0	26,979.2	4.8	150,416.4	5.6	30,420.9	7.8	212,612.0	2.8	87,376.0	8.2	265,998.0	4.0	118,139.5
std	1.1	5641.0	3.8	121,306.0	1.3	12,885.4	0.1	1192.2	0.0	0.0	0.0	0.0	0.0	0.0

Table 9. Forecast Accuracy Distribution for All Trials for *total deaths* in South Africa.

	MES-LSTM		LSTM		LGB		DeepAR		VARMAX		SARIMAX		MLR	
	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE	sMAPE	RMSE
mean	4.9	446.2	7.8	7902.4	5.9	2773.4	10.4	8515.8	10.7	10,806.6	16.1	15,438.9	6.5	6276.2
std	1.2	106.0	5.3	5578.8	2.2	1259.7	0.1	61.2	0.0	0.0	0.0	0.0	0.0	0.0

MES-LSTM and SARIMAX results generally being on opposite sides of the spectrum is also evidenced by the box and whisker plots in Figure 5.

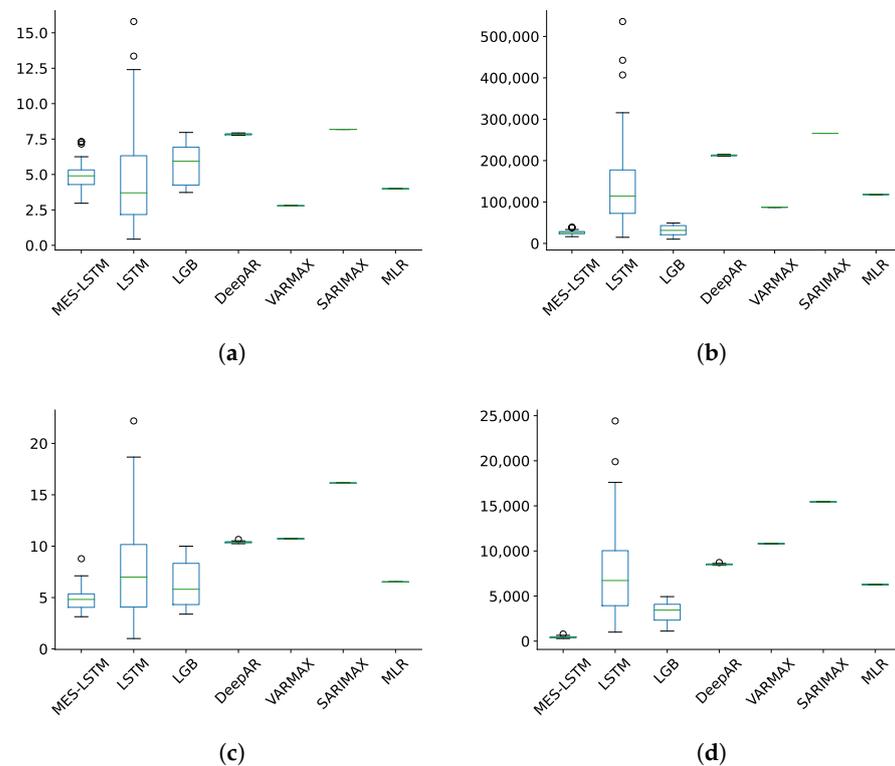


Figure 5. Forecast Accuracy Distribution Boxplots for All Trials in South Africa. (a) sMAPE for total cases. (b) RMSE for total cases. (c) sMAPE for total deaths. (d) RMSE for total deaths.

The results from the LSTM are far more variable for each trial than MES-LSTM. LSTM seems competitive to our model for some trials but shows little consistency for multiple independent runs. The LSTM's relatively poor and oft-inconsistent performance with COVID-19 modeling agrees with observations by other researchers (as detailed in Section 1.1).

When viewed together with the bar graphs presented in Figure 6, one possible interpretation of all the forecast results is a ranking of the best performance for South Africa in increasing order as SARIMAX, VARMAX, LSTM, DeepAR, MLR, LGB, and MES-LSTM. MES-LSTM shows consistent outperformance over (or at least competitiveness to) the benchmarks. Our model also presents results with a tight distribution, second in terms of tightness to DeepAR (from the probabilistic forecasts models). Even in instances where outliers are present in the forecast distribution of our model, these outliers are still very close to the core distribution. This tendency reaffirms the robustness of our model when it comes to producing accurate forecasts.

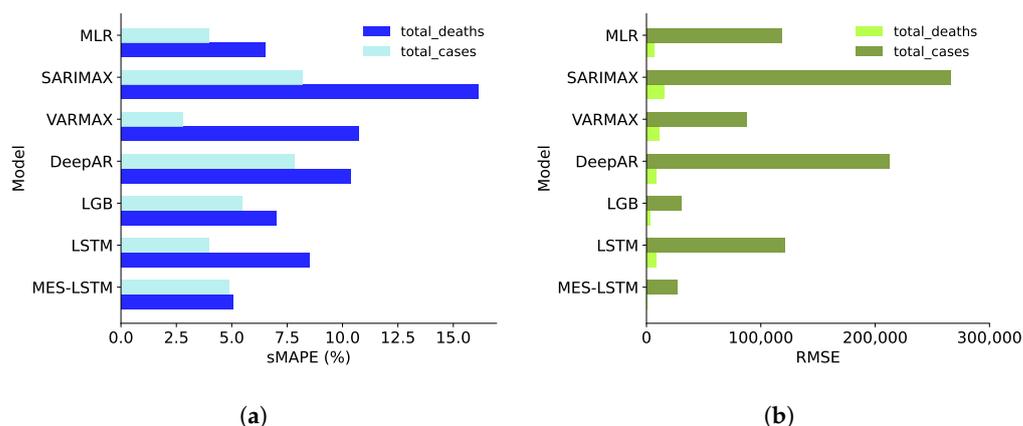


Figure 6. Forecast Accuracy in South Africa. (a) sMAPE. (b) RMSE.

4.4. Prediction Interval Performance for South Africa

We turn our attention to the prediction interval accuracy for South Africa. Tables 10 and 11 show the prediction interval summary statistics at the $\alpha = 0.05$ level of significance for our respective predictands. We note that the distribution for MES-LSTM is tight and this characteristic persists throughout all significance levels.

Examining the bar graphs in Figure 7, we note that even though our model's coverage score is outperformed in some instances, it is not by a great margin. Furthermore, our model has the benefit of the lowest MIS, indicating the tightest prediction intervals. The second benefit is consistency, i.e., where other models may only have decent coverage for one predictand, our model has a consistent level of coverage across all the predictands. Moreover, we can depend on the significance level, with stricter levels leading to marginally higher coverage. This coverage-alpha dependence from MES-LSTM is important as it also speaks to the robustness of our model. In contrast, DeepAR for instance, goes from almost complete coverage ($\alpha = 0.05$, $\alpha = 0.1$) for *total cases* to none ($\alpha = 0.02$) and MLR presents close to perfect coverage for *total cases* throughout. This does not instill much trust in the statistical method's uncertainty quantification.

Overall, all models perform worse for the SADC region than for South Africa. The worse performance is due to more variability introduced in the data and thus a higher level of uncertainty. This variability can be seen when comparing the Coverage scores from Figure 7 to those we presented in Figure 4.

As an additional point, we perform a one-sided t-test to check whether or not the distribution of our model's forecast and prediction interval results are significantly better than those produced by the benchmark models. Concisely, we compare MES-LSTM, to each of the other models in turn. The t-test results are presented in Tables 12–15 truncated to three decimal places. Our null hypothesis is H_0 : The benchmark models produce forecasts that are more accurate and prediction intervals superior to those produced by MES-LSTM.

Table 10. Prediction Interval Accuracy Distribution for All Trials for *total cases* in South Africa ($\alpha = 0.05$).

	MES-LSTM		LSTM		LGB		DeepAR		VARMAX		SARIMAX		MLR	
	MIS	Coverage	MIS	Coverage	MIS	Coverage	MIS	Coverage	MIS	Coverage	MIS	Coverage	MIS	Coverage
mean	193,075.6	80.7	348,860.7	97.7	1,418,713.2	86.6	239,009.3	100.0	740,131.9	100.0	1,061,162.0	100.0	828,447.2	100.0
std	13,881.3	11.2	76,242.6	13.7	0.0	30.2	1674.9	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Table 11. Prediction Interval Accuracy Distribution for All Trials for *total deaths* in South Africa ($\alpha = 0.05$).

	MES-LSTM		LSTM		LGB		DeepAR		VARMAX		SARIMAX		MLR	
	MIS	Coverage	MIS	Coverage	MIS	Coverage	MIS	Coverage	MIS	Coverage	MIS	Coverage	MIS	Coverage
mean	5785.9	79.8	10,713.1	97.4	97,311.6	84.7	9589.2	17.6	63,626.1	48.1	33,910.2	100.0	26,489.6	100.0
std	1714.1	15.1	3448.0	15.6	0.0	28.5	723.9	6.5	0.0	0.0	0.0	0.0	0.0	0.0

Table 12. Student’s t-test for Forecast RMSE (H_0 : Benchmark Forecasts Outperform MES-LSTM).

	Total Cases						Total Deaths					
	LSTM	LGB	DeepAR	VARMAX	SARIMAX	MLR	LSTM	LGB	DeepAR	VARMAX	SARIMAX	MLR
statistic	-6.014	-1.448	-190.477	-63.342	-250.673	-95.605	-7.906	-10.891	-389.954	-578.058	-836.515	-325.283
p-value	0.000	0.077	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Table 13. Student’s t-test for Forecast sMAPE (H_0 : Benchmark Forecasts Outperform MES-LSTM).

	Total Cases						Total Deaths					
	LSTM	LGB	DeepAR	VARMAX	SARIMAX	MLR	LSTM	LGB	DeepAR	VARMAX	SARIMAX	MLR
statistic	0.270	-1.883	-15.455	12.180	-17.383	5.616	-3.216	-2.422	-27.328	-29.110	-56.140	-8.166
p-value	0.606	0.032	0.000	1.000	0.000	1.000	0.001	0.009	0.000	0.000	0.000	0.000

Table 14. Student’s t-test for Prediction Interval MIS (H_0 : Benchmark Prediction Intervals Outperform MES-LSTM).

	Total Cases						Total Deaths					
	LSTM	LGB	DeepAR	VARMAX	SARIMAX	MLR	LSTM	LGB	DeepAR	VARMAX	SARIMAX	MLR
statistic	-5.328	-25.791	-2.851	-22.095	-34.220	-25.119	-2.299	-37.924	-10.070	-40.222	-20.029	-14.501
p-value	0.000	0.000	0.004	0.000	0.000	0.000	0.013	0.000	0.000	0.000	0.000	0.000

Table 15. Student’s t-test for Prediction Interval Coverage (H_0 : Benchmark Prediction Intervals Outperform MES-LSTM).

	Total Cases						Total Deaths					
	LSTM	LGB	DeepAR	VARMAX	SARIMAX	MLR	LSTM	LGB	DeepAR	VARMAX	SARIMAX	MLR
statistic	-1.8770	1.6106	0.2267	-2.7459	-2.7459	-2.7459	-2.8639	-0.2274	8.3710	3.5521	-3.5613	-3.5613
p-value	0.9666	0.0567	0.4110	0.9952	0.9952	0.9952	0.9968	0.5895	0.0000	0.0006	0.9994	0.9994

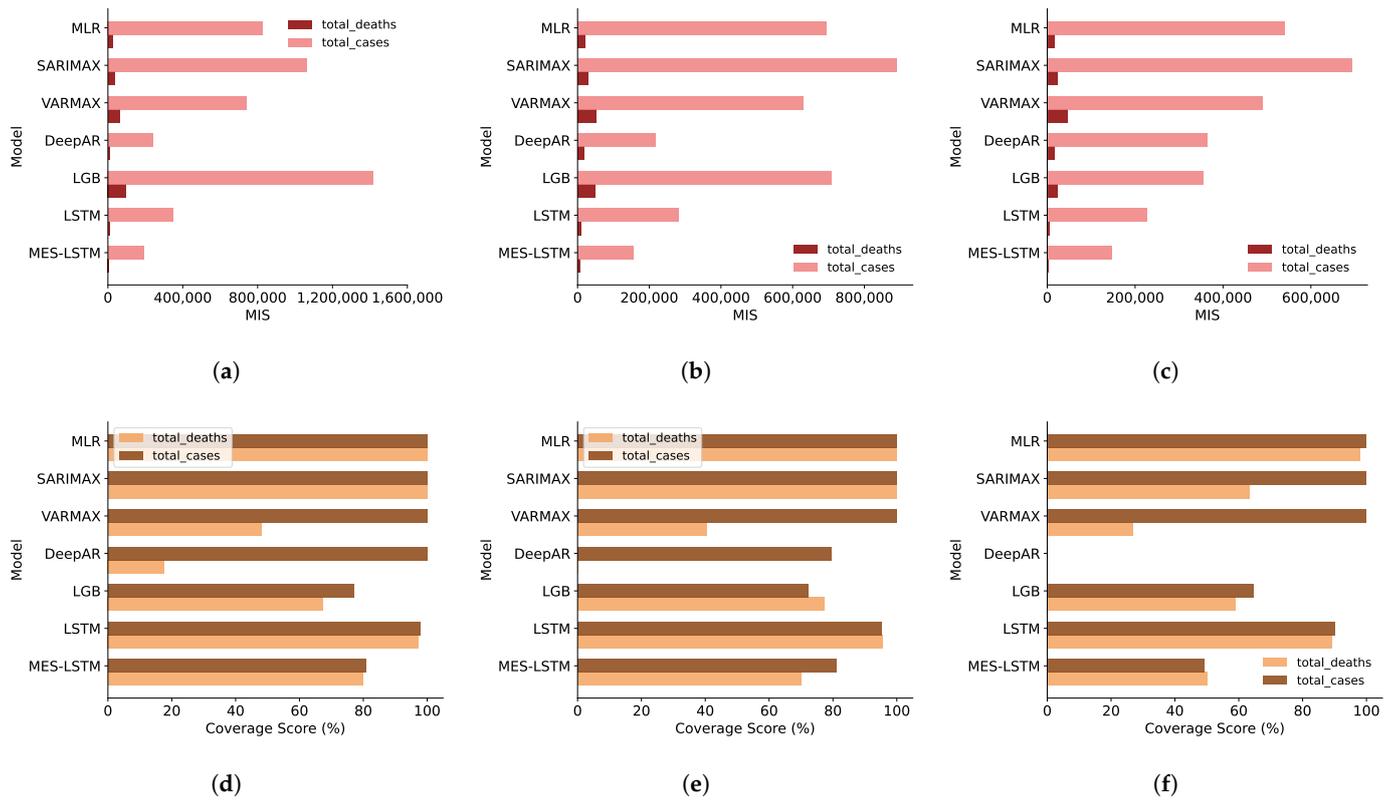


Figure 7. Prediction Interval Accuracy in South Africa. (a) MIS ($\alpha = 0.05$). (b) MIS ($\alpha = 0.1$). (c) MIS ($\alpha = 0.2$). (d) Coverage Score ($\alpha = 0.05$). (e) Coverage Score ($\alpha = 0.1$). (f) Coverage Score ($\alpha = 0.2$).

We note in terms of forecasting accuracy (Tables 12 and 13) in each instance that the null hypothesis w.r.t. both predictands is rejected at the $\alpha = 0.01$ level of significance, except for LSTM, LGB, VARMAX, and MLR. This may seem contrary to previously presented results in this section, but the box plots in Figure 5 explain this. Although the core and lower distributions may overlap (the box and bottom whisker) for the distributions in questions compared to MES-LSTM, the higher levels of inaccuracy are only reported for the benchmarks, i.e., the upper whiskers for the benchmarks peak higher than MES-LSTM.

For the prediction interval performance (Tables 14 and 15) we note that we are able to reject the null hypothesis at the $\alpha = 0.01$ level of significance for MIS in all instances. Finally, we are unable to conclude that our model’s Coverage Score is not outperformed except perhaps by DeepAR and VARMAX for *total deaths*. Again, the t-test results are consistent with the results previously discussed.

We also conduct a Diebold-Mariano (DM, [55]) test to compare the out-of-sample predictive skill of MES-LSTM to the benchmark models (Table 16). The test is configured to use Mean Absolute prediction Error (MAPE, [56]) and the null hypothesis is H_0 : There is no significant difference in the accuracy of the competing forecasts. We reject the null hypothesis w.r.t. both predictands at the $\alpha = 0.01$ level of significance.

Table 16. Diebold-Mariano test for Forecast accuracy (H_0 : Benchmark Forecasts Outperform MES-LSTM).

	Total Cases						Total Deaths					
	LSTM	LGB	DeepAR	VARMAX	SARIMAX	MLR	LSTM	LGB	DeepAR	VARMAX	SARIMAX	MLR
statistic	-19.817	-277.776	-144.279	-12.665	-16.160	-8.048	-5.567	-374.711	-112.500	-7.066	-11.817	-43.647
p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

We conclude this chapter with a deeper look into the effects of introducing more variability on our model’s performance.

4.5. Effects of Variability on Model Performance

Our methodology suggests that the introduction of more variability into the data impacts both forecasting accuracy and prediction interval construction. We focus here specifically on MES-LSTM. We note from Figures 8 and 9 that in most instances, the countries in which the predictions are least (most) accurate are also the countries in which the prediction intervals are widest (narrowest). This direct correlation reinforces the consistency of our model. We have normalized the MIS in Figure 9 before plotting the maps for ease of interpretation.

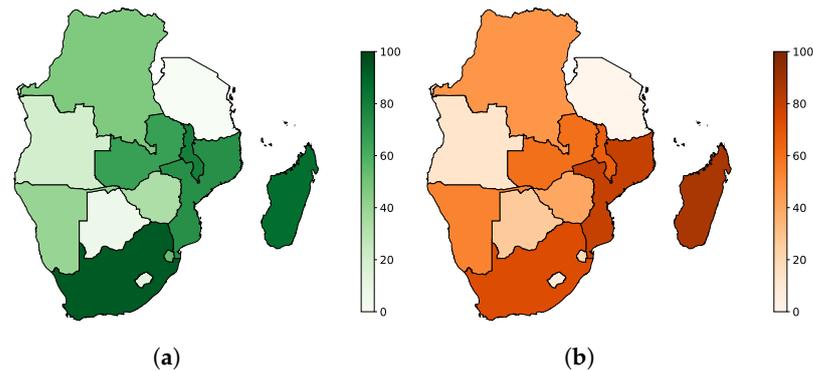


Figure 8. MES-LSTM Forecast Accuracy Ranked for Each Country in the SADC Region. (a) sMAPE for total cases. (b) sMAPE for total deaths.

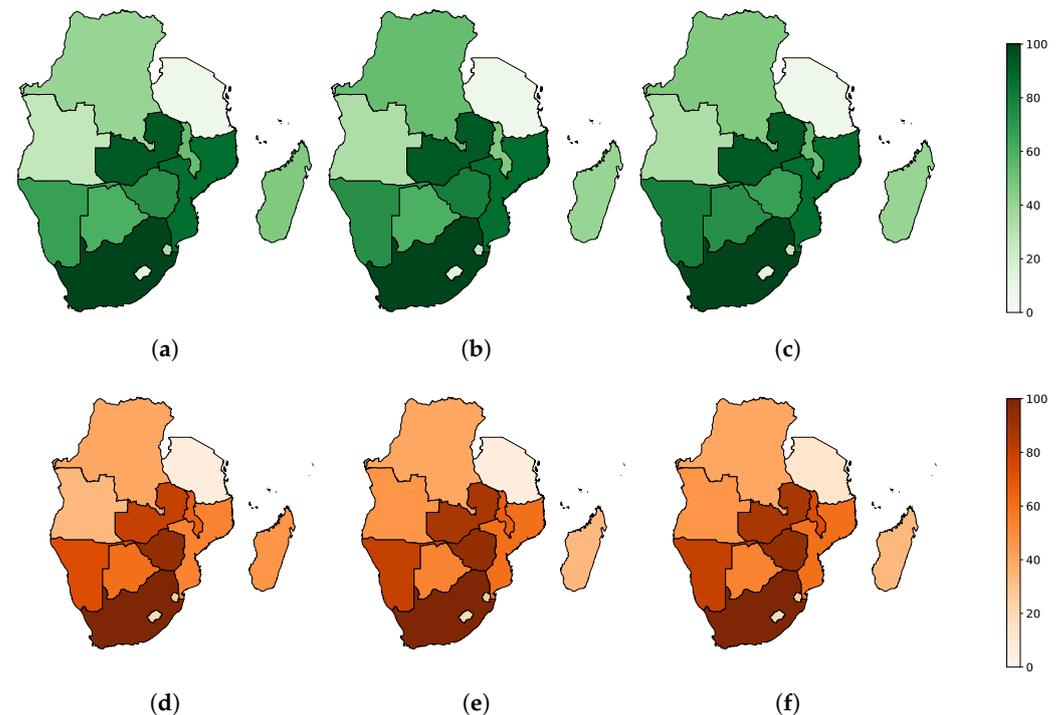


Figure 9. MES-LSTM Prediction Interval Accuracy (normalized MIS) Ranked for Each Country in the SADC Region. (a) total cases ($\alpha = 0.05$). (b) total cases ($\alpha = 0.1$). (c) total cases ($\alpha = 0.2$). (d) total deaths ($\alpha = 0.05$). (e) total deaths ($\alpha = 0.1$). (f) total deaths ($\alpha = 0.2$).

In terms of Coverage, MES-LSTM is in instances outperformed by the benchmark methods. This Coverage is an area where the model can be improved. However, there are some crucial areas where the model improves on the skill of its counterparts. Our methodology suggests that the hybrid MES-LSTM is indeed able to outperform statistical methods and deep learning techniques at both forecasting tasks and prediction interval

construction for morbidity and mortality data with exogenous factors. In terms of the overall prediction error, there is a significant improvement over the benchmark models considered. Our model reports forecast consistently within a tight range for multiple independent trials. We also note that MES-LSTM offers the narrowest prediction intervals for all the predictands for all geographical regions at all the significance levels considered.

5. Conclusions

We introduce a hybrid model, MES-LSTM, for multivariate prediction and forecast uncertainty quantification and apply it to morbidity and mortality data with exogenous factors. The univariate counterpart, Smyl's ES-RNN [2], has been shown to perform well in the univariate case, outperforming both pure machine learning and pure statistical methods. We hypothesise that our multivariate extension also outperforms statistical and machine learning models at both forecasting tasks and constructing prediction intervals. With the methodology presented and the aggregated datasets considered, MES-LSTM generally improves upon the skill of its classical probabilistic and pure deep learning counterparts.

MES-LSTM shows consistent outperformance with forecast accuracy and the MIS of the prediction intervals constructed at all significance levels considered. There remains room for improvement when it comes to the coverage of the prediction intervals.

In this paper we mostly limit our attention to the multivariate setting (except in Section 2.1 where we use the univariate exposition as a building block before introducing our model). Future work may include running our model on univariate datasets as well. The benchmarks can also be applied with ease since both SARIMAX and VARMAX are univariate models if no exogenous inputs are declared, MLR is fit using OLS, and gradient boosting techniques can also be applied to univariate data.

Applying the kind of techniques discussed in this paper to varied data with expedience is still a major limitation in related research. For example, the M5 Competitions although arguably among the most important forecast competitions globally, only considers retail units for one supermarket chain. Future work may also include applying our model to more multivariate datasets from a broader cross-section of industries and applications beyond morbidity and mortality modeling.

We motivate our choice of benchmark models in Section 2.5, but we also consider future work comparing MES-LSTM against more deep learning models such as convolutional, attention and transformer models. We also consider future applications where an adaptation or extension of our model can be applied, such as in multivariate anomaly detection. Explainability and interpretability of our model and its results is also an avenue worth considering for future work.

Supplementary Materials: The following supporting information can be downloaded at: https://github.com/zulucomputer/MES_LSTM, code repository: MES-LSTM.

Author Contributions: Conceptualization, T.M. and T.L.v.Z.; methodology, T.M. and T.L.v.Z.; software, T.M.; validation, T.M.; formal analysis, T.M. and T.L.v.Z.; investigation, T.M. and T.L.v.Z.; resources, T.M.; data curation, T.M.; writing—original draft preparation, T.M.; writing—review and editing, T.M. and T.L.v.Z.; visualization, T.M. and T.L.v.Z.; supervision, T.L.v.Z.; project administration, T.M.. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are openly available in Our World in Data repository at <https://github.com/owid/covid-19-data/tree/master/public/data> (accessed on 23 October 2021), references [53,54].

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Lim, B.; Arik, S.Ö.; Loeff, N.; Pfister, T. Temporal Fusion Transformers for interpretable multi-horizon time series forecasting. *Int. J. Forecast.* **2021**, *37*, 1748–1764. [[CrossRef](#)]
2. Smyl, S. A hybrid method of exponential smoothing and recurrent neural networks for time series forecasting. *Int. J. Forecast.* **2020**, *36*, 75–85. [[CrossRef](#)]
3. Hyndman, R.J.; Koehler, A.B.; Snyder, R.D.; Grose, S. A state space framework for automatic forecasting using exponential smoothing methods. *Int. J. Forecast.* **2002**, *18*, 439–454. [[CrossRef](#)]
4. Jaeger, H. *The “Echo State” Approach to Analysing and Training Recurrent Neural Networks*; GMD Report 148; GMD—German National Research Institute for Computer Science: Hanover, Germany, 2001.
5. Hochreiter, S.; Schmidhuber, J. Long Short-Term Memory. *Neural Comput.* **1997**, *9*, 1735–1780. [[CrossRef](#)]
6. Hochreiter, S.; Bengio, Y.; Frasconi, P. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies. In *A Field Guide to Dynamical Recurrent Neural Networks*; Kolen, J.F., Kremer, S.C., Eds.; IEEE Press: Piscataway, NJ, USA, 2001.
7. Makridakis, S.; Spiliotis, E.; Assimakopoulos, V. The M4 Competition: 100,000 time series and 61 forecasting methods. *Int. J. Forecast.* **2020**, *36*, 54–74. [[CrossRef](#)]
8. Redd, A.; Khin, K.; Marini, A. Fast ES-RNN: A GPU Implementation of the ES-RNN Algorithm. *arXiv* **2019**, arXiv:1907.03329.
9. Beeram, S.R.; Kuchibhotla, S. Time Series Analysis on Univariate and Multivariate Variables: A Comprehensive Survey. In *Communication Software and Networks*; Satapathy, S.C., Bhateja, V., Ramakrishna Murty, M., Gia Nhu, N., Kotti, J., Eds.; Springer: Singapore, 2021; pp. 119–126.
10. Bharathi Priya, C.; Arulanand, N. Univariate and multivariate models for Short-term wind speed forecasting. *Mater. Today Proc.* **2021**. [[CrossRef](#)]
11. Olkin, I.; Sampson, A. Multivariate Analysis: Overview. In *International Encyclopedia of the Social and Behavioral Sciences*; Smelser, N.J., Baltes, P.B., Eds.; Pergamon: Oxford, UK, 2001; pp. 10240–10247.
12. Jones, R.H. Exponential Smoothing for Multivariate Time Series. *J. R. Stat. Soc. Ser. B Methodol.* **1966**, *28*, 241–251. [[CrossRef](#)]
13. Enns, P.G.; Machak, J.A.; Spivey, W.A.; Wroblewski, W.J. Forecasting Applications of an Adaptive Multiple Exponential Smoothing Model. *Manag. Sci.* **1982**, *28*, 1035–1044. [[CrossRef](#)]
14. Trigg, D.W.; Leach, A.G. Exponential Smoothing with an Adaptive Response Rate. *OR* **1967**, *18*, 53–59. [[CrossRef](#)]
15. Harvey, A.C. Analysis and Generalisation of a Multivariate Exponential Smoothing Model. *Manag. Sci.* **1986**, *32*, 374–380. [[CrossRef](#)]
16. Pfeffermann, D.; Allon, J. Multivariate exponential smoothing: Method and practice. *Int. J. Forecast.* **1989**, *5*, 83–98. [[CrossRef](#)]
17. Tan, F. Regression analysis and prediction using LSTM model and machine learning methods. *J. Phys. Conf. Ser.* **2021**, *1982*, 012013. [[CrossRef](#)]
18. Hu, Y.; O’Donncha, F.; Palmes, P.; Burke, M.; Filgueira, R.; Grant, J. A spatio-temporal LSTM model to forecast across multiple temporal and spatial scales. *arXiv* **2021**, arXiv:2108.11875.
19. Kirbaş, I.; Sözen, A.; Tuncer, A.D.; Kazancıoğlu, F.Ş. Comparative analysis and forecasting of COVID-19 cases in various European countries with ARIMA, NARNN and LSTM approaches. *Chaos Solitons Fractals* **2020**, *138*, 110015. [[CrossRef](#)]
20. Ibrahim, M.; Jemei, S.; Wimmer, G.; Hissel, D. Nonlinear autoregressive neural network in an energy management strategy for battery/ultra-capacitor hybrid electrical vehicles. *Electr. Power Syst. Res.* **2016**, *136*, 262–269. [[CrossRef](#)]
21. Chandra, R.; Jain, A.; Chauhan, D.S. Deep learning via LSTM models for COVID-19 infection forecasting in India. *arXiv* **2021**, arXiv:2101.11881.
22. Chimmula, V.K.R.; Zhang, L. Time series forecasting of COVID-19 transmission in Canada using LSTM networks. *Chaos Solitons Fractals* **2020**, *135*, 109864. [[CrossRef](#)]
23. Shahid, F.; Zameer, A.; Muneeb, M. Predictions for COVID-19 with deep learning models of LSTM, GRU and Bi-LSTM. *Chaos Solitons Fractals* **2020**, *140*, 110212. [[CrossRef](#)] [[PubMed](#)]
24. Chung, J.; Gülçehre, Ç.; Cho, K.; Bengio, Y. Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling. In Proceedings of the NIPS 2014 Deep Learning and Representation Learning Workshop, Montreal, QC, Canada, 12 December 2014.
25. Mathonsi, T.; van Zyl, T.L. Multivariate Anomaly Detection based on Prediction Intervals Constructed using Deep Learning. *arXiv* **2021**, arXiv:2110.03393.
26. Hu, M.J.C.; Root, H.E. *Application of the Adaline System to Weather Forecasting*; Technical Report 6775-1; Stanford Electronic Laboratories: Stanford, CA, USA, 1964.
27. Mathonsi, T.; v. Zyl, T.L. Prediction Interval Construction for Multivariate Point Forecasts Using Deep Learning. In Proceedings of the 2020 7th International Conference on Soft Computing Machine Intelligence (ISCMI), Stockholm, Sweden, 14–15 November 2020; pp. 88–95.
28. Oreshkin, B.N.; Carpov, D.; Chapados, N.; Bengio, Y. N-BEATS: Neural basis expansion analysis for interpretable time series forecasting. *arXiv* **2020**, arXiv:1905.10437.
29. Olivares, K.G.; Challu, C.; Marcjasz, G.; Weron, R.; Dubrawski, A. Neural basis expansion analysis with exogenous variables: Forecasting electricity prices with NBEATSx. *arXiv* **2021**, arXiv:2104.05522.
30. Makridakis, S.; Spiliotis, E. The M5 Competition and the Future of Human Expertise in Forecasting. *Foresight Int. J. Appl. Forecast.* **2021**, *60*, 33–37.

31. Makridakis, S.; Spiliotis, E.; Assimakopoulos, V. The M5 competition: Background, organization, and implementation. *Int. J. Forecast.* **2021**. [[CrossRef](#)]
32. Wen, Y.; Vicol, P.; Ba, J.; Tran, D.; Grosse, R. Flipout: Efficient Pseudo-Independent Weight Perturbations on Mini-Batches. In Proceedings of the International Conference on Learning Representations, Vancouver, BC, Canada, 30 April–3 May 2018.
33. Blundell, C.; Cornebise, J.; Kavukcuoglu, K.; Wierstra, D. Weight Uncertainty in Neural Networks. In Proceedings of the 32nd International Conference on Machine Learning—Volume 37. JMLR.org, 2015, ICML'15, Lille, France, 7–9 July 2015; pp. 1613–1622.
34. Joyce, J.M., Kullback-Leibler Divergence. In *International Encyclopedia of Statistical Science*; Springer: Berlin/Heidelberg, Germany, 2011; pp. 720–722. [[CrossRef](#)]
35. Dillon, J.V.; Langmore, I.; Tran, D.; Brevdo, E.; Vasudevan, S.; Moore, D.A.; Patton, B.; Alemi, A.A.; Hoffman, M.; Saurous, R. TensorFlow Distributions. *arXiv* **2017**, arXiv:1711.10604.
36. Gal, Y.; Ghahramani, Z. Dropout As a Bayesian approximation: Representing Model Uncertainty in Deep Learning. In Proceedings of the 33rd International Conference on Machine Learning. JMLR.org, ICML'16, New York, NY, USA, 19–24 June 2016; Volume 48, pp. 1050–1059.
37. Davison, A.C.; Hinkley, D.V. *Bootstrap Methods and Their Application*; Cambridge University Press: New York, NY, USA, 2013.
38. Hesterberg, T. What Teachers Should Know about the Bootstrap: Resampling in the Undergraduate Statistics Curriculum. *Am. Stat.* **2014**, *69*, 371–386. [[CrossRef](#)]
39. Lever, J.; Krzywinski, M.; Altman, N. Points of Significance: Model selection and overfitting. *Nat. Methods* **2016**, *13*, 703–704. [[CrossRef](#)]
40. Akaike, H. Information theory and an extension of the maximum likelihood principle. In *Second International Symposium on Information Theory*; Petrov, B.N.; Csaki, F., Eds.; Akadémiai Kiado: Budapest, Hungary, 1973; pp. 267–281.
41. Matthews, D.E., Multiple Linear Regression. In *Encyclopedia of Biostatistics*; American Cancer Society: Atlanta, GA, USA, 2005; Chapter 5, pp. 119–133.
42. Kingma, D.P.; Ba, J. Adam: A Method for Stochastic Optimization. In Proceedings of the 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, 7–9 May 2015.
43. Makridakis, S.; Hibon, M. The M3-Competition: Results, conclusions and implications. *Int. J. Forecast.* **2000**, *16*, 451–476. [[CrossRef](#)]
44. Koehler, A. Commentaries on the M3-Competition. *Int. J. Forecast.* **2001**, *17*, 537–584.
45. Goodwin, P.; Lawton, R. On the asymmetry of the symmetric MAPE. *Int. J. Forecast.* **1999**, *15*, 405–408. [[CrossRef](#)]
46. Gneiting, T.; Raftery, A.E. Strictly Proper Scoring Rules, Prediction, and Estimation. *J. Am. Stat. Assoc.* **2007**, *102*, 359–378. [[CrossRef](#)]
47. Hannan, E.J.; Deistler, M. The Statistical Theory of Linear Systems. *Econom. Theory* **1992**, *8*, 135–143.
48. Arunraj, N.; Ahrens, D.; Fernandes, M. Application of SARIMAX Model to Forecast Daily Sales in Food Retail Industry. *Int. J. Oper. Res. Inf. Syst.* **2016**, *7*, 1–21. [[CrossRef](#)]
49. Salinas, D.; Flunkert, V.; Gasthaus, J.; Januschowski, T. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *Int. J. Forecast.* **2020**, *36*, 1181–1191. [[CrossRef](#)]
50. Ke, G.; Meng, Q.; Finley, T.; Wang, T.; Chen, W.; Ma, W.; Ye, Q.; Liu, T.Y. LightGBM: A highly efficient gradient boosting decision tree. *Adv. Neural Inf. Process. Syst.* **2017**, *30*, 3146–3154.
51. Makridakis, S.; Spiliotis, E.; Assimakopoulos, V. The M5 Accuracy Competition: Results, Findings and Conclusions. Available online: https://www.researchgate.net/publication/344487258_The_M5_Accuracy_competition_Results_findings_and_conclusions (accessed on 23 October 2021).
52. Makridakis, S.; Spiliotis, E.; Assimakopoulos, V.; Chen, Z.; Gaba, A.; Tsetlin, I.; Winkler, R. The M5 Uncertainty competition: Results, findings and conclusions. *Int. J. Forecast.* **2021**. [[CrossRef](#)]
53. Mathieu, E.; Ritchie, H.; Ortiz-Ospina, E.; Roser, M.; Hasell, J.; Appel, C.; Giattino, C. A global database of COVID-19 vaccinations. *Nat. Hum. Behav.* **2021**, *5*, 947–953. [[CrossRef](#)]
54. Hasell, J.; Mathieu, E.; Beltekian, D.; Macdonald, B.; Giattino, C.; Ortiz-Ospina, E.; Roser, M.; Ritchie, H. A cross-country database of COVID-19 testing. *Sci. Data* **2020**, *7*, 345–347. [[CrossRef](#)]
55. Diebold, F.; Mariano, R. Comparing Predictive Accuracy. *J. Bus. Econ. Stat.* **1995**, *13*, 253–263.
56. Mean Absolute Percentage Error. In *Encyclopedia of Production and Manufacturing Management*; Swamidass, P.M., Ed.; Springer: Boston, MA, USA, 2000; pp. 462–462.