Predicting Soil Moisture for Improved Environmental Sustainability: A Multivariate Time-Series Forecasting Approach using Machine Learning

LUCAS ALBERT HARTMANS, University of Twente, The Netherlands

Soil moisture plays an essential role in the overall health of newly planted vegetation. Hence, it is crucial to develop precise soil moisture forecasting techniques that facilitate precision irrigation, effectively mitigating the risk of vegetation succumbing to drought. Recent research compared the performance of different machine learning techniques and found it to be an effective forecasting approach. This research was about extending this existing research on machine learning approaches by assessing the performance of a Long Short-Term Memory Recurrent Neural Network using multivariate time-series data of soil moisture sensors which were subject to manual irrigation complemented with meteorological data from weather stations within the Netherlands. The performance was evaluated using MAE, MSE, RMSE and MAPE and showed compelling results, with an average MAPE of 3.41% for a 1-day horizon, 5.85% for a 3-day horizon, and 10.09% for a 7-day horizon.

Additional Key Words and Phrases: Soil moisture, prediction, forecasting, machine learning, LSTM, multivariate, timeseries, RNN.

1 INTRODUCTION

Drastic changes are needed in order to reduce green house gasses to limit global warming to 1.5 degrees Celsius [4]. As a result many governmental policy makers aim, besides many other measurements, to increase canopy coverage (i.e. proportion of ground surface covered by vegetation material) of municipalities to 25-30%. Increasing canopy coverage as a means to limit greenhouse gases is sought to be important because at the one hand trees absorb CO2 through photosynthesis and at the other hand provide shade cooling down cities between 10 and 15 degrees Celsius [5, 12].

To meet the goal of higher canopy coverage, municipalities are increasingly planting new trees. In the Netherlands 1.6. million new trees have been planted in 2020 alone, twice as much as in previous years [13]. Importantly, newly planted trees require extensive health maintenance within the first three years to survive. 90% of these newly planted trees that do not survive within these three years suffered form too dry or wet soil moisture [10].

To help monitoring the soil moisture, company ConnectedGreen B.V. offers wireless moisture sensors from which the data can be viewed within their web portal. Currently over 3000 of their sensors are being used in the Netherlands and help green managers to make data-driven decisions about whether a tree requires additional watering.

However, the data only gives insight on historical and current soil moisture values which makes it hard to make the right, most sustainable and efficient decisions on when to water the trees.

To improve the decision making process of green managers on the irrigation scheduling of trees, accurate soil moisture forecasting is desired such that likely future developments of soil moisture influenced by meteorological factors [6] can be taken into account.

First soil moisture forecasting attempts were made using physical models of weather forecast models that based their predictions on calculations of the simulations. These calculations however did not yield high accuracy when validated with real observations. Possibly because of the lack of real-time and historical data of soil moisture. Later and more recent research, after the rise of IoT [11], approached the problem as a time-series problem and applied statistical methods as well as machine learning methods. The statistical methods did not perform well because of the non-linear characteristics of soil moisture and the machine learning approaches either used datasets of rural areas that were not subject to any manual watering or did not use multivariate data such as meteorological factors.

This research aims to fill the gap of training and evaluating the performance of a Long Short-Term Memory (i.e. LSTM) Recurrent Neural Network on forecasting soil moisture using multivariate data that includes meteorological variables and a sensor dataset that was subject to manual watering. The performance will be evaluated by calculating the Mean Absolute Error (i.e. MAE) and the performance is expected to be higher than previous attempts given the multivariate approach and quality of the dataset.

2 PROBLEM STATEMENT

Currently, green managers are already basing their decisions on watering vegetation based on the data of their soil moisture sensors. However, this data only gives an impression of the historical and current soil moisture without incorporating the future. Without having a forecast of the soil moisture values, vegetation are getting inaccurate irrigation (e.g. too much, late or little additional watering) besides inefficient watering rounds that are being driven as a result. Therefore, to solve this problem, accurate forecasting of soil moisture sensor data is needed in order to enable in-time adequate anticipation of moisture change. Recent research has shown promising forecasting results by utilizing a Long Short-Term Memory (i.e. LSTM) artifical neural network, however the dataset used included only data of sensors placed in rural areas that were not subject to manual watering. Furthermore, the LSTM model was trained solely using uni-variate data and did not incorporate any (historical) meteorological data on which soil moisture is highly dependent [8].

2.1 Objectives and Goals

Given the problem statement and limitations of recent research presented above, this paper has the following objective:

- Develop predictive LSTM model to forecast soil moisture sensor data using multivariate time-series data.

The main objective of this paper entails the following goals:

TScIT 39, July 8, 2023, Enschede, The Netherlands

(1) Identify what features are most informative to include with training the LSTM model
(2) Assess the accuracy of the LSTM model in forecasting soil moisture data at different depths and forecasting time-steps
(3) Determine most optimal hyper-parameters of the LSTM model to maximize the accuracy of soil moisture forecasting
(4) Investigate the effect of different forecasting horizons and sensor depths on the accuracy of the model
(5) Draw conclusions and limitations on the effectiveness of utilizing LSTM models for soil moisture forecasting

2.2 Research Question

The problem statement and the objective of this research leads to the following research question:

- How accurate can we forecast future soil moisture sensor data with multivariate time series data utilizing a LSTM model?

The following sub-questions will guide us in answering the main research question:

(1) What features are most informative to include as input for training the LSTM model?
(2) What is the effect of the forecasting horizon and depth of sensor on the accuracy of the LSTM model
(3) How does the choice of hyper-parameters affect the accuracy of the LSTM model at different depths and forecasting time-steps?

3 RELATED WORK

This section summarizes related work done on forecasting or predicting soil moisture from the early past to the present.

In the early past, soil moisture prediction was approached by empirical models and simplistic models. Statistical analysis such as regression techniques were used to find relationships between soil moisture and meteorological variables. This research showed us the influence of meteorological factors such as precipitation on soil moisture [14].

With the advancements of remote sensing technologies in the 1980s, researchers began to predict soil moisture using satellite-based microwave radiation sensors that measured radiation emitted by the earth's surface which was found to be influenced by soil moisture [7]. However, this prediction was limited to soil surface only.

With the rise of Internet of Things (i.e. IoT), abundance of granular historical soil moisture data was available. This opened doors for new research utilizing machine learning techniques using this data [11].

Acharya et al. compared different machine learning models for predicting soil moisture and concluded that machine learning models can perform well in predicting soil moisture. In his comparison, Recurrent Neural Networks were not taken into account and the dataset that was used contained soil moisture values in 16-day intervals and was not subject to manual irrigation [1]

More recently, Singh utilized a Long Short-Term Memory Recurrent Neural Network (i.e. LSTM-RNN) to predict soil moisture. This time-series based approach showed how to overcome the exploding/vanishing gradients problem using LSTM and had promising results. However, the dataset used was highly specific to one region in India and the model did not include any meteorological data making it univariate [8].

Khalil utilized artificial neural networks and used meteorological data as well, however the evaluation metrics showed significant deviation between test and train suggesting overfitting of the data. Furthermore, the input only went back in history for one time step causing the loss of potential temporal dependencies [3].

4 METHODOLOGY AND APPROACH

To answer the research question we approached the problem as a supervised time-series machine learning problem. We complemented the soil moisture data with meteorological data making our approach multivariate. We trained and evaluated the LSTM model at different depths (15, 30, 60 and 90cm) and forecasting horizons (1, 3 and 7 days) using predefined evaluation metrics. The more detailed approach is elaborated below.

4.1 Datasets

For training the LSTM model, two datasets have been used:

(1) Historical sensor data of ConnectedGreen B.V. (i.e. sensor dataset).
(2) Historical weather-station data of KNMI (i.e. weather dataset)

The sensor dataset included historical hourly soil moisture readings of more than 3000 sensors placed around the Netherlands from 2017 to 2023. Next to the soil moisture value per timestamp, the dataset also included general information such as soil type, sensor depth, location and vegetation type.

The weather dataset included historical daily meteorological readings of 51 different weather stations in the Netherlands from 1951 until present. The meteorological variables included 39 variables such as precipitation amount, mean temperature and relative atmospheric humidity.

4.2 Data pre-processing

In order to use the datasets as input for training the LSTM model, several pre-processing steps have been taken.

Since the weather dataset has daily values whereas the sensor dataset has hourly values we needed to sample down the sensor dataset to daily intervals in order to make them compatible. This was achieved by taking the last hourly value of each day.

Then we merged the two datasets by combining all variables of both datasets that correspond to the same timestamp. To make sure the weather data is most accurate for the sensor in question, we calculated which of the 51 weather stations had the least distance to the given sensor using the haversine formula. Once merged, the data needed to be cleaned. Missing values were replaced with the mean value of the previous and next step and erroneous data has been replaced with zero.

To select the most important features, pearson correlation coefficients has been calculated between all meteorological variables and the soil moisture variable. Variables without any significant correlation will then be omitted from the data to yield higher accuracy and decrease computation time making it more sustainable.

Furthermore, to train the LSTM model, all variables needed to be numerical. Therefore we encoded all non-numerical values such as timestamps and wind-direction using dedicated techniques to convert them to numerical values.

Finally, Min-Max scaling is applied to normalize all numerical values between zero and one in order to ensure the same magnitude for all input features. Data has been split between train and test with a ratio of 80:20 and various window sizes will be used to see how it affects the accuracy.

4.3 Machine Learning Architecture

Since the data that we are using is sequential and multivariate, we wanted to train a recurrent neural network model (ie. RNN). However, classic vanilla RNN's suffer from vanishing/exploding gradients when dealing with bigger window sizes making it less suitable for large sequential data. That is why we utilized a LSTM model since it solves this problem by introducing gating mechanisms and memory cells.

We used the architecture of a classic LSTM model consisting of three gates (Input gate, Forget gate & Output gate). The input and output layer had a 3-dimensional shape of *(samples, timesteps, features)* and *(1, forecastingHorizon, features)* respectively. The hyperparameters such as hidden layers, batch size, window size will has been evaluated and optimized using grid search.

4.4 Experiments

4.4.1 Feature importance. The first experiment was about identifying the most significant features for training our model. To determine the importance and consequently which features can be dropped we employed two complementary techniques to determine their importance.

First, we computed the Pearson correlation coefficient to determine the correlations between the timeseries data from the weather station and the moisture sensor timeseries data.

Next, based on the correlations, we selected the top 5 features having the highest negative or positive relationship, used it for training a single dummy machine learning model and applied our second feature selection technique: Integrated gradients. Integrated gradients, a technique introduced by Sundarajan et al., assesses the contribution of each feature to the predicted outcome by reversing the loss process through integration [9]. By using this technique, we gained deeper insights into feature importance and their impact on accuracy, thus aiding in addressing SQ. Also this experiment was a prerequisite for the next experiment as the result of this experiment decides what final features will be fitted into the machine learning model.

4.4.2 Experiment 2: Hyperparameter optimization. In the second experiment, the goal was to help addressing research questions SQ2 and SQ3 by implementing and executing a custom gridsearch technique and computing and comparing our proposed metrics for each search. Our analysis focused on sensor depth-based scenarios, systematically exploring 108 scenarios per sensor depth (15cm, 30cm, 60cm, 90cm) to evaluate the impact of varying hyperparameters and conditions. In total this resulted in 423 evaluated scenarios. The table below illustrates the parameters considered for gridsearch along with the corresponding range of potential values. Notably, we quickly discovered that the evaluation metrics of models with two layers performed significantly better compared to those with one or three layers. Therefore, we limited the gridsearch to two layers, reducing the number of combinations by at least 50%. Following the recommendations of Ibrahim Kandel, we limited our investigations on smaller batch sizes and hidden sizes, as they are known to improve training speed and accuracy [2]. Additionally, we gradually increased the window size in increments of 5, up to a maximum of 30 days (equivalent to one month). Each search was constrained to a maximum of 600 training epochs; nevertheless, the model's training would terminate if the evaluation metrics of the current epoch exceeded those of the previous epoch, thereby preventing the model against overfitting.

Table 1. (Hyper)parameters and Values

Parameter	Possible Values
forecasting days	[1, 3, 7]
number of layers	[2]
window size	[5, 10, ..., 30]
hidden size	[50, 75, 100]
batch size	[16, 32]

4.5 Evaluation Metrics

In order to answer our main research question on how accurate the final model is, we evaluated the performance based on several conventional forecasting evaluation metrics for all depths and forecasting horizons. The following metrics were calculated to evaluate the performance of the model:

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i|$$

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2$$

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}$$

$$\text{MAPE} = \frac{1}{n}\sum_{i=1}^{n}\left|\frac{y_i - \hat{y}_i}{y_i}\right| \times 100$$

5 RESULTS

5.1 Experiment 1

Table 2 displays a filtered subset of variables, showcasing their Pearson correlation coefficients limited to those having at minimum a weak positive or negative correlation. Out of 27 considered variables, 14 matched this criterion. It can be observed that the daily humidity (UG) has the highest positive correlation with moisture at 0.2665 while the sunshine duration (SQ) has the highest negative correlation at -0.184.

Table 2. Feature correlation

Variable	Correlation
	Moisture
Daily Temperature (TG)	0.0903
Min Temperature (TN)	0.14878
Min Temperature 10cm (T10N)	0.2043
Sunshine duration (SQ)	-0.184
Sunshine percentage (SP)	-0.1367
Global radiation (Q)	-0.1289
Daily Precipitation (RH)	0.1946
Max. Hourly precipitation (RHX)	0.0905
Min. Visibility (VVN)	-0.115
Mean cloud cover (NG)	0.1848
Daily humidity (UG)	0.2665
Max humidity (UX)	0.1426
Min humidity (UN)	0.1678
Potential evatransporation (EV24)	-0.1341

The top 5 variables that have the highest negative or positive correlation were the minimum temperature at 10cm (T10N), duration of sunshine (SQ), daily precipitation (RH), mean cloud coverage (NG) and the daily humidity (UG). When using these variables as features within our trained dummy model the feature attributions calculated by the integrated gradients can be seen in Table 3.

Table 3. Feature Attributions (IG)

Feature	Attribution (IG)
Daily Humidity (UG)	0.029203194
Daily Precipitation (RH)	0.019523483
Min Temperature 10cm (T10N)	-0.016551758
Sunshine Duration (SP)	-0.027794981
Mean Cloud Coverage (NG)	-0.021663292
Moisture (M)	0.619978998

5.2 Experiment 2

Tables 4 and 5 present the evaluation metrics obtained for each sensor depth and forecasting horizon. In Table 4, only the hyper-parameters that yielded the most accurate metrics, achieved through gridsearch, are displayed, while Table 5 shows the hyper-parameters that performed worst.

Table 4 reveals that the forecasting of soil moisture exhibited the highest performance for a sensor depth of 90cm, whereas the lowest performance was observed for a sensor depth of 60cm. Irrespective of the sensor depth, soil moisture forecasting achieved an average MAPE of 3.41% for a 1-day horizon, 5.85% for a 3-day horizon, and 10.09% for a 7-day horizon.

Interestingly, When looking at the hyper-parameters that yielded the best results the batch size (batch_size) did not seem to have a consistent impact on the forecasting performance across different combinations. However, besides the forecasting horizon, The evaluation metrics are most significantly influenced by the window size (window_size) and hidden size (hidden_size). Upon comparing the most optimal window size and hidden size with the least optimal ones, as shown in Table 5, a substantial decrease in performance is evident. In fact, when comparing the best- and worst-case window and hidden sizes of a 7-day forecasting horizon with a sensor depth of 90cm, a substantial difference of 531.87% becomes evident.

The optimal window size (window_size) for achieving the best evaluation metrics varied between 5 and 20 for smaller sensor depths of 15cm and 30cm. On the other hand, for bigger sensor depths of 60cm and 90cm, larger window sizes (window_size) ranging between 20 and 30 were found to be necessary. Notably, the forecasting horizon (forecasting_days) did not seem to have a noteworthy correlation with the most optimal window size (window_size).

The hidden size (hidden_size) achieved the best evaluation metrics between 75 and 100 while a hidden size (hidden_size) of 50 contributed to the least optimal evaluation metrics several times.

In Figure 1 and Figure 2 we can see a comparison of a plotted visualization of a 1-day forecast scenario for a near-surface sensor at 15cm sensor depth and a 90cm deep placed sensor. It is evident that the 90cm sensor seen in Figure 2 yields a better performance than the near-surface 15cm sensor seen at Figure 1

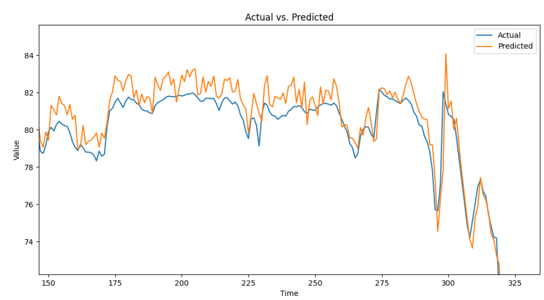

Fig. 1. Predicted vs. Actual soil moisture 1-day forecast for 15cm sensor depth with MAE of 3.96%.

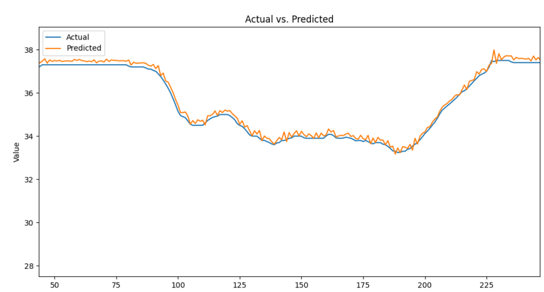

Fig. 2. Predicted vs. Actual soil moisture 1-day forecast for 90cm sensor depth with MAE of 0.66%.

Table 4. Table that shows the gridsearch results of the evaluation metrics under different hyperparameters filtered to the best performance

forecasting_days	depth	window_size	hidden_size	batch_size	MAE	MSE	RMSE	MAPE
1	15	5	75	16	1.605845809	6.259361267	2.501871586	3.985815048
1	30	20	75	32	1.245525002	5.976701736	2.444729328	3.618539572
1	60	30	75	32	0.249674723	0.240912452	0.490828335	5.371599197
1	90	25	50	16	0.226022452	0.198213428	0.445211679	0.662917137
3	15	10	100	16	2.733232975	18.77821159	4.33338356	7.598301888
3	30	5	100	16	1.986030221	14.99341583	3.872133255	6.055431366
3	60	30	100	16	0.428185552	0.735471725	0.857596457	8.172030449
3	90	20	75	16	0.538294196	1.08403337	1.041169286	1.563376904
7	15	15	75	32	4.404176235	34.46481323	5.870674133	12.94023705
7	30	5	100	32	3.352068186	33.10631561	5.753808975	10.31196976
7	60	30	50	16	0.754276216	1.715686679	1.309842229	14.63491631
7	90	30	75	16	0.856117189	1.946182609	1.395056486	2.486559868

Table 5. Table that shows the gridsearch results of the evaluation metrics under different hyperparameters filtered to the worst performance

forecasting_days	depth	window_size	hidden_size	batch_size	MAE	MSE	RMSE	MAPE
1	15	30	50	32	2.338097095	9.288647652	3.0477283	7.254351139
1	30	25	75	16	1.688043118	8.732964516	2.955158949	5.40770483
1	60	20	75	32	0.522857547	0.42932412	0.655228317	15.21760082
1	90	10	50	32	0.70194608	2.663289547	1.631958842	2.434599638
3	15	30	75	32	4.13811636	40.56227875	6.368852139	16.48276901
3	30	25	50	16	3.372938633	26.21930504	5.120479107	11.10694885
3	60	20	75	16	0.576828241	0.825701356	0.908681095	23.62647438
3	90	10	50	32	2.483380556	34.86305237	5.904494286	8.060541153
7	15	30	50	16	10.38335323	288.1854553	16.97602654	41.59179688
7	30	30	50	16	6.523112297	78.6312561	8.867426872	25.85913849
7	60	20	50	16	1.126405478	3.846907616	1.96135354	43.90224838
7	90	10	50	32	4.398735046	76.59194183	8.751682281	13.2251997

6 DISCUSSION

To answer our first subquestion (SQ1) about what features are most informative to include as input for training the LSTM model we can say the following things:

The overall Pearson correlation coefficient between the meteorological variables from the weather station and the soil moisture dataset is relatively low, measuring less than 0.2. However, it is important to note that the relationship between these variables may not be easily determinable based on correlation alone. This could be because the influence of meteorological variables on soil moisture might have a delay or be dependent upon the duration of specific conditions.

Features such as temperature, precipitation, and humidity, which are known to have a relationship with soil moisture based on domain-specific and contextual knowledge, do exhibit a low but noticeable correlation. This finding further strengthens our argument regarding the potential delay or duration-dependent relationship between these variables.

Integrated gradients technique revealed that a correlation between a feature and moisture does not guarantee an improvement in LSTM accuracy when using that variable. Comparing the top 5 correlated features from Table 2 with the feature attribution in Table 3, we observe that only two of these features significantly contribute to the model's outcome. However, the contribution of these features is relatively small, almost negligible, as the model heavily relies on historical soil moisture, which outweighs the influence of daily humidity by a factor of 21.

On one hand, this phenomenon can be attributed to the initial weak correlation between soil moisture and the top 5 meteorological variables. On the other hand, the inclusion of redundant features may introduce noise and unnecessary complexity by duplicating information already captured by the soil moisture variable, potentially adversely impacting the model's performance.

The evaluation metrics and accuracy of our model primarily rely on historical soil moisture, with the inclusion of features known to be correlated with it showing minimal improvement in performance. This raises the question of whether multivariate soil moisture forecasting provides substantial benefits compared to univariate soil moisture forecasting. To gain further insights, future experiments should be conducted by solely considering soil moisture as the feature, allowing for a direct comparison of evaluation metrics.

By analyzing the results based on the forecasting horizon and sensor depth at Table 4 we can get a better understanding of the effects of these variables on the model's performance that can answer our second sub research question (SQ2):

In table 4, a significant exponential decline in the model's performance can be seen as the forecasting horizon was increased. On average, when extending the forecasting horizon from one day to seven days, a considerable decrease of 3.4 times in MAPE was observed independently from the depth of the sensor. This decline can be attributed to the fact that as the forecasting horizon grows larger, a larger portion of the window size consists of predicted values that already deviate from the actual values due to inherent biases. Based on our conducted experiments, we assume that further increasing the forecasting horizon would result in a further exponential decrease in performance. This is because, at a certain point, the forecasting would rely solely on previous predictions that could already be inaccurate. In other words, the forecasting horizon has a high negative correlation with the models performance.

When evaluating the performance of sensors based on their depth, a notable trend can be seen: the deepest sensor, located at 90cm, consistently outperforms sensors placed at shallower depths which can also be seen when comparing Figure 1 with Figure 2. This disparity can be attributed to several key observations. Firstly, soil moisture at greater depths show higher consistency and lower levels of noise due to reduced exposure to evaporation and less influence of both manual and natural irrigation when compared to sensors more close to the surface. Additionally, natural factors such as root water uptake and capillary action contribute to the enhanced consistency observed in sensors placed at greater depths as it moves the water to the surface.

The increased consistency of the soil moisture data of deeper sensors holds significant implications for the effectiveness of LSTM models. LSTM models achieve better performances on consistent patterns and relationships within the data that are needed to make accurate forecasts. As the soil moisture becomes more consistent at greater depths, the LSTM model is better equipped to generalize and make reliable predictions.

Therefore, a positive correlation can be established between the depth of the sensor and the performance of the LSTM model. As the sensor depth increases, the consistency of soil moisture improves, facilitating easier forecasting and ultimately enhancing the model's overall performance as can be seen in the evaluation metrics in Table 4.

When comparing Table 4 and Table 5 it is evident that choosing the hyper-parameters has a large impact on the evaluation metrics. Choosing the right versus choosing the worst hyper-paramaters resulted in a significant difference of up to 531.87% . Mainly the choice of hidden size and window size leads to this difference.

Notably, our analysis, which contributes to the third sub research question (SQ3) indicates that the deeper sensors (60cm and 90cm) had superior performance with larger window sizes, while the shallower sensors required smaller window sizes. Our argument regarding increased consistency in moisture for deeper sensors explains why bigger window sizes are required as this consistency implies less changes. That is why bigger window sizes are needed to capture patterns that emerge in the long-term.

On the other hand, shallower placed sensors, which are inherently more prone to noise which can be seen in Figure 1, often tend to perform better with smaller window sizes. This might be because the smaller window size reduced the amount of noise and the model can adapt to a more localized and recent subset of the data which is more prone to local variations such as precipitation and humidity.

Analyzing the performance of Table 4 in relation to the hidden size, it becomes apparent that deeper sensors achieve the best performance with lower hidden sizes, whereas shallower sensors need higher hidden sizes to achieve their best performance. In general, when the data contains intricate patterns and complexity, a higher hidden size is needed. This gives the LSTM model the capacity to store the intricate and complex relationships between input and output variables. Applying this to our sensor forecasting context, we can confirm that due to the higher noise levels and complex nature of soil moisture development near the surface, a higher hidden size is needed for achieving the most optimal performance compared to deeper placed sensors.

While our LSTM model can achieve high performance for deeper sensors due to the data being less noisy and more consistent, alternative approaches specifically targeting the performance of near-surface sensors could be explored.

In addition, our model is based on a conventional vanilla LSTM architecture, which can impact its performance. To further enhance the accuracy and effectiveness of soil moisture forecasting, it would be valuable to conduct experiments using diverse LSTM architectures and compare their performance against the classic vanilla architecture. These experiments could achieve intriguing insights and potential significant improvements in the field of soil moisture prediction.

7 CONCLUSION

Accurately forecasting soil moisture is a crucial challenge to address, as it can greatly assist green managers and landscape professionals in optimizing their irrigation processes and ensuring the best health of their green.

Our model represents a significant advancement in addressing this challenge by achieving remarkable accuracy in forecasting soil moisture across various sensor depths for up to seven days. By utilizing a LSTM model that combines historical soil moisture data and metereological information we made a substantial progression in this field.

Our model achieves compelling results, with an average MAPE of 3.41% for a 1-day horizon, 5.85% for a 3-day horizon, and 10.09% for a 7-day horizon. In other words, this means that, on average, our predictions deviate from the actual soil moisture values by only 3.41%, 5.85% or 10.09% respectively, depending on the forecasting horizon.

To achieve higher accuracy in soil moisture forecasting, further research should prioritize the prediction of near-surface sensors, considering that these sensors notably contributed to the decrease in our overall performance.

Finally, our research has contributed to the field of machine learning soil moisture forecasting by uniquely utilizing and evaluating a LSTM model with multivariate timeseries data to accurately forecast soil moisture. While previous research was limited to one-size depths and rural areas, this research has evaluted performance accross different depths, forecasting horizons and more diverse soil moisture sensors subject to manual irrigation.

REFERENCES

[1] Umesh Acharya, Aaron L. M. Daigh, and Peter G. Oduor. 2021. Machine Learning for Predicting Field Soil Moisture Using Soil, Crop, and Nearby Weather Station Data in the Red River Valley of the North. *Soil Systems* 5, 4 (Dec. 2021), 57. https://doi.org/10.3390/soilsystems5040057 Number: 4 Publisher: Multidisciplinary Digital Publishing Institute.

[2] Ibrahem Kandel and Mauro Castelli. 2020. The effect of batch size on the generalizability of the convolutional neural networks on a histopathology dataset. *ICT Express* 6, 4 (Dec. 2020), 312–315. https://doi.org/10.1016/j.icte.2020.04.010

[3] A. Khalil, M.K. Gill, and M. McKee. 2005. New applications for information fusion and soil moisture forecasting. *2005 7th International Conference on Information Fusion* (2005), 7 pp. https://doi.org/10.1109/ICIF.2005.1592050 Conference Name: 2005 7th International Conference on Information Fusion ISBN: 9780780392861 Place: Philadelphia, PA, USA Publisher: IEEE.

[4] Mingshuai Liu, Chao Huang, Long Wang, Yu Zhang, and Xiong Luo. 2020. Short-Term Soil Moisture Forecasting via Gaussian Process Regression with Sample Selection. *Water* 12, 11 (Nov. 2020), 3085. https://doi.org/10.3390/w12113085

[5] Zulkarnaen Rashid and Syed Al Junid. 2014. Trees' Cooling Effect on Surrounding Air Temperature Monitoring System: Implementation and Observation. *International Journal of Simulation Systems Science & Technology* 15 (April 2014), 70–77. https://doi.org/10.5013/IJSSST.a.15.02.10

[6] U. Rosenbaum, H. R. Bogena, M. Herbst, J. A. Huisman, T. J. Peterson, A. Weuthen, A. W. Western, and H. Vereecken. 2012. Seasonal and event dynamics of spatial soil moisture patterns at the small catchment scale. *Water Resources Research* 48, 10 (2012). https://doi.org/10.1029/2011WR011518 _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2011WR011518.

[7] Thomas J. Schmugge. 1983. Remote Sensing of Soil Moisture: Recent Advances. *IEEE Transactions on Geoscience and Remote Sensing* GE-21, 3 (July 1983), 336–344. https://doi.org/10.1109/TGRS.1983.350563 Conference Name: IEEE Transactions on Geoscience and Remote Sensing.

[8] Sukhwinder Singh, Sanmeet Kaur, and Parteek Kumar. 2020. Forecasting Soil Moisture Based on Evaluation of Time Series Analysis, S. N. Singh, R. K. Pandey, Bijaya Ketan Panigrahi, and D. P. Kothari (Eds.), Vol. 609. Springer Singapore, Singapore, 145–156. https://doi.org/10.1007/978-981-15-0313-9_11 Book Title: Advances in Power and Control Engineering Series Title: Lecture Notes in Electrical Engineering.

[9] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic Attribution for Deep Networks. (June 2017). https://doi.org/10.48550/arXiv.1703.01365 arXiv:1703.01365 [cs].

[10] TerraSpect. 2023. Efficiënt watergeven in tijden van droogte - TerraSpect. *TerraSpect* (May 2023). https://www.terraspect.nl/nieuws/efficient-watergeven-in-tijden-van-droogte

[11] Rodrigo Togneri, Diego Felipe dos Santos, Glauber Camponogara, Hitoshi Nagano, Gilliard Custódio, Ronaldo Prati, Stênio Fernandes, and Carlos Kamienski. 2022. Soil moisture forecast for smart irrigation: The primetime for machine learning. *Expert Systems with Applications* 207 (Nov. 2022), 117653. https://doi.org/10.1016/j.eswa.2022.117653

[12] Pantana Tor-ngern and Nisa Leksungnoen. 2020. Investigating carbon dioxide absorption by urban trees in a new park of Bangkok, Thailand. *BMC Ecology* 20, 1 (April 2020), 20. https://doi.org/10.1186/s12898-020-00289-4

[13] Luuk van der Lee. 2020. Nederland heeft er 1,6 miljoen bomen bij, en dat is nog maar het begin. *NOS* (July 2020). https://nos.nl/artikel/2339646-nederland-heeft-er-1-6-miljoen-bomen-bij-en-dat-is-nog-maar-het-begin

[14] Peter J. Wetzel and Robert H. Woodward. 1987. Soil Moisture Estimation Using GOES-VISSR Infrared Data: A Case Study with a Simple Statistical Method. *Journal of Applied Meteorology and Climatology* 26, 1 (Jan. 1987), 107–117. https://doi.org/10.1175/1520-0450(1987)026<0107:SMEUGV>2.0.CO;2 Publisher: American Meteorological Society Section: Journal of Applied Meteorology and Climatology.