

# Machine Learning-based Energy Optimisation in Smart City Internet of Things

Eric Samikwa, Jakob Schärer, Torsten Braun, Antonio Di Maio Institute of Computer Science, University of Bern Switzerland

# ABSTRACT

The deployment of Internet of Things (IoT) temperature sensors in urban areas is essential for the monitoring and understanding of the thermal environment. However, accurate temperature measurements can be compromised by factors such as direct sunlight, leading to overheating and inaccurate readings. We propose a Machine Learning-based approach that addresses this challenge by dynamically ventilating the sensor environment using small fans, enabling accurate and energy-efficient temperature measurements. This paper focuses on two interconnected problems: predicting steady-state temperature using a limited window of initial temperature measurements and investigating the impact of ventilation time. We employ various DNNs suitable for low-power IoT sensor devices to predict temperature using multivariate time series from different sensors and compare their accuracy. Furthermore, we highlight the tradeoff between prediction accuracy, which is correlated to the length of the observed input sequence, and energy consumption dependent on ventilation time. By adopting advanced prediction techniques, we can develop efficient IoT systems for accurate and energy-efficient environment monitoring in smart cities.

# **CCS CONCEPTS**

• Information systems → Sensor networks; • Computing methodologies → Neural networks; Distributed artificial intelligence; • Hardware → Temperature monitoring.

# **KEYWORDS**

Machine Learning, Internet of Things, Energy Optimisation, Smart Sensors, Temperature Monitoring.

#### **ACM Reference Format:**

Eric Samikwa, Jakob Schärer, Torsten Braun, Antonio Di Maio. 2023. Machine Learning-based Energy Optimisation in Smart City Internet of Things. In *The Twenty-fourth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing (MobiHoc '23), October 23–26, 2023, Washington, DC, USA.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3565287.3616527

# **1** INTRODUCTION

Deploying Internet of Things (IoT) devices, particularly temperature sensors, in urban areas has revolutionized our ability to monitor and understand the dynamic thermal environment. These sensors



This work is licensed under a Creative Commons Attribution International 4.0 License.

MobiHoc '23, October 23–26, 2023, Washington, DC, USA © 2023 Copyright held by the owner/author(s). ACM ISBN 978-1-4503-9926-5/23/10. https://doi.org/10.1145/3565287.3616527 play a crucial role in applications such as urban planning, building management, and environmental monitoring for smart cities [1]. However, their accuracy can be compromised by local factors such as direct sunlight, which can cause the sensors to overheat and yield inaccurate temperature readings. To address this challenge, we introduce a Machine Learning (ML)-based approach for ventilating the sensor environment using small fans, enabling accurate and energy-efficient temperature measurements.

This paper focuses on two interconnected problems that arise in the context of the IoT temperature sensors. Firstly, we aim to predict the final temperature at a steady state after an extended ventilation time, leveraging only a limited window of temperature and humidity measurements taken during the initial ventilation phase. This approach enables us to conserve energy by activating the fan for a shorter duration, thus reducing the overall power consumption of the sensor device.

Secondly, we delve into the issue of determining the optimal ventilation time that maximizes the amount of data observed and subsequently enhances the accuracy of temperature prediction at steady-state. This exploration needs to strike a balance between energy consumption and prediction accuracy. Longer ventilation times yield more comprehensive datasets that have the potential to yield more accurate predictions of the final, steady-state temperature [2]. Several time series analysis studies have shown that longer input sequence length allows Deep Neural Networks (DNN) models to capture enough historical context, trends, and seasonality in the data, enhancing the predictive capabilities [3]. Conversely, shorter ventilation times save energy but provide a more limited dataset and observation sequence length, compromising prediction accuracy. Thus, the ventilation period significantly affects the energy consumption for the sensor devices, which introduces a tradeoff between prediction accuracy and energy consumption.

Efficiently addressing these challenges holds significant practical implications. Optimizing energy consumption can extend the operational lifespan of IoT monitoring systems and reduce their environmental impact. Moreover, accurate temperature measurements are vital for various applications, including energy-efficient building management, climate change mitigation, and urban planning [4]. While previous studies have explored aspects of temperature prediction and energy optimization in IoT systems, few have focused on the interaction between these two aspects in the context of IoT temperature sensors in urban environments.

In this paper, we propose an ML-based approach for accurate and energy-efficient temperature monitoring in urban areas while offering valuable insights into optimizing energy efficiency and prediction accuracy. We utilize a range of DNNs suitable for low-power IoT sensor devices to predict temperature using a limited multivariate time series of initial temperature and humidity sensor values and assess their accuracies. Moreover, we emphasize the tradeoff

Eric Samikwa, Jakob Schärer, Torsten Braun, Antonio Di Maio

between prediction accuracy, which is affected by the length of the observed input sequence, and energy consumption, which is influenced by ventilation time. By adopting prediction techniques, we can develop more efficient environmental monitoring systems that achieve both accuracy and energy efficiency in smart city IoT.

The remainder of the present article is organized as follows. Section 2 presents the relevant related works. Section 3 presents the system model and describes the proposed method. Section 4 provides the evaluation. Section 5 concludes the paper.

# 2 RELATED WORKS

ML has been applied to environmental sensing in various studies, considering the tradeoff between information accuracy and energy expenditure. One study [5] explored the application of modern techniques, including ML, for enhancing reliability and reducing costs in environmental monitoring IoT. The authors [6] proposed a real-time, power-efficient environmental monitoring system and studied coverage quality and deployment mechanisms.

Temperature prediction and energy optimization in IoT systems have been widely studied, but limited research has specifically focused on the interaction between these aspects in the context of IoT temperature sensors in urban environments. Traditional approaches, such as autoregressive integrated moving average models and exponential smoothing methods, have shown reasonable accuracy but may not capture the complex dynamics of urban temperature variations [7]. ML algorithms were used with multi-temporal data capture changes in environmental parameters to understand climate variables [8, 9]. In the context of collaborative sensing and communication systems, distributed learning methods are proposed to improve the efficiency of IoT systems [10, 11].

More recently, ML techniques, particularly Deep Learning (DL) models, have gained attention for temperature prediction [4, 8]. However, these studies have primarily focused on general temperature prediction without considering the specific challenges in urban environments and the energy efficiency aspect. On the other hand, energy optimization in IoT systems has been a significant research area, aiming to prolong the operational lifespan of sensor devices and reduce energy consumption using various techniques.

While temperature prediction and energy optimization have been individually explored, their interaction in the context of improving sensor measurements in urban environments remains largely unexplored. The tradeoff between accuracy and energy consumption, specifically in relation to ventilation time, has not been investigated. Therefore, this paper aims to bridge this gap by investigating the interaction between these two aspects and proposing an approach to achieve accurate and energy-efficient temperature monitoring in urban environments.

# 3 MACHINE LEARNING-BASED ENERGY OPTIMISATION FOR SMART CITY TEMPERATURE SENSORS

#### 3.1 System Model

Our system contains a set  $\Phi$  of devices with constrained computation, communication, and energy resources, whose task is monitoring environmental variables over a large geographical extension



Figure 1: System Model.

periodically. Each device  $\phi \in \Phi$  is endowed with a set of *n* sensors that can measure environmental variables such as temperature and humidity. We assume that each device performs a measurement periodically and transmits it to a central server via a long-range wireless link (e.g., LoRa) for visualization and processing.

The sensors might be located in environmentally challenging situations, such as under direct sunlight or intense humidity, which might bias their measurements. To mitigate such potential biases, the device initiates a *ventilation* operation on the sensor at the moment of performing a measurement to disperse excess heat from the sensor casing. We assume that sensors are self-contained and cannot access the wired power grid of data connectivity, therefore, the fan is powered with the limited energy of the onboard battery. After a fixed ventilation period of  $t \in \mathbb{N}$  time slots, the device records the measurement and transmits it to the central server. We assume that the ventilation period t is sufficient to measure the environmental conditions accurately and without biases. Specifically, t is predetermined by a sensor bench-marking process in variable controlled lab environments, where the device is heated and subsequently ventilated until the sensor readings do not change further.

The sensor device is also able to record the sensor values during the ventilation period. Without loss of generality, we assume that a ventilation period is divided into time slots where a measurement is carried out. For each ventilation operation, the time slots are indexed starting from the beginning of ventilation, i.e., the first measurement is  $x_{1,i}$  and the last measurement is  $x_{t,i}$ .We denote  $X_i = (x_1, \ldots, x_t)_i \in \mathbb{R}^{n \times t}$  as the *measurement time-series* for the *i*-th measurement, and we assume that its values at every time slot during the ventilation period are recorded and saved on the device's local storage. Let us denote  $\mathbb{X} = (X_1, \ldots, X_m)$  as the set of all measurements collected by the sensor devices.

The central server periodically collects measurements from the sensor devices for processing and analysis via the long-range wireless network or manual collection. We index each measurement sent to the server by a device as  $y_i$ . Since the ventilation process is only aimed at improving the quality of the measurements, the sensor device does not need to send the entire measurement time series (i.e., recorded during the ventilation process) to the server. Thus only the last sensor measurements recorded at the end of the ventilation period are sent to the server, i.e.,  $y_i = x_{t,i}$  (last value of

Machine Learning-based Energy Optimisation in Smart City Internet of Things

Figure 2: Sensor Device.

the *measurement time-series*). Let us denote  $Y = (y_1, \ldots, y_m)$  as the set of all measurements collected by the server.

We assume that the device fan's average power consumption  $P(\phi)$  for device  $\phi$  is constant over time. Empirical results show that the device uses the majority (over 70%) of its energy for the ventilation process; therefore, our approach focuses on optimizing the ventilation energy. We define the average energy consumption for a single measurement sent to the server by device  $\phi$  as  $E(\phi) = t\tau P(\phi)$ , where  $\tau$  is the duration of one time slot in seconds. Figure 1 represents our model of the system. Each sensor device has a processor capable of executing small DNN models. Figure 2 shows a picture of the sensor node equipped with a small solar panel. The sensors, control module, and small fan are housed within a ceramic casing with a design commonly used for environmental sensors [12].

#### 3.2 **Problem Formulation**

Devices' local energy storage is a precious resource, as it determines the devices' lifetime and therefore impacts the frequency of battery replacement by operators. In this work, we aim to devise a solution to reduce the device's energy consumption by reducing the ventilation time *t*, while maintaining accurate measurement information about  $x_{t,i} = y_i$  at ventilation convergence. We propose using a multivariate time-series DNN model to predict the sensor measurements (e.g., the temperature) at the end of a ventilation period for the *i*-th measurement (i.e.,  $x_{t,i}$ ) by observing a smaller sample of temperature and humidity measurements (i.e.,  $(x_1, \ldots, x_{t'})_i$ , with t' < t), which consequently reduces energy consumption.

Several studies have shown that the dataset volume and the length of the observed time series (i.e., *measurement time-series*) have a significant impact on the prediction accuracy [2, 13]. When the observed part of the time series is short, the model may not have access to enough historical context to make accurate predictions. On the other hand, longer *input sequence length* allows the model to capture input sequence length, trends, and seasonality in the data, enhancing the predictive capabilities [3]. In such cases, the model learns from a larger historical data set and potentially uncovers more complex relationships within the time series. However, the ventilation period duration significantly increases the energy consumption  $E(\phi)$  for the sensor device  $\phi$ , which introduces a potential tradeoff between the measurements' prediction accuracy and the energy needed to perform the measurement.



Figure 3: 1-Dimensional Convolutional Neural Networks.



Figure 4: Feed Forward Neural Network.

We aim to build a function  $f(X_i)$  that takes in input a multivariate time series  $X_i = (x_1, ..., x_{t'})_i$ , with t' < t, and predicts the value  $x_{t,i} = y_i$  as  $\hat{y}_i = f(X_i)$ . To simplify notation, and without loss of generality, we focus on a simpler subproblem where the predicted measurement  $y_i$  is the last temperature value. Let us denote  $\hat{Y} = f(\mathbb{X})$  as the set of predicted measurements. We outline a regression problem where our goal is to minimize the average prediction loss  $\mathcal{L}(Y, \hat{Y})$  between the ground truth of collected measurements Y and their model prediction  $\hat{Y}$ . In our solution, we consider the root mean squared error (RMSE) as loss function  $\mathcal{L}(Y, \hat{Y}) = \sqrt{\frac{1}{m} \sum_{i=1}^{m} l(y_i, \hat{y}_i)} = \sqrt{\frac{1}{m} \sum_{i=1}^{m} (y_i - \hat{y}_i)^2}$ . RMSE measures the squared root of the mean of the deviation squares, which quantifies the difference between the predicted values and the actual ones. On the other hand, the energy consumption required to perform the measurement of the sensor values depends on the input sequence length (i.e., the period of ventilation for collecting the measurement time series). Thus, the energy  $E(\phi) = t' \tau P(\phi)$  is required to collect the multivariate time series  $X_i = (x_1, ..., x_{t'})_i$ .

In this study, we perform an offline parameter analysis to find the optimal ventilation time (i.e., the optimal input sequence length t' for the observed multivariate time series) to strike a balance between energy consumption and prediction accuracy.

# 3.3 DNN Prediction of Temperature Sensor Values Using Multivariate Time-Series

We utilize the DNNs, Long Short-Term Memory (LSTM), Feed Forward Neural Network (FFNN), and Convolutional Neural Networks (CNN), for temperature values prediction using time series. MobiHoc '23, October 23-26, 2023, Washington, DC, USA

Eric Samikwa, Jakob Schärer, Torsten Braun, Antonio Di Maio



Figure 5: Long-Short Term Memory Network.

LSTM is a type of recurrent neural network (RNN) architecture that is specifically designed to capture long-term dependencies and process sequential data (Figure 5). It overcomes the limitations of traditional RNNs, which struggle with long-term dependencies due to the vanishing gradient problem.LSTMs have a memory cell that can store information for extended periods, allowing them to capture and retain long-term dependencies in the input data. This is crucial for time series data, where the previous inputs can have a significant impact on future predictions. LSTMs can effectively handle multivariate time series inputs, where multiple variables or features are observed simultaneously at each time step by capturing complex relationships between different variables. LSTMs have fewer parameters compared to other deep learning architectures, making them computationally efficient and plausible for low-power IoT sensors that typically have limited computational resources [14, 15]. Moreover, LSTMs can process sequential data in an incremental and online manner, which is desirable for low-power sensors.

1D CNN is a type of neural network architecture commonly used for processing one-dimensional sequential data, such as time series [16]. While CNNs are widely known for their effectiveness in computer vision tasks, they can also be applied to time series analysis (Figure 3). 1D CNNs employ convolutional filters that slide over the input sequence, extracting local patterns and features. This localized feature extraction is advantageous for time series data, as it enables the network to identify meaningful patterns at different time scales. 1D CNNs can handle multivariate time series inputs, where multiple variables are observed simultaneously at each time step by capturing cross-variable interactions and extracting relevant features from the input. 1D CNN models can be deployed directly on IoT sensors or edge devices, eliminating the need for transmitting raw sensor data.

FFNN, also known as Multi-Layer Perceptrons (MLPs), is a neural network where information flows in one direction (Figure 4). FFNNs can capture complex nonlinear relationships between input features and output predictions. This is beneficial for time series data, as it often contains patterns and dependencies that can be effectively captured by the nonlinear activation functions in FFNN. FFNNs have been proven to be universal function approximators, meaning they can approximate any continuous function given enough hidden units. This property allows FFNNs to learn and model the underlying patterns in time series data, making them suitable for a wide range of time series prediction tasks [17]. FFNNs typically have a simpler structure compared to most DNN architectures, resulting in a smaller number of parameters making FFNNs computationally efficient, which is advantageous for low-power sensor devices.

# **4 EVALUATION**

# 4.1 Experimental Setup

We developed the ML models in Python, utilizing the deep learning frameworks Tensorflow and Keras. The key metric of the evaluation is the accuracy of the temperature prediction determined by the error of the regression output of the DNN models for variable input sequence length. We then visualize the accuracy and energy consumption tradeoff for the variable input sequence lengths. Furthermore, we compare the accuracy of the multivariate time series approach to a univariate time series approach. We released the source code and trained models on a public GitHub repository<sup>1</sup>.

# 4.2 Dataset

We collected measurements from a real-world deployment of environmental sensors in the city of Bern, Switzerland. Our proposed approach can be applied to determine the tradeoff between the accuracy of temperature measurements and reducing the energy consumption for a single sensor; hence, without loss of generality, the evaluation is conducted on a dataset from a single sensor. Overall, we acquired 3697 measurements, each long t = 138 s. To correct the measurements, we set the maximum ventilation duration of 138 seconds, during which the multivariate time series of humidity and temperature sensor values are recorded together with their corresponding timestamps. The sensor values are recorded at a fixed frequency. To promote our findings' reproducibility, we publicly release the collected dataset on Zenodo<sup>2</sup>.

Figure 6 shows a summary of the dataset used for the evaluation with different characteristics depending on the time of the day the measurements were carried out. Figure 6 (a) temperature variations without ventilation and after a 120 s ventilation for different hours of the day. The largest gap is observed around noon due to the sun's overheating action on the sensor casings. Figure 6 (b) shows sampled temperature time series (curves) recorded during the ventilation process for 138 seconds, with t = 138 (i.e., one record per second), at different times of the day. Figure 6 (c) shows some scaled temperature and humidity values for corresponding measurements. The data shows a correlation of the pattern between the two features that indicate interactions and dependencies between temperature and humidity. This pattern can be leveraged by a multivariate approach as complementary information for a more comprehensive and accurate understanding of the time series.

From this raw data, we created the training and test sets through data augmentation to simulate time series of different lengths. Namely, for each measurement, we generated 136 samples with the increasing length of measurement time-series, padding the residual time-series length with zeros until reaching a time-series length of 137. We randomly select 70% of the generated dataset for training, 15% for validation, and 15% for testing, shuffling the samples to avoid imbalances in the different parts of the dataset.

<sup>&</sup>lt;sup>1</sup>https://www.github.com/ricsamikwa/ml-iot-smartcitytemp

<sup>&</sup>lt;sup>2</sup>https://doi.org/10.5281/zenodo.8287290

Machine Learning-based Energy Optimisation in Smart City Internet of Things

MobiHoc '23, October 23-26, 2023, Washington, DC, USA



Figure 6: Temperature (a) variations without ventilation and after a 120 s ventilation for different hours of the day, (b) sampled temperature curves during ventilation, and (c) scaled values for corresponding temperature and humidity measurements.

#### 4.3 Results

Figure 7 shows the performance comparison of the selected baseline prediction models (LSTM, 1D-CNN, and FFNN), under varying input sequence lengths t' of measurement multivariate time-series on the test set. The plots indicate the root mean square error  $\mathcal{L}(Y, \hat{Y})$ for all samples in the test set for various input sequence length t', and the confidence intervals show the variability of the prediction error for each t'. Overall, the LSTM has a higher prediction accuracy compared to the 1D-CNN and the FFNN. Moreover, for lower value input sequence length t' the LSTM significantly outperforms the 1D-CNN, and FFNN). This can be attributed to the fact that LSTM networks are designed to capture long-term dependencies in time series effectively. Furthermore, LSTMs are good at variable-length sequences, where the number of time steps can vary. This is more prominent in our case since we experimented with variable input sequence length for the multivariate time series, and we introduced padding (with zeros), which might contribute to more noisy features for CNN and FFNN. Moreover, the performance of the CNN and FFNN significantly improves for longer input sequence length. The naive estimation takes the last temperature measurement at the end of a ventilation period t' as the actual value at the end of the curve. However, it is highly dependent on the specific range of temperature variability for the sensor. The naive estimation is further visualized in Figure 9.

Figure 8 shows the tradeoff between the prediction accuracy (based on the RMSE) and the energy consumption on the devices for variable input sequence lengths t'. The energy consumption is shown in normalized values  $(t'\tau P/t\tau P)$  as a proportion of the ventilation energy to observe the shorter input sequence of length t' instead of full ventilation for the duration t. The results indicate that longer input sequence length t' results in better prediction accuracy. On the other hand, the energy consumption grows proportionally to the length of the observed input sequence. This is a useful indication that can be used by the application policymaker to consider the tradeoff between energy consumption and accuracy. In accuracy-sensitive scenarios, longer input sequence lengths of the temperature and humidity time series can be utilized to achieve more accurate predictions of the measurements. On the other hand, in energy-sensitive scenarios, e.g., when the battery of the sensor is low, shorter input sequence lengths can be desirable for the prediction of the temperature values.



Figure 7: Prediction error (RMSE) for various DNN methods under variable observed input sequence lengths.



Figure 8: Accuracy and energy consumption tradeoff.

Figure 9 shows a visualization of the prediction accuracy for LSTM on a single ventilation curve compared with a naive approach and the full ventilation. The naive approach records the last temperature measurement at the end of a ventilation period t' as the actual value at the end of the curve. The predicted result is closer to the actual temperature reading while utilizing the same

MobiHoc '23, October 23-26, 2023, Washington, DC, USA

Eric Samikwa, Jakob Schärer, Torsten Braun, Antonio Di Maio



Figure 9: Sample prediction for t' = 50 and a naive estimation.



Figure 10: RMSE for multivariate and univariate LSTM.

amount of ventilation energy consumption as the naive approach. When the input sequence length t' approaches the end of the curve, the naive approach is the same as full ventilation for all 138 seconds.

Figure 10 shows the RMSE for the temperature prediction using multivariate times series ( $|x_i| = N = 2$ ), i.e., the model observes both the temperature and humidity values at each time step of the input sequence and univariate time series  $(|x_i| = N = 1)$ , i.e., the model observes only the temperature sequence for various input sequence lengths t'. The prediction error is lower for the multivariate LSTM compared to the univariate LSTM. This can be attributed to the ability of the multivariate approach to leverage complementary information from multiple relevant or correlated features, such as humidity measurements in this case.

#### 5 CONCLUSION

Our study demonstrates the feasibility of using time-series ML methods to predict temperature using a shorter time series than the target time, resulting in significant energy savings through reduced sensor ventilation time. We employed various DNNs suitable for low-power IoT sensor devices to predict temperature using a multivariate time series approach and compared their prediction accuracy. Moreover, we emphasized the tradeoff between prediction accuracy, which improves with increasing input sequences, and the energy consumption of the sensor device, which is associated with

ventilation time. These findings highlight the potential benefits of adopting these prediction techniques in diverse applications, facilitating the development of more efficient environmental monitoring systems for urban heat environments and smart cities. Furthermore, federated aggregation of model parameters can be applied to improve performance across different geographical locations with diverse weather conditions and seasonal variations.

### REFERENCES

- [1] S. Chen, L. Zhang, Y. Tang, C. Shen, R. Kumar, K. Yu, U. Tariq, and A. K. Bashir, "Indoor temperature monitoring using wireless sensor networks: a smac application in smart cities," Sustainable Cities and Society, vol. 61, p. 102333, 2020.
- Y. Hua, Z. Zhao, R. Li, X. Chen, Z. Liu, and H. Zhang, "Deep learning with long [2] short-term memory for time series prediction," IEEE Communications Magazine, vol. 57, no. 6, pp. 114–119, 2019.
- L. Yunpeng, H. Di, B. Junpeng, and Q. Yong, "Multi-step ahead time series forecasting for different data patterns based on lstm recurrent neural network," in 2017 14th web information systems and applications conference (WISA), pp. 305-310, IEEE, 2017
- [4] A. L. Diedrichs, F. Bromberg, D. Dujovne, K. Brun-Laguna, and T. Watteyne, 'Prediction of frost events using machine learning and iot sensing devices," IEEE Internet of Things Journal, vol. 5, no. 6, pp. 4589-4597, 2018
- M. A. Abid, N. Afaqui, M. A. Khan, M. W. Akhtar, A. W. Malik, A. Munir, J. Ahmad, and B. Shabir, "Evolution towards Smart and Software-Defined Internet of Things," AI, vol. 3, pp. 100-123, Mar. 2022. Number: 1 Publisher: Multidisciplinary Digital Publishing Institute.
- [6] C.-Y. Chang, C.-Y. Lin, G.-J. Yu, and C.-H. Kuo, "An energy-efficient hole-healing mechanism for wireless sensor networks with obstacles," Wireless Communications and Mobile Computing, vol. 13, no. 4, pp. 377-392, 2013. \_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/wcm.1106.
- [7] A. M. J. C. Wadoux, B. Minasny, and A. B. McBratney, "Machine learning for digital soil mapping: Applications, challenges and suggested solutions," Earth-Science Reviews, vol. 210, 2020.
- M. A. Haq, P. Baral, S. Yaragal, and B. Pradhan, "Bulk Processing of Multi-Temporal Modis Data, Statistical Analyses and Machine Learning Algorithms to Understand Climate Variables in the Indian Himalayan Region," Sensors, vol. 21, no. 21, 2021.
- T. Bouramtane, H. Hilal, A. T. Rezende-Filho, K. Bouramtane, L. Barbiero, S. Abra-[9] ham, V. Valles, I. Kacimi, H. Sanhaji, L. Torres-Rondon, D. D. d. Castro, J. d. C. V. Santos, J. Ouardi, O. E. Beqqali, N. Kassou, and M. Morarech, "Mapping Gully Erosion Variability and Susceptibility Using Remote Sensing, Multivariate Statistical Analysis, and Machine Learning in South Mato Grosso, Brazil," Geosciences, vol. 12, no. 6, 2022
- [10] H. Wu, Z. Zhang, C. Guan, K. Wolter, and M. Xu, "Collaborate edge and cloud computing with distributed deep learning for smart city internet of things," IEEE Internet of Things Journal, vol. 7, no. 9, pp. 8099–8110, 2020. E. Samikwa, A. Di Maio, and T. Braun, "Ares: Adaptive resource-aware split
- [11] learning for internet of things," Computer Networks, vol. 218, p. 109380, 2022.
- [12] D. T. Young, L. Chapman, C. L. Muller, X.-M. Cai, and C. Grimmond, "A low-cost wireless temperature sensor: Evaluation for use in environmental monitoring applications," Journal of Atmospheric and Oceanic Technology, vol. 31, no. 4, pp. 938-944, 2014.
- [13] A. Di Maio, M. Aghaei Dinani, and G. Rizzo, "The Upsides of Turbulence: Baselining Gossip Learning in Dynamic Settings," in The Twenty-fourth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing (MobiHoc '23), Washington, DC, USA, p. in press, ACM, New York, NY, USA, Oct. 2023.
- [14] D. Dennis, D. A. E. Acar, V. Mandikal, V. S. Sadasivan, V. Saligrama, H. V. Simhadri, and P. Jain, "Shallow rnn: accurate time-series classification on resource constrained devices," Advances in Neural Information Processing Systems, vol. 32, 2019
- [15] E. Samikwa, T. Voigt, and J. Eriksson, "Flood prediction using iot and artificial neural networks with edge computing," in 2020 International Conferences on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData) and IEEE Congress on Cybermatics (Cybermatics), pp. 234-240, IEEE, 2020.
- [16] W. Tang, G. Long, L. Liu, T. Zhou, J. Jiang, and M. Blumenstein, "Rethinking 1d-cnn for time series classification: A stronger baseline," arXiv preprint arXiv:2002.10061, pp. 1-7, 2020.
- J. J. Dabrowski, Y. Zhang, and A. Rahman, "Forecastnet: a time-variant deep feed-[17] forward neural network architecture for multi-step-ahead time-series forecasting," in Neural Information Processing: 27th International Conference, ICONIP 2020, Bangkok, Thailand, November 23-27, 2020, Part III 27, pp. 579-591, Springer, 2020.