

The Role of Artificial Intelligence in Early Moisture Risk Detection for Wooden Buildings

Kristina Piiarska
University of Primorska, Faculty of
Mathematics, Natural Sciences and
Information Technologies
Koper, Slovenia
kristina.piiarska@famnit.upr.si

Aleksandar Tošić
University of Primorska, Faculty of
Mathematics, Natural Sciences and
Information Technologies
Koper, Slovenia
InnoRenew CoE
Izola, Slovenia
aleksandar.tosic@upr.si
aleksandar.tosic@innorenew.eu

Pau Brunet
Forest Science and Technology
Centre of Catalonia
Solsona, Spain
pau.brunet@ctfc.cat

Abstract

The operation and maintenance of modern buildings increasingly relies on sensor networks that continuously collect large amounts of environmental and structural monitoring data. In wooden buildings, continuously monitoring these measurements remains challenging, allowing moisture-related problems to go unnoticed until costly intervention is required. This study investigates whether a fine-tuned Granite Time Series Foundation Model can support early moisture risk detection by forecasting future moisture development and providing automated alerts. A separate Granite model was fine-tuned for each building using moisture and environmental measurements collected from the InnoRenew CoE building in Slovenia and the Forestry HUB building in Spain. Performance was evaluated on unseen seasonal test data and compared with daily-naïve, weekly-naïve, Linear Regression, and Random Forest baselines. The Granite model mostly outperformed the naïve methods and Random Forest, while achieving MAE and RMSE comparable to Linear Regression. Despite the similar numerical accuracy, Granite better captured the temporal patterns and fluctuations present in the ground-truth measurements. The results demonstrate the potential of pre-trained time-series models for automated moisture monitoring and early risk alerts in wooden buildings.

Keywords

wood moisture, time-series forecasting, alert systems, building monitoring

1 Introduction

Wood has become an increasingly important construction material due to its renewable nature and widespread use in sustainable construction. As the number of wooden buildings continues to increase, ensuring their long-term durability and reliable operation becomes increasingly important. Among the various factors influencing the service life of wooden structures, moisture is one of the most critical. Elevated moisture levels can lead to fungal decay, mould growth, and other degradation processes that reduce structural performance and may ultimately require costly maintenance or repair.

To mitigate these risks, modern wooden buildings are increasingly equipped with sensor networks that continuously monitor wood moisture content together with environmental variables such as temperature and relative humidity. These systems generate large volumes of valuable information. However, collecting data alone does not guarantee effective preventive maintenance. In practice, building owners and facility managers do not continuously inspect thousands of sensor measurements or identify subtle temporal patterns that may indicate the early stages of moisture-related problems. Consequently, important developments may remain unnoticed despite being recorded by the monitoring infrastructure.

Artificial intelligence offers an opportunity to transform building monitoring from passive data collection into an active decision-support process. Recent advances in time-series forecasting enable learning complex temporal relationships from historical observations and forecasting future system behaviour. Such systems are not intended to replace human expertise, but instead to continuously analyse incoming measurements and automatically draw attention to situations that may require further inspection.

If successful, this approach could improve preventive maintenance by allowing building managers to identify developing moisture-related risks before visible damage occurs, reducing maintenance costs while improving building durability. Beyond wooden buildings, similar methodologies could support automated monitoring and early warning in other structural health monitoring applications where large volumes of sensor data are continuously collected, including bridges, historical buildings, industrial facilities, and other critical infrastructure.

In this paper, we investigate whether a fine-tuned Granite Time Series Foundation Model can support early moisture risk detection in wooden buildings through short-term forecasting of wood moisture content. A separate Granite model is fine-tuned and evaluated for each building: the InnoRenew CoE building in Slovenia and the Forestry HUB building in Spain. The forecasting performance is compared with statistical and machine-learning baseline methods to assess the suitability of this approach for practical building monitoring applications.

The remainder of this paper is organized as follows. Section 2 reviews related work on wood moisture monitoring, predictive maintenance, and time-series forecasting. Section 3 describes the study buildings, datasets, data preparation, forecasting models, and experimental methodology. Section 4 presents the experimental results and compares the proposed approach with the baseline methods. Finally, Section 5 concludes the paper by summarizing the main findings, discussing the limitations of the study, and outlining directions for future research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Information Society 2026, Ljubljana, Slovenia

© 2026 Copyright held by the owner/author(s).

2 Related Work

Previous studies have established the relationship between wood moisture and long-term durability, showing that prolonged exposure to elevated moisture levels increases the risk of fungal decay and other degradation processes [4]. Other works have focused on continuous monitoring of timber buildings using embedded sensors, wireless monitoring systems, and hygrothermal modelling to analyse moisture evolution under real operating conditions [8, 11]. More recently, machine learning has also been applied to predict moisture behaviour in timber structures, for example through Support Vector Regression models for long-term hygrothermal performance under future climate scenarios [2]. Together, these studies demonstrate the importance of continuous moisture monitoring for preserving the durability of wooden buildings. However, they primarily focus on moisture measurement, retrospective analysis, or long-term prediction rather than operational forecasting capable of automatically identifying developing moisture-related risks from continuously collected sensor data. In contrast, our work investigates whether a pre-trained time-series foundation model can forecast future moisture development and provide automated early warnings for building monitoring.

Artificial intelligence has become an increasingly important tool for predictive maintenance and forecasting in buildings. Bouabdallaoui et al. [3] proposed a machine-learning framework for predictive maintenance of building facilities, while Guo et al. [7] demonstrated that deep learning models can accurately forecast multivariate building energy consumption. Kramar and Alchakov [9] further showed that forecasting performance depends strongly on the ability of models to capture recurring seasonal patterns. Although these studies demonstrate the potential of artificial intelligence for analysing building-related time series, they either address different application domains or rely on task-specific forecasting models. Our work instead focuses on operational moisture monitoring in wooden buildings and investigates whether a pre-trained foundation model can be adapted through fine-tuning rather than developing a dedicated forecasting architecture from scratch.

Recent advances in time-series foundation models have introduced a new generation of forecasting models pre-trained on large and diverse collections of time-series data. Chronos reformulated forecasting as a language modelling problem [1], while Das et al. [5] proposed a decoder-only foundation model trained on heterogeneous datasets. Ekambaram et al. [6] introduced Tiny Time Mixers, demonstrating that lightweight pre-trained models can achieve competitive forecasting performance with substantially lower computational requirements. More recently, Park et al. [10] demonstrated that fine-tuning foundation models significantly improves forecasting performance for building energy systems compared with zero-shot inference. These studies confirm the strong potential of foundation models for time-series forecasting, yet their application to operational wood moisture forecasting remains largely unexplored. Building upon this recent progress, our work evaluates the IBM Granite Time Series Foundation Model for forecasting wood moisture content using long-term monitoring data collected from two wooden buildings and compares its performance with conventional statistical and machine-learning forecasting methods.

3 Methodology

3.1 Data Preparation

The experiments presented in this work were conducted using monitoring data collected from two wooden buildings: the InnoRenew CoE building in Slovenia and the Forestry HUB building in Spain. Both buildings are equipped with monitoring systems that continuously measure the condition of wood.

The preprocessing pipeline was designed to create a consistent forecasting dataset. After hourly aggregation, the measurements from each sensor were organised into independent multivariate time series. Wood moisture content was selected as the forecasting target, while temperature, relative humidity, absolute humidity, and temporal variables were used as model inputs. Since temporal variables are cyclical by nature, the month, hour of the day, and day of the year were encoded using sine and cosine transformations.

Each sensor time series was divided into four seasonal subsets corresponding to winter, spring, summer, and autumn. For every season, the final 96 hours were reserved as the test period, while the preceding 768 hours were used for validation and the remaining observations for training. To make better use of the limited seasonal data, the 512-hour input context required for testing was taken from the last 512 hours of the validation period immediately preceding the test horizon.

The available datasets also contained periods with missing observations. The InnoRenew CoE dataset was missing approximately two weeks of measurements during autumn, while the Forestry HUB dataset was missing almost one month of summer observations. Consequently, fewer samples were available for training and evaluation in these seasons, which affected the forecasting performance.

3.2 Granite Time Series Foundation Model

The primary forecasting model evaluated in this study is the IBM Granite Time Series Foundation Model (Granite TTM), a pre-trained deep learning model designed for multivariate time-series forecasting. Instead of being trained from scratch for a single application, Granite is first pre-trained on a large collection of time-series datasets and can then be adapted to a specific forecasting task through fine-tuning. This approach enables the model to benefit from previously learned temporal representations while requiring only a relatively small amount of task-specific training data. The experiments were performed using the publicly available IBM Granite Time Series Foundation Model (ibm-granite/granite-timeseries-ttm-r2, revision 512-96-ft-r2.1), which supports a context length of 512 hours and a prediction horizon of 96 hours.

3.3 Baseline Models

To evaluate the forecasting performance of the Granite model, its predictions were compared with four baseline approaches representing different levels of forecasting complexity: two persistence-based methods and two commonly used machine-learning algorithms.

The first baseline is the *daily-naive* model, which assumes that future moisture values follow the same pattern observed 24 hours earlier. Each prediction is obtained by repeating the corresponding observation from the previous day, providing a simple reference for short-term forecasting.

The second baseline is the *weekly-naïve* model, which predicts future values using observations recorded at the same time one week earlier. Compared with the daily-naïve approach, it captures longer repeating patterns associated with weekly environmental cycles.

The third baseline is a multivariate Linear Regression model, which estimates future wood moisture as a linear combination of the input variables. All input variables were standardized using the mean and standard deviation computed from the training data, and the same transformation was applied to the validation and test data to ensure consistent feature scaling.

The final baseline is a Random Forest regressor consisting of 500 decision trees with a maximum depth of 20 and a minimum leaf size of two samples. Random Forest models nonlinear relationships by combining the predictions of multiple decision trees, making it capable of capturing more complex interactions between wood moisture, environmental conditions, and temporal variables than a linear model.

All baseline models were evaluated using the same training, validation, and test datasets and generated forecasts for the same 96-hour prediction horizon. This ensured that differences in forecasting performance were attributable solely to the forecasting methods.

3.4 Experimental Design

Three experiments were conducted. The first experiment evaluated the Granite model against the daily-naïve, weekly-naïve, Linear Regression, and Random Forest baselines using the InnoRenew CoE dataset. The second experiment applied the same experimental methodology to the Forestry HUB dataset, allowing the forecasting performance to be evaluated under a different building and climatic environment. The third experiment evaluated whether the proposed forecasting system can detect an unexpected increase in wood moisture using an artificial moisture increase introduced into the test data.

For the first and second experiments, the Granite model was fine-tuned for 100 epochs, while validation was performed using a sliding-window approach with multiple overlapping input-target windows. The checkpoint achieving the lowest validation loss was selected for the final evaluation. The selected checkpoint was then evaluated on the final 96-hour test horizon using the last 512-hour historical context from the validation period. This procedure was repeated independently for winter, spring, summer, and autumn to evaluate forecasting performance across different environmental conditions.

For the third experiment, an artificial moisture increase was introduced only into the measured values of the final 96-hour winter test period. Since neither monitored building had abnormally high moisture levels during the observation period, no real moisture event suitable for evaluating the proposed alerting procedure was available. The artificial sequence was therefore generated by adding a smooth nonlinear drift based on a Smooth-step interpolation function together with small temporally correlated fluctuations to the original measurements, simulating a gradually developing moisture increase. At each forecast step, the absolute difference between the Granite prediction and the modified observation was calculated. A deviation alert was raised when this difference exceeded the selected threshold for at least three consecutive hours, while a separate alert was generated when the modified moisture value exceeded the predefined safety threshold.

3.5 Evaluation Metrics

Forecasting performance was evaluated using two widely adopted regression metrics: Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE). Both metrics quantify the difference between the predicted and observed wood moisture content over the 96-hour forecasting horizon. For each forecasting model, MAE and RMSE were computed separately for every season. In addition, overall MAE and RMSE values were obtained by averaging the results across all seasonal forecasting experiments, providing an overall comparison of model performance.

4 Results

Tables 1 and 2 summarize the overall forecasting accuracy of all evaluated methods, while Figure 1 presents the corresponding seasonal forecasts for both datasets. On the InnoRenew CoE dataset, Granite consistently outperformed both persistence-based baselines and achieved forecasting accuracy comparable to the classical machine learning models. Although Linear Regression obtained the lowest overall MAE and RMSE, the difference compared with Granite was small. The seasonal forecasts further confirm these results, showing that Granite consistently reproduced both the overall trend and short-term moisture variations across all seasons.




Figure 1: Seasonal forecasting results for the InnoRenew and ForestryHub datasets.

The ForestryHub dataset proved more challenging due to the higher variability of the measured wood moisture. Granite produced accurate forecasts during winter and autumn and remained competitive during spring. The main exception was summer, where the measured moisture dropped abruptly immediately after the beginning of the forecasting horizon (Figure 1), creating a




Figure 2: Artificial moisture increase experiment.

pattern that was not present in the preceding input context. This season also contained the smallest amount of available training data, which likely contributed to the higher prediction errors. Consequently, Granite ranked behind the persistence-based baselines in the overall MAE and RMSE reported in Table 2. Nevertheless, the seasonal forecasts indicate that Granite followed the overall evolution of the moisture signal more closely than the baseline methods, which often repeated previously observed patterns or gradually diverged from the measured trend.

Table 1: Overall forecasting performance on the InnoRenew CoE dataset.

Method	MAE	RMSE
Granite	0.0488	0.0643
Daily naive	0.1137	0.1243
Weekly naive	0.1740	0.2149
Linear Regression	0.0483	0.0627
Random Forest	0.0572	0.0745

Table 2: Overall forecasting performance on the ForestryHub dataset.

Method	MAE	RMSE
Granite	0.1551	0.2205
Daily naive	0.1538	0.2123
Weekly naive	0.1522	0.1879
Linear Regression	0.1784	0.3135
Random Forest	0.2055	0.2810

The artificial moisture increase experiment further demonstrated the practical applicability of the proposed forecasting framework for building monitoring. As illustrated in Figure 2, the modified measurements gradually deviated from the Granite forecast and eventually exceeded the predefined safety threshold. By continuously comparing predicted and measured moisture values, the monitoring procedure successfully detected the developing moisture increase and generated an alert before the end of the forecasting horizon. The figure highlights the first detected alert, which indicates the moment when the predefined alert conditions were first satisfied.

5 Conclusion

This work investigated the applicability of the Granite TinyTimeMixer foundation model for short-term wood moisture forecasting using two independent datasets collected from wooden buildings. The proposed approach was evaluated against naive forecasting methods and classical machine learning models using

seasonal forecasting experiments with a 96-hour prediction horizon. The experimental results showed that Granite consistently ranked among the best-performing forecasting methods across both datasets. On the InnoRenew CoE dataset, its forecasting accuracy closely matched the best-performing classical machine learning model, while on the ForestryHub dataset it remained competitive despite the higher variability of the measurements. In addition to the numerical evaluation, the seasonal forecast curves showed that Granite generally reproduced the temporal evolution of wood moisture more faithfully than the baseline methods. The artificial moisture increase experiment showed that the proposed approach can detect abnormal moisture development and issue an alert when predefined conditions are met.

Overall, the results indicate that pretrained foundation models represent a promising alternative for wood moisture forecasting and that their predictions can be effectively combined with threshold-based monitoring to support the early identification of abnormal moisture behaviour. Future work should investigate training on larger datasets, including additional buildings and longer observation periods, to further improve the robustness and generalization of forecasting.

References

- [1] Abdul Fatir Ansari et al. 2024. Chronos: learning the language of time series. *Transactions on Machine Learning Research*. <https://arxiv.org/abs/2403.07815>.
- [2] Naman Bansal, Maurice Deho, and Michael Lacasse. 2021. Application of Support Vector Regression to the prediction of the long-term impacts of climate change on the moisture performance of wood frame and massive timber walls. *Buildings*, 11, 5, 188. doi:10.3390/buildings11050188.
- [3] Yassine Bouabdallaoui, Zoubeir Lafhaj, Pascal Yim, Laure Ducoulombier, and Belkacem Bennadji. 2021. Predictive maintenance in building facilities: a machine learning-based approach. *Sensors*, 21, 4, 1044. doi:10.3390/s21041044.
- [4] Christian Brischke and Gry Alfredsen. 2020. Wood-water relationships and their role for wood susceptibility to fungal decay. *Applied Microbiology and Biotechnology*, 104, 3781–3795. doi:10.1007/s00253-020-10479-1.
- [5] Abhimanyu Das, Weihao Kong, Rajat Sen, and Yichen Zhou. 2024. A decoder-only foundation model for time-series forecasting. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*. PMLR. <https://arxiv.org/abs/2310.10688>.
- [6] Vijay Ekambaram, Arindam Jati, Pankaj Dayama, Sumanta Mukherjee, Nam H. Nguyen, Wesley M. Gifford, Chandra Reddy, and Jayant Kalagnanam. 2024. Tiny Time Mixers (TTMs): fast pre-trained models for enhanced zero/few-shot forecasting of multivariate time series. In *Advances in Neural Information Processing Systems (NeurIPS)*. NeurIPS Foundation. <https://arxiv.org/abs/2401.03955>.
- [7] Jing Guo, Penghui Lin, Limao Zhang, Yue Pan, and Zhonghua Xiao. 2023. Dynamic adaptive encoder-decoder deep learning networks for multivariate time series forecasting of building energy consumption. *Applied Energy*, 350, 121803. doi:10.1016/j.apenergy.2023.121803.
- [8] Steven Kordziel, Samuel V. Glass, Shiling Pei, Samuel L. Zelinka, and Paulo Cesar Tabares-Velasco. 2018. Moisture monitoring and modeling of mass timber building systems. In *Proceedings of the World Conference on Timber Engineering (WCTE)*. World Conference on Timber Engineering. <https://api.semanticscholar.org/CorpusID:198425188>.
- [9] Vadim Kramar and Vasily Alchakov. 2023. Time-series forecasting of seasonal data using machine learning methods. *Algorithms*, 16, 5, 248. doi:10.3390/a16050248.
- [10] Young-Jin Park, François Germain, Jing Liu, Ye Wang, Toshiaki Koike-Akino, Gordon Wichern, Navid Azizan, Christopher Laughman, and Ankush Chakrabarty. 2025. Probabilistic forecasting for building energy systems using time-series foundation models. (2025). arXiv: 2506.00630 [Cs.LG]. doi:10.48550/arXiv.2506.00630.
- [11] Maryam Shirmohammadi, Marcus Strang, Paola Leardini, Kim Baber, and Adam Faircloth. 2025. Moisture monitoring of mass timber building – study of condition variation and building environment design. In *Proceedings of the World Conference on Timber Engineering (WCTE)*, 2589–2598. doi:10.52202/080513-0316.