

Article

Not peer-reviewed version

Spatio-Temporal Graph Autoencoder for Sensor Data Reconstruction in Vineyard Microclimate Monitoring

[Filippo Costanti](#)*, [Irene Cappelli](#), [Monica Bianchini](#), [Ada Fort](#)

Posted Date: 23 April 2026

doi: 10.20944/preprints202604.1685.v1

Keywords: denoising autoencoder; environmental monitoring; gated recurrent unit with decay; graph sample and aggregate; internet of things; time series reconstruction



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Spatio-Temporal Graph Autoencoder for Sensor Data Reconstruction in Vineyard Microclimate Monitoring

Filippo Costanti * , Irene Cappelli, Monica Bianchini and Ada Fort

Department of Information Engineering and Mathematics, University of Siena, Siena, Italy

* Correspondence: filippo.costanti@unisi.it

Abstract

Continuous monitoring of microclimatic variables is essential for precision viticulture and data-driven decision support systems. However, agricultural sensor networks are frequently affected by missing data due to hardware failures, communication issues, or maintenance interruptions. In this work, we propose a spatio-temporal graph-based autoencoder for reconstructing missing temperature and relative humidity time series collected from a five-node vineyard sensor network over a two-year period. The model combines a GRU-D-based temporal encoder with a GraphSAGE spatial module, enabling the joint exploitation of temporal dynamics and inter-node spatial correlations. Experimental results on real-world data show that the proposed approach achieves accurate reconstruction under realistic missing-data conditions. For moderate corruption levels ($p = 0.3$), the model attains reconstruction losses of 0.003 for temperature and 0.005 for humidity using short temporal windows ($L = 36 \sim 3$ h), corresponding to MAE values below 0.03°C and 0.1%, respectively. Even at higher corruption levels ($p = 0.7$), performance remains stable, with losses below 0.008 and 0.011, and MAE values within 0.05°C and 0.17%. The results highlight a trade-off between temporal context and reconstruction stability: shorter windows yield better performance under moderate corruption, while longer windows ($L = 144 \sim 12$ h) improve robustness under extreme data loss ($p = 0.9$), reducing temperature reconstruction loss from 0.027 to 0.021 and MAE from 0.133° to 0.226° . Additionally, temperature is consistently reconstructed more accurately than humidity, reflecting its smoother dynamics and stronger spatial coherence.

Keywords: denoising autoencoder; environmental monitoring; gated recurrent unit with decay; graph sample and aggregate; internet of things; time series reconstruction

1. Introduction

Precision viticulture has progressively transitioned from empirical management practices to data-driven decision frameworks supported by distributed sensing infrastructures [1]. Within vineyard ecosystems, microclimatic variables such as air temperature, relative humidity (RH), solar radiation, and precipitation directly influence plant physiology, evapotranspiration rates, pathogen development, and fruit quality [2]. Since these processes are highly sensitive to local environmental fluctuations, continuous and spatially distributed monitoring is essential to capture the fine-scale variability that characterizes vineyard terrains.

Wireless sensor networks (WSNs) based on low-power internet of things (IoT) technologies have become fundamental for smart agriculture systems [3]. By enabling long-term deployment of sensing nodes across heterogeneous plots, these platforms generate high-resolution time series that can feed predictive models for irrigation scheduling, disease risk estimation, and yield optimization. In this context, the value of the sensing infrastructure is not limited to data acquisition alone; rather, it lies in its integration within a complete digital pipeline that includes storage, preprocessing, modeling, and decision support.

However, real-world IoT deployments are inherently exposed to reliability constraints. Battery-powered nodes operating under outdoor conditions may experience energy instabilities, sensor degradation, or hardware faults. At the communication layer, low-power wide-area network (LPWAN) protocols such as long range wide-area network (LoRaWAN) introduce duty-cycle limitations, and latency variability. These phenomena may result in irregular sampling, missing observations, and node-specific anomalies in the collected datasets. In terms of IoT reliability, these aspects should be regarded as inherent characteristics of distributed, resource-constrained systems rather than exceptional events [4].

Although these issues may appear primarily infrastructural, originated at the sensing and communication layers, they have direct implications for data-driven modeling. Both classical statistical models and modern machine learning (ML) approaches are sensitive to the temporal structure of the input data. Irregular sampling and missing observations, if not explicitly modeled, can alter the effective dependency structure of the series, leading to biased inference, degraded forecast accuracy, or reduced training stability. These effects are particularly relevant when temporal alignment across multiple sensing nodes is required as in multi-node vineyard deployments. Indeed, inconsistent timestamps or heterogeneous data gaps across nodes can hinder comparability and the detection of physically meaningful microclimatic differences. Consequently, inadequate handling of missing or irregular data does not merely reduce the dataset size, it alters the statistical and dynamical properties of the signals, with direct effect on the prediction accuracy [5].

Sensor data reconstruction therefore represents a critical interface between IoT system reliability and ML robustness. Its role extends beyond simple gap filling: it ensures that the statistical properties of the reconstructed signals remain coherent with the physical dynamics of vineyard microclimate [6]. Temperature and relative humidity, for instance, exhibit smooth diurnal cycles governed by radiative forcing and atmospheric mixing processes. Reconstruction strategies must preserve these dynamics without introducing artificial oscillations or attenuating variability. From a modeling perspective, reconstruction directly influences feature quality, temporal dependencies, and spatial comparability across nodes. Inadequate handling of missing data may lead deep learning models to learn artifacts rather than true environmental patterns. Conversely, physically informed interpolation and consistent temporal alignment can enhance the stability of training procedures and improve predictive accuracy.

In this work, we address missing data reconstruction using real-field air temperature and RH measurements collected from a five-node vineyard sensor network deployed under field conditions over a two-year period [7]. To jointly exploit temporal dynamics and spatial correlations among neighboring sensors, we propose a spatio-temporal graph-based autoencoder. The model combines a GRU-D module to capture temporal dependencies under irregular sampling and missing observations, with a GraphSAGE-based spatial encoder operating on the graph representation of the sensor network. Experimental results support the view that, in distributed agricultural monitoring systems, reliable data reconstruction benefits from explicitly integrating network topology and environmental dynamics within a unified learning framework.

The paper is organized as follows. Section 2 reviews related work on missing data reconstruction and spatio-temporal modeling in environmental sensor networks. Section 3 describes the vineyard sensing infrastructure, the dataset collection and preprocessing, and the proposed spatio-temporal graph autoencoder. Section 4 presents the experimental results while Section 5 discusses the obtained outcomes. Finally, Section 6 draws the conclusions and outlines directions for future research.

2. Related Works

Recent research on sensor network modeling has progressively moved toward structure-aware learning architectures that explicitly incorporate both spatial topology and temporal dynamics into deep neural models. In this context, spatio-temporal graph learning frameworks have emerged as an effective paradigm for representing dependencies among geographically distributed sensors and irregular temporal observations.

Spatio-Temporal Graph Neural Networks (STGNNs) extend classical time-series models by integrating graph structures that encode spatial relationships among nodes. Early architectures such as the Spatio-Temporal Graph Convolutional Network (STGCN) [8], the Diffusion Convolutional Recurrent Neural Network (DCRNN) [9], and Graph WaveNet [10] demonstrated that combining graph-based spatial modeling with temporal learning mechanisms significantly improves forecasting accuracy in networked time-series data. These models typically integrate graph convolutional operators with recurrent modules or temporal convolutional layers in order to jointly capture spatial correlations and temporal dependencies [11].

Compared with traditional sequence-based models, spatio-temporal embeddings provide a richer representation of environmental sensing systems. In many real-world monitoring scenarios, measurements collected at different locations are not independent but are governed by spatially correlated physical processes such as atmospheric circulation, hydrological flows, or soil moisture diffusion. By embedding sensor observations into a graph representation, STGNNs explicitly model these spatial interactions, allowing information to propagate across neighboring nodes and enabling the network to capture both local and global spatial patterns. This joint spatial-temporal representation is particularly advantageous for environmental forecasting tasks, where the evolution of a variable at a given location depends not only on its historical behavior but also on the conditions of nearby regions. As a result, STGNN-based models have demonstrated strong performance in a variety of environmental monitoring applications, including air quality forecasting (e.g., PM_{2.5} prediction) [12], soil moisture estimation [13,14], and hydrological forecasting systems [15,16].

In parallel to spatio-temporal graph models, deep learning models based on temporal convolutional architectures have also been explored for environmental data reconstruction in sensor network applications. In particular, WaveNet-inspired models leveraging dilated causal convolutions have been proposed to infer unmeasured environmental variables from correlated meteorological observations, effectively capturing non-linear dependencies and long-range temporal patterns in vineyard monitoring datasets [17]. These approaches demonstrate that predictive reconstruction of environmental quantities can be interpreted as an implicit missing data recovery task, highlighting the potential of data-driven sequence modeling techniques in real-world agricultural sensing systems.

More recently, the use of self-supervised learning techniques have been explored, to improve the robustness of spatio-temporal models when dealing with noisy or incomplete sensor data. In particular, masked autoencoder strategies have gained attention as a general framework for representation learning. Originally introduced for visual representation learning [18], masked autoencoders have subsequently been extended to spatio-temporal data by reconstructing partially masked observations in both spatial and temporal domains [19]. Within the graph learning domain, several works have proposed masked graph autoencoder architectures that combine graph message passing with self-supervised reconstruction tasks. For instance, the Spatio-Temporal Graph Masked Autoencoder (STGMAE) framework introduces a generative self-supervised paradigm in which both node features and graph structures are partially masked during training. The model learns to reconstruct the missing information, enabling more robust latent representations for noisy and sparse spatio-temporal datasets [20]. These approaches have shown promising results in urban sensing, mobility analysis, and environmental monitoring scenarios [21,22].

However, handling missing observations remains a major challenge in sensor network. Traditional imputation techniques, such as statistical interpolation or simple forward-filling strategies, often fail to capture the temporal dynamics of irregularly sampled environmental data. For this reason, several studies have proposed learning-based approaches in which neural networks are used to reconstruct or infer missing information directly from the temporal structure of the data. Recurrent architectures, in particular, have proven effective for modeling temporal dependencies and identifying anomalous patterns in meteorological time series [23].

A widely adopted approach for handling missing data in time series is the GRU-D framework, which introduces time-dependent decay mechanisms to model the influence of past observations as a

function of the elapsed time since the last measurement [24]. In this formulation, the contribution of previously observed values progressively decreases as the time gap increases, while statistical baselines derived from historical data are used to guide the reconstruction of missing values. More recently, these temporal decay concepts have been extended to dynamic graph learning frameworks, where missing data may affect not only node features but also the underlying graph structure. In such settings, hybrid models combining graph reconstruction mechanisms with temporal sequence modeling have been proposed to handle incomplete graph snapshots caused by communication failures or sensor outages [25,26].

These developments highlight the growing relevance of graph-based deep learning frameworks for modeling complex sensor networks characterized by spatial dependencies, temporal irregularity, and incomplete observations. However, the joint modeling of physical spatial topology, temporal irregularity, and fault-induced missing data remains a largely open problem in distributed environmental sensor network learning. Motivated by these considerations, the present work introduces a spatio-temporal graph-based autoencoder designed for the reconstruction of missing environmental observations in vineyard sensor networks. The proposed framework explicitly combines temporal decay-aware sequence modeling with spatial message passing over a graph reflecting the physical topology of the sensing infrastructure. By jointly learning temporal dynamics within individual sensor streams and spatial dependencies among neighboring nodes, the model aims to provide robust and physically consistent reconstruction of temperature and humidity time series under realistic conditions of data loss and irregular sampling.

3. Materials and Methods

3.1. Hardware Platform for Environmental Data Collection

The datasets considered in this study were acquired using a custom-developed wireless sensing platform for precision agriculture applications. The hardware architecture and validation procedures have been extensively described in previous contributions [7,17,23,27]; therefore, only a concise overview of the system configuration and deployment is provided.

The platform is composed of five autonomous, battery-powered sensing nodes designed for long-term operation in outdoor vineyard environments. The system was engineered with particular attention to low power consumption, reduced implementation cost through commercial off-the-shelf components, and modular design, enabling flexible and scalable deployment strategies.

Each sensing node integrates multiple environmental sensors for microclimatic monitoring. Air temperature and RH are measured by a digital SHT30 sensor (Sensirion), selected for its stability and compatibility with ultra-low-power embedded systems. Solar irradiance is estimated through cadmium sulfide (CdS) light-dependent resistors (Advanced Photonics), whose radiometric behavior under field conditions was characterized in previous studies [7]. Leaf wetness is detected using resistive interdigitated probes (MH-RD type) positioned to sense moisture conditions on both upper and lower leaf surfaces. Prior electrical characterization of the resistive sensors allowed impedance variations to be mapped into a discrete wetness index used for subsequent data analysis. Precipitation is monitored through an integrated tipping-bucket rain gauge, enabling site-specific rainfall measurements. The processing unit is based on an STM32L073RBT6 ultra-low-power microcontroller, supporting duty-cycled acquisition and communication to extend battery lifetime. Wireless connectivity is provided by a HopeRF RFM95x LoRa transceiver operating in the 868 MHz ISM band. A schematic representation of the node architecture is reported in Figure 1.

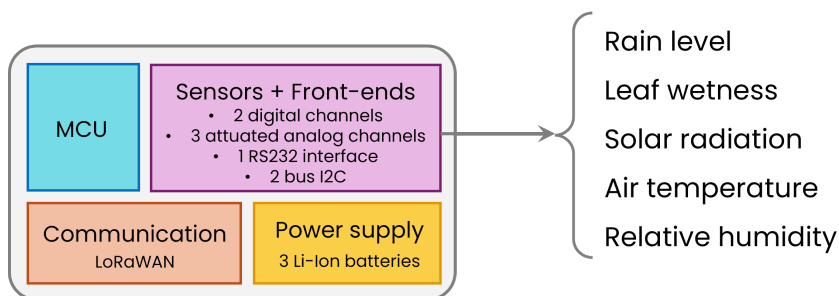


Figure 1. Schematic representation of the building blocks composing the realized sensor node.

Five wireless sensors were deployed at the Istituto Tecnico Agrario B. Ricasoli (Siena, Italy), as shown in Figure 2, enabling continuous monitoring of vineyards microclimatic conditions over an extended observation window spanning from July 2023 to May 2025. The sensors' readings were transmitted every five minutes according to a periodic sampling scheme optimized for energy efficiency and communication reliability. The packets were then collected by an outdoor gateway installed on the main farm building and forwarded to a ChirpStack LoRaWAN network server for storage and management. Remote access, visualization, and inspection of the acquired measurements are provided through a Grafana-based dashboard interface. While the sensing infrastructure monitored several environmental parameters, the analysis presented in this work is restricted to the air temperature and RH time series. Representative temporal profiles for these quantities are illustrated in Figure 3 for about two weeks of test. The temperature series exhibit the expected diurnal oscillations, with peak values occurring during daytime and minima during nighttime, while RH shows an inverse trend. The five nodes display a high degree of temporal coherence, with only minor differences attributable to site-specific environmental conditions and variability across vineyard locations. Indeed, although the sensing units were deployed within a relatively confined geographical area (inter-node distance < 500 m), they were installed in vineyard plots characterized by different elevations and solar exposures, resulting in moderate microclimatic variability. At the same time, the spatial arrangement of the sensing nodes can induce similarities in the temporal evolution of the recorded variables, as sensors located closer to each other are likely to experience comparable environmental conditions. This observation suggests that spatial information can be exploited to support data-driven modeling of the sensor measurements. To this end, the sensing network is represented as a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where each vertex $v \in \mathcal{V}$ corresponds to a sensor node and edges $e \in \mathcal{E}$ encode relationships among sensors. In particular, the graph topology is defined according to a spatial relationship due to geographical proximity, by connecting each node to its two nearest neighbors based on the straight-line distance between sensor locations. The corresponding graph is reported in Figure 2.




Figure 2. Geographical map of the sensor network deployed in the chosen test site and corresponding graph representation defined according to a spatial proximity relationship (i.e., each node is connected to its two nearest neighbors in terms of straight-line distance).

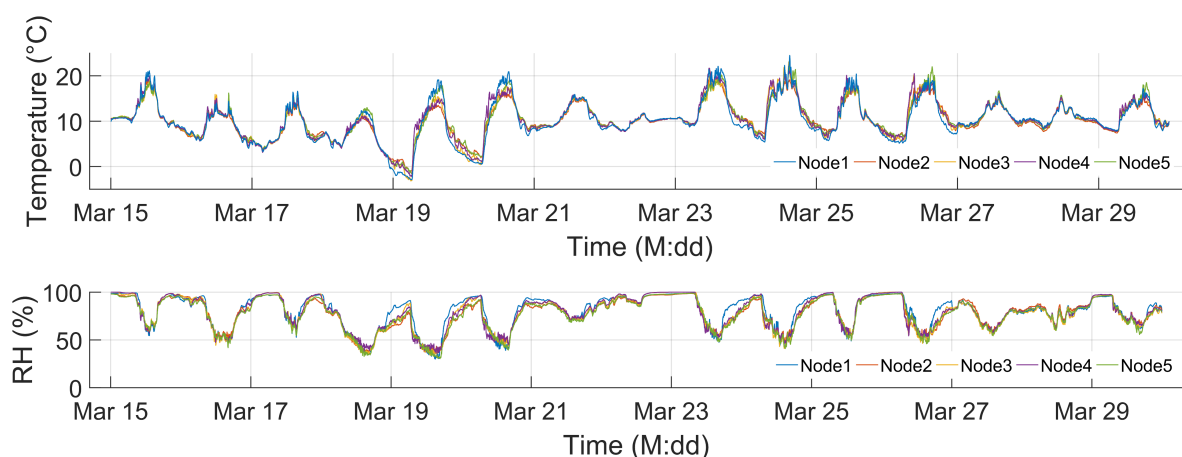


Figure 3. Temperature and RH temporal evolution for the five nodes (as per legend) during 15 days of tests in March 2025.

3.2. Neural Network Architecture

A Spatio-Temporal Graph Neural Network Autoencoder (STGNN-AE) is developed to reconstruct multivariate time series collected from interconnected sensing units and naturally represented as a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. The network architecture is shown in Figure 4.

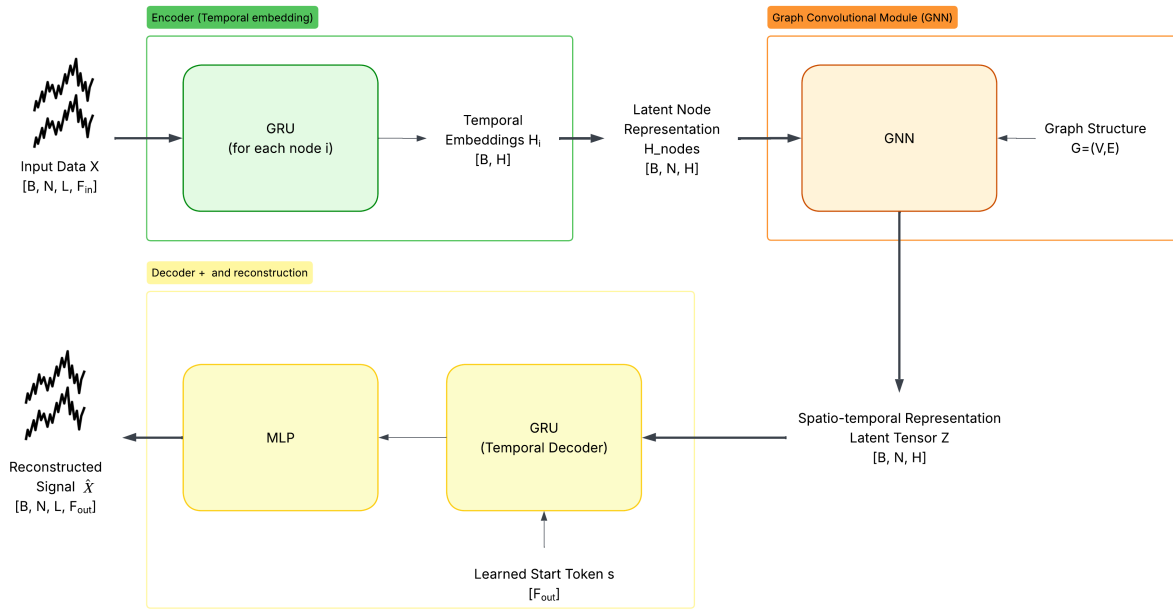


Figure 4. Network structure.

Specifically, temporal patterns are first encoded independently at the node level in order to extract compact representations of local dynamics. These node-wise embeddings are then combined through graph-based message passing, allowing spatial dependencies to be incorporated into a unified latent representation. In this way, both local temporal evolution and cross-node interactions are simultaneously captured.

Temporal Node Encoder

The input to the model is organized as a four-dimensional tensor $\mathbf{X} \in \mathbb{R}^{B \times N \times L \times F_{in}}$, where B denotes the batch size, $N = |\mathcal{V}|$ the number of sensor nodes, L the temporal window length, and F_{in} the number of measured features at each timestep.

For each node $i \in \mathcal{V}$, the corresponding time series $\mathbf{x}_i = \{x_{i,1}, \dots, x_{i,L}\}$ is processed independently by a temporal encoder ϕ_{enc} based on a Gated Recurrent Unit (GRU) [28]. GRU is adopted for its ability to efficiently model sequential dependencies, while mitigating vanishing gradient issues. Through the sequential update of hidden states over $t \in [1, L]$, temporal information is progressively aggregated into a compact latent representation.

At each timestep, two gating mechanisms regulate the information flow. The update gate z_t controls the extent to which the candidate hidden state contributes to the current hidden state, effectively balancing between retaining past information and incorporating new content. The reset gate r_t , instead, modulates the contribution of the previous hidden state in the computation of the candidate activation, enabling the model to selectively forget irrelevant past information.

The recurrent dynamics is defined as:

$$\begin{aligned} z_t &= \sigma(W_z \cdot [h_{t-1}, x_{i,t}] + b_z) \\ r_t &= \sigma(W_r \cdot [h_{t-1}, x_{i,t}] + b_r) \\ \tilde{h}_t &= \tanh(W_h \cdot [r_t \odot h_{t-1}, x_{i,t}] + b_h) \\ h_t &= (1 - z_t) \odot h_{t-1} + z_t \odot \tilde{h}_t \end{aligned}$$

where b_z, b_r, b_h are bias vectors and W_z, W_r, W_h are learnable weight matrices. After processing the entire temporal window, the final hidden state $\mathbf{h}_L \in \mathbb{R}^H$ is retained as the latent temporal embedding

for node i , denoted as \mathbf{H}_i . This embedding summarizes the temporal behavior of the sensor over the considered time horizon and serves as the input to the subsequent spatial modeling stage.

Graph Convolutional Module

To explicitly model the spatial relationships among sensors, a Graph Neural Network (GNN) module is incorporated into the architecture. In this work, the spatial component of the model is implemented using the GraphSAGE framework [29], an inductive graph representation learning method designed to efficiently propagate information across neighboring nodes. GraphSAGE learns node representations by aggregating information from the local neighborhood of each node and combining it with the node own features. This inductive formulation allows the model to generalize to unseen nodes or graph structures, which is particularly advantageous in dynamic sensor networks where node availability may change over time. Formally, at layer $l + 1$, the representation of node v is updated according to:

$$h_v^{(l+1)} = \sigma\left(\mathbf{W} \cdot \text{concat}\left(h_v^{(l)}, \text{AGG}\left(\{h_u^{(l)} \mid u \in \mathcal{N}(v)\}\right)\right)\right)$$

where $\mathcal{N}(v)$ denotes the set of neighbors of node v , \mathbf{W} is a learnable weight matrix, and $\sigma(\cdot)$ is a nonlinear activation function. The aggregation operator $\text{AGG}(\cdot)$ is permutation-invariant (e.g., mean or pooling), ensuring robustness with respect to the ordering of neighboring nodes. In this work, mean aggregation is adopted due to its stability and robustness to noisy measurements.

Spatio-Temporal Fusion in Latent Space

Following temporal encoding, the node-level embeddings are arranged into $\mathbf{H}_{nodes} \in \mathbb{R}^{B \times N \times H}$, with $\mathbf{H}_i \in \mathbb{R}^{B \times H}$ denoting the representation of node i . These embeddings are then processed through m layers of graph message passing, yielding the spatially enriched latent tensor \mathbf{Z} :

$$\mathbf{Z} = \text{GNN}(\mathbf{H}_{nodes}, \mathcal{E})$$

This operation enables information exchange among spatially connected nodes, allowing each latent representation to integrate neighborhood context. As a consequence, spatial dependencies among sensors are explicitly modeled within the latent space. The resulting tensor \mathbf{Z} jointly captures temporal dynamics, learned by the GRU encoder, and spatial correlations—modeled through graph message passing. These enriched embeddings are subsequently used to initialize and condition the decoder during the reconstruction process.

No additional graph operations are performed during decoding. This design choice reduces computational complexity while preserving the essential spatial structure encoded in the latent representation, thus ensuring suitability for practical sensor network deployments.

Temporal Decoding and MLP Projection

Sequence reconstruction is carried out by a decoder ψ_{dec} implemented as a GRU. The decoder is initialized from the spatially enriched latent representation \mathbf{Z} , which compactly captures both temporal and spatial information.

Let $\mathbf{S} \in \mathbb{R}^{B \times L \times F_{out}}$ be a tensor obtained by replicating the learned start token $\mathbf{s} \in \mathbb{R}^{F_{out}}$ along the temporal dimension. The decoder dynamics can then be written as:

$$\mathbf{H}_{dec} = \text{GRU}(\mathbf{S}, \mathbf{Z})$$

Through this process, the decoder progressively reconstructs the temporal signal over the entire prediction horizon. The resulting sequence of hidden states \mathbf{H}_{dec} is finally projected back into the original feature space by means of a Multi-Layer Perceptron (MLP), producing the reconstructed signal:

$$\hat{\mathbf{X}} = \text{MLP}(\mathbf{H}_{dec})$$

This projection ensures that the output $\hat{\mathbf{X}}$ matches the dimensionality of the input data, enabling direct comparison between the reconstructed and observed sensor measurements.

Loss Function

To improve reconstruction performance, a composite loss function is employed that combines a conventional point-wise error term with a correlation-based component. Specifically, the proposed objective integrates the Mean Absolute Error (MAE), which penalizes absolute deviations between predicted and target values, with a Pearson correlation loss designed to promote similarity in temporal trends. Correlation-based objectives have previously demonstrated effectiveness in tasks such as neural rendering and depth reconstruction [30,31], where augmented Pearson loss formulations have been shown to enhance local structural consistency between predicted and ground-truth representations. By jointly minimizing amplitude discrepancies and maximizing linear correlation, the proposed loss encourages accurate signal reconstruction while preserving the underlying temporal dynamics.

Given the predictions $\hat{\mathbf{X}}$ and the targets \mathbf{X} , the Pearson correlation coefficient is computed along the temporal dimension after mean-centering both signals. The correlation term encourages alignment in both shape and temporal dynamics, regardless of differences in scale. The corresponding loss component is defined as:

$$\mathcal{L}_{\text{Pearson}} = 1 - \text{mean}(\rho)$$

where ρ denotes the Pearson correlation coefficient computed across the time axis. The total loss is then obtained as the weighted combination:

$$\mathcal{L}_{\text{total}} = \alpha \mathcal{L}_{\text{MAE}} + (1 - \alpha) \mathcal{L}_{\text{Pearson}}$$

where $\alpha \in [0, 1]$ controls the trade-off between amplitude accuracy and temporal trend alignment. This formulation allows the model to simultaneously minimize reconstruction error while preserving the temporal structure of the signals, which is particularly important in sensor-based time series where relative dynamics may carry significant information.

Denoising Training Strategy

The proposed model is trained to reconstruct clean signals from deliberately corrupted inputs. This training strategy encourages the latent representation to capture the intrinsic structure of the data rather than memorizing individual observations, thereby improving robustness to noise and missing measurements.

Input corruption is performed through a structured masking procedure applied to the input tensor \mathbf{X} . For each element in the batch, a single node u is randomly selected from the set of available nodes. A temporal mask vector is then sampled according to a Bernoulli distribution with probability p , which determines the timesteps to be corrupted.

For each selected timestep, the masking operation is applied simultaneously across all feature channels of the chosen node. The corresponding entries are replaced with a constant sentinel value s (set to -4.0 in our implementation), representing missing or corrupted measurements, while all other nodes and unmasked timesteps remain unchanged. This procedure forces the model to infer the corrupted values by exploiting both temporal context and spatial information from neighboring nodes.

The complete data corruption procedure is summarized in Algorithm 1.

Algorithm 1 Structured Node-wise Temporal Masking**Require:** Input batch $\mathbf{X} \in \mathbb{R}^{B \times N \times L \times F_{in}}$, masking probability p , sentinel value s **Ensure:** Corrupted batch $\tilde{\mathbf{X}}$

```

1:  $\tilde{\mathbf{X}} \leftarrow \mathbf{X}$ 
2: for  $b = 1, \dots, B$  do
3:   Sample node index  $n \sim \mathcal{U}(1, N)$ 
4:   Sample mask vector  $\mathbf{m} \in \{0, 1\}^T$  with  $P(m_t = 1) = p$ 
5:   for  $t$  such that  $m_t = 1$  do
6:      $\tilde{\mathbf{X}}[b, n, t, :] \leftarrow s$ 
7:   end for
8: end for
9:
10: return  $\tilde{\mathbf{X}}$ 

```

This node-wise and time-selective corruption strategy simulates realistic missing-data scenarios in sensor networks, where individual sensor nodes may experience temporary data gaps due, for example, to radio transmission issues, or sensors signal degradation. By training the model to recover the original signal from these partially masked inputs, the architecture learns to exploit both temporal continuity and spatial correlations across nodes, thereby improving generalization and reconstruction performance.

3.3. Dataset Preprocessing

The raw measurements generated by each node were transmitted via LoRaWAN and stored at the server side as discrete uplink events. Each event contains the reception timestamp, the unique node identifier, and the associated sensor readings. This packet-based structure guarantees internal temporal consistency among variables acquired within the same transmission cycle. However, the five nodes operate autonomously and are not synchronized. Although configured for a nominal sampling interval of five minutes, the effective inter-arrival times exhibit variability. Consequently, the raw datasets consist of irregularly spaced time series that cannot be directly compared across nodes. In addition, each node may experience distinct failure mechanisms, including sensor-specific malfunctions, power instabilities, or communication-related packet losses intrinsic to LPWAN operation. Consequently, the raw datasets consist of irregularly spaced and node-dependent time series that cannot be directly compared across sensing locations.

Therefore, prior to model development, the pre-processing phase was designed to ensure temporal comparability and data reliability. First, the five independent sequences were mapped onto a common temporal reference grid to obtain aligned time series suitable for cross-node analysis. Subsequently, data quality control procedures were applied. Time intervals corresponding to verified sensor or node anomalies were removed, even when numerical readings were present but considered unreliable. To avoid temporal leakage while preserving seasonal variability, the dataset was split using a month-wise stratified strategy. Each monthly segment was divided chronologically into training (50%), validation (20%), and test (30%) subsets. This approach ensures that sliding windows do not cross discontinuous temporal segments while maintaining representative seasonal patterns in all subsets. Temporal discontinuities were analyzed to distinguish short-term transmission losses from prolonged outages. Short-duration gaps were reconstructed using a GRU-D-inspired imputation strategy [24].

GRU-D Imputation and Time-Dependent Baseline

Missing values were handled using a GRU-D-based imputation mechanism that combines temporal decay with a time-dependent baseline. Specifically, the baseline term $\mu_f(t)$ was defined as the mean value of feature f conditioned on both the month and the hour associated with each timestamp, i.e., $\mu_f(\text{month}, \text{hour})$. Each timestamp was converted into its corresponding month (1–12) and hour (0–23), and feature-wise averages were computed across all observed samples within each month-hour combination, yielding 288 temporal bins (12 months \times 24 hours). For a given timestamp t , when

sufficient observations were available in the corresponding bin, the baseline $\mu_f(t)$ was assigned as the mean value of that bin.

To ensure robustness under sparse data conditions, a hierarchical fallback strategy was adopted. If the month-hour bin contained no observations, the algorithm reverted to the hourly mean aggregated across all months. If this was also unavailable, the monthly mean was used. As a final fallback, the global feature mean over the entire dataset was applied.

Missing values were then reconstructed using an exponential decay formulation:

$$x_{\text{imp}}(t, f) = \gamma(t, f) \cdot x_{\text{last}}(t, f) + (1 - \gamma(t, f)) \cdot \mu_f(t)$$

where $x_{\text{last}}(t, f)$ denotes the most recent observed value and $\gamma(t, f) = \exp(-\lambda\Delta t(t))$ modulates the influence of past observations according to the elapsed time $\Delta t(t)$ since the last valid measurement. The decay rate λ was determined using a fixed half-life criterion of 4 hours, ensuring that the contribution of past observations decreases progressively while avoiding excessively long memory effects. This strategy enables the model to exploit seasonal and diurnal regularities while progressively attenuating outdated information.

In addition, cyclic time-of-day encodings (sine and cosine components) were derived from the timestamps to capture daily periodicity. For each timestamp t , the final feature vector was constructed as:

$$[x_{\text{imp}}(t, f), \Delta t(t), \sin(\frac{2\pi \cdot \text{hour}(t)}{24}), \cos(\frac{2\pi \cdot \text{hour}(t)}{24})]$$

Following feature construction, the processed time series were transformed into samples using a sliding-window approach. Each input window had length L (lookback), and the prediction horizon was set to one step ahead. Different strides were adopted across splits to control window overlap: a high overlap in the training set to increase sample density, reduced overlap in validation, and no overlap in the test set to prevent information leakage. Finally, windows with insufficient true data availability were removed by enforcing a minimum mask-coverage threshold, discarding sequences in which less than 50% of the target values were observed. The data are subsequently normalized using z-score normalization, where each feature is standardized by subtracting the mean and dividing by the standard deviation computed on the training set.

4. Experimental Results

The training procedure was designed to reconstruct a single sensor variable at a time. In particular, the model was trained alternatively to reconstruct either temperature or humidity measurements, allowing the reconstruction performance to be evaluated independently for the two sensing modalities.

The entire framework was implemented in *PyTorch*, enabling efficient training and experimentation with GPU acceleration. All experiments were conducted on a workstation equipped with an NVIDIA GeForce RTX 4090.

After an initial hyperparameter optimization phase performed using the Optuna framework [32], the most promising model configuration was identified. The selected hyperparameters are summarized in Table 1. Optuna was employed to efficiently explore the hyperparameter space and identify the combination of architectural and training parameters that minimized the validation loss.

Table 1. Hyperparameter selected by Optuna framework.

Parameter	Value
latent space dimension	128
layers of GRU	2
layers of GNN	2
drop out	0.0
learning rate	$2 \cdot 10^{-4}$
batch size	32
loss weight (α)	0.6

The hyperparameter search included the dimension of the latent representation h , the number of layers in the GRU encoder module, the number of message-passing layers in the GNN component, and the weighting parameter α used in the loss function. The decoder uses a single GRU layer, while the final projection module is implemented as an MLP composed of three layers with ReLU activations.

The input window length L was set to two values, $L = 36$ and $L = 144$, corresponding to temporal spans of 3 hours and 12 hours, respectively. Subsequent experiments were conducted using both values of L in order to assess the model under different temporal contexts.

Once the optimal hyperparameter configuration was identified, the model was further evaluated using a k -runs training strategy with $k = 10$, in order to mitigate the influence of random initialization and training stochasticity. In each run, the model was trained using the corruption procedure described in Algorithm 1, which applies structured node-wise temporal masking to the input signals as part of the denoising training paradigm. The same hyperparameter configuration was, therefore, employed multiple times using different random seeds. The corruption probability during training was fixed at $p = 0.5$. The model achieving the best test performance was selected as the definitive configuration used for subsequent experiments. Results of the k -runs experiments are reported in Table 2.

Table 2. 10-runs results.

Sensor	Window length (L)	Reconstruction Loss (mean \pm std)
Temperature	36 (3 h)	0.0020 ± 0.0002
Temperature	144 (12 h)	0.0082 ± 0.0003
Humidity	36 (3 h)	0.0024 ± 0.0002
Humidity	144 (12 h)	0.0121 ± 0.0004

In addition to the quantitative evaluation, a qualitative assessment of the model's reconstruction capability was performed on representative samples from the test set. Figs. 5 and 6 present examples of signal reconstruction for temperature and humidity sensors, respectively, using a window length $L = 36$. In these examples, one node is partially corrupted according to the masking strategy (i.e., $p = 0.5$), while the neighboring nodes provide contextual information to the model. The plots illustrate how the proposed spatio-temporal architecture is able to accurately reconstruct the missing measurements by jointly exploiting the temporal dynamics of the corrupted node and the spatial information propagated through the sensor network.

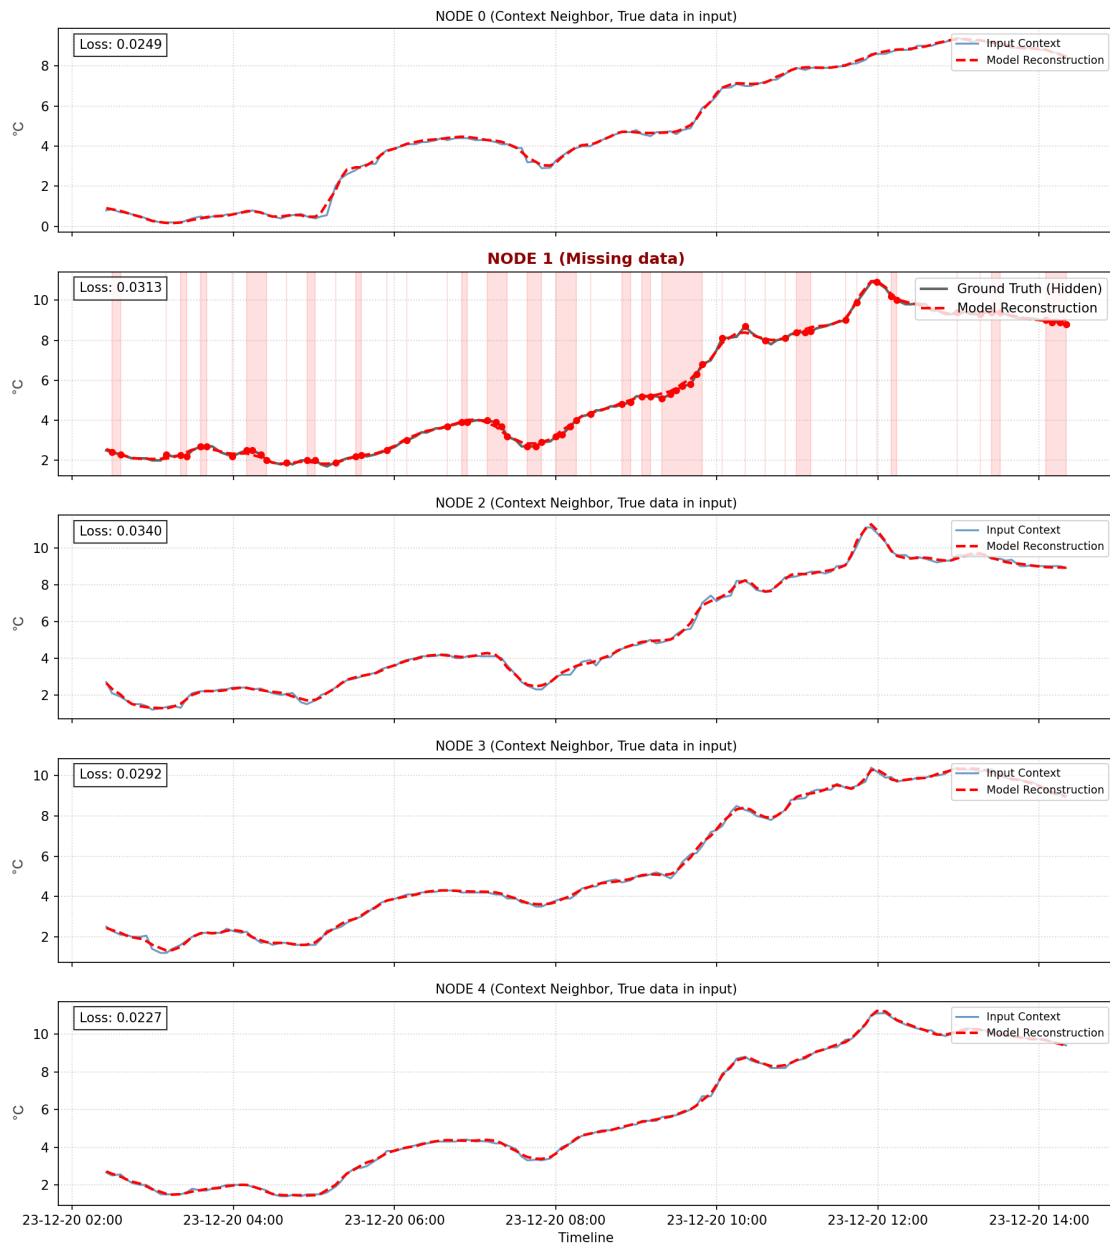


Figure 5. Example of the temperature signal reconstruction for $L = 36$ as per legend. The shaded regions indicate masked intervals where the model must infer the missing values and red dots are the corresponding input points in the ground truth signal. Neighboring nodes provide contextual information to support the reconstruction.

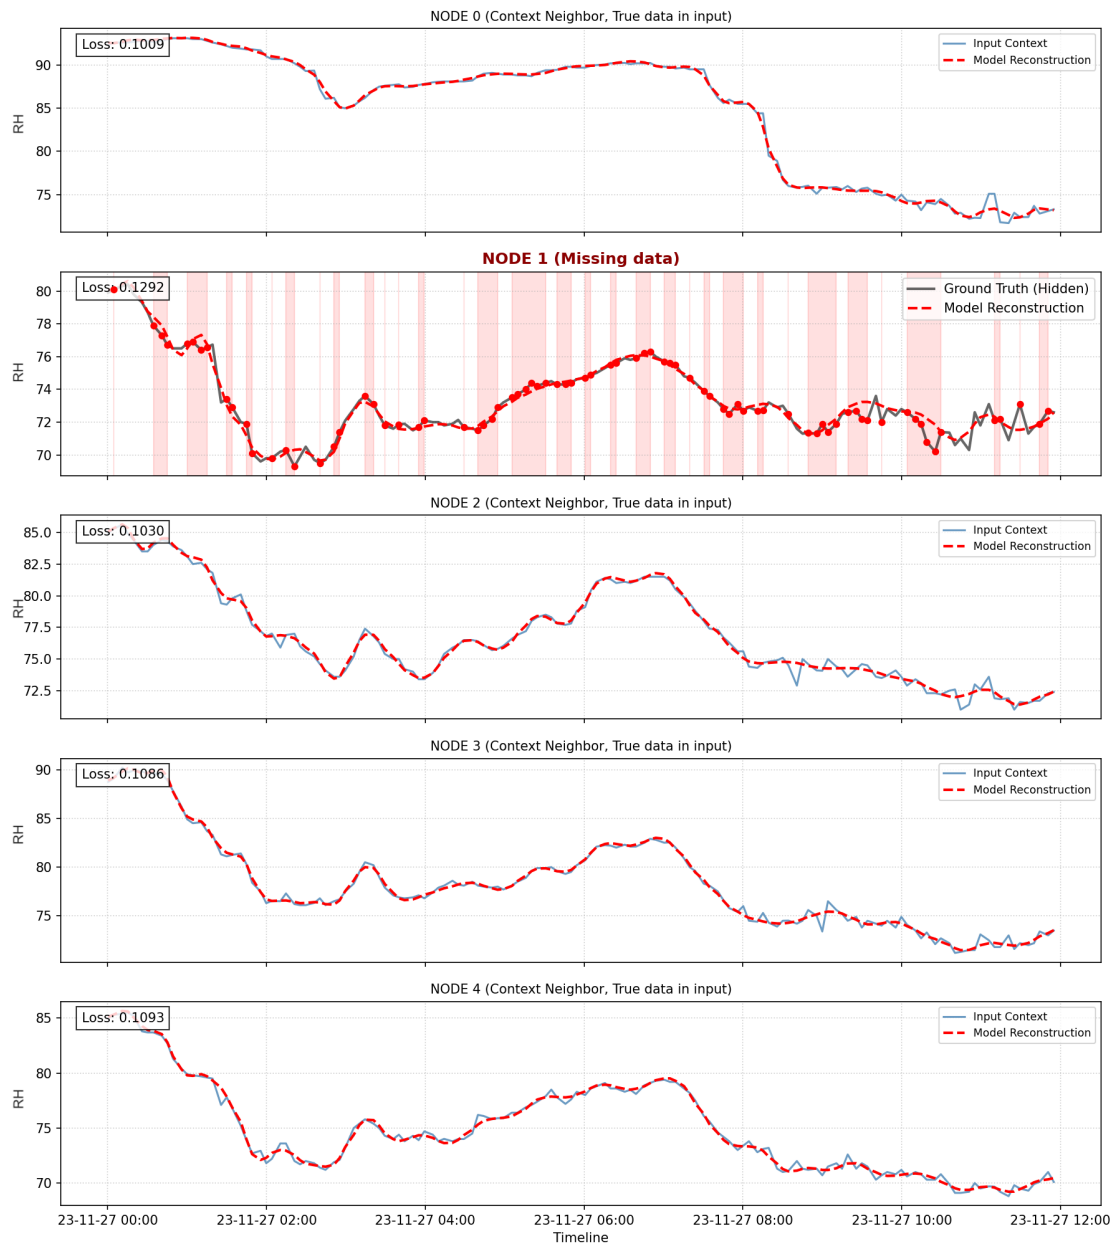


Figure 6. Example of the humidity signal reconstruction for $L = 36$ as per legend. The shaded regions indicate masked intervals where the model must infer the missing values and red dots are the corresponding input points in the ground truth signal. Neighboring nodes provide contextual information to support the reconstruction.

To further assess the robustness of the learned representation, the selected model was evaluated on the test set under varying levels of input corruption, simulating realistic scenarios in which different portions of the signal may be missing or corrupted. Specifically, the structured node-wise temporal masking procedure described in Algorithm 1 was applied using corruption probabilities $p \in \{0.1, 0.3, 0.7, 0.9\}$. This analysis enables the evaluation of the model's ability to reconstruct signals under increasing levels of missing or corrupted measurements. The corresponding results are reported in Table 3 and Figure 7.

Performance was measured using the same loss function adopted during training, ensuring consistency between the training objective and the evaluation metric.

Table 3. Reconstruction performance as a function of different corruption probabilities p . The table shows the training loss and the non-normalized MAE, expressed in degrees Celsius ($^{\circ}\text{C}$) for temperature and in percentage (%) for humidity.

Temperature				Humidity			
Window L	Noise p	Loss	MAE [$^{\circ}\text{C}$]	Window L	Noise p	Loss	MAE [%]
36	0.1	0.002	0.018	36	0.1	0.003	0.074
	0.3	0.003	0.022		0.3	0.005	0.095
	0.5	0.002	0.028		0.5	0.002	0.122
	0.7	0.008	0.042		0.7	0.011	0.169
	0.9	0.027	0.133		0.9	0.031	0.425
144	0.1	0.008	0.099	144	0.1	0.011	0.432
	0.3	0.008	0.101		0.3	0.011	0.437
	0.5	0.002	0.104		0.5	0.012	0.448
	0.7	0.009	0.115		0.7	0.013	0.484
	0.9	0.021	0.226		0.9	0.025	0.783




Figure 7. Reconstruction loss across varying corruption probabilities p .

5. Discussion

The experimental results provide several insights into the behavior of the proposed spatio-temporal reconstruction framework under realistic missing-data conditions.

Table 3 and Figure 7 report the reconstruction performance of the model under different corruption levels, illustrating the relationship between the corruption probability and the reconstruction loss. Unsurprisingly, the reconstruction loss increases monotonically with the corruption probability p , confirming the expected degradation in performance as the amount of missing information grows. However, as shown in Table 3, the model maintains stable performance under moderate corruption levels ($p \leq 0.7$), demonstrating its ability to recover missing information by leveraging both temporal continuity and spatial dependencies across neighboring sensing nodes.

This robustness is further validated by the non-normalized MAE values, which provide a clear measure of the model's precision in physical units. For the temperature, the reconstruction error

remains remarkably low: considering $L = 36$, it ranges from $0.018\text{ }^{\circ}\text{C}$ ($p = 0.1$) to a maximum of $0.133\text{ }^{\circ}\text{C}$ at $p = 0.9$. When the window size is increased to $L = 144$, the model starts from a higher baseline ($0.099\text{ }^{\circ}\text{C}$) but maintains a stable error profile, peaking at $0.226\text{ }^{\circ}\text{C}$ under extreme masking conditions. A similar behavior is observed for humidity, where the MAE scales from 0.074% to 0.425% for $L = 36$, and reaches 0.783% for $L = 144$ at the highest corruption level.

A key observation emerging from the experimental results concerns the role of spatial information in stabilizing the reconstruction process. Even when a significant portion of the temporal series of a node is masked, the model is still able to infer plausible signal dynamics by exploiting information propagated through the graph structure. This suggests that the latent representations learned by the GraphSAGE module effectively encode correlations between nodes induced by shared environmental conditions. In practical terms, this confirms that geographically proximate sensors can provide meaningful contextual information for reconstructing missing microclimatic measurements. The comparison between the two temporal window configurations ($L = 36$ and $L = 144$) reveals an additional trade-off between temporal context and reconstruction stability. Shorter windows, corresponding to approximately three hours of observations, consistently yield lower reconstruction losses than longer twelve-hour windows for $p \leq 0.7$. This result may be explained by the quasi-smooth and locally predictable evolution of temperature and humidity signals over short horizons. Conversely, longer temporal contexts introduce increased variability due to changes in site-specific microclimatic effects. As a consequence, the reconstruction task becomes more challenging, especially when large portions of the signal are corrupted.

Interestingly, the degradation pattern observed at high masking probabilities ($p = 0.9$) highlights the limits of spatial compensation mechanisms. When nearly all temporal information from the corrupted node is removed, the model increasingly relies on neighboring nodes, whose dynamics, although correlated, are not identical. This leads to larger reconstruction errors and indicates that spatial dependencies alone are insufficient to fully recover node-specific signal characteristics under extreme data loss conditions. Such behavior is consistent with the physical heterogeneity of vineyard environments, where elevation, solar exposure, and local canopy structure can induce persistent microclimatic differences.

Furthermore, in the extreme corruption scenario ($p = 0.9$), the configuration with $L = 144$ achieves slightly better reconstruction performance than the one with $L = 36$. This outcome may be explained by the increased temporal context available to the model when longer input windows are used. Under very high masking levels, the amount of reliable information from the corrupted node becomes extremely limited, reducing the effectiveness of short-term temporal dependencies. In this condition, the reconstruction process benefits from a wider temporal receptive field, which provides additional historical data and allows the model to better exploit slow-varying environmental trends such as diurnal cycles and gradual atmospheric transitions. Furthermore, longer windows enable the latent representation to capture more global signal regularities, partially compensating for the reduced availability of local temporal information. As a consequence, the model can leverage broader temporal patterns to stabilize the reconstruction, leading to a slightly improved performance compared to shorter windows under extreme data loss.

A further insight can be obtained by comparing the reconstruction performance achieved for both temperature and RH signals. As visible in Figure 7, humidity consistently exhibits higher reconstruction losses than temperature across all masking probabilities and temporal window configurations. This behavior can be attributed to both physical and modeling-related factors. From a physical perspective, air temperature in outdoor environments typically evolves according to smooth and strongly driven diurnal cycles governed by solar radiation and atmospheric heat exchange processes. Such dynamics are relatively predictable over short temporal horizons and tend to exhibit high spatial coherence across nearby sensing locations. Relative humidity, on the other hand, is influenced by a more complex interplay of processes, including evapotranspiration, canopy wetness, soil moisture conditions, and localized airflow patterns. These factors may introduce higher short-term variability and reduce spatial

synchrony among nodes, thereby increasing the intrinsic difficulty of the reconstruction task. From a modeling standpoint, this may limit the effectiveness of both temporal extrapolation and spatial message passing. When large portions of the temporal evolution are masked, the model relies more heavily on neighboring nodes. However, weaker spatial correlation in humidity dynamics reduces the reliability of this contextual information, leading to larger reconstruction errors compared to temperature.

The qualitative examples reported in Figs. 5 and 6 further support these quantitative findings. The reconstructed series preserve the smooth temporal evolution and the relative amplitude variations of the ground-truth signals, even across consecutive masked intervals. This suggests that the composite loss function, combining MAE minimization with a correlation-based objective, effectively encourages the model to learn both point-wise accuracy and global trend consistency. In environmental sensing applications, this property is particularly desirable, as preserving diurnal dynamics and relative fluctuations is often more informative than minimizing instantaneous errors alone. Another important aspect concerns the denoising training paradigm adopted in this work. By systematically exposing the model to structured node-wise temporal masking during training, the network learns to infer missing observations from partially corrupted inputs, resulting in improved robustness at inference time.

However, despite the promising reconstruction performance, some limitations should be considered when interpreting the results. The relatively small size of the sensor network may limit the variety of spatial dependency patterns observed during training. As a consequence, the graph-based representations learned by the model may capture only a restricted set of inter-node interactions, reflecting the relatively strong spatial coherence of the monitored microclimatic conditions. Moreover, the graph topology is static and defined solely based on geographical proximity, which may not fully capture the complexity of environmental interactions among nodes. Real environmental interactions are often governed by time-varying processes that are not strictly distance-dependent and that cannot be adequately represented by a fixed graph topology (e.g., transient cloud coverage, canopy growth dynamics, terrain-induced airflow, topographical barriers). Consequently, the spatial message-passing mechanism may not always reflect the true instantaneous coupling between sensing locations. Furthermore, although the temporal masking provides an effective and realistic approximation of communication losses or temporary sensor faults, it may not fully reproduce the diversity of failure mechanisms encountered in real deployments. For instance, long-term sensor drifts or correlated outages affecting multiple nodes simultaneously may generate reconstruction scenarios that differ from the corruption patterns simulated during training. Exploring more heterogeneous and physically informed corruption models could therefore improve the robustness of the learned representations. Finally, the experimental evaluation focuses on reconstructing individual variables independently. Although this design simplifies the analysis and allows clearer interpretation of the reconstruction behavior for temperature and RH, environmental processes are inherently multivariate. Joint reconstruction of multiple correlated variables could enable the model to exploit additional cross-feature dependencies, potentially improving robustness under severe data loss conditions. Future work could therefore investigate fully multivariate spatio-temporal reconstruction frameworks capable of capturing both inter-node and inter-variable interactions.

6. Conclusions

In this work, a spatio-temporal graph-based autoencoder has been proposed for the reconstruction of missing environmental data in vineyard sensor networks. The model integrates temporal sequence modeling through a GRU-D encoder with spatial message passing based on the GraphSAGE framework, enabling the joint exploitation of temporal dynamics and inter-node correlations.

The experimental evaluation conducted on real-world vineyard data demonstrates that the proposed approach provides robust reconstruction performance under realistic missing-data conditions. The results confirm that, even when 90% of the data is corrupted, the network effectively exploits correlations between nodes to reconstruct missing values with high accuracy, keeping average errors

consistently below 0.3 °C for temperature and 1% for humidity. This behavior is particularly relevant in vineyard monitoring scenarios, where short-term data losses are common due to communication instability or temporary sensor malfunction. The ability to provide reliable reconstructions from sparse observations ensures the continuity of data streams necessary for precise agro-climatic modeling and decision support systems.

The analysis further highlights a trade-off between temporal context and reconstruction accuracy. Shorter temporal windows yield improved performance in standard conditions due to the smoother and more locally predictable behavior of environmental signals. Conversely, longer temporal contexts become advantageous under extreme data loss scenarios, where broader temporal patterns and slow-varying dynamics provide additional information to compensate for the lack of local observations. Additionally, the comparison between temperature and RH reconstruction reveals that the physical characteristics of the monitored variables significantly impact model performance. Temperature, characterized by smoother dynamics and higher spatial coherence, is reconstructed more accurately than humidity, which exhibits higher variability and weaker inter-node correlation.

To conclude, the proposed framework represents a promising solution for improving data reliability in distributed agricultural monitoring systems, with potential implications for applications such as disease prediction, irrigation management, and anomaly detection. Future research directions may include the development of adaptive graph construction strategies capable of modeling time-varying spatial dependencies, the extension to multivariate reconstruction frameworks that jointly exploit cross-variable interactions, and the validation of the approach on larger and more heterogeneous sensor networks. These extensions may further improve the generalization capability and practical applicability of spatio-temporal graph-based models in real-world environmental monitoring scenarios.

Author Contributions: Conceptualization, F.C. and I.C.; methodology, F.C. and I.C.; software, F.C.; validation, F.C., I.C., A.F. and M.B.; formal analysis, F.C.; investigation, F.C., I.C., A.F. and M.B.; resources, A.F. and M.B.; data curation, I.C. and A.F.; writing—original draft preparation, F.C. and I.C.; writing—review and editing, M.B. and A.F.; visualization, F.C.; supervision, M.B. and A.F.; project administration, A.F.; funding acquisition, A.F. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Acknowledgments: This study was carried out within the Agritech National Research Center and received funding from the European Union Next-GenerationEU (PIANO NAZIONALE DI RIPRESA E RESILIENZA (PNRR) – MISSIONE 4 COMPONENTE 2, INVESTIMENTO 1.4 – D.D. 1032 17/06/2022, CN00000022). This manuscript reflects only the authors' views and opinions, neither the European Union nor the European Commission can be considered responsible for them.

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

Abbreviations

The following abbreviations are used in this manuscript:

IoT	Internet of Things
WSN	Wireless Sensor Network
LPWAN	Low-Power Wide-Area Network
LoRaWAN	Long Range Wide-Area Network
RH	Relative Humidity
ML	Machine Learning
GNN	Graph Neural Network
STGNN	Spatio-Temporal Graph Neural Network
STGCN	Spatio-Temporal Graph Convolutional Network
DCRNN	Diffusion Convolutional Recurrent Neural Network
STGMAE	Spatio-Temporal Graph Masked Autoencoder
GRU	Gated Recurrent Unit
GRU-D	Gated Recurrent Unit with Decay
LSTM	Long Short-Term Memory
MLP	Multi-Layer Perceptron
MAE	Mean Absolute Error
SHT3x	Sensirion Temperature and Humidity Sensor Series
CdS	Cadmium Sulfide
ISM	Industrial, Scientific and Medical
MCU	Microcontroller Unit
ReLU	Rectified Linear Unit
AE	Autoencoder
DAE	Denosing Autoencoder

References

1. Wolfert, S.; Ge, L.; Verdouw, C.; Bogaardt, M.J. *Agricultural systems* **2017**, *153*, 69–80. <https://doi.org/10.1016/j.agsy.2017.01.023>.
2. Rouxinol, M.I.; Martins, M.R.; Barroso, J.M.; Rato, A.E. Wine grapes ripening: A review on climate effect and analytical approach to increase wine quality. *Applied Biosciences* **2023**, *2*, 347–372. <https://doi.org/10.3390/applbiosci2030023>.
3. Mowla, M.N.; Mowla, N.; Shah, A.F.M.S.; Rabie, K.M.; Shongwe, T. Internet of Things and Wireless Sensor Networks for Smart Agriculture Applications: A Survey. *IEEE Access* **2023**, *11*, 145813–145852. <https://doi.org/10.1109/ACCESS.2023.3346299>.
4. Mansouri, T.; Sadeghi Moghadam, M.R.; Monshizadeh, F.; Zareravasan, A. IoT data quality issues and potential solutions: a literature review. *The Computer Journal* **2023**, *66*, 615–625. <https://doi.org/10.1093/comjnl/bxab183>.
5. Decorte, T.; Mortier, S.; Lembrechts, J.J.; Meysman, F.J.; Latré, S.; Mannens, E.; Verdonck, T. Missing value imputation of wireless sensor data for environmental monitoring. *Sensors* **2024**, *24*, 2416. <https://doi.org/10.3390/s24082416>.
6. Choi, C.; Jung, H.; Cho, J. An ensemble method for missing data of environmental sensor considering univariate and multivariate characteristics. *Sensors* **2021**, *21*, 7595. <https://doi.org/10.3390/s21227595>.
7. Cappelli, I.; Parri, L.; Tani, M.; Vignoli, V.; Fort, A. Pervasive Monitoring in the Context of Precision Agriculture: Using Low-Cost LDR Sensors for Solar Intensity Measurement. In Proceedings of the 2024 IEEE International Instrumentation and Measurement Technology Conference (I2MTC), 2024, pp. 1–6. <https://doi.org/10.1109/I2MTC60896.2024.10560913>.
8. Yu, B.; Yin, H.; Zhu, Z. Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting. In Proceedings of the IJCAI, 2018. <https://doi.org/10.5555/3304222.3304273>.
9. Li, Y.; Yu, R.; Shahabi, C.; Liu, Y. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In Proceedings of the ICLR, 2018. <https://doi.org/10.48550/arXiv.1707.01926>.
10. Wu, Z.; Pan, S.; Long, G.; Jiang, J.; Zhang, C. Graph WaveNet for Deep Spatial-Temporal Graph Modeling. In Proceedings of the IJCAI, 2019. <https://doi.org/10.5555/3367243.3367303>.
11. Longa, A.; Lachi, V.; Santin, G.; Bianchini, M.; Lepri, B.; Lio, P.; Scarselli, F.; Passerini, A. Graph neural networks for temporal graphs: State of the art, open challenges, and opportunities. *arXiv preprint arXiv:2302.01018* **2023**. <https://doi.org/10.48550/arXiv.2302.01018>.

12. Chabalala, V.; Rudolph, C.; Mosala, K.; Nkadimeng, E.K.; Mosomane, C.; Mathaha, T.; Basu, P.; Mahboob, M.A.; Kong, J.; Bragazzi, N.; et al. Spatiotemporal Graph Neural Networks for PM 2.5 Concentration Forecasting. *Air* **2026**, *4*, 2. <https://doi.org/10.3390/air4010002>.
13. Pan, Z.; Xu, L.; Chen, N. Combining graph neural network and convolutional LSTM network for multistep soil moisture spatiotemporal prediction. *Journal of Hydrology* **2025**, *651*, 132572. <https://doi.org/10.1016/j.jhydrol.2024.132572>.
14. Tuo, Y.; Wirthensohn, M.; Disse, M. Spatio-Temporal Graph Neural Networks for Soil Moisture Drought Forecasting: Adaptability, Predictability, and Interpretability. In Proceedings of the AGU Fall Meeting Abstracts, 2023, Vol. 2023, *AGU Fall Meeting Abstracts*, pp. H43J–2219.
15. Akkala, A.; Boubrahimi, S.F.; Hamdi, S.M.; Hosseinzadeh, P.; Nassar, A. Spatio-Temporal Graph Neural Networks for Streamflow Prediction in the Upper Colorado Basin. *Hydrology* **2025**, *12*. <https://doi.org/10.3390/hydrology12030060>.
16. Feng, J.; Sha, H.; Ding, Y.; Yan, L.; Yu, Z. Graph convolution based spatial-temporal attention LSTM model for flood forecasting. In Proceedings of the 2022 International Joint Conference on Neural Networks (IJCNN). IEEE, 2022, pp. 1–8. <https://doi.org/10.1109/IJCNN55064.2022.9892371>.
17. Costanti, F.; Cappelli, I.; Ceroni, E.G.; Bianchini, M.; Fort, A. Foliar Wetness Prediction Using Sensor Network Data and WaveNet-Based Deep Learning Models. In Proceedings of the 2025 IEEE International Conference on Metrology for eXtended Reality, Artificial Intelligence and Neural Engineering (MetroXRINE), 2025, pp. 1036–1041. <https://doi.org/10.1109/MetroXRINE66377.2025.11340571>.
18. He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; Girshick, R. Masked autoencoders are scalable vision learners. In Proceedings of the Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2022, pp. 16000–16009. <https://doi.org/10.1109/CVPR52688.2022.01553>.
19. Feichtenhofer, C.; Li, Y.; He, K.; et al. Masked autoencoders as spatiotemporal learners. *Advances in neural information processing systems* **2022**, *35*, 35946–35958. <https://doi.org/10.5555/3600270.3602875>.
20. Jin, G.; Liang, Y.; Fang, Y.; Shao, Z.; Huang, J.; Zhang, J.; Zheng, Y. Spatio-temporal graph neural networks for predictive learning in urban computing: A survey. *IEEE transactions on knowledge and data engineering* **2023**, *36*, 5388–5408. <https://doi.org/10.1109/TKDE.2023.3333824>.
21. Zhang, Q.; Gao, X.; Wang, H.; Huang, D.; Yiu, S.M.; Yin, H. HGAurban: Heterogeneous Graph Autoencoding for Urban Spatial-Temporal Learning. In Proceedings of the Proceedings of the 34th ACM International Conference on Information and Knowledge Management, 2025, pp. 4139–4148. <https://doi.org/10.1145/3746252.3761383>.
22. Zhao, F.; Cao, X.; Zhao, J.; Duan, Y.; Yang, X.; Zhang, X. Masked graph autoencoder-based multi-agent dynamic relational inference model for trajectory prediction. *Neurocomputing* **2025**, *634*, 129922. <https://doi.org/10.1016/j.neucom.2025.129922>.
23. Costanti, F.; Cappelli, I.; Fort, A.; Ceroni, E.G.; Bianchini, M. LSTM-based Siamese Networks for Fault Detection in Meteorological Time Series Data. In Proceedings of the 2024 IEEE International Conference on Metrology for eXtended Reality, Artificial Intelligence and Neural Engineering (MetroXRINE), 2024, pp. 906–911. <https://doi.org/10.1109/MetroXRINE62247.2024.10796514>.
24. Che, Z.; Purushotham, S.; Cho, K.; Sontag, D.; Liu, Y. Recurrent neural networks for multivariate time series with missing values. *Scientific reports* **2018**, *8*, 6085. <https://doi.org/10.1038/s41598-018-24271-9>.
25. Tang, H.; Yang, H.; Zhang, W. DAHG: A Dynamic Augmented Heterogeneous Graph Framework for Precipitation Forecasting with Incomplete Data. *Information* **2025**, *16*, 946. <https://doi.org/10.3390/info16110946>.
26. Sasal, L.; Busby, D.; Hadid, A. Tempokgat: A novel graph attention network approach for temporal graph analysis. In Proceedings of the International Conference on Neural Information Processing. Springer, 2024, pp. 212–226. https://doi.org/10.1007/978-981-96-6579-2_15.
27. Dimitri, G.M.; Cappelli, I.; Scarselli, F.; Fort, A.; Gori, M. Graph Neural Networks for Missing Data Imputation in Time Series from Meteorological Sensors. In Proceedings of the 2024 IEEE International Conference on Metrology for eXtended Reality, Artificial Intelligence and Neural Engineering (MetroXRINE), 2024, pp. 1242–1247. <https://doi.org/10.1109/MetroXRINE62247.2024.10796616>.
28. Chung, J.; Gulcehre, C.; Cho, K.; Bengio, Y. Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling. *arXiv preprint arXiv:1412.3555* **2014**. <https://doi.org/10.48550/arXiv.1412.3555>.
29. Hamilton, W.; Ying, Z.; Leskovec, J. Inductive representation learning on large graphs. *Advances in neural information processing systems* **2017**, *30*. <https://doi.org/10.5555/3294771.3294869>.

30. Lence, A.; Granese, F.; Fall, A.; Hanczar, B.; Salem, J.E.; Zucker, J.D.; Pifti, E. ECGrecover: a deep learning approach for electrocardiogram signal completion. In Proceedings of the Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1, 2025, pp. 2359–2370. <https://doi.org/10.1145/3690624.3709405>.
31. Liang, Y.; Zhang, Y.; Chen, F.; Lu, J.; Lin, Z. Decoding Speech Envelopes from Electroencephalogram with a Contrastive Pearson Correlation Coefficient Loss. *arXiv preprint arXiv:2601.20542* 2026. <https://doi.org/10.48550/arXiv.2601.20542>.
32. Akiba, T.; Sano, S.; Yanase, T.; Ohta, T.; Koyama, M. Optuna: A next-generation hyperparameter optimization framework. In Proceedings of the Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining, 2019, pp. 2623–2631. <https://doi.org/10.1145/3292500.3330701>.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.