International Journal of Civil Engineering
https://doi.org/10.1007/s40999-026-01231-3

RESEARCH PAPER

Monitoring Building Settlements Using Space–Time Cubes and Geospatial AI

Luigi Barazzetti[1] · Mattia Previtali[1] · Fabio Roncoroni[2]

Received: 31 August 2024 / Revised: 28 December 2025 / Accepted: 22 March 2026

Abstract
This paper explores the application of space–time cubes (S–T cubes) and geospatial artificial intelligence (GeoAI) for monitoring vertical settlements in structures. Although S–T cubes are not commonly employed in this type of structural analysis, they enable the storage and visualization of multi-temporal monitoring data, offering an effective framework to represent differential settlement over time within a defined spatial domain, including at the scale of individual structural elements. The study further integrates GeoAI techniques for predictive analysis aimed at detecting discontinuities, leveraging the temporal datasets organized within S–T cubes. We employ an index that combines accuracy across multiple prediction steps with least-squares statistics of adjusted data. This provides users with a rapid diagnostic to verify the effectiveness of forecasts before deeper analysis. Results from three datasets—each representing a monitoring project with distinct characteristics—demonstrate that machine-learning-based forecasts remain reliable for at least three prediction steps ahead. This level of consistency is sufficient to support the detection of discontinuities in new monitoring campaigns, with a precision on the order of ± 0.2–0.3 mm.

Keywords GeoAI · Monitoring · Prediction · Simulation · Space–time cube · Structures · Subsidence

1 Introduction

1.1 Overview of Settlement Monitoring Methods

Activities in structural monitoring involve identifying and measuring differential movements among load-bearing elements of buildings and infrastructure. Advances in instrumentation and methods have enhanced these activities, enabling the precise detection of geometric variations, such as displacements and rotations. Periodic assessments conducted at specific time intervals can reveal any changes that may have occurred and, where necessary, support strategic planning for maintenance or extraordinary consolidation to ensure structural safety [1, 2].

Currently, a variety of methods are employed in the surveying and geomatics sectors to monitor structures both periodically and continuously. These methods include movement detection and measurement using total stations or GNSS (global navigation satellites systems), radar interferometry (ground-based and satellite-based), and photogrammetric or vision metrology techniques. Representative examples can be found in [3–8]. However, this list is not exhaustive, as these methods are frequently integrated with data from inclinometers, clinometers, accelerometers, crack monitoring devices, linear variable displacement transducers (LVDTs), and other instruments [9, 10].

The identification and quantification of settlements are fundamental for engineers and architects to ensure the long-term stability and integrity of buildings and infrastructure. Settlement monitoring measures and analyzes how the ground supporting a structure moves over time, playing a critical role in assessing the stability and performance of structures across various scales. Settlement may result from

✉ Luigi Barazzetti
luigi.barazzetti@polimi.it

Mattia Previtali
mattia.previtali@polimi.it

Fabio Roncoroni
fabio.roncoroni@polimi.it

[1] Department of Architecture, Built Environment and Construction Engineering (ABC), Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milan, Italy

[2] Polo territoriale di Lecco, Politecnico di Milano, Via Previati 1c, 23900 Lecco, Italy

Published online: 04 May 2026

factors such as soil consolidation, changes in groundwater levels, and other geotechnical conditions. Two primary types of settlement are generally recognized: uniform settlement, where the entire structure subsides evenly, and differential settlement, where varying rates of subsidence across a structure can cause damage. The latter necessitates special attention due to uneven stress distributions and relative movements that may compromise structural safety.

Settlement monitoring may be conducted continuously using automated sensors or periodically through targeted measurement campaigns. Data collected through monitoring must be analyzed to identify trends and discontinuities. Advanced analysis techniques, including statistical approaches, geospatial analysis, and modeling, are frequently employed to support a comprehensive understanding. Unlike methods commonly used for monitoring settlements in urban areas—such as GNSS satellite positioning [11, 12], SAR interferometry [13–15], or a combination of both [16, 17]—the approach presented in this article operates at a finer scale. It focuses on individual buildings and their structural components, encompassing both interior and exterior environments.

One method employed to assess vertical differential settlements is high-precision geometric leveling. Typically, a series of benchmarks is strategically placed on load-bearing elements such as columns or pillars. A high-precision level enables the measurement of elevation differences with sub-millimeter accuracy. By implementing redundant measurement schemes based on multiple closed loops and processing the data using the least squares adjustment method, it is possible to calculate point elevations with a precision of approximately ± 0.1 mm. Monitoring points consist of stable benchmarks where leveling rods can be positioned to minimize repositioning errors. One benchmark is designated as the reference point, and its elevation is held constant over time, relative to which the elevations of the other points are calculated.

To support the interpretation of settlement trends and patterns, a monitoring protocol must include effective data visualization tools such as charts, graphs, and reports. Advanced analytics can be applied to anticipate potential issues and support informed decisions regarding interventions or structural consolidation. In the case of settlement monitoring, a GIS-based method that integrates geospatial visualization with numerical data is particularly valuable for conducting precise analyses focused on buildings and their architectural elements. The approach proposed in this paper employs a space–time cube to store, visualize, and analyze geospatial time-series data.

Space–time cubes are commonly used in GIS environments to provide a three-dimensional representation of spatiotemporal data. They combine spatial and temporal dimensions, allowing for a deeper understanding of dynamic processes. A space–time cube typically includes two axes for geographic coordinates and a third axis for time [18]. These cubes have been applied to analyze the motion of objects and spatial changes anchored in time [19]. In the context of a leveling network, the same benchmarks are measured periodically. Their spatial positions remain constant over time, forming a cube in which each defined location is associated with variation data derived from least squares estimates and stored as a single attribute for each time interval.

The integration of space–time cubes into GIS has emerged as a promising method for visualizing and analyzing spatiotemporal data, with applications spanning numerous domains. The concept of the space–time cube was originally introduced by Hägerstrand to represent spatial dimensions horizontally and time vertically [20]. Since its inception, it has been adapted for use in a wide range of contexts, demonstrating its versatility and efficacy in handling complex spatiotemporal datasets. In GIS, space–time cubes have been used for population density studies, weather forecasting (e.g., temperature and precipitation), and movement tracking. Additional applications include historical landscape analysis, traffic flow evaluation, and the development of epidemic warning systems. Notably, the technique has supported environmental studies, such as identifying hotspots of marine mammal strandings along coastlines. Applications span hotspot detection, trend analysis, and trajectory data visualization, highlighting the adaptability of space–time cubes for diverse analytical tasks. For further examples, the reader is referred to [21–35].

In this manuscript, we assess the applicability of space–time cubes in structural monitoring, with particular emphasis on tracking settlement variations in buildings and infrastructure at the architectural scale rather than the environmental scale. Our objective is to explore the potential of this method to improve the understanding of structural behavior by analyzing differential displacements recorded at monitoring points located on key load-bearing components.

The structured data organization enabled by the space–time cube is integrated with geospatial processing techniques and enhanced by artificial intelligence (AI), particularly machine learning (ML) and deep learning (DL) approaches. These tools are effective in analyzing time-series data, supporting insight generation, predictive modeling, and pattern recognition. A wide array of forecasting methods exists for time-series analysis, from traditional approaches such as ARIMA and Exponential Smoothing State Space Models (ETS) to contemporary machine learning methods like decision trees and support vector machines. Deep learning models—including Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks—have proven particularly capable of capturing sequential

dependencies. Ensemble methods combine multiple models to improve accuracy. Beyond forecasting, AI methods also facilitate anomaly detection by training models to identify irregularities or outliers. The incorporation of AI into time-series analysis thus represents a comprehensive approach to uncovering patterns, predicting future behaviors, and extracting actionable insights from temporal datasets.

This manuscript aims to evaluate and explore a variety of AI techniques through the analysis of time-series data stored within space–time cubes and integrated into a GIS environment. Several case studies are presented to illustrate the advantages and limitations of the proposed methodologies. Real monitoring datasets, obtained through high-precision geometric leveling, are used and require tailored visualization strategies to effectively interpret the geospatial characteristics of the monitoring process. This aspect is critical in the context of structural monitoring applications.

2 Time-Series of Vertical Settlements in Space–Time Cubes

This section presents the selected case studies and outlines the procedures used to generate the corresponding space–time cubes. The selection of case studies considered a range of structural characteristics, including geometry, observed vertical movements, and overall structural behavior. The section begins by introducing the general workflow, then describes the specific features of each case study. Finally, it examines and discusses how the monitoring datasets are organized within the space–time cube framework.

2.1 Methodological Workflow

This paper explores the use of space–time cubes, geospatial analysis, and machine and deep learning methods for monitoring vertical settlements. Constructing these datasets involves acquiring observations through periodic on-site measurement campaigns. These observations are subsequently post-processed to derive elevation values and their associated precision using least squares adjustment. However, the paper primarily focuses on the creation of space–time cubes and the extraction of supplementary information, and therefore does not provide a detailed discussion of this crucial preliminary phase.

The space–time cube is based on NetCDF (Network Common Data Form), a set of software libraries and machine-independent data formats that support the creation, access, and sharing of array-oriented scientific data. It is widely adopted for the dissemination of scientific data due to its self-describing, portable, scalable, appendable, shareable, and archivable characteristics. NetCDF is supported by the Unidata Program Center and provides interfaces for several programming languages, including C, C++, Java, and Python.

The continuity of time series data is of fundamental importance, particularly in the context of periodic measurements collected over extended periods. In the scenarios examined in this paper, data collection spanned several decades. Ensuring continuity in these time series was achieved through the consistent maintenance of monitoring systems, which at times required the replacement or integration of monitoring points. Any disruptions caused by the replacement of equipment—such as upgrades to more advanced data acquisition tools—were addressed by estimating the differences between measurements recorded with the old and new equipment.

Missing data or the loss of specific points is sometimes unavoidable. Restoration work may temporarily block access, while occasional benchmark damage can create discontinuities. When such gaps cannot be measured directly, they are estimated through interpolation from neighboring points, with all issues and corrective actions recorded in the monitoring reports. Although infrequent, these cases must be accounted for, and the next sections provide further details on the missing-data procedures adopted in the proposed case studies.

When organizing data within space–time cubes, several statistical approaches can be used to recover missing values during their creation. Temporary gaps—often due to restricted access—can be interpolated, and although null values may be stored, some processing methods require uninterrupted time series. For this reason, space–time cubes must be continuously updated as new monitoring campaigns provide additional measurements.

However, permanent discontinuities—such as the installation of a new monitoring system or the replacement of monitoring points—cannot be handled automatically during cube generation. In these cases, numerical corrections are applied directly to the time-series data before producing the final space–time cube. Because space–time datasets are typically structured as tables, these adjustments can be made at the data-storage stage, ensuring that discontinuities are properly managed before further analysis.

2.2 Description of the Case Studies

The case studies were selected to represent a variety of structural movement patterns and to evaluate the effectiveness of the proposed methods in revealing inconsistencies in the results. These case studies comprise datasets spanning several decades, with movements documented and analyzed in technical reports dedicated to structural monitoring. This approach allows for a preliminary interpretation

of the observed data, followed by a comparative analysis of the outcomes achieved through the integration of space–time cubes and artificial intelligence-based processing techniques, offering an alternative method for result validation.

This article investigates three distinct case studies located in Milan: (i) the Meazza San Siro Stadium, situated in the western part of the city; (ii) the Catholic University of Milan; and (iii) the Milan Cathedral, both located in the historical city center. The geometric leveling methodologies were designed according to the geometry and construction techniques of each structure, taking into account the expected types of movements during the planning phase. The three case studies are illustrated in Fig. 1 and are briefly described as follows:

- *Case 1*. Curved Beam of the Meazza San Siro Stadium: The analyzed beam section has a curved geometry and supports the third ring of the stadium. This box-type beam, with hollow interiors made of pre-stressed concrete and reinforced by transverse elements, is periodically monitored using high-precision leveling measurements across 19 benchmarks distributed along the longitudinal profile. Additional benchmarks placed on the beam sides enable rotation calculations. The reference benchmark is located at one end, on the large pillar with the spiral staircase. Although precision may decrease at the end opposite the benchmark due to the configuration of three closed loops, the overall movement estimation precision—after least squares adjustment—is typically better than $\pm 0.1\,\mathrm{mm}$, owing to the improved acquisition geometry enabled by the three consecutive loops. Measurements have been conducted annually since 2001, and the space–time cube includes data collected consistently during the summer and autumn seasons, ensuring comparability under similar environmental conditions. It is worth noting that monitoring the curved beam is part of a broader system extending from the base of the pillars (ground level) to the stadium's metal roof. A dynamic control system is also in place, although it is not discussed in this paper [36, 37].
- *Case 2*. Courtyards at the Catholic University of Milan: The two courtyards consist of two sets of 44 columns, each equipped with base benchmarks to detect vertical displacements. The acquisition geometry includes two large interconnected rings, subdivided into sub-rings, each covering a quarter of the courtyard. The rings are connected via a reference benchmark located on a

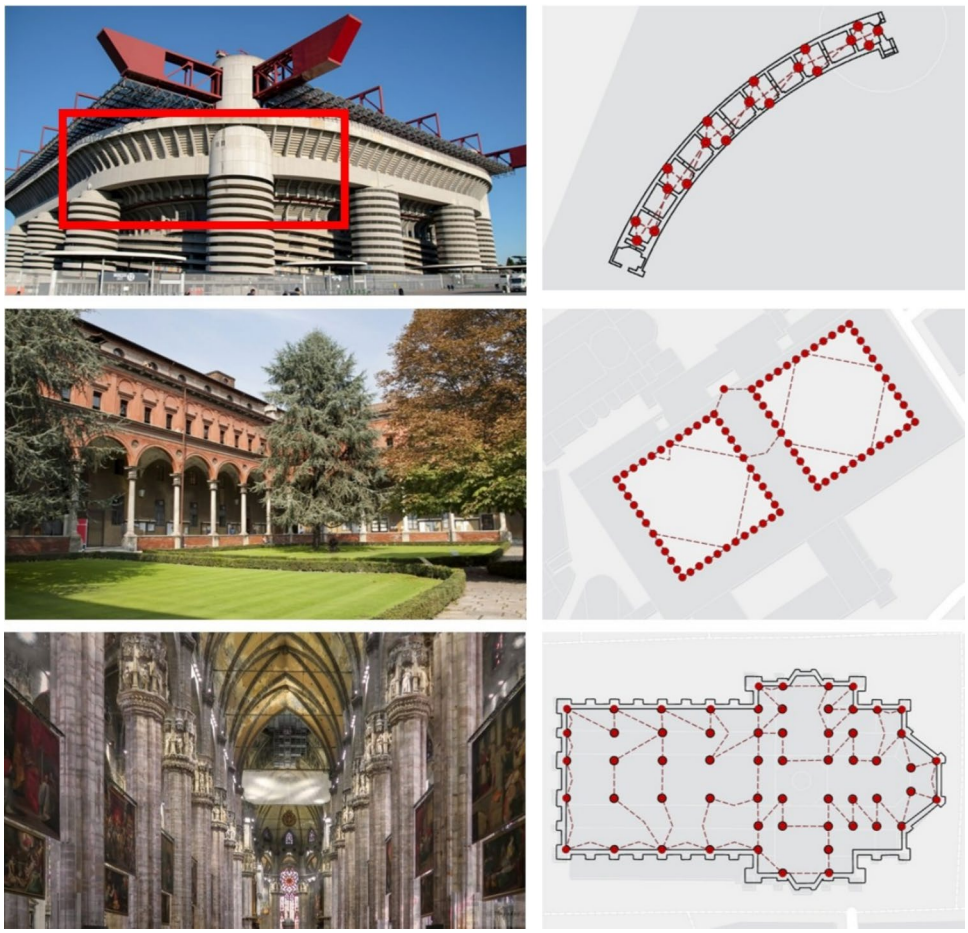

Fig. 1 The three case studies considered in this paper: curved beam of the Meazza San Siro Stadium (top), courtyards at the Catholic University of Milan (center), and Duomo di Milano (bottom). Figures also show point distribution and network geometry

load-bearing wall between the courtyards, with an additional point measured in the interconnecting corridor. Least squares adjustment of the observations generally yields a precision better than ± 0.1–$0.2\,\mathrm{mm}$. The dataset used in this analysis includes measurements conducted consistently during the winter season from 1992 to 2022.

- *Case 3*. Duomo di Milano: The monitoring system for the Milan Cathedral includes leveling measurements initiated in the early 1960 s, prompted by subsidence-related movements that required continuous monitoring, which remains ongoing. For additional details on the triggering factors of the subsidence phenomenon, see [38]. This paper focuses on measurements recorded since 1970. Unlike the previous cases, the benchmark distribution spans the entire interior of the cathedral, forming a spatially extensive configuration. In this context, leveling measurements are part of a larger monitoring system. For further technical insights, refer to [39]. The current analysis builds upon preliminary work presented in [40], where leveling data were combined with space–time cubes. In this manuscript, the same data are partially reused and processed using machine learning and deep learning techniques.

As shown in Fig. 1, the distribution of monitoring points varies across the three case studies, resulting in distinct spatial configurations. The first example (Meazza Stadium) exhibits a predominantly longitudinal development. The second case (Catholic University) features a linear configuration that closes into large interconnected loops with associated sub-rings. The third case (Duomo di Milano) presents a more extensive spatial configuration with a dense network of sub-rings.

The identification and interpretation of structural movements depend heavily on the spatial distribution of the benchmarks and the structural characteristics of each site. Consequently, interpreting results requires careful attention not only to temporal changes (i.e., evaluating variations over time at individual columns) but also to the relative spatial positioning of different columns. The following sections will further explore how movements are computed and interpreted using spatial analyses and visualizations obtained through space–time cubes within a GIS environment.

Regarding missing data, the proportions observed in the three case studies are very limited relative to the total dataset, making the adopted interpolation strategy both simple and effective. Case Study 1 contains no missing values, while Case Study 3 shows only about 0.4% missing data. In these two case studies, the interpolation error introduced during the recovery of missing values is not considered influential. Case Study 2 represents the least favorable scenario, with approximately 6% missing data, mostly concentrated in 2 years during which measurements were not performed. However, because measurements are available immediately before and after this period, linear interpolation is considered sufficient to reconstruct the missing values and to ensure the completeness required for generating the space–time cube and for performing subsequent analyses.

2.3 Creation of Space–Time Cubes with Time-Series of Differential Settlements

The planimetric position (East, North) of control points, identified by permanent benchmarks, remains constant in leveling measurements used to detect local settlements. In the space–time cube, the planimetric coordinates are fixed, while the Z-axis represents time. Variations are indicated by differences in elevation values computed at different epochs. It is important to note that the space–time cube comprehensively captures these relative changes. A single benchmark serves as a reference, making all detected changes relative (differential) to this established point. Consequently, the variations can be both positive and negative.

A 3D visualization of the space–time cubes for the different case studies is presented in Fig. 2. The figure employs a color scale to represent the magnitude of movements recorded during various epochs, arranged along the vertical axis. The value at ground level corresponds to the initial measurement epoch, while the topmost value represents the most recent observation. This form of 3D visualization differs from conventional methods, such as time-series plots with time on the x-axis and measured values on the y-axis. Although Fig. 2 effectively illustrates the distribution of maximum and minimum variations, it does not provide a clear understanding of the temporal evolution leading to these values. Therefore, it does not replace traditional graphs but rather complements them. Its primary function is to confirm the successful construction of the space–time structure and to support subsequent data extraction for further analysis, as elaborated in the following sections.

The space–time cubes used in this study contain no empty cells; in other words, values for all measurement epochs are available for each point. Achieving continuous measurements in complex environments—such as the various areas of the case studies affected by periodic maintenance over several decades—poses significant challenges. A crucial preliminary step before generating the cubes involved reconstructing data archives for epochs where direct field measurements were not feasible.

In the case of the Meazza Stadium beam, complete measurement data were available for each epoch, thanks to the closed and protected nature of the beam, which preserved the installed benchmarks. However, for the courtyards at the Catholic University of Milan, interpolation was required

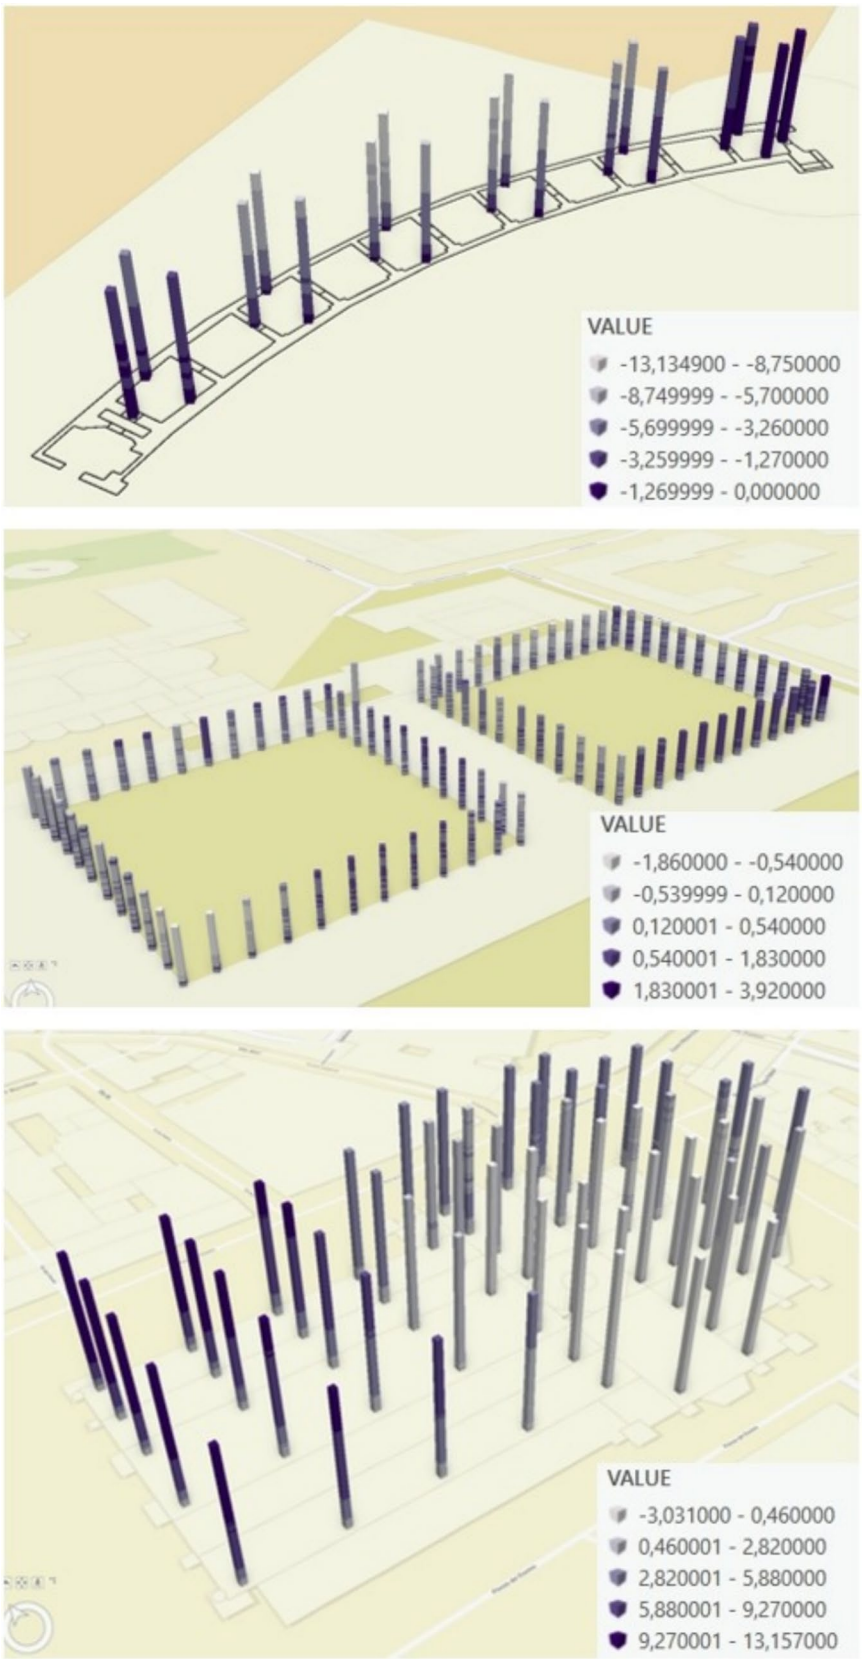

Fig. 2 3D visualization of the space–time cubes of the different case studies

for the years 1994 and 1998 due to missing measurements. Data continuity was restored for one point in the southwest courtyard following a restoration intervention that led to the replacement of the benchmark. The missing epochs were reconstructed through linear interpolation based on existing time series values. For the Milan Cathedral, measurements from 1969–1970 were selected, preserving the configuration of points still in use today. Reconstructing the archive in this case was more complex, owing to ongoing maintenance and conservation activities.

One limitation of using a space–time cube as the sole data storage structure lies in the difficulty of associating supplementary information, including non-numeric data, which is essential for contextualizing the results. Interpretation therefore requires an understanding of the broader problem, beyond the values recorded. Accordingly, the space–time cube and the techniques proposed in this study should not be regarded as a standalone solution for interpreting structural behavior. Instead, when integrated with data analysis methods—including GeoAI—the space–time cube becomes a highly valuable tool that enhances interpretation through targeted analysis and geospatial visualization. This approach complements existing methodologies commonly employed in monitoring operations.

As previously noted, the visualization in Fig. 2 does not provide a complete understanding of the vertical displacements experienced by the structures. Therefore, further analyses based on the space–time cubes are introduced in the following sections, with a focus on the generation of new information and the application of specific visualization techniques. The objective is to assist in the interpretation of structural behavior, a task that should always be carried out by experienced professionals.

3 Monitoring and Data Processing Based on Space–Time Cubes

In this section, the space–time cube is employed for processing and visualizing settlement information. The objective is to produce numerical results and visual representations that facilitate user interpretation, particularly in the context of georeferenced measurements. These methods integrate seamlessly with traditional reports generated for structural health monitoring, thereby enhancing the understanding of processes that involve both spatial and temporal components.

3.1 Trend Analysis

Analyzing trends in time-series data is essential for identifying movements or changes occurring over extended periods. This process involves examining the overall direction of the data—either upward or downward—based on its chronological progression. A thorough analysis can be conducted using the space–time cube method in combination with the Mann-Kendall statistic, which assesses the rank correlation between observed values and their temporal order. In this approach, the value of each data point within a specified time interval (referred to as a 'bin') is compared to the value in the subsequent bin. A positive score (+ 1) is assigned if the subsequent value is higher, indicating an upward trend; a negative score (− 1) is assigned if the value is lower, indicating a downward trend; and a neutral score (0) is recorded if the values are equal, implying no change. The scores from all pairwise comparisons are then aggregated.

Ideally, in the absence of a significant trend, the cumulative score will approximate zero, suggesting that increases and decreases are balanced over time. However, to evaluate whether any deviation from zero is statistically significant—indicating a genuine trend rather than random variation—the cumulative score is analyzed in relation to the variance of the bin values, the number of tied scores, and the total number of bins. This comprehensive evaluation facilitates the identification and statistical validation of trends within the time-series data.

An important consideration in trend analysis is that differential displacements stored in space–time cubes are always calculated relative to a designated reference benchmark. Consequently, when visualizing trends and interpreting results, it is essential to understand that these measurements are relative. For example, in the case of the curved beam of the San Siro Stadium, all points exhibit a downward trend with 99% confidence (Fig. 3). It should be noted, however, that trend analysis does not provide information on the magnitude of the displacements. To address this, users may plot the time-series data directly or employ color-based visualizations, as shown in Fig. 2.

The dataset from the Catholic University displays greater variability in the behavior of individual columns. By examining groups of adjacent columns—particularly those located along the sides and corners of the courtyards, as illustrated in Fig. 3—distinct settlement patterns become evident. In contrast, the dataset from the Cathedral of Milan exhibits a more homogeneous color distribution. Notably, the transept area displays negative trends, in contrast to the nave and apse, suggesting a localized settlement tendency in that portion of the cathedral.

3.2 Change Detection: Automated Discontinuity Analysis in Time-Series of Differential Settlements

While an overall evaluation of a time series can provide insight into general trends, individual data points may

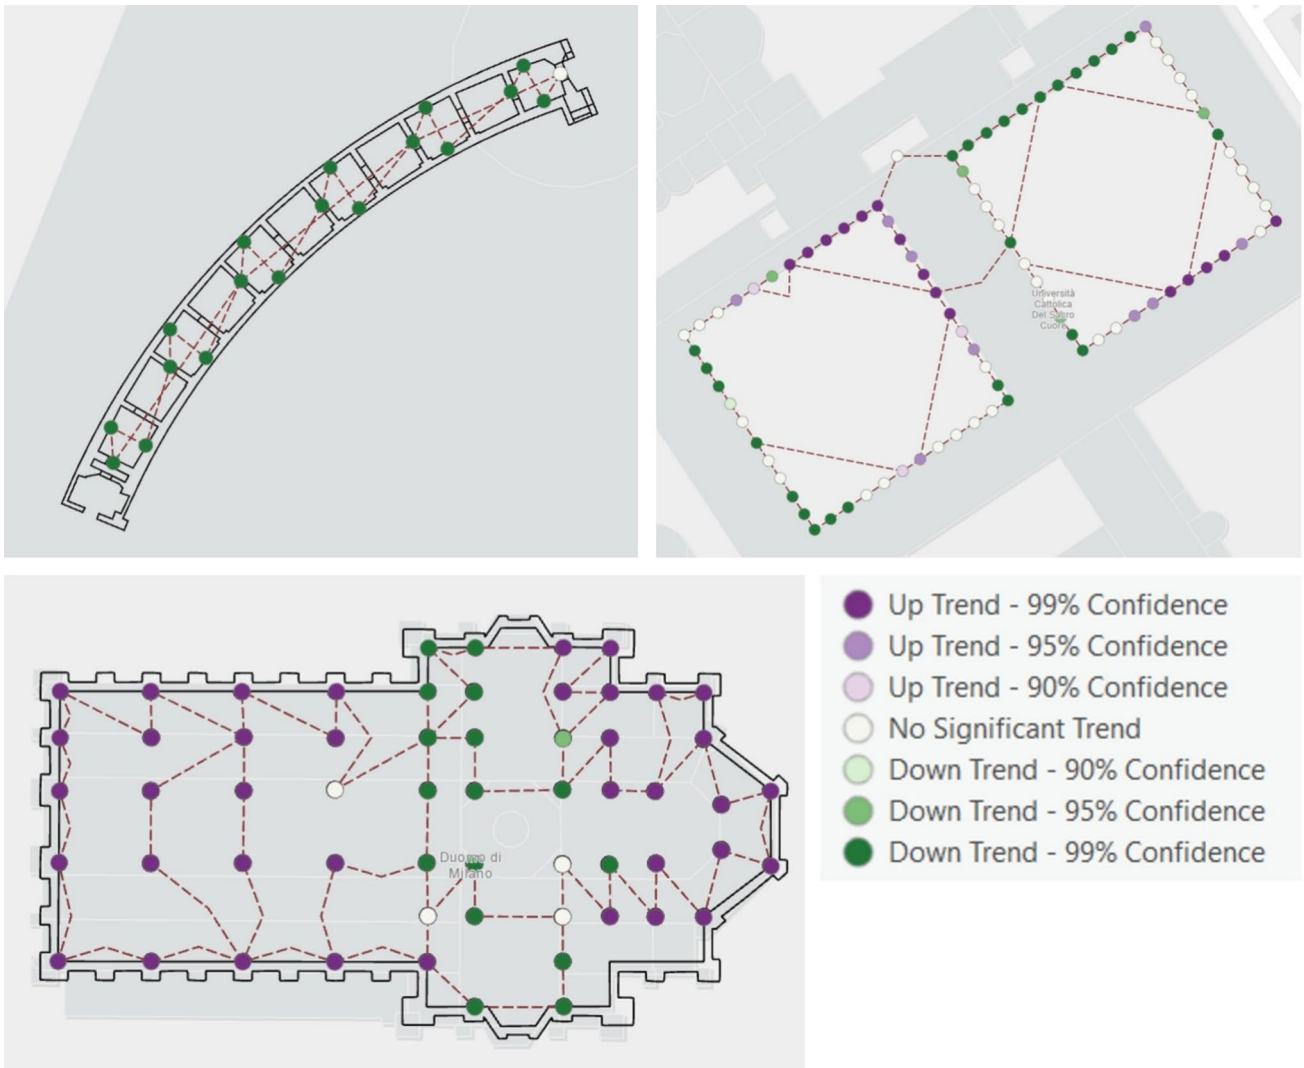

Fig. 3 The three case studies and trend visualization with variable confidence levels

exhibit different behaviors over time. For instance, certain measurement epochs may reveal an upward trend followed by a downward trend. In such cases, identifying the various trends within the dataset and examining their relationships with similar behaviors becomes essential. To perform this type of analysis, we employed discontinuity detection techniques.

Discontinuity detection is used to identify changes in the statistical properties of a time series within a space–time cube. Each detected change point partitions the time series into segments, where the values within each segment exhibit similar statistical characteristics (Fig. 4). The types of discontinuities considered include the following:

- *Mean value:* detects shifts in the mean of the analyzed variable. It is assumed that the data values follow a normal distribution, with all time steps sharing the same standard deviation. The mean value is constant within each segment and changes at the identified change points.
- *Standard deviation:* identifies changes in the standard deviation of the analyzed variable, assuming a normal distribution with a constant mean across all time steps. The standard deviation remains stable within each segment and shifts at the change points.
- *Linear trend (slope):* detects modifications in the linear trend. The data are assumed to follow a normal distribution with a mean value defined by a straight line. Both the slope and intercept are constant within each segment and change at the breakpoints.
- *Mean of count variables:* identifies shifts in the mean value of count data, assuming a Poisson distribution. The mean remains constant within each segment and changes at the detected change points.

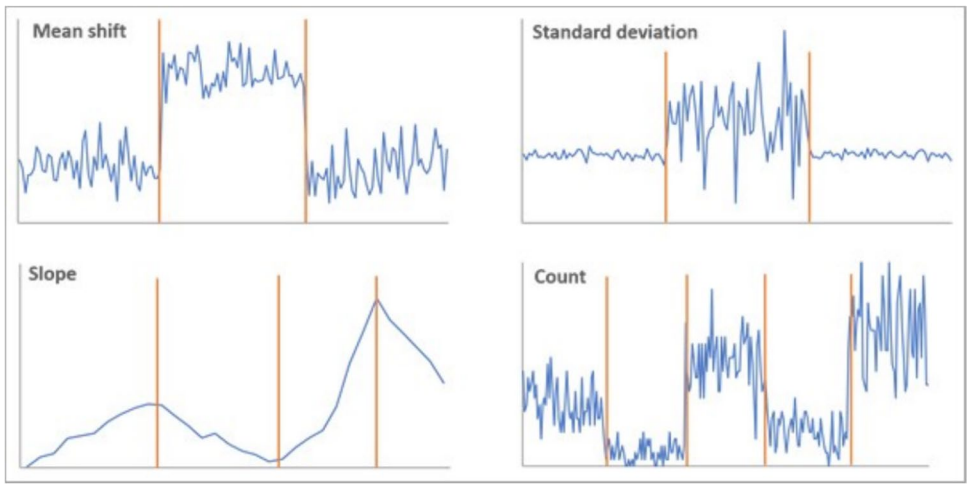

Fig. 4 The different types of feature changes analyzed for discontinuity detection.Figure adapted from https://pro.arcgis.com/

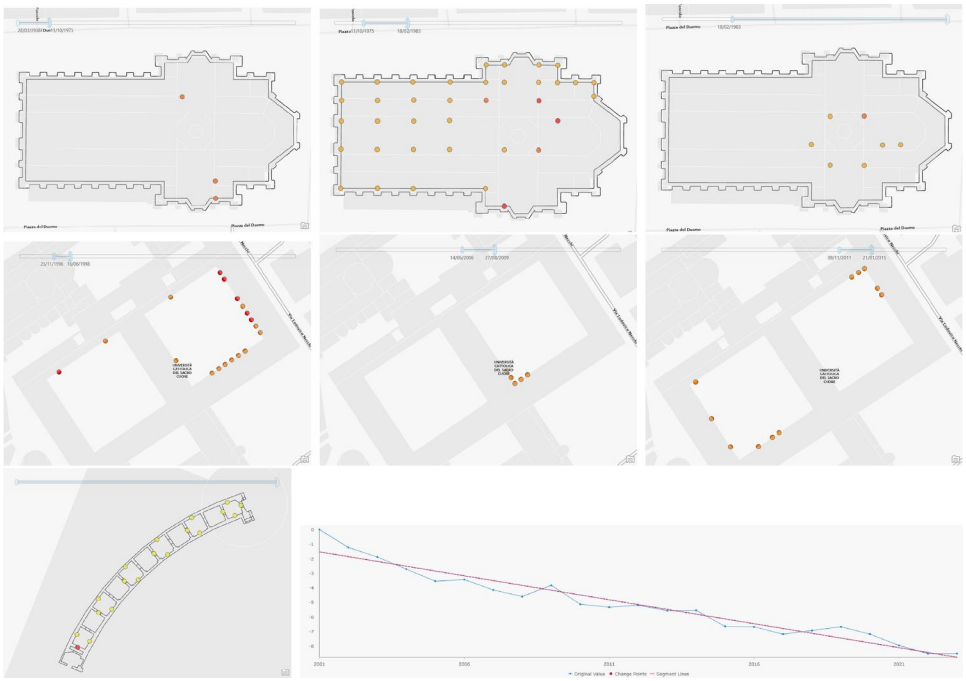

Fig. 5 Results of linear trend discontinuity detection for the three presented case studies: Duomo di Milano (top), Catholic University of Milan (middle), and San Siro Stadium (bottom)

The outputs of the discontinuity detection process include the total number of changes in each time series of the space–time cube and the time at which each change occurs. In our analysis, we prioritized the mean value and linear trend as the primary features for detection. The mean value helps identify abrupt shifts, while the linear trend highlights changes in directional movement (e.g., upward vs. downward trends).

The results of linear trend discontinuity detection are shown in Fig. 5. A notable finding emerged from clustering the data by the timing of slope changes in the Duomo di Milano dataset (Fig. 5, top). Specifically, only three pillars showed a trend change before 1975, while 37 pillars exhibited a change between 1975 and 1983, and only seven pillars have shown a change since 1983. This temporal distribution provides insights into the evolution of structural behavior over time. It is noteworthy that the upward trend persisted until 1983, after which the behavior of the pillars remained substantially stable.

A similar analysis was carried out for the Catholic University dataset (Fig. 5, middle). In this case, slope changes occurred predominantly in three periods: (i) 1996–1998; (ii) 2006–2009; and (iii) 2011–2015. No significant changes were detected outside these intervals.

No discontinuities were detected in the San Siro Stadium dataset; instead, a consistent downward trend was identified (Fig. 5, bottom).

Results from mean value discontinuity detection are presented in Fig. 6, where data are visualized based on the number of detected changes. For the Catholic University dataset

(Fig. 6, top left), the majority of points do not exhibit significant changes (shown in yellow). Pillars identified with mean value discontinuities correspond closely with those displaying changes in linear behavior, suggesting that trend shifts significantly impacted the mean displacement values.

The results for the San Siro Stadium (Fig. 6, top right) appear more scattered and less interpretable, indicating that mean value analysis may not be well-suited to this dataset. In contrast, the Duomo di Milano dataset (Fig. 6, bottom) reveals that most displacements involve points located near the façade, confirming localized settlement patterns.

3.3 Clustering Structural Elements Based on Their Time-Series

As previously noted, partitioning space–time cube locations provides valuable insights into the differential behavior of monitoring points, which is crucial for evaluating the structural performance of the system. To achieve this, we conducted time-series clustering on the data obtained from the monitoring network.

Time-series clustering involves identifying locations within a space–time cube that exhibit similar temporal patterns based on shared time-series attributes and organizing them into distinct clusters. Three specific criteria were employed to categorize these clusters:

- *Similarity of Values Across Time*: benchmarks within the same cluster exhibit comparable values over time;
- *Tendency to Increase and Decrease Simultaneously*: benchmarks within the same cluster display proportional trends in value changes;
- *Display of Similar Periodic Patterns Across Time*: benchmarks within the same cluster share recurring temporal behaviors (i.e., cycles).

Clustering based on the similarity of values over time is the most direct method. It effectively groups time series that maintain similar values across their entire timelines. Similarity is computed using the Euclidean distance between corresponding values at each time step.

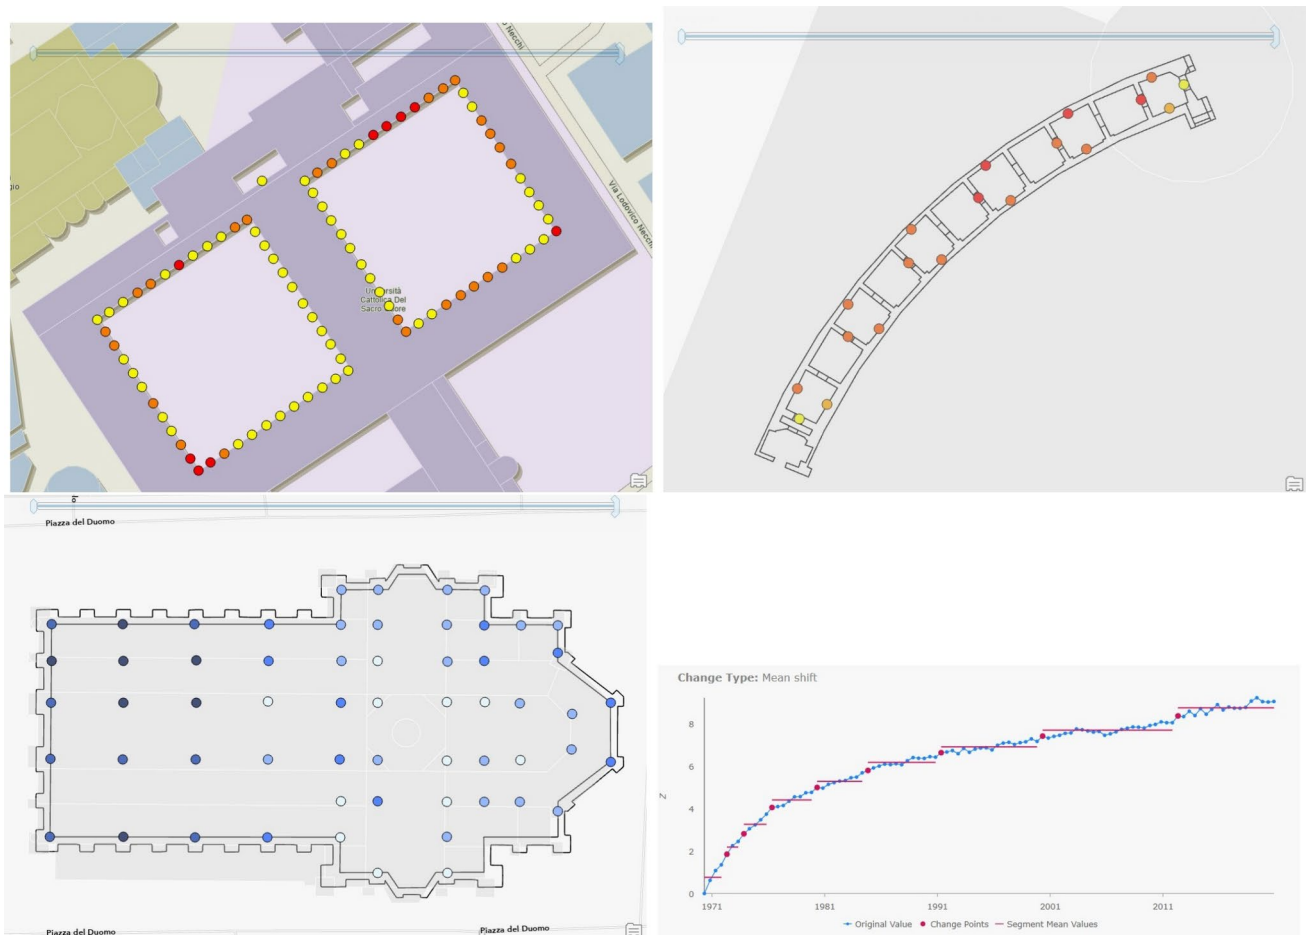

Fig. 6 Results of mean value discontinuity detection for the three presented case studies: Catholic University of Milan (top left), San Siro Stadium (top right), and Duomo di Milano (bottom)

Clustering based on synchronous increases or decreases over time groups time series exhibiting consistent, proportional trends, even if their absolute values differ. For instance, a time series with values (1, 2, 3, 4) and another with (10, 20, 30, 40) would be clustered together, as both show similar growth patterns. This method evaluates similarity using statistical correlation across time steps.

Clustering based on similar periodic patterns is applied to time series that display smooth, repeating cycles. Each time series is decomposed into a set of basis functions representing dominant signals using Fourier decomposition. Each basis function has an associated weight indicating the prominence of that signal. The dissimilarity between two time series is calculated as the sum of squared differences between the corresponding weights of their basis functions. Time series with similar dominant frequencies and oscillatory behavior are thus grouped together.

Optionally, some characteristics of these periodic patterns—such as starting times, durations, or magnitudes—can be ignored to focus solely on shape similarity. For example, two time series are considered similar if they share similar periodic durations, regardless of amplitude or alignment.

This analysis results in the partitioning of the space–time cube into clusters, with an associated average time series for each cluster. The optimal number of clusters is determined using the pseudo-*F* statistic. Clustering results for the three case studies are presented in Figs. 7, 8, and 9.

Figure 7 illustrates the clustering results for the Duomo di Milano using the three approaches: value-based, correlation-based, and Fourier-based clustering. While the first two methods identify well-separated clusters, Fourier clustering reveals three main patterns: a group oscillating around zero (blue), a group with a slightly increasing trend (red), and a group with a more pronounced increasing trend (green). These results show similarities with those obtained through hotspot detection (see the following section).

Figure 8 shows clustering results for the Catholic University case study using value and Fourier methods. The results identify clusters with broadly similar trends, as well as one distinctive cluster (in purple) representing a unique benchmark with a different behavioral pattern.

For the San Siro Stadium case (Fig. 9), clustering results from the value and Fourier methods are nearly identical, while the correlation-based method produces slightly different groupings. All methods reveal a similar pattern in which the beam's endpoints remain nearly stable, while differential behavior is observed between the inner and outer sections of the beam.

3.4 A Novel Method to Assist User's Interpretation: Hotspot and Coldspot in Time Series of Different Settelements

The space–time cube is employed to identify significant hotspots and coldspots using the Getis-Ord Gi* statistic. This spatial statistical method detects areas of high (hotspots) or low (coldspots) values by analyzing the spatial distribution of the data. The Gi* statistic is calculated for each spatial bin by considering its neighboring features, generating a *z*-score that indicates statistical significance. High positive *z*-scores correspond to hotspots—dense clusters of significantly high values—while high negative *z*-scores correspond to coldspots—clusters of significantly low values.

The *z*-scores and associated *p*-values allow for the identification of spatial clustering in high or low-value features. This analysis incorporates each feature's context by comparing it with its neighbors. For a feature to be classified as a significant hotspot, it must not only exhibit a high value itself but also be surrounded by other high-value features. This determination is made by comparing the local sum of a feature and its neighbors to the global sum across all features. A statistically significant deviation from the expected local sum, unlikely to occur by chance, results in a significant *z*-score.

The accuracy of the analysis relies heavily on how well spatial interactions among features are modeled. The objective is to reflect the natural spatial relationships among monitored elements. Given that the locations stored within space–time cubes may have a non-uniform spatial distribution, a strategy was implemented to ensure a minimum number of neighboring points for each feature. To this end, the *k*-nearest neighbors method was selected, with each feature evaluated using at least eight neighbors. This approach ensures consistency across the analysis, particularly in scenarios where feature density varies spatially, and is well-suited for situations where maintaining a fixed number of neighbors is prioritized over a fixed spatial scale.

In the case of the Duomo di Milano, as shown in the bottom-right panel of Fig. 10, hotspot analysis reveals a clear and consistent pattern across different sections of the cathedral. Notably, the southern transept area exhibits relative movements that contrast with those observed in the nave, presbytery, and northern transept. These results are consistent with long-term subsidence processes that began in the mid-twentieth century and have been extensively documented in prior works such as [38, 40]. This phenomenon necessitated significant restoration work on the columns during the 1980 s.

The dataset from the Catholic University reveals distinct motion patterns affecting different portions of the courtyards (Fig. 10, bottom-left). In particular, movement is

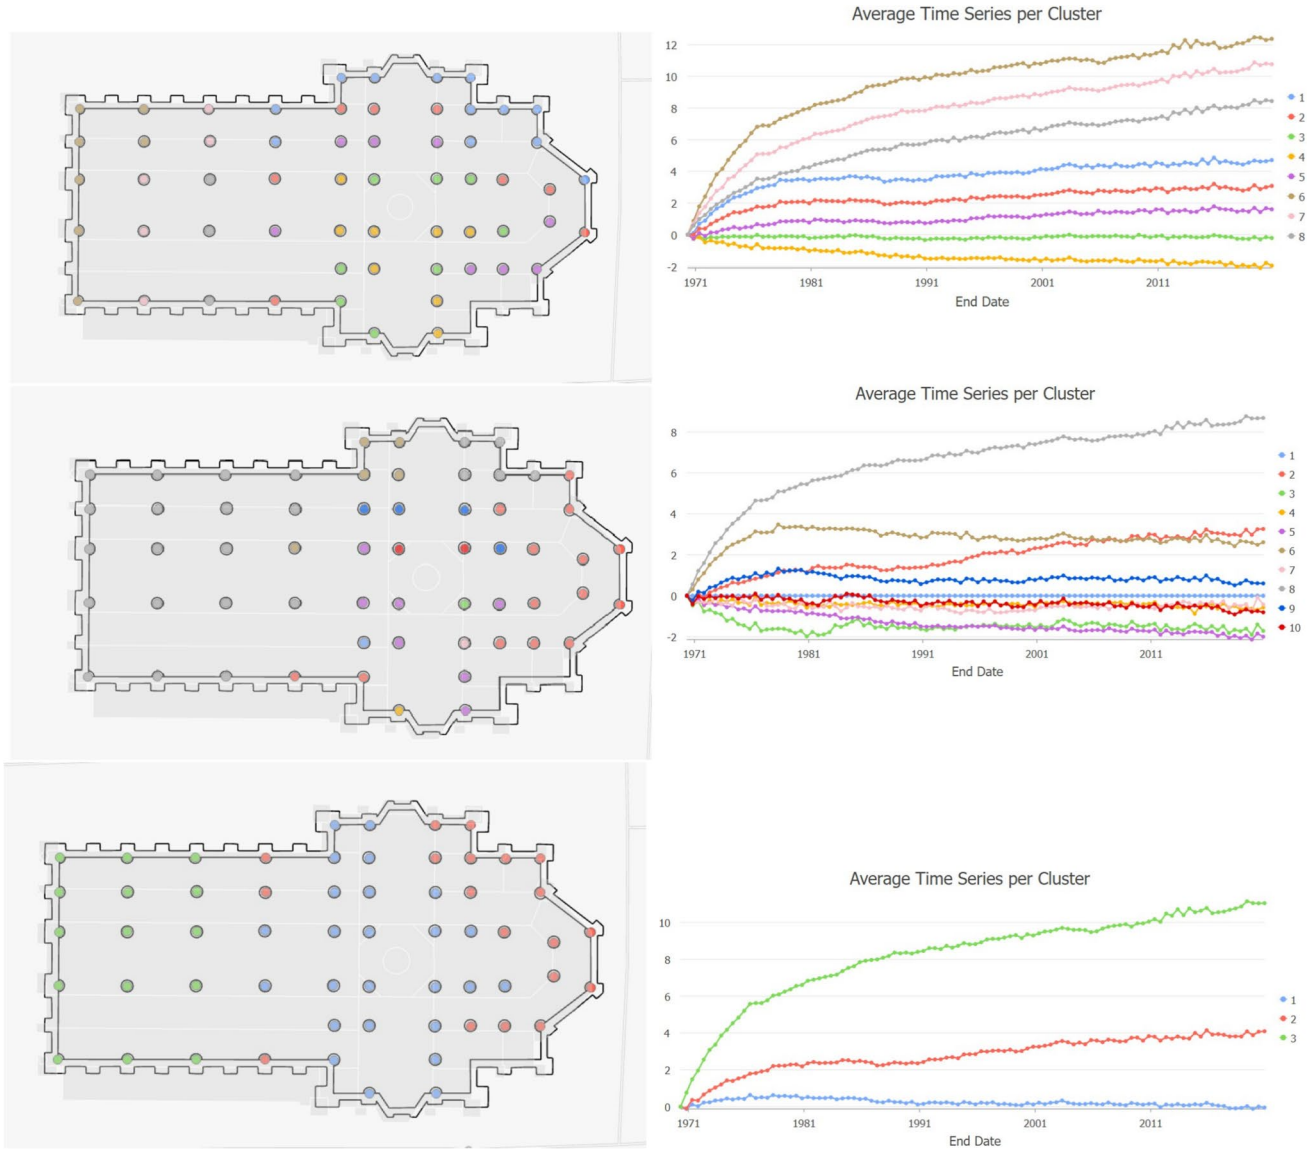

Fig. 7 Time-series clustering results for Duomo di Milano: clusters (left) and average time series per cluster (right) obtained using value-based (top), correlation-based (middle), and Fourier-based (bottom) clustering methods

more concentrated at the corners than along the sides. This observation aligns with manual interpretations derived from multiple years of monitoring campaigns. Interestingly, some areas that were previously considered relatively stable have shown slight movements in recent years, potentially linked to increasing temperatures compared to earlier decades.

Lastly, the San Siro Stadium dataset, which consists of measurements distributed along a linear geometry, illustrates that hotspot analysis does not merely reflect upward or downward displacement (Fig. 10, top). The alternating red and blue colors should not be directly interpreted as corresponding to movement direction alone, as the analysis integrates spatiotemporal patterns holistically. Although the beam is supported by two columns, the structural behavior is asymmetrical, influenced by the non-symmetrical layout of the columns supporting the stadium's third ring.

In summary, the combination of space–time cube visualization and hotspot analysis does not directly quantify the magnitude of observed movements. However, it provides a valuable exploratory tool that enables users to generate detailed graphs by accessing individual components of the space–time cube via the GIS interface. Interpretation of results remains complex and must be informed by prior knowledge of each specific case study.

Fig. 8 Time-series clustering results for the Catholic University: clusters (left) and average time series per cluster (right) using value-based (top) and Fourier-based (bottom) clustering methods

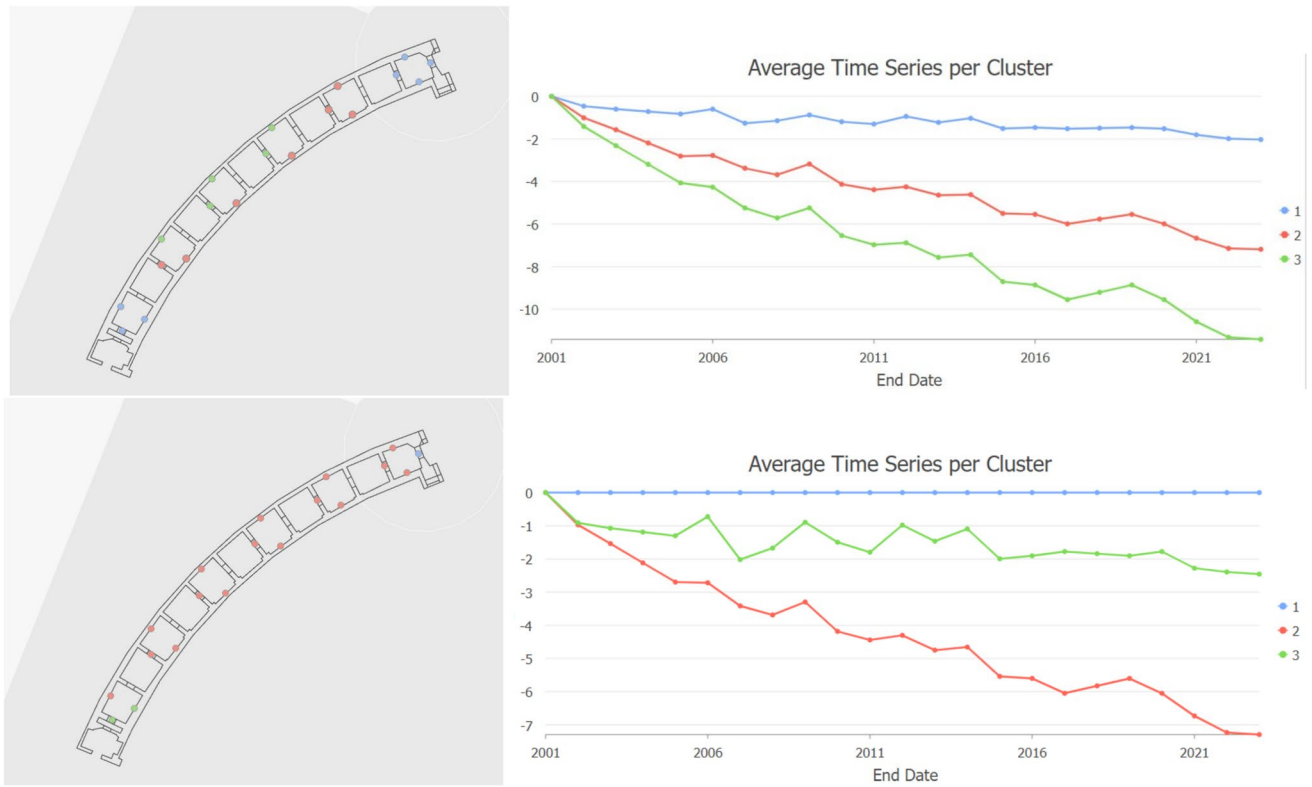

Fig. 9 Time-series clustering results for San Siro Stadium: clusters (left) and average time series per cluster (right) using value- and Fourier-based methods (top), and correlation-based method (bottom)

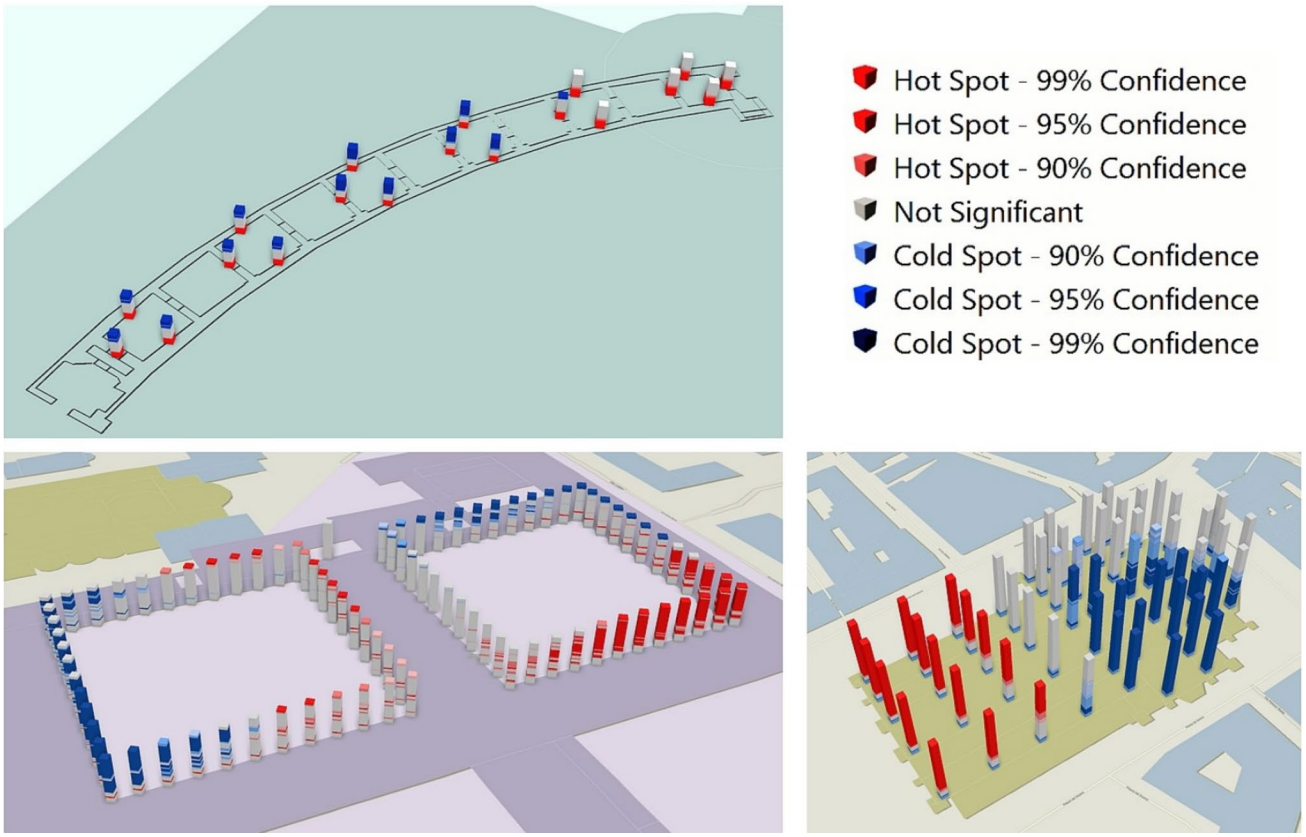

Fig. 10 Results of hotspot and coldspot detection using the Getis-Ord Gi* statistic

4 Geospatial Artificial Intelligence for the Analysis of Differential Settlements

4.1 Forecasting Using Space–Time Cubes and Machine Learning

We address the task of time-series forecasting utilizing the space–time cube to predict future values. Our focus is not on long-term forecasting (e.g., multiple years), but rather on short-term forecasts spanning a maximum of 2–3 years. These forecasts serve as reference datasets for future monitoring activities, enabling a comparison between forecasted and on-site measured values. Typically, variations are computed using the initial epoch (total variations) and the latest available epoch (partial variations). The proposed approach extends this by integrating forecasted values derived from the time series.

The method described in this section supports simultaneous processing and visualization of all measured points. It combines the space–time cube with a machine learning forecasting technique based on the forest-based method [41], which incorporates future time steps into the existing cube. Two models are constructed: a forecast model to predict future values, and a validation model to assess forecast accuracy.

The forecast model builds a forest using all time-series values and recursively predicts future values. After each forecasted time step, the predicted value is fed into the model to generate the next step. The forecast root mean square error (RMSE) is computed as the squared difference between the forecast and the actual time series values. The validation model operates similarly but excludes a number of final time steps from training. These excluded values serve as the reference to compare predicted outcomes, yielding the validation RMSE, which is generally considered a more accurate indicator of model performance.

Various tests were conducted on the proposed datasets, allowing the user to define the number of forecasted steps (set to 2 years in this case) and the number of steps excluded for validation. Forecasting was performed with 1, 3, and 5 years of excluded data. In the experiments, the random forest model was configured with 100 trees, which provided a stable balance between predictive accuracy and computational efficiency. The training data are built by constructing time-windowed pairs of explanatory and dependent variables. Any remaining portion of the training data is then used to optimize the model parameters.

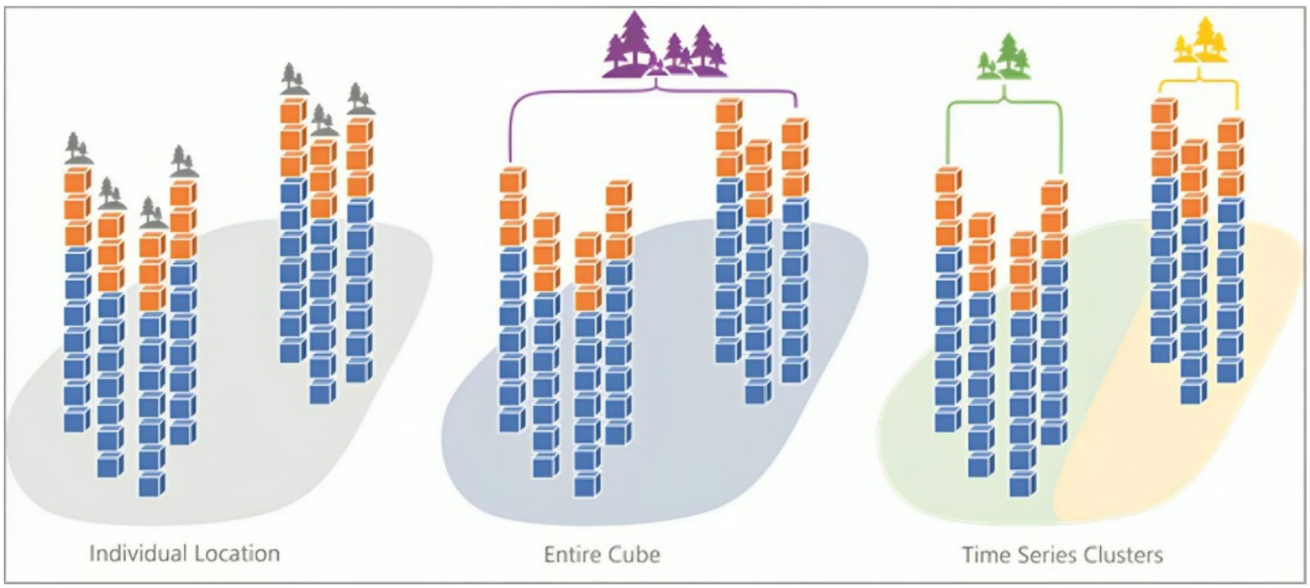

Fig. 11 Different processing strategies for time-series forecasting using random forest. Figure adapted from www.esri.com

Table 1 Statistics for the curved beam of San Siro stadium using forest-based forecast

San Siro					
Individual location					
Excluded time steps	Type	Max	Mean	Median	Std. dev.
5	Forecast RMSE	0.46	0.26	0.28	0.12
	Validation RMSE	1.92	1.01	1.08	0.52
3	Forecast RMSE	0.42	0.25	0.26	0.11
	Validation RMSE	2.21	1.15	1.25	0.61
1	Forecast RMSE	0.45	0.26	0.27	0.12
	Validation RMSE	0.94	0.55	0.57	0.25
Entire space–time cube					
Excluded time steps	Type	Max	Mean	Median	Std. dev.
5	Forecast RMSE	0.40	0.19	0.19	0.09
	Validation RMSE	1.84	0.72	0.81	0.43
3	Forecast RMSE	0.20	0.21	0.21	0.09
	Validation RMSE	2.21	0.49	0.29	0.55
1	Forecast RMSE	0.39	0.19	0.21	0.09
	Validation RMSE	1.13	0.24	0.09	0.31
Clustered time-series					
Excluded time steps	Type	Max	Mean	Median	Std. dev.
5	Forecast RMSE	0.33	0.18	0.18	0.08
	Validation RMSE	1.85	0.81	0.73	0.47
3	Forecast RMSE	0.36	0.18	0.20	0.08
	Validation RMSE	2.25	1.06	1.03	0.57
1	Forecast RMSE	0.32	0.18	0.21	0.08
	Validation RMSE	1.15	0.60	0.56	0.29

The RMSE values are in millimeters

For the San Siro curved beam and the Catholic University of Milan, which feature annual measurements, the number of excluded steps equals the number of excluded years. For the Cathedral of Milan, which has biannual measurements, the number of excluded steps is twice the number of excluded years.

Additional processing considered different time-series clustering strategies (see Fig. 11). In the individual location analysis, each time series is modeled independently. Alternatively, a model is constructed using the entire space–time cube. Finally, clustering-based models are built separately for each identified cluster.

Forecasting results are presented in Tables 1, 2, and 3 for the respective case studies. RMSE values are given in millimeters. The results exhibit minor differences across the three strategies (individual series, entire cube, and clustering). However, for the San Siro curved beam, the validation RMSE is significantly higher, ranging from ± 0.3 to 0.6 mm. This exceeds the height precision typically achieved through least squares adjustment (about ± 0.1 mm), indicating that the forecasted values lack sufficient accuracy for practical use in this case.

In contrast, the Catholic University and Cathedral of Milan datasets show significantly better performance, with validation RMSE values ranging between ± 0.08 and 0.15 mm even over a 5-year period. These findings suggest that the forest-based method yields reliable forecasts that could serve as supplementary references for upcoming monitoring campaigns.

The inferior performance observed for San Siro remains unexplained and merits further investigation. One hypothesis is that the relatively short time series (only 19 years) lacks sufficient data to train a robust model. Additionally, some measurements were taken during seasonal transitions (summer to autumn), potentially introducing variability due

Table 2 Statistics for the Catholic University using forest-based forecast

Catholic University					
Individual location					
Excluded time steps	Type	Max	Mean	Median	Std. dev.
5	Forecast RMSE	0.16	0.09	0.09	0.02
	Validation RMSE	0.60	0.28	0.26	0.12
3	Forecast RMSE	0.17	0.09	0.09	0.02
	Validation RMSE	0.68	0.31	0.29	0.14
1	Forecast RMSE	0.16	0.09	0.09	0.02
	Validation RMSE	0.74	0.19	0.17	0.14
Entire space–time cube					
Excluded time steps	Type	Max	Mean	Median	Std. dev.
5	Forecast RMSE	0.16	0.07	0.07	0.02
	Validation RMSE	0.61	0.30	0.28	0.12
3	Forecast RMSE	0.17	0.07	0.07	0.02
	Validation RMSE	0.74	0.30	0.27	0.16
1	Forecast RMSE	0.16	0.07	0.07	0.02
	Validation RMSE	0.50	0.15	0.13	0.12
Clustered time-series					
Excluded time steps	Type	Max	Mean	Median	Std. dev.
5	Forecast RMSE	0.18	0.07	0.06	0.02
	Validation RMSE	0.70	0.29	0.26	0.13
3	Forecast RMSE	0.16	0.07	0.06	0.02
	Validation RMSE	0.71	0.30	0.28	0.16
1	Forecast RMSE	0.17	0.07	0.06	0.02
	Validation RMSE	0.51	0.16	0.14	0.12

The RMSE values are in millimeters

Table 3 Statistics for the Cathedral of Milan using forest-based forecast

Cathedral of Milan					
Individual location					
Excluded time steps	Type	Max	Mean	Median	Std. dev.
5	Forecast RMSE	0.09	0.05	0.05	0.01
	Validation RMSE	0.39	0.21	0.20	0.07
3	Forecast RMSE	0.09	0.05	0.06	0.01
	Validation RMSE	0.52	0.23	0.21	0.10
1	Forecast RMSE	0.09	0.05	0.06	0.01
	Validation RMSE	0.35	0.17	0.16	0.08
Entire space–time cube					
Excluded time steps	Type	Max	Mean	Median	Std. dev.
5	Forecast RMSE	0.09	0.05	0.05	0.01
	Validation RMSE	0.47	0.23	0.20	0.10
3	Forecast RMSE	0.09	0.05	0.05	0.01
	Validation RMSE	0.51	0.19	0.17	0.08
1	Forecast RMSE	0.09	0.05	0.05	0.01
	Validation RMSE	0.37	0.17	0.15	0.09
Clustered time-series					
Excluded time steps	Type	Max	Mean	Median	Std. dev.
5	Forecast RMSE	0.09	0.05	0.05	0.01
	Validation RMSE	0.59	0.22	0.20	0.09
3	Forecast RMSE	0.09	0.05	0.05	0.01
	Validation RMSE	0.52	0.20	0.18	0.09
1	Forecast RMSE	0.09	0.05	0.05	0.01
	Validation RMSE	0.35	0.16	0.14	0.08

The RMSE values are in millimeters

to environmental conditions. Future studies may consider incorporating environmental variables into the model to improve forecast reliability.

4.2 From Geospatial Machine Learning to Deep Learning: Advantages and Limitations

Time-series analysis via geospatial deep learning involves applying advanced neural network architectures to extract patterns from temporal geospatial data. This area has gained considerable attention due to the increasing availability of sensor-based data and the development of deep learning tools tailored for temporal and spatial applications.

Deep learning models offer advantages over traditional machine learning by automatically learning hierarchical representations from raw data, thereby capturing complex relationships in time-series without the need for handcrafted features. Architectures such as recurrent neural networks (RNNs) and long short-term memory networks (LSTMs) are particularly suited to model temporal dependencies. These models are effective for handling univariate, multivariate, and spatiotemporal data, capable of capturing both short- and long-term trends. However, deep learning models require large datasets to generalize effectively and may underperform when time series are sparse. In contrast, traditional models like linear regression or decision trees offer more interpretability, allowing for greater insight into the forecasting logic, while deep learning models tend to behave as black-box systems.

In this work, we implemented a deep learning approach based on the diverse trends and seasonal components observed in datasets such as the Duomo di Milano. Traditional statistical models may struggle when data exhibit non-uniform trends and seasonality, while deep learning models—given sufficient training data—can adapt to such complexity and provide accurate forecasts.

The forecasting process using the space–time cube involves two main steps: (i) training a deep learning model on the available data, and (ii) forecasting future values at each spatial location. The framework supports multiple deep learning models, including Fully Connected Networks (FCNs), LSTMs, InceptionTime, ResNet, and ResCNN, all of which can be applied to univariate or multivariate time series. In this study, a univariate model was used with elevation variation as the input variable. Model parameters are optimized over a maximum of 20 epochs, although early

stopping is enabled by default and automatically interrupts training when the validation loss does not improve over five consecutive epochs, preventing unnecessary computation and overfitting. The batch size used during training defaults to 64.

The deep-learning architectures available in the tool differ significantly in depth and internal structure. The Fully Convolutional Network (FCN) is the simplest model, composed of three sequential 1-D convolutional layers followed by a global average-pooling layer. ResNet employs a deeper configuration based on three residual blocks, each containing multiple convolutional layers combined with batch-normalization and nonlinear activations; the residual connections further increase the effective depth of the network before the final global pooling stage. InceptionTime adopts an even more modular architecture, built from a series of inception modules that operate in parallel with different kernel sizes and include max-pooling branches; skip connections are inserted after every third module. Because the number of inception modules may vary internally, the total layer count depends on the number of modules used in the default implementation. ResCNN combines ideas from convolutional and residual networks, using several convolutional layers and skip connections to form a moderately deep architecture. The LSTM model is structurally simpler, typically consisting of a single recurrent hidden layer with 100 units by default, followed by the output layer.

As in the machine learning experiments, time series were split into training and validation sets. The last 10% of time steps were reserved for validation. The training process included a maximum of 20 epochs, as further increases did not improve model performance. The most influential hyperparameter was the sequence length, representing the number of prior time steps used for prediction. For the Duomo dataset, the minimum meaningful sequence length was set to one season (2 time steps). Forecasting performance was evaluated using RMSE for both forecast and validation sets. The RMSE values, measured in millimeters, are reported in Tables 4, 5, and 6.

For the Duomo dataset, we observed significant RMSE differences across methods. LSTM and InceptionTime achieved the best results, with LSTM validation RMSE ranging from ± 0.25 to 0.43 mm and a standard deviation of approximately ± 0.1 mm. Validation RMSE was generally twice that of the forecast RMSE, and in some cases (e.g., ResCNN), 3–4 times higher.

The Catholic University dataset showed more consistent RMSE values across different models, with validation RMSE ranging from ± 0.23 to 0.33 mm and a standard deviation of about ± 0.2 mm. As with the Duomo dataset, longer sequence lengths improved model performance.

Table 4 Deep learning statistics for the Cathedral of Milan

Cathedral of Milan						
LSTM						
Sequence length	Type	Max	Mean	Median	Std.dev.	
20	Forecast RMSE	0.02	0.24	0.15	0.15	0.04
	Validation RMSE	0.05	0.55	0.27	0.25	0.12
4	Forecast RMSE	0.08	0.48	0.22	0.18	0.09
	Validation RMSE	0.24	0.63	0.43	0.43	0.10
2	Forecast RMSE	0.04	0.35	0.18	0.15	0.06
	Validation RMSE	0.19	0.70	0.42	0.42	0.12
Inception						
Sequence length	Type	Max	Mean	Median	Std.dev.	
20	Forecast RMSE	0.04	0.49	0.25	0.18	0.13
	Validation RMSE	0.14	1.44	0.63	0.54	0.37
4	Forecast RMSE	0.06	1.46	0.37	0.19	0.38
	Validation RMSE	0.14	4.32	1.08	0.37	1.34
2	Forecast RMSE	0.07	0.34	0.21	0.22	0.04
	Validation RMSE	0.05	3.06	0.54	0.40	0.52
ResCNN						
Sequence length	Type	Max	Mean	Median	Std.dev.	
20	Forecast RMSE	0.11	0.68	0.31	0.26	0.15
	Validation RMSE	0.15	1.83	0.79	0.71	0.38
4	Forecast RMSE	0.13	1.25	0.38	0.20	0.35
	Validation RMSE	0.10	7.11	1.86	0.96	2.11
2	Forecast RMSE	0.13	0.56	0.31	0.30	0.11
	Validation RMSE	0.15	2.20	1.07	1.12	0.58

The RMSE values are in millimeters

Table 5 Deep learning statistics for the Catholic University

Catholic University						
LSTM						
Sequence length	Type	Max	Mean	Median	Std.dev.	
10	Forecast RMSE	0.07	0.67	0.17	0.16	0.07
	Validation RMSE	0.04	1.70	0.34	0.35	0.21
4	Forecast RMSE	0.00	0.52	0.19	0.19	0.05
	Validation RMSE	1.15	0.28	0.27	0.17	
2	Forecast RMSE	0.07	0.36	0.23	0.23	0.04
	Validation RMSE	0.05	0.66	0.27	0.24	0.14
ResCNN						
Sequence length	Type	Max	Mean	Median	Std.dev.	
10	Forecast RMSE	0.03	0.62	0.14	0.14	0.06
	Validation RMSE	0.04	1.99	0.34	0.34	0.24
4	Forecast RMSE	0.14	0.48	0.22	0.22	0.05
	Validation RMSE	0.10	0.92	0.37	0.32	0.18
2	Forecast RMSE	0.05	0.57	0.27	0.26	0.06
	Validation RMSE	0.05	0.85	0.31	0.29	0.16
FCN						
Sequence length	Type	Max	Mean	Median	Std.dev.	
10	Forecast RMSE	0.03	0.57	0.14	0.13	0.06
	Validation RMSE	0.04	1.93	0.34	0.31	0.26
4	Forecast RMSE	0.01	0.73	0.18	0.17	0.07
	Validation RMSE	0.01	1.01	0.32	0.31	0.20
2	Forecast RMSE	0.05	0.84	0.27	0.26	0.10
	Validation RMSE	0.05	2.63	0.37	0.33	0.32

The RMSE values are in millimeters

Table 6 Deep learning statistics for the San Siro Stadium

San Siro Stadium						
LSTM						
Sequence length	Type	Max	Mean	Median	Std.dev.	
7	Forecast RMSE	0.09	0.90	0.43	0.43	0.21
	Validation RMSE	0.03	1.60	0.38	0.19	0.43
4	Forecast RMSE	0.29	0.61	0.39	0.36	0.09
	Validation RMSE	0.06	0.79	0.36	0.26	0.23
2	Forecast RMSE	0.14	0.82	0.46	0.48	0.19
	Validation RMSE	0.09	1.32	0.41	0.28	0.34
Inception						
Sequence length	Type	Max	Mean	Median	Std.dev.	
7	Forecast RMSE	0.27	0.65	0.38	0.36	0.10
	Validation RMSE	0.06	1.07	0.49	0.40	0.32
4	Forecast RMSE	0.47	0.96	0.63	0.55	0.15
	Validation RMSE	0.05	1.25	0.61	0.74	0.38
2	Forecast RMSE	0.55	1.82	0.82	0.63	0.37
	Validation RMSE	0.05	3.42	0.62	0.40	0.81
FCN						
Sequence length	Type	Max	Mean	Median	Std.dev.	
7	Forecast RMSE	0.13	0.66	0.35	0.32	0.15
	Validation RMSE	0.03	1.26	0.57	0.41	0.39
4	Forecast RMSE	0.39	0.61	0.47	0.45	0.06
	Validation RMSE	0.48	1.11	0.75	0.71	0.15
2	Forecast RMSE	0.60	1.33	0.84	0.88	0.20
	Validation RMSE	0.34	1.95	1.17	1.24	0.48

The RMSE values are in millimeters

The San Siro dataset exhibited the highest RMSE values ($\pm$ 0.4–1.2 mm) and the least variability among methods. Forecast and validation RMSE values were similar, suggesting limited benefit from deep learning. The reduced number of time steps (19 epochs) likely constrained the models' ability to generalize.

In general, deep learning predictions produced RMSE values comparable to or slightly worse than those from machine learning. Deep learning excels with large training datasets and can model complex dependencies, but in data-sparse scenarios, both machine learning and deep learning encounter similar challenges in capturing underlying patterns. This results in comparable forecast performance despite the greater complexity of deep learning architectures.

4.3 Quality Assessment of Forecasted Time-Series Values and Discontinuity Detection for New Monitoring Epochs

An essential task in any new monitoring campaign is to verify that the newly acquired measurements remain sufficiently close to previous observations collected under comparable environmental and operational conditions. When a structure exhibits similar (and typically small) displacements across different years, and the measurements are taken in analogous acquisition conditions, this behaviour generally suggests that no significant trend is developing and that the structure is undergoing a periodic or quasi-periodic motion. Traditionally, this verification is performed by directly comparing the most recent measured values with those from the previous epoch, following the temporal resolution established in the monitoring plan.

However, relying solely on the last available measurements ignores the information contained in the entire time series. Using historical observations together with forecasted values and their associated statistics provides a more robust and statistically rigorous mechanism for detecting anomalies. Forecasts include an estimate of their uncertainty, which makes it possible to detect discontinuities using statistically meaningful thresholds (for example, twice the forecast standard deviation), rather than simply checking whether the latest measurement deviates from the previous one. Nonetheless, to ensure meaningful comparisons, the precision of the forecasted values must be compatible with the precision of the least-squares adjusted measurements; otherwise, the forecasted values cannot serve as a reliable reference dataset.

To provide a quantitative assessment of this compatibility, we introduce a simple index that compares the accuracy of the forecasts with the intrinsic precision of the monitoring network obtained through least-squares adjustment. The index is defined as the ratio between the mean value of the forecasted RMSE over p prediction steps and the reference standard deviation σ_0 derived from the least-squares adjustment. The proposed index is therefore computed as $\frac{1}{p}\sum_{k=1}^{p} \mathrm{RMSE}_k/\sigma_0$.Values of the index close to 1 indicate that the accuracy of the forecasts is consistent with the precision of the adjusted monitoring measurements, whereas significantly larger values suggest that the forecasted values may not be sufficiently reliable for discontinuity detection. Conversely, values substantially smaller than 1 may indicate problems or inconsistencies in the acquisition or adjustment of the new monitoring measurements.

The computation of σ_0 follows standard least-squares theory. Let v be the vector of residuals and P the weight matrix of the observations. The a posteriori variance factor (or reference standard deviation) is computed as $\sigma_0^2 = v^T P v/(n-u)$, where n is the total number of observations and u the number of estimated unknown parameters. The variance–covariance matrix of the adjusted parameters, used to assess the precision of the estimated displacements, is then given by $C_x = \sigma_0^2 \left(A^T P A\right)^{-1}$.

For levelling networks, both σ_0 and the forecast RMSEs have the same physical units (typically millimetres), which makes the ratio I easy to interpret as a global quality indicator describing the balance between measurement precision and forecasting accuracy.

If the index is much larger than 1, this indicates that forecasted values are considerably less accurate than the measured and adjusted values, and should therefore be used with caution when checking for discontinuities. If the index is significantly smaller than 1, it may instead highlight issues in the acquisition or least-squares adjustment of the monitoring data. Values close to 1 are generally desirable, as they suggest that the forecasted values and the adjusted measurements exhibit comparable precision, enabling reliable detection of deviations during new monitoring campaigns.

The three monitoring datasets analysed in the present work produced values of 5.1, 1.1, and 1.46 for Case Studies 1, 2, and 3 respectively (computed as the mean over all prediction steps).

The first dataset showed a particularly small value of σ_0 owing to the very short distances between the monitored points, which resulted in highly precise least-squares adjustments. Combined with the limited number of epochs available for training the forecasting model, this situation reduced the accuracy of the predictions and led to a relatively large value of I. This suggests that, for this specific dataset, traditional comparison techniques based solely on adjusted measurements may be preferable for discontinuity detection. By contrast, the other two case studies produced values close to the theoretically expected range, indicating

a good balance between measurement precision and forecasting accuracy. In these cases, the index confirms that forecasted values can be confidently used as a reference to support anomaly detection in future monitoring campaigns.

5 Conclusion

This paper presented the application of space–time cubes for storing and processing monitoring measurements related to differential settlements in constructions, detected through high-precision geometric leveling. The space–time structure stores absolute variations relative to a conventionally fixed benchmark, enabling a consistent representation of temporal deformation patterns.

The space–time cube offers not only an effective storage solution but also a platform for visualizing and further analyzing the acquired data. It complements traditional time-series plots and conventional displacement visualizations (e.g., vector fields in 2D or 3D). After constructing the space–time cube, users can interact with the data through both static and dynamic visualizations. While these methods do not replace the well-established graphical representations commonly used in technical reports, they provide an intuitive and accessible means to explore the dataset, supporting a deeper understanding of structural behavior.

Once populated with data from multiple monitoring campaigns, the space–time cube enables advanced processing operations such as discontinuity detection and hotspot analysis. These tools offer automated, scalable mechanisms for identifying temporal or spatial anomalies across numerous time series. The resulting numerical and visual outputs assist users in identifying behavioral trends and deviations, making the space–time cube a practical solution for processing complex, multi-location monitoring datasets.

Moreover, the space–time cube framework supports seamless integration with geospatial artificial intelligence (GeoAI) techniques and GIS-based processing. The context of settlement monitoring is particularly well-suited to geospatial representation, given that monitored phenomena typically occur in localized areas represented by fixed measurement benchmarks. In this study, we evaluated both machine learning and deep learning forecasting techniques for detecting trends and discontinuities. Results indicate that the forest-based machine learning method generally outperformed deep learning models, particularly in scenarios involving limited training data. This can be attributed to its simpler setup, reduced parameter space, and robustness under constrained data conditions—characteristics especially advantageous when the number of epochs is below 130, as in our case.

Nonetheless, deep learning methods may be better suited for applications involving larger datasets or more complex temporal dynamics. Therefore, testing and comparing multiple forecasting approaches remains essential to identify the most suitable method for a specific dataset. In future applications, combining different forecasting models or averaging their outputs could further enhance predictive accuracy and robustness.

Finally, we assess forecasting performance through a composite index that integrates multi-step temporal prediction accuracy with least-squares statistics of the adjusted data. This index provides users with an immediate diagnostic to verify whether the forecasts are reliable before proceeding with further analysis. Across the three datasets examined—each representing monitoring projects with distinct characteristics—machine-learning forecasts remain stable for at least three steps ahead. This predictive horizon is sufficient to support outlier detection in new monitoring campaigns by comparing observed measurements with the corresponding forecasted values.

In conclusion, the integration of space–time cubes with machine learning and deep learning techniques offers a powerful framework for structural health monitoring. It supports efficient data organization, enables advanced visualization, and facilitates intelligent trend analysis and forecasting, contributing to a more comprehensive and interpretable understanding of structural behavior over time.

6 Article Highlights

A novel approach to effectively store and process datasets from structural monitoring, ensuring precision and efficiency in structural monitoring operations related to settlements.

Exploiting artificial intelligence to analyze georeferenced time-series, utilizing space–time cubes for comprehensive monitoring and prediction.

Demonstrating practical implementations and simulations in real-world scenarios with thorough validation and interpretation of results, testing reliability and applicability.

Funding Open access funding provided by Politecnico di Milano within the CRUI-CARE Agreement.

Declarations

Conflict of interest The authors declare that they have no conflict of interest related to this work.

References

1. Farrar CR, Worden K (2007) An introduction to structural health monitoring. Philos Trans R Soc A Math Phys Eng Sci 365(1851):303–315. https://doi.org/10.1098/rsta.2006.1928
2. Worden K, Farrar C, Manson G (2007) The fundamental axioms of structural health monitoring. Proc R Soc A Math Phys Eng Sci 463:1639–1664. https://doi.org/10.1098/rspa.2007.1834
3. Moretto S, Bozzano F, Mazzanti P (2021) The role of satellite InSAR for landslide forecasting: limitations and openings. Remote Sens. https://doi.org/10.3390/rs13183735
4. Zhang B, Chang L, Stein A (2021) Spatio-temporal linking of multiple SAR satellite data from medium and high resolution RADARSAT-2 images. ISPRS J Photogramm Remote Sens 176:222–236. https://doi.org/10.1016/j.isprsjprs.2021.04.005
5. Shen N, Chen L, Liu J, Wang L, Tao T, Wu D, Chen R (2019) A review of global navigation satellite system (GNSS)-based dynamic monitoring technologies for structural health monitoring. Remote Sens 11:9. https://doi.org/10.3390/rs11091001
6. Lazecky M, Hlavacova I, Bakon M, Sousa JJ, Perissin D, Patricio G (2017) Bridge displacements monitoring using space-borne x-band SAR interferometry. IEEE J Sel Top Appl Earth Obs Remote Sens 10(1):205–210. https://doi.org/10.1109/JSTARS.2016.2587778
7. Giordano PF, Turksezer ZI, Previtali M, Limongelli MP (2022) Damage detection on a historic iron bridge using satellite DInSAR data. Struct Health Monit 21(5):2291–2311. https://doi.org/10.1177/14759217211054350
8. Lindner G, Schraml K, Mansberger R, Hübl J (2016) UAV monitoring and documentation of a large landslide. Appl Geomat 8(1):1–11. https://doi.org/10.1007/s12518-015-0165-0
9. Pehlivan H, Bayata H (2016) Usability of inclinometers as a complementary measurement tool in structural monitoring. Struct Eng Mech 58:1077–1085. https://doi.org/10.12989/sem.2016.58.6.1077
10. Lekidis V, Tsakiri M, Makra K, Karakostas C, Klimis N, Issam S (2004) Evaluation of dynamic response and local soil effects of the Evripos cable-stayed bridge using multi-sensor monitoring systems. Eng Geol 79:43–59. https://doi.org/10.1016/j.enggeo.2004.10.015
11. Bian H-F, Zhang S-B, Zhang Q-Z, Zheng N-S (2014) Monitoring large-area mining subsidence by GNSS based on IGS stations. Trans Nonferrous Met Soc China 24(2):514–519. https://doi.org/10.1016/S1003-6326(14)63090-9
12. Bitelli G, Bonsignore F, Pellegrino I, Vittuari L (2015) Evolution of the techniques for subsidence monitoring at regional scale: the case of Emilia-Romagna region (Italy). Proc Int Assoc Hydrol Sci 372:315–321. https://doi.org/10.5194/piahs-372-315-2015
13. Osmanoğlu B, Sunar F, Wdowinski S, Cabral-Cano E (2016) Time series analysis of InSAR data: methods and trends. ISPRS J Photogramm Remote Sens 115:90–102. https://doi.org/10.1016/j.isprsjprs.2015.10.003 (Theme issue 'State-of-the-art in photogrammetry, remote sensing and spatial information science')
14. Pepe A, Calò F (2017) A review of interferometric synthetic aperture radar (InSAR) multi-track approaches for the retrieval of earth's surface displacements. Appl Sci. https://doi.org/10.3390/app7121264
15. Bokhari R, Shu H, Tariq A, Al-Ansari N, Guluzade R, Chen T, Jamil A, Aslam M (2023) Land subsidence analysis using synthetic aperture radar data. Heliyon 9(3):14690. https://doi.org/10.1016/j.heliyon.2023.e14690
16. Grgić M, Bender J, Bašić T (2020) Estimating vertical land motion from remote sensing and in-situ observations in the Dubrovnik area (Croatia): a multi-method case study. Remote Sens. https://doi.org/10.3390/rs12213543
17. Fabris M, Battaglia M, Chen X, Menin A, Monego M, Floris M (2022) An integrated InSAR and GNSS approach to monitor land subsidence in the Po River Delta (Italy). Remote Sens. https://doi.org/10.3390/rs14215578
18. Bogucka EP, Jahnke M (2018) Feasibility of the space–time cube in temporal cultural landscape visualization. ISPRS Int J Geo-Inf. https://doi.org/10.3390/ijgi7060209
19. Putrenko V, Pashvnska N, Nazarenko S (2018) Data mining of network events with space–time cube application, pp 79–83. https://doi.org/10.1109/DSMP.2018.8478437
20. Li X, Coltekin A, Kraak M-J (2010) Visual exploration of eye movement data using the space–time-cube, vol 6292, pp 295–309. https://doi.org/10.1007/978-3-642-15300-6_21
21. Xu D, Zhang Q, Ding Y, De Zhang A (2021) Spatiotemporal pattern mining of drought in the last 40 years in China based on the SPEI and space–time cube. J Appl Meteorol Climatol 60(9):1219–1230. https://doi.org/10.1175/JAMC-D-21-0049.1 (Cited by: 6; All Open Access, Bronze Open Access)
22. Osman A, Owusu AB, Adu-Boahen K, Atamey E (2023) Space–time cube approach in analysing conflicts in Africa. Soc Sci Humanit Open. https://doi.org/10.1016/j.ssaho.2023.100557 (Cited by: 0; All Open Access, Gold Open Access)
23. Hernández-Lara OG, Díaz-GarayúaJR, Butler KA (2022) COVID-19 deaths in México: a spatiotemporal analysis, vol 1, pp 949–961 (2022). https://doi.org/10.1007/978-3-030-94350-9_51 (Cited by: 0)
24. Ogneva-Himmelberger Y, Haynes M (2022) Using space–time cube to analyze trends in adverse birth outcomes and maternal characteristics in Massachusetts, USA. GeoJournal 87(4):2491–2504. https://doi.org/10.1007/s10708-021-10382-w
25. Arco E, Ajmar A, Cremaschini F, Monaco C (2021) Spatio temporal data cube applied to AIS containerships trend analysis in the early years of the belt and road initiative—from global to local scale, vol 43, pp 71–78. https://doi.org/10.5194/isprs-archives-XLIII-B4-2021-71-2021
26. Wu P, Meng X, Song L (2022) Identification and spatiotemporal evolution analysis of high-risk crash spots in urban roads at the microzone-level: using the space–time cube method. J Transp Saf Secur 14(9):1510–1530. https://doi.org/10.1080/19439962.2021.1938323 (Cited by: 7)
27. çalışkan M, Anbaroğlu B (2023) Space time cube analytics in QGIS and python for hot spot detection. SoftwareX. https://doi.org/10.1016/j.softx.2023.101498 (Cited by: 0; All Open Access, Gold Open Access)
28. Song B, Wang Y, Li F (2021) The visualization representation of space–time-path in the space–time-cube, vol 906 (2021). https://doi.org/10.1088/1755-1315/906/1/012030 (Cited by: 0; All Open Access, Bronze Open Access)
29. Mo C, Tan D, Mai T, Bei C, Qin J, Pang W, Zhang Z (2020) An analysis of spatiotemporal pattern for COVID-19 in China based on space–time cube. J Med Virol 92(9):1587–1595. https://doi.org/10.1002/jmv.25834
30. Putrenko V, Pashvnska N, Nazarenko S (2018) Data mining of network events with space–time cube application, pp 79–83. http

s://doi.org/10.1109/DSMP.2018.8478437 (Cited by: 0; All Open Access, Green Open Access)
31. Ma Z, Bai L, Yan L (2020) Spatiotemporal data and spatiotemporal data models. Stud Comput Intell 894:1–18 (2020) https://doi.org/10.1007/978-3-030-41999-8_1 (Cited by: 3)
32. Krishnan P, Aggarwal P, Mridha N, Bajpai V (2019) Spatio-temporal changes in wheat crop cultivation in India 42:385–395. https://doi.org/10.5194/isprs-archives-XLII-3-W6-385-2019
33. Ssin SY, Walsh JA, Smith RT, Cunninqham A, Thomas BH (2019) GeoGate: correlating geo-temporal datasets using an augmented reality space–time cube and tangible interactions, pp 210–219. https://doi.org/10.1109/VR.2019.8797812 (Cited by: 27)
34. Bogucka EP, Jahnke M (2018) Feasibility of the space–time cube in temporal cultural landscape visualization. ISPRS Int J Geo-Inf. https://doi.org/10.3390/ijgi7060209 (Cited by: 11; All Open Access, Gold Open Access, Green Open Access)
35. Bogucka EP, Jahnke M (2017) Space–time cube—a visualization tool for landscape changes [space–time cube—ein visualisierungswerkzeug zur darstellung landschaftlicher veränderungen]. Kartogr Nachr 67(4):183–191 (Cited by: 4)
36. Cigada A, Moschioni G, Vanali M, Caprioli A (2009) The measurement network of the San Siro Meazza stadium in Milan: origin and implementation of a new data acquisition strategy for structural health monitoring: dynamic testing of civil engineering structures series. Exp Tech 34:70–81. https://doi.org/10.1111/j.1747-1567.2009.00536.x
37. Turrisi S, Cigada A, Zappa E (2021) The new monitoring system of the roof of the G. Meazza stadium in Milan. Structural 234:1–15. https://doi.org/10.12917/Stru234.07
38. Passano C (1988) Il Duomo Rinato. Storia e Tecnica del Restauro Statico dei Piloni del Tiburio del Duomo di Milano, p 136. Diakronia, Milano
39. Barazzetti L, Canali F, Torre SD, Gentile C, Previtali M, Roncoroni F (2022) Monitoring the cathedral of Milan: an archive with more than 50 years of measurements. In: Gervasi O, Murgante B, Hendrix EMT, Taniar D, Apduhan BO (eds) Computational science and its applications—ICCSA 2022. Springer, Cham, pp 575–590
40. Barazzetti L, Previtali M, Roncoroni F (2022) Visualization and processing of structural monitoring data using space–time cubes. Computational science and its applications—ICCSA 2022: 22nd international conference. Malaga, Spain, July 4–7, 2022, Proceedings, Part II. Springer, Berlin, Heidelberg, pp 19–31
41. Breiman L (2001) Random forests. Mach Learn 45(1):5–32. https://doi.org/10.1023/A:1010933404324