ECEE 2026
18th European Conference on Earthquake Engineering

Prototype System Architecture for SONATA: Sensor-Driven Statistical Tools to Evaluate and Manage Safety in the Built Environment

Davit Shahnazaryan[1], Marco Todaro[1], Mauricio Guamán[1], Volkan Ozsarac[1], Daniele Sivori[1], Gerard J. O'Reilly[1],

1- Centre for Training and Research on Reduction of Seismic Risk (ROSE Centre), Scuola Universitaria Superiore IUSS di Pavia, Pavia 27100, Italy

davit.shahnazaryan@iusspavia.it

Abstract

Italy's geographical position makes it particularly vulnerable to seismic hazards, a risk compounded by the widespread presence of structures built before the introduction of modern building codes. Non-structural elements, which comprise much of the built environment, are especially susceptible to repeated earthquake damage. Their failure threatens occupant safety, drives up repair costs, and delays the restoration of functionality underscoring the urgent need for more effective seismic risk management strategies.

The SONATA project addresses this challenge through an integrated, end-to-end architecture that combines distributed sensor networks, AI-based damage detection, and advanced probabilistic modeling into a unified workflow. The system begins with heterogeneous sensor networks and data from experimental campaigns, where non-structural elements are characterised under controlled conditions. Raw vibration and vision-based sensor data undergo preliminary filtering and validation before secure transmission to cloud-based infrastructure comprising time-series databases and scalable object storage optimised for high-velocity continuous monitoring. Distributed computing resources support the parallel processing of these large-scale datasets.

Deep learning algorithms through analysis of periodic sensor data are deployed for automated damage detection and identification of structural anomalies during and after seismic events. Risk models are dynamically updated by integrating data from both experimental campaigns and continuous field monitoring, enabling rapid post-event assessment and informed decision-making. The key innovation of this architecture lies in bridging experimental testing, structural health monitoring, and probabilistic risk assessment through a scalable, coherent data pipeline. By combining continuous monitoring with AI-driven damage characterisation, the SONATA project provides a replicable workflow for enhancing the resilience of the built environment ultimately improving occupant safety, reducing repair costs, and accelerating post-earthquake recovery.

Keywords: Seismic risk, Non-structural elements, Sensor networks, Experimental testing, Artificial intelligence.

1- INTRODUCTION

Use of structural health monitoring (SHM) has increased significantly over the last decade [1, 2]. The technology comprises a continuous or periodic process of damage detection and characterisation through the integration of sensors, data acquisition systems, and computational algorithms that track changes in structural behaviour over time. By enabling real-time or near real-time assessment of structural condition, SHM offers significant advantages over traditional inspection methods [3], including reduced reliance on manual inspections, earlier detection of damage or deterioration, lower long-term maintenance costs, and improved safety for both occupants and emergency responders. These properties make SHM particularly valuable in post-earthquake scenarios [3–8], where rapid condition assessment is critical for informing decisions on continued occupancy and repair prioritisation. Despite this progress, the field faces persistent challenges. Sensor deployment costs remain substantial, data volumes are large and heterogeneous, distinguishing damage-induced signals from environmental and operational variability is technically demanding, and the absence of standardised methodologies and labelled damage datasets continues to constrain the development and validation of automated damage detection algorithms. The framework presented in this paper directly targets these limitations, proposing an integrated architecture that addresses data management, automated damage detection, and probabilistic model updating within a single coherent workflow.

Historically, SHM has been developed and deployed with a primary focus on structural components, targeting load-bearing elements such as beams, columns, foundations, and connections whose failure poses an immediate threat to structural integrity. This orientation is most prominent in bridge and transport infrastructure monitoring [9], where the consequences of undetected deterioration are severe and the economic case for continuous monitoring is well established. Platforms such as MODISAT [10] exemplify this trajectory, integrating satellite-based displacement measurement technologies including GNSS and SAR interferometry with dynamic sensor networks, operational modal analysis [11, 12], and finite element modelling into unified GIS environments designed to assess the condition and safety of bridges. More recently, artificial intelligence-based methods have been incorporated into SHM frameworks to automate damage detection and classification. Deep learning approaches, including convolutional neural networks (CNN) and encoder-decoder architectures, have demonstrated promise for detecting and localising structural damage from both vibration signals and images [13–15]. Vision-based SHM, in particular, has emerged as a complementary approach to vibration-based monitoring: rather than inferring damage indirectly from changes in dynamic properties, camera-based systems are used to identify surface-level damage features such as cracks, spalling, and rebar exposure directly from images [16, 17]. This gives a more complete picture of damage than vibration data alone.

Non-structural elements (NSE)s, however, encompass a broad range of building elements including partition walls, ceilings, cladding, mechanical and electrical systems, and building components that, while not part of the primary load-bearing system, can not only have consequences regarding the safety of the building's occupants, but also account for the majority of repair costs following seismic events and play a decisive role in a building's ability to maintain functionality after an earthquake. Despite this economic and operational significance, NSEs have remained largely outside the scope of conventional SHM frameworks, reflecting both the historical prioritisation of life-safety over functionality and the practical difficulties of monitoring elements that are numerous, heterogeneous, and distributed throughout a building. Kanda et al. [6] demonstrated an important precedent using the "q-NAVI" SHM system for post-earthquake survey of 26 buildings in Japan and deriving fragility functions for NSE. This case notwithstanding, dedicated monitoring of NSEs remains the exception rather than the rule, and the gap between structural and non-structural SHM represents a critical blind spot in current seismic risk management practice.

Recent advances in sensor technology and artificial intelligence (AI) offer new ways to address longstanding challenges in seismic risk assessment [4, 13, 14, 17]. This work exploits these developments to build more comprehensive and responsive monitoring. The approach uses two complementary sources of data: experimental testing campaigns, from which fragility functions for NSEs are derived under controlled conditions, and in-situ sensor measurements, through which those functions are continuously refined to reflect real-world boundary conditions, installation variability, and cumulative damage histories that laboratory-derived models cannot fully capture. Smart sensor networks thus serve a dual purpose, providing response data during seismic events and supplying the observational evidence required for ongoing model updating.

A significant challenge, however, arises in the post-earthquake phase. While accelerometric data may be available throughout a seismic event, visual information about damage state is typically accessible only after the shaking has ceased, making it difficult to reconstruct the progression of damage across discrete damage states during the event itself. This temporal gap limits the precision with which fragility models can be updated for individual damage states. To address this, the proposed framework integrates camera-based sensor networks capable of capturing image data in real-time during seismic events, enabling direct observation of damage accumulation and identification of damage state transitions as they occur rather than retrospectively. The continuous influx of heterogeneous sensor and image data necessitates robust statistical methods for model learning and updating. Bayesian updating is adopted as the core probabilistic framework, through which prior fragility models established from experimental testing and numerical simulation are revised as new observational evidence accumulates, yielding estimates that grow progressively more reliable over the operational lifetime of the monitored components.

2- SONATA

This section will discuss the high-level architecture of the SONATA framework. The architecture represents the preliminary components envisaged for the framework. The system is organised into the following layers:

- Distributed sensor network and edge processing: heterogeneous accelerometer and camera arrays deployed at the monitored buildings

- Real-time data routing and event streaming: a centralised Apache Kafka cluster organising data into topic-partitioned streams. The cluster is a central hub where applications send data that other applications can read. It allows continuous, fault-tolerant data flow in real-time, making it useful for monitoring systems and streaming analytics.
- Data ingestion, archiving, and signal processing
- Probabilistic risk assessment and decision support: Bayesian updating of NSE fragility functions using a combination of vibration-based and vision-based damage observations, a risk engine and an alert service
- Real-time visualisation and stakeholder communication

2-1- DISTRIBUTED SENSOR NETWORK AND EDGE PROCESSING

SONATA architecture uses a distributed sensor network and its associated edge processing infrastructure at its base level. Heterogeneous sensor arrays are deployed at each monitored building, combining accelerometers for measuring floor accelerations and drift demands with camera units for capturing real-time visual information during and after seismic events. The number of camera units required is not fixed on a per-storey basis but is governed by the placement of the accompanying accelerometers and by the location and distribution of the target NSEs, so that only those storeys and zones containing components of interest need to be instrumented while storeys without relevant NSEs may be omitted. The associated level of effort is further reduced because the approach is designed to exploit surveillance infrastructure already present in many buildings rather than to mandate a bespoke installation. Recent studies have demonstrated that consumer-grade indoor surveillance cameras installed for security purposes can be repurposed for the earthquake-induced damage assessment of both structural and non-structural systems [18, 19]. Where such coverage already exists, the incremental effort is largely confined to camera calibration and integration of the resulting data streams rather than to new hardware deployment. Sensors generate continuous data streams whose volume and velocity preclude transmission of raw signals to centralised servers without prior filtering and compression. To address this, each deployment node has an edge compute container that pre-processes data locally. Raw sensor signals pass through validation checks for out-of-range values, timing gaps, and transmission errors before any data leaves the node. A precision sampler applies event-triggered decimation, preserving full-resolution records when accelerations exceed defined thresholds and reverting to reduced rates during quiet periods. An alert detector monitors incoming data in real-time and sends immediate notifications if shaking exceeds a threshold, enabling quick post-event triage without waiting for upstream processing.

Two channels are used to transfer pre-processed data. Continuous waveform data are written locally in MiniSEED format [20], a standardised binary format widely adopted in seismological practice for its compact Steim-2 compression and compatibility with downstream analysis tools. These files are simultaneously forwarded to a SeedLink Ringserver, enabling real-time waveform streaming to any connected client. Event records, alerts, and operational metrics are sent via the message queuing telemetry transport (MQTT) protocol [21], a lightweight publish-subscribe messaging protocol suited for unreliable network conditions. A bridge service collects these messages and feeds them into Apache Kafka [22], a central message system where downstream applications can access them. This dual-protocol approach separates continuous waveform archiving from event-driven messaging, ensuring that the two data flows can be independently scaled, monitored, and replicated without interference. Figure 1 summarises the edge layer data flow.

2-2- REAL-TIME DATA ROUTING AND EVENT STREAMING

Sensor data, alerts, validation errors, and operational metrics generated at the edge layer are ingested into a centralised Apache Kafka cluster, which serves as the backbone of the SONATA data pipeline. Kafka is a distributed event streaming platform designed for high-throughput, fault-tolerant, and ordered message delivery across multiple concurrent consumers [22]. Its architecture is well suited to the demands of a continuous SHM deployment, where multiple sensor streams must be ingested simultaneously, processed by independent services in parallel, and retained for retrospective analysis without any single consumer blocking or affecting another.

Data are organised into a set of dedicated topics, each containing a distinct message type and configured with a retention period appropriate to its operational role. For example, a topic `shm.sensor_readings` carries raw and pre-processed waveform data from all active stations and is retained for a short period of seven days, reflecting the high volume of continuous acquisition. In contrast, alert or anomaly topics carry event-triggered notifications and anomaly detection outputs respectively, both retained for a longer period (90 days) to support post-event review and model validation. Predictions and risk events are retained for a year given its relevance to long-term seismic risk records. All topics are partitioned by station identifier, ensuring that messages from each sensor node are processed in order and that computational load is distributed evenly across the broker cluster. Multiple independent consumer groups read from the same topics in parallel without interference, allowing the waveform archiver, database ingestor, anomaly detection service, and machine learning inference service to each maintain their own read position and processing rate. This decoupled architecture means that a failure or slowdown in any one consumer does not propagate to others, and new analytical services can be introduced by adding a consumer group without modification to the upstream pipeline.

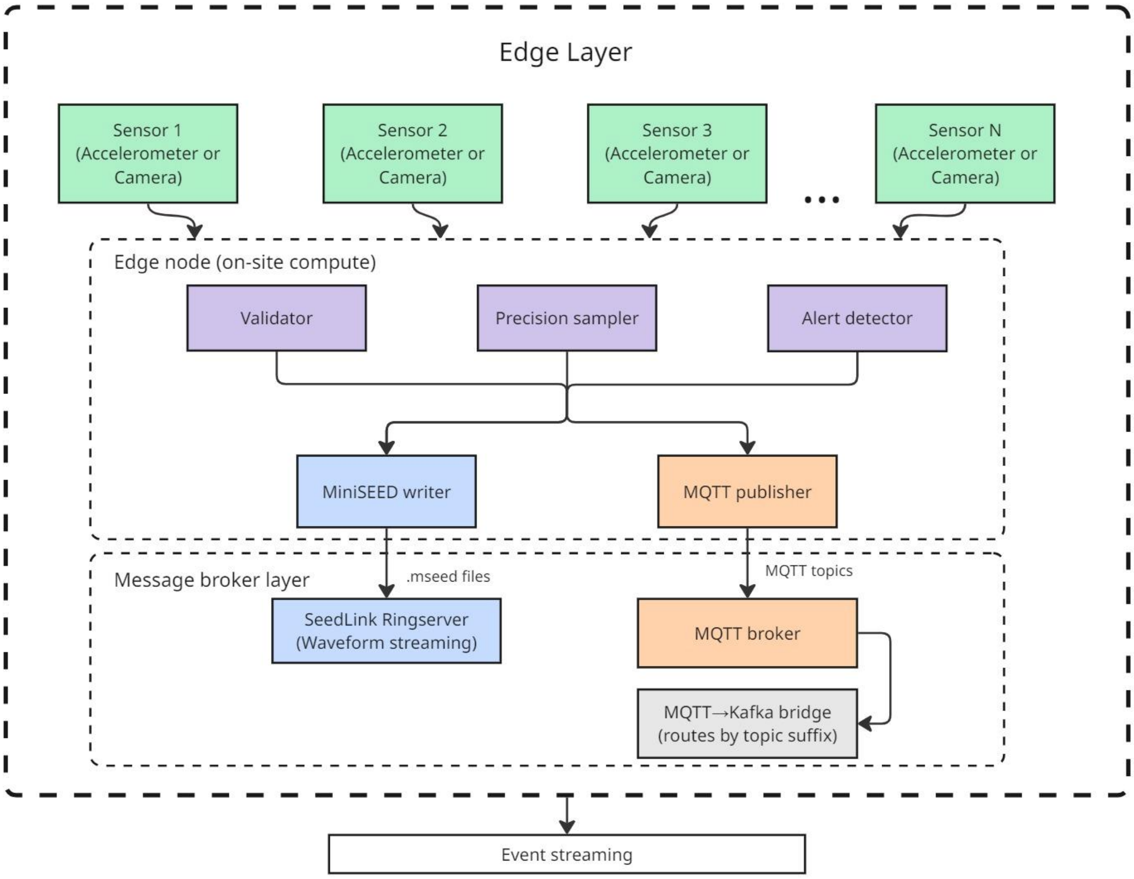

Figure 1- Edge layer data flow

2-3- DATA INGESTION, ARCHIVING, AND SIGNAL PROCESSING

Data arriving at the Kafka cluster are consumed in parallel by four independent services, each responsible for a distinct aspect of ingestion, archiving, and processing as displayed in Figure 2. This separation of concerns ensures that long-term storage, real-time database ingestion, anomaly detection, and machine learning (ML) inference operate at their own rates without contention. Additionally, the system foresees static datasets of fragility functions obtained from past experimental campaigns of NSEs that act as baselines for further fragility function inference.

The first service concerns waveform archiving, where continuous data are written to disk in the Standard for the Exchange of Earthquake Data (SEED) archive structure, commonly referred to as the SDS archive. Files are stored in a hierarchical directory tree organised by year, network code, station code, and channel code, following the convention established by the IRIS data management center [20]. Each file contains 512-sample MiniSEED records encoded with Steim-2 compression [20], providing an efficient binary representation suitable for long-term retention and subsequent reprocessing with standard seismological tools. This archive constitutes the primary research-grade record of all sensor activity and is intended to remain accessible indefinitely, independently of any upstream database or processing service.

Time-series database ingestion service stores decoded sensor readings, alerts, validation errors, and operational metrics into a TimescaleDB instance, which is a time-series extension of PostgreSQL optimised for high-frequency append-heavy workloads. Data are stored in hypertables partitioned into time chunks, with sensor readings partitioned at one-hour intervals and all other tables at one-day intervals. This partitioning strategy allows the database engine to prune irrelevant chunks during queries, significantly reducing scan cost for time-bounded analytical queries. Continuous aggregate views are maintained over the sensor readings hypertable, precomputing summary statistics at configurable time resolutions to support dashboard rendering and threshold-based alerting without repeated full-resolution scans.

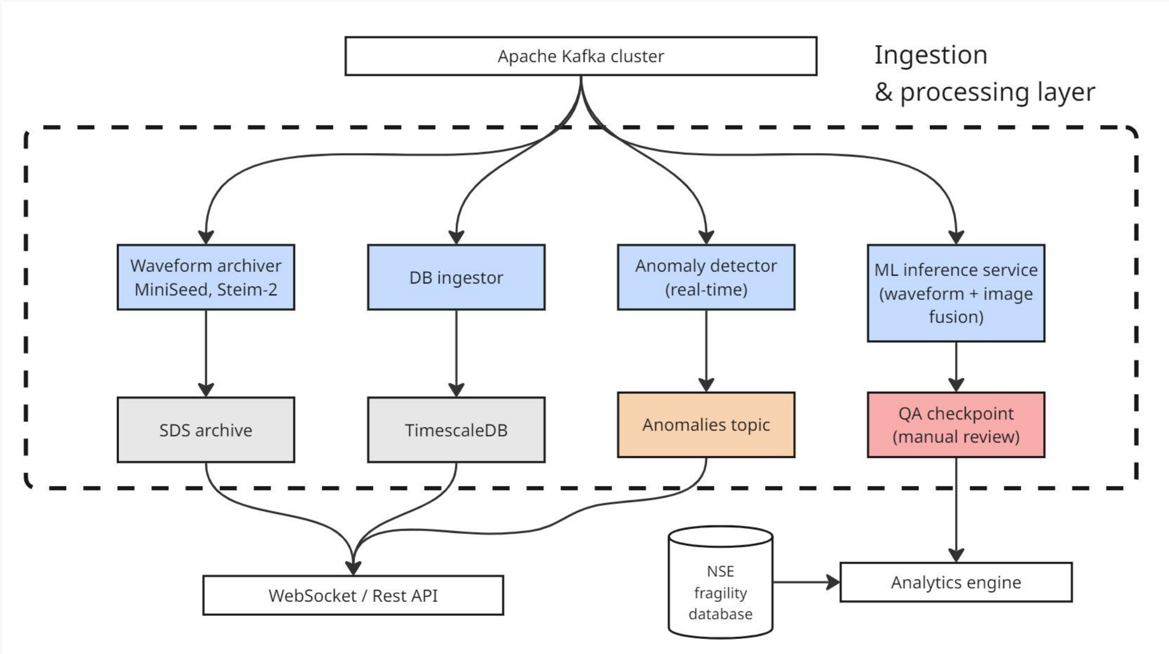

Figure 2- Data ingestion and processing layer of the SONATA framework

The next service is for real-time anomaly detection based on incoming data stream using Kafka Streams. Detected anomalies are published into the relevant topic, where they become available for downstream consumers including the risk engine and the alert routing service. Operating as a stream processor rather than a batch job allows anomaly detection to produce outputs within milliseconds of data arrival, a requirement for any system intended to support near-real-time post-earthquake assessment.

The last service is the machine learning inference service, which acts as an asynchronous inference service that subscribes independently to sensor readings topic and applies trained deep learning models to incoming waveform and image data. The service combines accelerometric measurements with camera-based visual observations of damage state, combining the two complementary information sources that individually provide incomplete characterisation of NSE condition. Inference outputs are published to a Kafka, where a fragility function updater uses them to refine existing fragility models via Bayesian updating or, in cases where no prior fragility models exist, to construct new ones from accumulated observations. The pipeline foresees manual intervention at defined quality assurance (QA) checkpoints, where domain experts review, accept, or reject model outputs before they are used to update fragility parameters, ensuring that automated predictions do not propagate unchecked errors into the risk assessment layer. This service operates asynchronously to avoid introducing latency into the upstream pipeline, and model weights can be updated without service interruption as new labelled data become available from experimental campaigns and field observations.

2-4- PROBABILISTIC RISK ASSESSMENT AND DECISION SUPPORT

The analytics and decision layer translates the continuous stream of sensor measurements, anomaly signals, and ML predictions into probabilistic risk estimates for decision support. It consists of three interacting components: the fragility function updater, the risk engine, and the alert routing service.

The fragility function updater consumes outputs from the machine learning inference service and maintains a continuously revised set of fragility models for each monitored NSE typology. Bayesian inference is adopted as the updating framework, in which a prior probability distribution over fragility parameters, derived from experimental testing campaigns or numerical simulation, is revised as new observational evidence arrives from the deployed sensor network. This formulation is particularly well suited to the SHM context because it accommodates sparse and irregularly timed observations, propagates uncertainty explicitly rather than collapsing it into a point estimate, and allows prior knowledge from laboratory-derived fragility functions to be retained and progressively refined rather than discarded as field data accumulate. Where no prior fragility model exists for a given NSE typology, the updater constructs an initial model from the accumulated field observations alone. All updates are subject to manual quality assurance review before revised parameters are committed to the fragility database, preventing automated inference from propagating errors into the risk assessment layer.

Vision-based sensor networks provide a direct observational channel that accelerometric data alone cannot offer. During and immediately following a seismic event, image data are processed by the machine learning inference service to identify and classify visible damage states on monitored NSEs. The system is designed to recognise a range of damage indicators relevant to the target component typologies, including surface cracking, deformation, and partial or full collapse of partition walls and ceiling systems. Beyond structural damage, the vision component is capable of detecting indirect damage signatures such as water stains, pooling, or wet surface patterns on floors and walls, which serve as observable proxies for piping system damage that would otherwise be invisible to accelerometers. This capacity to infer damage to hidden mechanical and plumbing systems from their surface manifestations substantially extends the diagnostic reach of the monitoring framework beyond what vibration-based methods can achieve.

The risk engine is a downstream service to the fragility functions which combines updated fragility functions with real-time engineering demand parameters derived from sensor measurements and experimental campaigns to compute the probability that each monitored NSE has exceeded each defined damage state. These component-level probabilities are aggregated across the building inventory to produce a system-level risk metric reflecting the expected repair cost, functionality loss, and occupancy disruption associated with the current seismic event.

The risk engine in combination with alert routing service can be used for rapid post-event inspections and prioritisation to allow optimised usage of expert resources and safety tagging of buildings in a regional scale. This allows operators to define graduated response protocols that distinguish between conditions requiring precautionary inspection, restricted occupancy, and immediate evacuation. Apart from inspection prioritisation, targeted notifications may be delivered to relevant stakeholders and is designed to integrate with external emergency management systems. The routing logic implements a quick-pass mode for low-latency notification during active seismic events and a high-accuracy mode for detailed post-event reporting once full inference results are available.

Figure 3 provides a sample of camera-documented damage patterns together with accelerometric data, which are subsequently used for inference of NSE fragility functions. Essentially, the vibration data together with vision-based predictions and prior fragility functions is used to update the fragility functions of the NSE (updated).

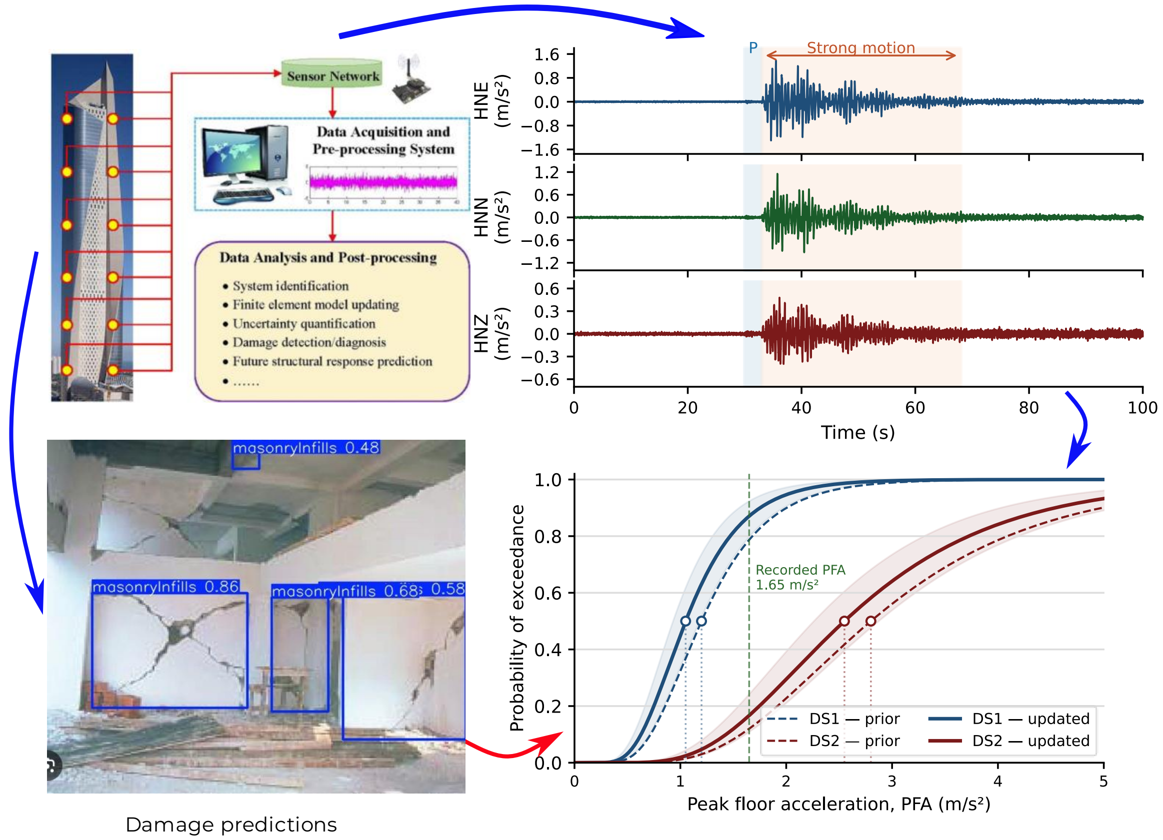

Figure 3- Vision-based damage detection from camera and vibration data and NSE fragility function inference

2-5- REAL-TIME VISUALISATION AND STAKEHOLDER COMMUNICATION

The presentation layer exposes the outputs of the SONATA framework to end users through a web-based dashboard served via a representational state transfer (REST) application programming interface (API) for asynchronous data retrieval and a WebSocket connection for live streaming updates. The dual communication model ensures that high-frequency data such as waveforms and real-time anomaly flags are delivered with minimal latency, while lower-frequency outputs such as risk estimates, prediction trends, and historical records are fetched on demand without burdening the streaming channel.

The dashboard is designed to provide six functional views: live waveform monitoring, anomaly detection status, a building portfolio risk heatmap, prediction trends, historical analysis, and event replay. The event replay capability leverages Kafka's configurable topic retention to re-stream any recorded event through the processing pipeline, allowing operators and researchers to rerun analyses with updated models against historical data. The conceptual dashboard layout is illustrated in Figure 4; full population of the interface with experimental and field data is reserved for future work. The dashboard is organised around eight panel types: live three-component waveform streaming, a normalised anomaly score derived from rolling RMS of accelerometric data, a colour-coded alert log, a building portfolio risk heatmap expressing component-level exceedance probabilities, fragility function visualisation showing prior and Bayesian-updated posterior curves, prediction trends across successive seismic events, a historical record of peak floor accelerations, and an event replay panel that leverages Kafka topic retention to re-stream any recorded event through the processing pipeline. The layout and panel composition shown are indicative and will evolve as the system moves from conceptual design toward experimental deployment.

Figure 4- Schematic layout of the SONATA real-time monitoring dashboard

Finally, beyond live monitoring, the REST API exposes endpoints that allow users to query and visualise the current state of the fragility functions for any monitored NSE typology, enabling direct inspection of how prior laboratory-derived models have been updated by field observations over time. The API additionally supports on-demand inference, accepting user-supplied accelerometric records and images as inputs and returning damage state probability estimates computed by the deployed machine learning models. This capability allows researchers and engineers to interrogate the system outside of live monitoring contexts, for example to assess the predicted performance of a specific component typology under a hypothetical ground motion scenario or to reprocess archived sensor data against a newly updated fragility model.

3- SUMMARY

This paper presented the SONATA framework, a scalable data pipeline architecture for seismic structural health monitoring (SHM) with an explicit focus on non-structural elements, a category of building components that dominates post-earthquake repair costs and governs functional recovery yet has remained largely outside the scope of conventional SHM systems. The framework integrates heterogeneous sensor networks combining accelerometers and camera units, a distributed edge processing layer, high-throughput event streaming, time-series data archiving, and machine learning inference into a single workflow. A key innovation of the proposed architecture is the combination of vibration and vision-based data streams to support real-time tracking of NSE damage state transitions during seismic events, enabling inference of fragility functions from field observations rather than relying exclusively on laboratory-derived models. This addresses a fundamental limitation of existing approaches, in which fragility models cannot account for the variability in real-world installation conditions and cumulative damage histories that characterise components in the field.

The architecture addresses several persistent challenges in operational SHM: lack of standardised data management workflows, delays between shaking and visual damage assessment, absence of automated damage detection for non-structural elements, and disconnect between lab-derived experimental fragility models and field measurements. By including quality assurance checkpoints where domain experts review model outputs, the framework balances automation with the oversight needed for safety-critical applications.

4- ACKNOWLEDGMENT

The work presented here has been developed within the framework of the "Sensor-driven statistical tools to evaluate risks and manage safety in the built environment" (SONATA) project, which is a Starting Grant awarded by the Italian Science Fund (Fondo Italiano per la Scienza) at IUSS Pavia under grant number FIS2023-03215.

5- REFERENCES

[1] Celebi, M.; Seismic Instrumentation of Buildings, Open-file Report, 2000

[2] Kaya, Y.; Ventura, C.; Huffman, S.; Turek, M.; Columbia smart infrastructure monitoring system, Canadian Journal of Civil Engineering, 2018, 579–588

[3] Giordano, P.F.; Iacovino, C.; Quqa, S.; Limongelli, M.P.; The value of seismic structural health monitoring for post-earthquake building evacuation, Bulletin of Earthquake Engineering, 2022, ISSN 1570-761X, https://doi.org/10.1007/s10518-022-01375-2

[4] Dondi, P.; Gullotti, A.; Inchingolo, M.; Senaldi, I.; Casarotti, C.; Lombardi, L.; Piastra, M.; Post-earthquake structural damage detection with tunable semi-synthetic image generation, Engineering Applications of Artificial Intelligence, Volume 147, 2025, 110302, ISSN 09521976, https://doi.org/10.1016/j.engappai.2025.110302

[5] Bragato, P.L.; Boaga, J.; Capotosti, G.; Comelli, P.; Parolai, S.; Rossi, G.; Siracusa, H.; Ziani, P.; Zuliani, D.; Implementing a dense accelerometer network in Veneto (NE Italy): a support for rapid earthquake impact assessment, Bulletin of Earthquake Engineering, Volume 23, 2025, 1859–1884, ISSN 1570-761X, https://doi.org/10.1007/s10518-025-02133-w

[6] Kanda, K.; Nakashima, M.; Suzuki, Y.; Ogasawara, S.; "q-NAVI": A case of market-based implementation of structural health monitoring in Japan, Earthquake Spectra, Volume 37, 2021, 160–179, ISSN 8755-2930, https://doi.org/10.1177/8755293020935884

[7] Fukutomi, Y.; O'Reilly, G.; Nakashima, M.; Kanda, K.; Kusaka, A.; Ogasawara, S.; Three benefits of "Caution (Yellow Tag)" in SHM-driven condition assessment of buildings: Eight years experience with market-based SHM, Earthquake Spectra, Volume 41, 2025, 979–994, ISSN 8755-2930, https://doi.org/10.1177/87552930241296377

[8] Poggi, V.; Scaini, C.; Moratto, L.; Peressi, G.; Comelli, P.; Bragato, P.L.; Parolai, S.; Rapid Damage Scenario Assessment for Earthquake Emergency Management, Seismological Research Letters, Volume 92, 2021, 2513–2530, ISSN 0895-0695, https://doi.org/10.1785/0220200245

[9] Karakostas, C.; Quaranta, G.; Chatzi, E.; Zülfikar, A.C.; Çetindemir, O.; Roeck, G. De; Döhler, M.; Limongelli, M.P.; Lombaert, G.; Apaydın, N.M.; Pakrashi, V.; Papadimitriou, C.; Yeşilyurt, A.; Seismic assessment of bridges through structural health monitoring: a state-of-the-art review, Bulletin of Earthquake Engineering, Volume 22, 2024, 1309–1357, ISSN 1570-761X, https://doi.org/10.1007/s10518-023-01819-3

[10] Dacarro, F.; Bolognini, D.; Curone, D.; Aster, B.; Lorusso, E.; Silvestri, D.; Lunghi, F.; Famà, A.; Eucentre, F.; Ferrata, V.; Ferrata, V.; MODISAT - Satellite technology-derived displacement measurements for smart monitoring of bridges, Earth Resources and Environmental Remote Sensing/GIS Applications XVI, 2025

[11] Quqa, S.; Landi, L.; Diotallevi, P.P.; Seismic structural health monitoring using the modal assurance distribution, Earthquake Engineering & Structural Dynamics, Volume 50, 2021, 2379–2397, ISSN 0098-8847, https://doi.org/10.1002/eqe.3451

[12] Dolce, M.; Nicoletti, M.; Sortis, A. De; Marchesini, S.; Spina, D.; Talanas, F.; Osservatorio sismico delle strutture: the Italian structural seismic monitoring network, Bulletin of Earthquake Engineering, Volume 15, 2017, 621–641, ISSN 1570-761X, https://doi.org/10.1007/s10518-015-9738-x

[13] Malekloo, A.; Ozer, E.; AlHamaydeh, M.; Girolami, M.; Machine learning and structural health monitoring overview with emerging technology and high-dimensional data source highlights, Structural Health Monitoring, Volume 21, 2022, 1906–1955, ISSN 1475-9217, https://doi.org/10.1177/14759217211036880

[14] Reuland, Y.; Martakis, P.; Chatzi, E.; A Comparative Study of Damage-Sensitive Features for Rapid Data-Driven Seismic Structural Health Monitoring, Applied Sciences, Volume 13, 2023, 2708, ISSN 2076-3417, https://doi.org/10.3390/app13042708

[15] Eltouny, K.A.; Liang, X.; Bayesian-optimized unsupervised learning approach for structural damage detection, Computer-Aided Civil and Infrastructure Engineering, Volume 36, 2021, 1249–1269, ISSN 10939687, https://doi.org/10.1111/mice.12680

[16] Zou, D.; Zhang, M.; Bai, Z.; Liu, T.; Zhou, A.; Wang, X.; Cui, W.; Zhang, S.; Multicategory damage detection and safety assessment of post-earthquake reinforced concrete structures using deep learning, Computer-Aided Civil and Infrastructure Engineering, Volume 37, 2022, 1188–1204, ISSN 10939687, https://doi.org/10.1111/mice.12815

[17] Mirzazade, A.; Popescu, C.; Gonzalez-Libreros, J.; Blanksvärd, T.; Täljsten, B.; Sas, G.; Semi-autonomous inspection for concrete structures using digital models and a hybrid approach based on deep learning and photogrammetry, Journal of Civil Structural Health Monitoring, Volume 13, 2023, 1633–1652, ISSN 2190-5452, https://doi.org/10.1007/s13349-023-00680-x

[18] Zhou, J.; Huo, L.; Huang, C.; Yang, Z.; Li, H.; Feasibility Study of Earthquake-Induced Damage Assessment for Structures by Utilizing Images from Surveillance Cameras, Structural Control and Health Monitoring, Volume 2024, 2024, ISSN 1545-2255, https://doi.org/10.1155/2024/4993972

[19] Wen, W.; Zhang, C.; Hu, J.; Guo, J.; Zhai, C.; Zhou, B.; Automatic monitoring method for seismic response of building structures and equipment based on indoor surveillance cameras, Mechanical Systems and Signal Processing, Volume 224, 2025, 112220, ISSN 08883270, https://doi.org/10.1016/j.ymssp.2024.112220

[20] IRIS Data Management Center; SEED Reference Manual: SEED Format Version 2.4, 2012

[21] OASIS; MQTT Version 5.0: OASIS Standard, 2019

[22] Kreps, J., Narkhede, N., and Rao, J.; Kafka: A distributed messaging system for log processing, Proceedings of the 6th International Workshop on Networking Meets Databases (NetDB), 2011