

CONTEXT-AWARE FOOTFALL INTELLIGENCE FOR HISTORIC DISTRICTS VIA DIGITAL TWINS

PER-SENSOR MACHINE LEARNING AND SPATIO-TEMPORAL GRAPH NETWORKS IN THE REGEN PROJECT

ASEL VILLANUEVA-MERINO¹, JOSE LUIS IZKARA², MARIA VIVAR², SILVIA URRÁ-URIARTE¹

¹TECNALIA, BASQUE RESEARCH AND TECHNOLOGY ALLIANCE (BRTA), PARQUE CIENTÍFICO Y TECNOLÓGICO DE BIZKAIA, EDIFICIO 700, 48160 DERIO (BIZKAIA), SPAIN.

²UNIVERSIDAD DE DEUSTO, UNIBERTSITATE ETORBIDEA 24, 48007 BILBAO (BIZKAIA), SPAIN.

KEYWORDS	ABSTRACT
Urban digital twin; Sustainable smart cities; Footfall prediction ; Graph neural networks; Urban planning; Contextual data; Low-cost sensing.	Managing historic districts requires accurate footfall forecasting, yet they often rely on sparse, low-cost sensor networks. We present an operational footfall-intelligence architecture deployed in Laredo's historic center (Spain) using 13 sensors under the EU REGEN project. The system combines a context-enriched tabular engine (Random Forest: RMSE 5.53, R^2 0.87) for interpretable single-step nowcasting with a graph neural network (T-GCN: MAE 6.17 at 1 h) to capture spatial flow redistribution and mitigate noise. Integrated into a 3D urban digital twin, this hybrid coupling of contextual machine learning with spatio-temporal graphs yields actionable intelligence for constrained heritage management.

1. Introduction

Historic city centres are among the most difficult urban environments to manage (Colavitti, 2018). They concentrate heritage value, residential life, commercial activity and, in most cases, tourism pressure, in a compact, often topographically constrained street network. Public space utilization in these districts is dynamic, influenced by a convergence of temporal and physical variables. The number of people present on a given street at a given hour is not a static figure: it is shaped by the calendar (working days, weekends, school terms), by local events and public holidays, by the weather, and by the physical geometry and accessibility of the streets themselves. Such insights provide municipalities with a critical framework for predictive planning. Knowing this number several hours or days in advance — and understanding why it will rise or fall — enables concrete action: planning pedestrian safety, cleaning and waste collection, hospitality and retail staffing, the organisation of cultural events, and the balance between economic vitality and residents' quality of life.

This paper addresses a guiding question: How can the integration of contextual data, complementary AI models and a digital twin turn sparse, noisy sensor counts into actionable footfall intelligence for managing a living historic city centre? We answer it in the concrete setting of the Puebla Vieja of Laredo, the historic-district demonstrator of the REGEN project, an EU-funded urban-regeneration initiative. Designated a Historic-Artistic Site in 1970, La Puebla Vieja is characterised by its grand historic residences and narrow, winding streets. While the district faces long-term population decline and accessibility issues, it experiences intense surges in tourism during the summer months. This creates a critical need to balance the growth of leisure activities, such as bars and restaurants, with the preservation of residential quality of life. In the context of the REGEN project, a local digital twin has been developed integrating physical, semantic and dynamic (sensor) data to support diagnosis, simulation and prediction (Evans et al., 2019; Ketzler et al., 2020).

Three observations motivate our integrated design. First, context dominates in a tourism-influenced town: raw per-sensor counts are noisy and only weakly self-explanatory; models that ignore calendar, weather and street geometry miss the main sources of variance. Second, streets are topographically coupled: Classical tabular models treat each sensor as an independent island, failing to exploit spatial dependencies across feeding alleys or access points along popular routes (e.g., the Camino de Santiago). Conversely, while spatio-temporal Graph Neural Networks (GNNs) capture this spatial coupling, standard architectures designed for dense vehicular highways risk over-smoothing and losing local specificity when applied to data-sparse, event-driven pedestrian graphs. Third, operational needs span multiple time scales: municipal staff need both immediate situational awareness (nowcasting) and multi-hour forecasts ahead of crowd surges, on a platform they can interpret without data-science expertise. No single algorithm addresses all three; the contribution is their combination within one deployable system.

Our central thesis is that a historic district is best served by an integrated footfall-intelligence stack built on a shared context layer, with two complementary prediction engines and a digital twin that makes outputs interpretable for non-technical users. We frame the contribution at three levels:

1. A context-enriched data layer: That fuses sensor time-series with calendar, open weather services and geospatial KPIs, and routes contextual features to each AI engine according to its role — reproducible through an automated retraining pipeline.
2. Two complementary AI engines in one platform: Context-enriched tabular ML (Machine Learning) for capturing repeating patterns and interpretable nowcasting, paired with spatio-temporal GNNs (GCN, T-GCN) over a correlation-based sensor graph to mitigate sensor noise, adapt to sudden flow redistributions, and enable multi-hour network-wide forecasting — whose combined coverage exceeds what either engine provides alone.
3. An operational digital-twin deployment: Consisting of a 3D urban digital twin and a complementary interactive 2D map view that turns complex model outputs into actionable

decisions for municipal staff by linking predictions to terrain slope, accessibility, events and weather. Ultimately, AI in the city only matters if end users can use and understand it.

The remainder of the paper is organised as follows. Section 2 reviews related work. Section 3 describes the study area, the low-cost sensor network and the contextual features. Section 4 presents the integrated system architecture, how context flows into each algorithm, and the two prediction engines. Section 5 reports experiments, contextual correlation analysis and operational integration. Section 6 discusses the value of integration, real municipal benefits and limitations. Section 7 concludes.

2. Related Work

2.1. Footfall forecasting and contextual learning

Footfall and pedestrian-flow prediction support public-space management, crowd safety, and retail analytics in smart-city applications (Zhang et al., 2017). In historic centres, sensor-based visitor monitoring has been used to address overtourism and plan public space (Zubiaga et al., 2019). However, most forecasting studies still focus on large, commuting-oriented traffic networks rather than small heritage districts characterized by sparse counting infrastructure and highly event-driven patterns. Recent applied work in Spanish municipalities—such as short-horizon demand forecasting using machine learning on municipal sensor data for bike-sharing (Angarita-Zapata et al., 2025) and AI-based parking occupancy prediction from urban sensors using explainable recurrent architectures (Domenech et al., 2025)—demonstrates that low-cost sensing pipelines can effectively power operational dashboards. Nonetheless, heritage districts with only a handful of counters remain significantly under-represented in the scientific literature.

Modeling and forecasting pedestrian footfall within these constrained environments has undergone a paradigm shift in recent years, driven by the availability of IoT sensor networks and advancements in applied artificial intelligence. Initially, the literature relied on classical statistical approaches like seasonal autoregressive integrated moving average (SARIMA) models, which continue to serve as strong baselines due to their accuracy in capturing stable, short-term seasonal patterns (Toshev et al., 2025). However, these linear models exhibit severe limitations when dealing with the non-linear, stochastic, and highly volatile nature of pedestrian flows, especially under the influence of tourism dynamics or mass events.

To capture these non-linearities, classical approaches often model each location as an independent time series, leveraging feature-based machine learning with calendar and weather covariates (Murcio & Wang, 2025). Within this tabular paradigm, tree-based ensemble models such as Random Forest (Breiman, 2001) and XGBoost (Chen & Guestrin, 2016) remain highly robust baselines when predictions are driven by complex interactions among temporal, calendar, and environmental features. This exogenous context—comprising weather, holidays, and local events—is the primary driver of footfall variance in tourism-dominated areas, where the explanatory role of calendar and meteorology changes dynamically across seasons and day types (Aultman-Hall et al., 2009). Furthermore, pedestrian movement is inherently constrained by the physical environment; spatial-analysis methods, such as space syntax, have successfully linked street-network structure and urban geometry to observed ground-level footfall density (Özkan Özbek et al., 2022).

To move beyond isolated per-sensor predictions, more recent spatio-temporal deep learning methods have emerged to jointly model space and time through recurrent or convolutional network architectures, as summarized in recent surveys of this fast-moving field (Jin et al., 2024; Yin et al., 2022). While these sophisticated architectures excel at capturing network-wide dynamics in dense, high-frequency transit networks, they often struggle in data-sparse, noisy environments that deploy only a handful of sensors.

Conversely, while context-enriched tabular and statistical approaches are highly robust and interpretable for long-term planning, they suffer from a critical methodological limitation: they treat each sensor as an independent island. By failing to exploit the intrinsic spatial coupling of

the urban network, they cannot adapt when pedestrian flows unexpectedly redistribute across streets (e.g., during street closures, sudden crowd surges, or unpredicted route changes).

2.2. Spatio-Temporal Learning and Graph Neural Networks

To overcome isolated per-sensor predictions and capture global city-wide dynamics, contemporary urban computing relies on spatio-temporal deep learning methods that jointly model spatial and temporal dependencies, as summarized in several comprehensive reviews of this field (Jin et al., 2024; Yin et al., 2022). Since urban streets form an intricate network rather than a regular grid, graph neural networks (GNNs) have become the standard approach for traffic and mobility forecasting. Under this paradigm, the sensing infrastructure is abstracted as a non-circular, non-Euclidean graph, where sensors represent nodes and physical or functional connections constitute the edges. The foundations of graph convolutional networks (GCNs), formally detailed by Kipf and Welling (2017), enable efficient propagation and aggregation of structural information across the network nodes.

In mobility forecasting, representing the urban sensor network as a graph allows for the explicit modeling of spatial dependencies among monitoring points. To capture the evolution of traffic flows, these spatio-temporal architectures couple spatial propagation with temporal dynamics models. Prominent examples of this evolution include diffusion convolutional recurrent neural networks such as DCRNN (Li et al., 2018), pure convolutional frameworks like STGCN (Yu et al., 2018), the use of dilated temporal convolutions in Graph WaveNet (Wu et al., 2019), and the integration of gated recurrent units in T-GCN (Zhao et al., 2020). These methods are particularly relevant when sensor data are sparse or noisy, as coupling with spatially correlated neighbors stabilizes predictions at poorly observed locations or those with incomplete metrics, a typical scenario in low-cost IoT deployments within historic districts.

However, a critical disconnection exists between theoretical developments and their practical application: the vast majority of unified architectures and benchmarks published in the literature—such as the LibCity ecosystem (Wang et al., 2023)—assume the existence of dense and homogeneous urban road networks, validating their results almost exclusively on large-scale vehicular transit systems like METR-LA or PeMS (Li et al., 2018). Small-scale pedestrian sensor graphs with heavily constrained counting infrastructure remain comparatively understudied. Furthermore, because these dynamic networks are strongly conditioned by massive crowd events, weather fluctuations, and tourist dynamics (Dong et al., 2025; Kiros et al., 2026), pure GNNs risk losing local specificity due to the over-smoothing phenomenon if they are not adequately enriched with exogenous context. This limitation justifies the development of our proposed approach, which hybridizes the spatial regularization capabilities of GCNs with the contextual robustness of a tabular univariate engine.

Urban digital twins increasingly serve as the interface through which planning and operations staff access live sensor data, contextual urban information and model outputs (Evans et al., 2019; Ketzler et al., 2020; Villanueva-Merino et al., 2024). Local Digital Twin frameworks that connect sensor streams to municipal decision-making through standardised urban data platforms (Jara et al., 2025) align closely with the REGEN architecture, and low-cost edge sensing for urban flow monitoring on public infrastructure (Glória, 2025) illustrates the feasibility of sparse-deployment regimes comparable to ours. In footfall intelligence, this delivery layer matters as much as offline accuracy: predictions must be linked to geography, accessibility and local context if they are to support municipal action.

Our work contributes an operational case study of an integrated system that combines (a) a shared contextual data layer, (b) complementary tabular and graph spatio-temporal engines, and (c) deployment through a 3D digital twin and a 2D graph dashboard — applied to the Puebla Vieja of Laredo, a small, low-cost, event-driven historic district under-represented in the literature.

3. Study Area and Data

3.1. Study area and low-cost instrumentation

The study area is the Puebla Vieja (old town) of Laredo, Cantabria, Spain (approx. lat 43.3933, lon -3.3973), declared a Historic-Artistic Site in 1970 and the historic-district demonstrator of the REGEN project (Evans et al., 2019; Villanueva-Merino et al., 2024). It is a compact heritage quarter shaped by steep topography, with critical accessibility constraints, a need for building rehabilitation, demographic decline, and an economy concentrated in hospitality and nightlife. Visitor patterns are strongly seasonal and event-driven. Cantabria's historic-enclave tourism and territorial identity — exemplified by heritage-led branding around Gaudí's legacy in neighbouring Comillas (Rodríguez Cánovas & Guzmán Rincón, 2026) — situates Laredo within a broader regional pattern of culture-driven visitor flows. Managing such a living heritage quarter also raises questions of how public space, conservation and community life intersect (Nomeikaite, 2019), and of how citizen participation can support sustainable regeneration in consolidated historic centres (Mejía Gómez, 2025; Yilmaz et al., 2025).

A deliberate design choice is scale and cost: the entire district is covered by thirteen street-level access points — a fraction of the sensor density typical in large-city traffic studies, yet sufficient, when combined with contextual enrichment and graph-based learning, to support operational footfall intelligence. This low-cost instrumentation is a strength of the case study: it demonstrates what a municipality can deploy without a city-wide sensing budget, and it forces methods that work with noisy, sparse data rather than dense, clean streams.

In such a setting, raw footfall counts are difficult to interpret in isolation: a low count on a steep side street may reflect limited accessibility rather than lack of interest, while a peak on a flat main artery may be amplified by event routing. The REGEN local digital twin — a CesiumJS 3D scene with building models, terrain elevation and geolocated sensors, together with a complementary 2D map view — is designed to make this context visible alongside the AI outputs (Figure 4).

3.2. Sensors and target variable

Thirteen access points are deployed across the district. They count pedestrians on most streets and vehicles on a subset (e.g. C. San Francisco, C. del Revellón, C. Santa María). Representative access points include C. El Medio, C. San Martín, C. San Marcial, C. Ruamayor, Cuesta del Infierno, C. Puerta de San Lorenzo, C. Santa María and C. San Francisco (with separate pedestrian and vehicle counters on some streets).

Each sensor records, per time step, the number of people entering in one direction (inputs) and in the opposite direction (outputs); a derived net = inputs – outputs captures directional balance. Data are stored as the access_data measurement in InfluxDB, tagged by centre, access_id and access_name, and aggregated at 15-minute and hourly resolution. The prediction target is the footfall count (inputs and/or outputs) per sensor per time step.

The sensors were placed to monitor accessibility at distinct entry points to the old town, not for joint forecasting. This motivates a data-driven (correlation-based) rather than purely distance-based graph (Section 4.3), and it explains why graph models that share information across correlated streets add value beyond per-sensor tabular models.

3.3. Contextual and exogenous features — the shared foundation

Context is the shared foundation of the integrated system, not an optional add-on. The REGEN platform enriches raw sensor counts with a unified feature set that both AI engines consume (directly or via the graph node signals):

- Temporal: hour, weekday, month, plus cyclical sin/cos encodings of hour, minute and weekday.
- Calendar context: is_holiday (Spanish/local holidays) and is_event (local events), built from curated JSON calendars for Laredo, 2025.

- Weather context (Open-Meteo API, no key required): temperature_2m, apparent_temperature, precipitation, weather_code, wind_speed_10m, cloud_cover. Historical values are fetched from the Open-Meteo archive for model training; forecast values are fetched at inference time so that upcoming conditions inform same-day predictions displayed in the digital twin.
- Urban / geospatial context: sensor coordinates (lat/lon), street type (fclass), is_main_street, distance-to-centre, street geometry (length, maxspeed, slope). The Urban KPI score is a per-sensor composite indicator derived from street-level attributes (security, slope-based accessibility and distance-based metrics) and exposed both as a training feature and as a layer in the digital twin. Slope and accessibility KPIs are particularly important in the Puebla Vieja: pedestrian flows are structurally biased away from steep segments.
- Derived: net = inputs – outputs.

The tabular ML engine (Section 4.2) consumes the full feature set, including weather. The graph-model experiments reported in Section 5 were trained on the consolidated imputed dataset, which includes temporal, calendar and urban-geometry features; weather is part of the live platform and the tabular branch, and its integration into the graph retraining pipeline is ongoing work.

3.4. Data quality and preprocessing

The deployment of this type of sensor network in historic urban settings is relatively uncommon due to the complexities of the built environment. In this specific case, the resulting data series are intermittent and contain gaps caused by connectivity failures. These issues were most pronounced during the summer and early autumn, leading to periods of sensor downtime. Additionally, the missing values were stored as zeros rather than nulls, which made it impossible to distinguish the real absence of pedestrians from a communication error. To fix this, intervals of more than a day containing only zeros were classified as missing data.

Beyond missing data handling, a small number of outliers were removed: instances where the sensors captured thousands of pedestrians within a single 15-minute interval were reported to the manufacturer, and they were confirmed to be errors.

Gaps were addressed through a hierarchical imputation procedure applied in order of priority: (1) same-timestamp values from the previous year, to preserve seasonal patterns for long gaps; (2) linear interpolation between adjacent observations; (3) forward fill; (4) backward fill; (5) sensor-specific hourly average; (6) sensor-specific monthly average; (7) average across correlated neighbouring sensors, to propagate spatial patterns; and, in the rare case none of the above were available, (8) the value was left at zero. Imputed values for July and August were additionally scaled by a factor of 1.2, derived from historical municipal tourism occupancy baselines, to reflect the higher pedestrian activity expected during peak tourist season and missing 15-minute timestamps were reconstructed before imputation to preserve a regular time index.


For the temporal graph model, we additionally restrict the training window to timestamps where all selected sensors are simultaneously active, ensuring a well-formed spatio-temporal tensor. Missing-data handling and imputation are explicit limitations discussed in Section 6.

4. Methodology

4.1. Integrated system architecture and context flow

Figure 1 depicts the integrated REGEN footfall-intelligence stack. Raw counts enter InfluxDB and are enriched with contextual features (Section 3.3) through an automated pipeline. The enriched observations feed two complementary engines (see Table 1) that answer different operational questions:

Figure 1. REGEN footfall-prediction system architecture.



Source(s): Own elaboration, 2026.

Table 1. Roles of the contextual tabular and graph-based prediction engines.

Engine	Primary role	Typical use in Laredo
Contextual tabular ML (Section 4.2)	Interpretable single step nowcasting; encodes repeating seasonal and calendar patterns	"How busy will C. El Medio be next hour given tomorrow is a holiday and rain is forecast?" planning cleaning, retail staffing, long-term pattern review in the digital twin
Spatio-temporal GNNs (Section 4.3)	Cross-street information sharing; multi-hour network-wide forecasts; robustness for noisy or data-poor sensors	"A festival is starting — will the crowd redistribute from the main artery to side streets we did not expect?" proactive crowd and security planning

Source(s): Own elaboration, 2026.

Both engines write to a shared model registry consumed by the real-time backend. The digital twin is the common delivery and interpretation layer: it geolocates each sensor, overlays

predictions, and displays contextual KPIs (slope, accessibility, events) so that municipal staff can relate a forecast to the physical environment.

This architecture is the paper's core methodological contribution: rather than a comparative benchmark between algorithms, it presents a purpose-built system integration where contextual features flow through a shared pipeline, allowing each prediction engine to address the specific operational gaps left open by the other.


4.2. Engine A — Per-sensor tabular ML with context

With thirteen sparsely observed access points and strongly seasonal, event-driven demand, the first operational question is local and explanatory: given tomorrow's calendar, weather forecast and the fixed geometry of a street, how busy will this access point be in the next hour? A per-sensor contextual regressor answers that question without requiring a complete, simultaneous observation tensor across all nodes — which is often unavailable under intermittent connectivity. Training one model per sensor also respects heterogeneous baselines (steep alleys vs. flat main corridors) that a single pooled model would blur.

For each `access_id`, observations are enriched with the shared context layer (Section 3.3): cyclical and calendar encodings (hour, weekday, month, `is_holiday`, `is_event`); Open-Meteo weather at training and inference time (temperature_2m, precipitation, cloud_cover, wind speed, ...); and geospatial attributes (`is_main_street`, street type, distance-to-centre, street length, slope). This set also includes urban KPIs used both for training and in the digital twin: a composite Urban KPI score per sensor, together with its constituent indicators (security, slope-based accessibility and distance-based metrics). Numerical features are standardised (StandardScaler); categorical fields are ordinal-encoded where needed. In exploratory pipelines, TruncatedSVD (10 components) was tested as a compact representation before regression; the deployed nowcasting path retains the full contextual feature set so that contributions remain interpretable in the digital twin. The target is the one-step footfall count (e.g. outputs); data are split 80/20 per sensor with a fixed random seed. Three regressors are evaluated:

- Linear Regression — interpretable linear baseline.
- Random Forest (Breiman, 2001) — `n_estimators = 200`, `max_depth = 20`, `min_samples_split = 2`.
- XGBoost — `learning_rate = 0.3`, `max_depth = 5`, `n_estimators = 200`, `eval_metric = rmse`.

Feature importances from the Random Forest (Figure 2), pooled across the thirteen sensors, show that the composite Urban KPI score and core temporal variables (hour, month) dominate single-step nowcasting. Weather features (temperature, wind, cloud cover) and finer-grained geometry/KPI terms contribute secondary but interpretable effects, whereas calendar flags (holiday, local event) rank lower on average — consistent with event effects being concentrated on few days rather than shaping routine baselines. This reading supports Engine A's role: encoding contextual drivers of repeating footfall patterns and delivering fast, interpretable single-step forecasts inside the digital twin. Its structural limitation remains per-sensor isolation: it cannot borrow strength from correlated neighbours or forecast several hours ahead across the whole network.




4.3. Engine B — Spatio-temporal Graph Neural Networks

Why graphs in Laredo. With only thirteen sensors, some nodes are intermittently noisy or data-poor. GNNs address this by propagating information along edges between functionally coupled streets: a sensor with a gap or spike can be regularised by its correlated neighbours. This is the primary motivation for the graph engine in our low-cost deployment — not benchmark leaderboard performance, but making a sparse network usable for crowd management.

Graph construction. Nodes are sensors. We connect sensors whose footfall time series are correlated: we pivot the series by sensor, compute the Pearson correlation matrix on the target signal, and add an undirected edge between any pair with correlation ≥ 0.6 , yielding the adjacency matrix A . This data-driven graph captures functional coupling — a main street and its feeding alleys, or sensors aligned with popular walking routes — that Euclidean distance would miss.

The resulting adjacency is shown in Figure 3. GNN model (GCN). A two-layer Graph Convolutional Network (GCNConv, PyTorch Geometric) performs node-level regression of footfall. Configuration: hidden dimension ≈ 106 , GELU activation, dropout ≈ 0.49 , Adam (lr ≈ 0.0103 , weight_decay $\approx 3.7 \times 10^{-5}$), ~ 300 epochs, MSE loss. Features and targets are standardised; nodes are split 70/15/15 train/validation/test with a fixed permutation saved to disk. Horizon: one step. Evaluated on held-out nodes, the GCN tests whether spatial message passing improves prediction for sensors that generalise poorly in isolation.

Figure 3. Sensor graph: correlation-based adjacency over the thirteen access points of the Puebla Vieja.



Source(s): Own elaboration, 2026.

T-GCN model. A Temporal GCN couples a custom GraphConvolution layer with a GRU. At each time step the GCN encodes the graph signal; the GRU models the temporal sequence over an input window $T_{seq} = 6$; a linear layer predicts the next step ($pred_steps = 1$, supporting horizons up to 6 h). Cyclical temporal features (hour and weekday sin/cos) are appended to node signals. Training uses Huber loss ($\delta = 1.0$), Adam ($lr \approx 0.0143$), $batch_size = 128$, ~ 283 epochs, per-node max normalisation, and a strictly chronological 70/15/15 split. The T-GCN is the engine for multi-hour, network-wide forecasting — the capability needed when a crowd surge or unexpected redistribution is developing across several streets at once. In deployment, the same correlation-based graph is rendered on an operational 2D dashboard (Figure 8) where staff inspect network-wide forecasts, sensor rankings and per-node prediction panels overlaid on the district map.

Hyperparameter optimisation. GCN and T-GCN hyperparameters were obtained through a systematic Bayesian-style (Optuna-style) search.

4.4. Integration with Local Digital Twin

In heritage districts with sparse sensing, model outputs are difficult to act upon unless they are linked to the physical and functional structure of the place. The Local Digital Twin is therefore treated as an integral component of the architecture: it is the common interaction and interpretation layer through which both prediction engines are exposed to municipal users.

The twin comprises a CesiumJS 3D scene of La Puebla Vieja — building stock, terrain elevation and geolocated sensors (Figure 4) — and a complementary 2D map view for network-level monitoring (Figure 8). Both engines write to a shared registry consumed by a real-time backend; the twin then geolocates each access point and overlays live counts, historical series and short-horizon forecasts together with the contextual drivers used at inference (calendar

type, weather and street geometry). The design choice is deliberate: Engine A supports street-level, context-conditioned nowcasting that staff can inspect sensor by sensor, while Engine B supports network-wide forecasts that must be read against the coupled street graph rather than in isolation.

Figure 4. REGEN Local Digital Twin with contextual ML prediction for C. Ruamayor.



Source(s): Own elaboration, 2026.

Equally important, the twin is not limited to plotting predictions. It combines model outputs with geospatial context layers (including WMS and related municipal overlays) — street typology, walkability/accessibility indicators and event-related information — so that forecasts can be interpreted against urban form and use. Terrain elevation and slope are retained in the 3D scene because orography is a first-order constraint in Laredo: steep segments systematically condition accessibility and baseline footfall, and a predicted peak or trough is only actionable when read against that topography. In methodological terms, the twin closes the pipeline: shared context feeds the engines; the engines return forecasts; and the twin returns those forecasts to geography, elevation and planning layers in a form usable by non-technical operators. Operational behaviour of this delivery layer is reported in Section 5.5.

4.5. Reproducibility and MLOps

The deployment regime — thirteen low-cost sensors, intermittent connectivity and occasional anomalous or missing readings — motivates an operational pipeline that is reproducible, resilient to data-quality issues and inexpensive to refresh. Rather than relying on a single high-capacity model, the system maintains two complementary engines that are comparatively cheap to retrain: contextual per-sensor tabular models for interpretable nowcasting, and graph-based models that propagate information across functionally coupled streets when individual nodes are noisy or data-poor. This dual-engine choice is therefore also an MLOps choice: coverage of distinct operational questions under a constrained sensing budget, with continuous refresh affordable for a small municipality.

Reproducibility is enforced through fixed random seeds and fixed train/validation/test splits persisted to disk. An automated retraining pipeline re-extracts data from InfluxDB, rebuilds the correlation-based sensor graph, retrains both engines on demand and performs data-freshness checks before model promotion. Trained artefacts — weights, scalars, metrics and hyperparameters — are serialised into a model registry consumed by the backend that feeds the Local Digital Twin, keeping live predictions aligned with the most recent validated training run.

5. Experiments and Results

This section reports a single-district operational case study in Laredo's Puebla Vieja. The goal is to characterise how the integrated system performs and what value the combination of engines and context provides — not to crown a single model winner. We evaluate all models on held-out test data using MAE, RMSE and R^2 . Classical models are evaluated per sensor and aggregated globally; the GCN on held-out nodes; the T-GCN on the chronological test window across all nodes, with metrics on de-normalised counts.


5.1. Performance of each engine

Table 2 reports per-engine performance on the consolidated imputed dataset (414,071 records across 13 sensors), using the hyperparameters of Section 4.3. Because the engines address different operational tasks, evaluation protocols are task-appropriate rather than identical: classical models are assessed as single-step, per-sensor regressions with an 80/20 split; the GCN as single-step node-level regression on held-out test nodes; and the T-GCN as single-step prediction on a chronological test window across all nodes, with metrics on de-normalised counts. Lower MAE/RMSE and higher R^2 indicate better fit. Multi-step horizon error for the T-GCN (1–6 h) is shown in Figure 6; agreement between predicted and observed footfall is illustrated in Figure 5.


Table 2. Footfall-prediction performance by engine (test set).

Model	RMSE	R^2	MAE	Horizon	Role in integrated system
Linear Regression	11.00	0.49	4.94	1 step	linear contextual baseline
Random Forest	5.53	0.87	1.72	1 step	contextual nowcasting engine
XGBoost	5.57	0.87	1.90	1 step	contextual nowcasting engine
GNN (GCN)	8.39	0.69	4.81	1 step	spatial sharing for noisy/poor sensors
T-GCN	10.12	0.74	6.17	1 step (chronological)	multi-hour network-wide engine

Source(s): Own elaboration, 2026.

Figure 5. Predicted vs observed footfall on the chronological test window (Random Forest and T-GCN).

Source(s): Own elaboration, 2026.

Figure 6. T-GCN prediction error by forecast horizon (1–6 h).*h-step MAE from autoregressive rollout on held-out test tensor · 1 h: 6.17 · 6 h: 8.55 people/h*

Source(s): Own elaboration, 2026.

Random Forest achieves the strongest single-step contextual nowcasting (RMSE 5.53, R^2 0.87, MAE 1.72), confirming that footfall in Laredo is governed by non-linear interactions among temporal, calendar and geometry features. Linear regression (R^2 0.49) shows that context alone, without non-linear modelling, is insufficient — whereas context with Random Forest is highly effective.

The graph engines report higher absolute error under their stricter evaluation protocols, but they solve different tasks: the GCN generalises to held-out nodes by borrowing strength from correlated streets, while the T-GCN forecasts the entire network chronologically up to six hours

ahead (autoregressive MAE from 6.17 at 1 h to 8.55 at 6 h; Figure 6). Table 2 should therefore be read alongside the capability matrix in Table 1, not as a single ranking of interchangeable models.

5.2. Complementary capabilities in practice

The integrated system is designed around two operational scenarios that illustrate why both engines are needed:

Scenario A — Repeating patterns and planning (tabular engine + digital twin). A municipal planner reviews expected footfall on a typical summer Saturday along flat main corridors, using contextual features (weekend, season, weather forecast, `is_main_street`, slope) rendered in the 3D twin. The tabular engine's strength is encoding why baseline levels differ across day types and making that visible alongside terrain and accessibility KPIs.

Scenario B — Crowd surge and unexpected redistribution (T-GCN + sensor graph). During a local festival, counts rise on C. Santa María while a correlated side street — linked in the sensor graph (Figure 3) but not on the expected parade route — also climbs. A per-sensor model on the main artery alone would miss the coupled surge; the T-GCN, propagating signals across the graph, forecasts the network-wide evolution several hours ahead, supporting proactive steward and security deployment.

Scenario C — Sparse or degraded sensing (GCN). When one access point is intermittently offline or noisy (e.g. connectivity gaps on a steep side street), a per-sensor model degrades locally. The GCN still produces usable forecasts by borrowing strength from correlated neighbours on the sensor graph (Figure 3), keeping network coverage under the same low-cost instrumentation.

5.3. Contextual drivers across day types

To justify context as a first-class contribution — and to show why a single model without contextual enrichment would fail — we stratify the consolidated dataset (414,071 records, November 2024–December 2025) into representative day types and compute Pearson correlations between contextual features and footfall (outputs), and between inter-sensor footfall series (the same structure that defines the sensor graph, but computed per regime). Weather variables are merged from the Open-Meteo archive (hourly, Laredo coordinates); all other features come from the consolidated imputed dataset. The regime-stratified correlations are shown in Figure 7.

Regimes analysed: winter (Dec–Feb) vs. summer (Jun–Aug); weekday vs. weekend; event vs. non-event days (`is_event`).

Feature–footfall correlations shift materially across regimes (Figure 7):

Figure 7. Contextual and inter-sensor correlations by day-type regime.



Slope is consistently negative (winter $r = -0.15$; summer $r = -0.14$; weekend $r = -0.11$), confirming that steeper streets carry systematically lower counts regardless of season — a geometry effect the tabular engine must encode.

Temperature matters more on routine days than on event days: weekday $r = +0.18$, weekend $r = +0.16$, but event-day $r = +0.06$ only — consistent with visitor flows on festival days being calendar-driven rather than weather-driven.

Directional balance (‘net’) is more explanatory on high-variance regimes (winter $r = -0.28$; weekend $r = -0.26$; event $r = -0.26$) than on routine non-event days ($r = -0.10$), signalling stronger ingress/egress asymmetry when tourism peaks.

Summer shows the densest coupling (16 strong edges), reflecting sustained co-movement along main corridors (C. El Medio ↔ C. San Martín, $r = 0.69$). Event days show fewer edges above threshold overall, but the strongest pairs shift: C. Santa María ↔ Cuesta del Infierno reaches $r = 0.83$ on event days (vs. 0.81 in summer), indicating crowd redistribution along routes activated by local festivities — precisely the situation where a network-wide T-GCN adds value over per-sensor nowcasting.

These regime shifts explain why a shared context layer feeding both engines matters: the tabular engine encodes the dominant features for each day type, while the graph engine must track how spatial coupling itself changes when crowds redistribute across the thirteen-sensor network.

5.4. Applications for Urban Planning and Local Economic Development

Beyond offline accuracy, the integrated stack is intended to support municipal decisions in a living historic centre, where economic recovery and infrastructure works must coexist with visitor peaks and accessibility constraints. In line with the complementary roles of the two engines (Table 1; Section 5.2), Engine A informs street-level, context-conditioned planning, while Engine B informs network-wide redistribution under events or disruptions. The digital

twin (Section 5.5) is the shared interface through which these forecasts become readable for non-technical staff.

A first application domain is local economic revitalisation. Vacant commercial premises remain a strategic concern in La Puebla Vieja. Contextual nowcasts from Engine A — combined with seasonal and calendar patterns already encoded in the shared feature layer (feast days, weekends, weather and street geometry) — can help planners relate expected pedestrian intensity to candidate locations and surrounding activity. In practice, this supports evidence-based guidance to entrepreneurs (which streets and time windows are more likely to sustain retail, hospitality or services) and prioritisation of strategic units when municipal layers on vacancy or ownership are available. The claim here is not that a full commercial-matching product is already deployed, but that footfall intelligence linked to context and place is a necessary input for such suitability analysis.

A second domain is infrastructure maintenance under constrained accessibility. Historic-centre works (e.g. undergrounding of utilities or pavement restoration) temporarily remove or degrade street segments and redistribute pedestrian flows. Network-wide forecasts from Engine B, read against the correlation-based sensor graph (Figure 3), help anticipate how activity may shift toward side streets when a main corridor is disrupted or when festivals concentrate demand. This matters in Laredo because functionally related streets need not be adjacent: access points that are geographically apart but both lie on the Camino de Santiago may still co-move, and a correlation-based graph is precisely intended to capture such non-Euclidean coupling when it appears in the data — without treating the Camino itself as a fully validated causal layer in the present study. The same foresight supports sequencing interventions to avoid peak summer and major event windows where possible, and to keep essential accessibility routes usable so that local businesses and residents are not cut off during construction. Again, full what-if closure simulation is a natural extension of the graph engine rather than a completed module of the present deployment; the operational contribution already available is multi-hour, network-level foresight under the same low-cost sensing regime.

Together, these applications illustrate how sparse footfall sensing becomes actionable urban intelligence only when contextual tabular nowcasting and graph-based network forecasts are delivered in a geographic planning frame — the focus of the next subsection.

5.5. Operational integration in the digital twin


Section 4.4 described the twin as the delivery and interpretation layer of the architecture. Here we report how that layer behaves in the deployed REGEN platform and what it enables in practice.

Trained models are served through a real-time backend (Strapi CMS with Socket.IO) that streams live counts and predictions to the CesiumJS twin and to a complementary 2D dashboard (React + Leaflet). In the 3D scene, operators can select an access point, inspect its historical series and overlay short-term contextual forecasts — including predicted entries, exits and peak hours — against terrain, building geometry and weather context (Figure 4). In operational use, this view is most informative for day-type and street-level planning: forecast differences between flat main corridors and steeper side streets become readable when slope, accessibility layers and calendar/weather context are shown on the same scene.

The 2D dashboard supports the complementary network-level task. The correlation-based sensor graph is drawn over the district map, together with global T-GCN summary metrics, ranked sensors by flow, and per-node panels with recent history and multi-hour forecasts (Figure 8). This view is where Engine B's contribution is observable in practice: coupled rises on linked streets, including secondary accesses activated under events, can be inspected as a network evolution rather than as isolated sensor spikes. Together, the twin and dashboard provide evidence that the integrated stack is usable beyond offline metrics: tabular nowcasts and graph forecasts are delivered in the geographic frame municipal staff already work with, linked to WMS/context layers and orography rather than presented as detached numerical outputs.

Broader urban digital-twin programmes are exploring natural-language interfaces over municipal data spaces to further lower the barrier for non-technical staff (Carmona-Martínez et al., 2025); such interfaces are outside the present deployment but consistent with the same delivery rationale.

Figure 8. Operational 2D graph prediction dashboard.



Source(s): Own elaboration, 2026.

6. Discussion

6.1. Why an integrated stack

The combination of contextual tabular ML, graph spatio-temporal learning and digital-twin delivery is deliberate. The tabular engine answers “what will happen on this street, given what we know about the day?” — the question behind cleaning schedules, retail staffing and seasonal planning. The graph engines answer “how will the whole network evolve, including streets we did not expect?” — the question behind crowd safety and event management. Neither subsumes the other; the platform runs both because municipal operations require both time scales and both spatial scopes.

Low cost is both a design constraint and a result. Thirteen sensors are few by traffic-forecasting standards, yet the system delivers operational predictions. Open-Meteo provides weather without proprietary API costs; the retraining pipeline keeps models fresh without a dedicated MLOps team. For other small historic districts, the transferable message is to invest in context and integration before investing in sensor density.

The sensor graph is itself a modelling decision. Correlation-based edges (threshold 0.6) capture functional coupling — shared corridors and event routes — better than distance alone in the Puebla Vieja. This is also a vulnerability: limited data and imputation can introduce spurious edges. Regime-stratified correlation analysis (Section 5.3) partially mitigates this by showing when coupling is stable versus event-specific.

6.2. Operational benefit

The clearest payoff is proactive management: a six-hour-ahead, context-aware, network-wide forecast that anticipates a peak or redistribution can support steward deployment, cleaning schedules, hospitality staffing and resident-friendly event planning. Sections 5.4–5.5 outline how these capabilities map onto local economic and infrastructure decisions and how they are exposed in the digital twin; formal validation with municipal operators remains future work. Strategic frameworks for embedding AI in urban governance (Fernández Añez, 2025) and for linking urban creativity with resilience (Ondiviela García et al., 2025) provide broader justification for deploying footfall intelligence as municipal infrastructure rather than as a standalone analytics experiment.

6.3. Lessons for practitioners

Four guidelines emerge for similar deployments:

- Build the context layer first — calendar, events, weather, geometry — and make it visible in the digital twin; in tourism-driven districts it often explains more variance than algorithm choice alone.
- Combine engines, do not pick one — contextual tabular ML for interpretable nowcasting and repeating patterns; GNNs for noisy sensors, spatial coupling and multi-hour network forecasts.
- Treat low sensor count as a problem GNNs are designed for — message passing compensates for sparse, noisy nodes when streets are functionally coupled.
- Deliver where decisions happen — the 3D twin and 2D map convert model output into municipal action; predictions staff cannot see have no operational value.

6.4. Limitations

Several limitations bound the present case study:

- Thirteen sensors and a single district — results may not transfer to larger networks; validation remains single-city.
- Sparse and missing data — imputation propagates into correlation graphs and temporal series; public-data quality barriers for predictive AI in municipal planning are well documented (Izkara et al., 2025).
- Correlation-based graph — edges reflect statistical association, not causal connectivity, and are sensitive to threshold and regime.
- Weather in graph training — used in the tabular engine and at inference; full integration into graph retraining is ongoing.
- Non-identical evaluation protocols — Table 2 reports per-engine metrics under task-appropriate splits; a unified chronological benchmark is future work.
- Regime-stratified correlations — computed on real data (Section 5.3, Figure 7), but event days cover only 33 calendar days in the study window.

7. Conclusions and Future Work

This paper shows how the integration of contextual data, complementary AI models and a digital twin can turn sparse, noisy sensor counts into actionable footfall intelligence for a living historic city centre. Applied to the Puebla Vieja of Laredo with only thirteen sensors, we designed, developed and deployed an integrated system — rather than an algorithm comparison — and reported it as an operational case study.

Three findings stand out. First, a shared context layer — calendar rules, scheduled events, weather and urban geometry — is the necessary baseline. It drives strong single-step nowcasting (Random Forest: RMSE 5.53 / R^2 0.87 / MAE 1.72), and its explanatory power shifts

across day types in ways that a context-free model would fail to resolve. Second, two complementary AI engines cover operational needs that no single paradigm meets: contextual tabular ensembles provide fast, interpretable forecasts for routine patterns, while spatio-temporal graph neural networks (GCN and T-GCN) mitigate sensor noise, capture spatial coupling across corridors and deliver multi-hour network forecasts (T-GCN autoregressive MAE up to 8.55 at 6 h). Third, delivery through an integrated 3D digital twin and 2D map — where terrain slope, accessibility and events are visible alongside predictions — converts model output into a spatial frame usable by non-technical municipal operators.

The lesson for similar historic centres is therefore not to search for the lowest RMSE in isolation, but to design an integrated, low-cost system in which context, complementary AI and a digital twin work together. The intended municipal value — from event-aware operations to local economic and infrastructure decisions (Sections 5.4–5.5) — follows from that integration; confirming it through structured operator validation is the next organisational step rather than a missing technical component of the present architecture.

Future work should deepen this operational trajectory without abandoning the sparse-sensing premise. A priority is transferability: validating the approach in other historic centres with different size, topography, visitor regimes and contextual data availability, while retaining low-cost instrumentation rather than assuming denser networks as a prerequisite. A second priority is to move from predictive to more prescriptive support — what-if scenarios such as temporary street closures or changes to event routing — building on the graph engine’s network foresight and on the planning use cases outlined in Section 5.4. Methodologically, promising extensions include adaptive or attention-based spatio-temporal graph models, fuller use of weather within graph retraining, longer horizons and richer targets (inputs/outputs/net), and tighter coupling between prediction, simulation and participation tools for historic-centre regeneration (Yilmaz et al., 2025). Across these directions, the aim remains the same: usable footfall intelligence for constrained heritage districts, not leaderboard performance on dense urban benchmarks.

8. Acknowledgements

This work was carried out as part of the REGEN project, which has received funding from the European Union under the Horizon Europe research and innovation programme (Grant Agreement No. 101123325).

Generative AI tools (Cursor) were used only to assist with language editing, consistency checks (figure/table numbering and cross-references) and formatting. All analyses, results, figures and scientific claims were produced and verified by the authors, who take full responsibility for the content.

References

- Colavitti, A. M. (2018). *Urban heritage management: Planning with history*. Springer.
<https://doi.org/10.1007/978-3-319-72338-9>
- Evans, S., Savian, C., Burns, A., & Cooper, C. (2019). *Digital twins for the built environment: An introduction to the opportunities, benefits, challenges and risks* [White paper]. Institution of Engineering and Technology. <https://www.theiet.org/media/8762/digital-twins-for-the-built-environment.pdf>
- Ketzler, B., Naserentin, V., Latino, F., Zangelidis, C., Thuvander, L., & Logg, A. (2020). Digital twins for cities: A state of the art review. *Built Environment*, 46(4), 547–573.
<https://doi.org/10.2148/benv.46.4.547>
- Zhang, J., Zheng, Y., & Qi, D. (2017). Deep spatio-temporal residual networks for citywide crowd flows prediction. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence* (pp. 1655–1661). AAAI Press. <https://doi.org/10.1609/aaai.v31i1.10735>

- Zubiaga, M., Izkara, J. L., Gandini, A., Alonso, I., & Saralegui, U. (2019). Towards smarter management of overtourism in historic centres through visitor-flow monitoring. *Sustainability*, 11(24), Article 7254. <https://doi.org/10.3390/su11247254>
- Angarita-Zapata, J. S., Torrent-Fontbona, F., Machón, B., & Benito Sanchez, F. (2025). Bike-sharing demand forecasting using machine learning. *Street Art & Urban Creativity*, 11(7), 329–342. <https://doi.org/10.62161/sauc.v11.6024>
- Domenech, F., Carot, V., Portales, C., & Coma Tatay, I. (2025). AI-based prediction models for urban parking availability: A case study of Valencia. *Street Art & Urban Creativity*, 11(7), 1–19. <https://doi.org/10.62161/sauc.v11.5990>
- Toshev, V., Karamitov, K., & Petrova-Antonova, D. (2025). Pedestrian movement prediction and pattern analysis. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-G-2025, 1447–1454. <https://doi.org/10.5194/isprs-archives-XLVIII-G-2025-1447-2025>
- Murcio, R., & Wang, Y. (2025). Identify optimal pedestrian flow forecasting methods in Great Britain retail areas: A comparative study of time series forecasting on a footfall dataset. *ISPRS International Journal of Geo-Information*, 14(2), Article 50. <https://doi.org/10.3390/ijgi14020050>
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). Association for Computing Machinery. <https://doi.org/10.1145/2939672.2939785>
- Aultman-Hall, L., Lane, D., & Lambert, R. R. (2009). Assessing impact of weather and season on pedestrian traffic volumes. *Transportation Research Record*, 2140(1), 35–43. <https://doi.org/10.3141/2140-04>
- Özkan Özbek, M., Ertürk, F., Çelebi, T. T., Kınacı, Y., & Yücel Caymaz, G. F. (2022). Interpretations and comparisons of pedestrian movement and land use activities in Kadıköy region using space syntax method. *Street Art & Urban Creativity*, 8(1), 48–61. <https://doi.org/10.25765/sauc.v8i1.585>
- Yin, X., Wu, G., Wei, J., Shen, Y., Qi, H., & Yin, B. (2022). Deep learning on traffic prediction: Methods, analysis and future directions. *IEEE Transactions on Intelligent Transportation Systems*, 23(6), 4927–4943. <https://doi.org/10.1109/TITS.2021.3054840>
- Jin, G., Liang, Y., Fang, Y., Shao, Z., Huang, J., Zhang, J., & Zheng, Y. (2024). Spatio-temporal graph neural networks for predictive learning in urban computing: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 36(10), 5388–5408. <https://doi.org/10.1109/TKDE.2023.3333824>
- Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. In *Proceedings of the 5th International Conference on Learning Representations*.
- Li, Y., Yu, R., Shahabi, C., & Liu, Y. (2018). Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. In *Proceedings of the International Conference on Learning Representations*.
- Yu, B., Yin, H., & Zhu, Z. (2018). Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence* (pp. 3634–3640). International Joint Conferences on Artificial Intelligence. <https://doi.org/10.24963/ijcai.2018/505>
- Wu, Z., Pan, S., Long, G., Jiang, J., & Zhang, C. (2019). Graph WaveNet for deep spatial-temporal graph modeling. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence* (pp. 1907–1913). International Joint Conferences on Artificial Intelligence. <https://doi.org/10.24963/ijcai.2019/264>
- Zhao, L., Song, Y., Zhang, C., Liu, Y., Wang, P., Lin, T., Deng, M., & Li, H. (2020). T-GCN: A temporal graph convolutional network for traffic prediction. *IEEE Transactions on Intelligent*

- Transportation Systems, 21(9), 3848–3858.
<https://doi.org/10.1109/TITS.2019.2935152>
- Wang, J., Jiang, J., Jiang, W., Han, C., & Zhao, W. (2023). Towards efficient and comprehensive urban spatial-temporal prediction: A unified library and performance benchmark. arXiv.
<https://doi.org/10.48550/arXiv.2304.14343>
- Kiros, A., Dorrheimer, J., Teshuva, N., & Cohen, A. (2026). Few-shot transfer learning for cross-city pedestrian level-of-service mapping using spatio-temporal graph models. *Urban Science*, 10(6), Article 334. <https://doi.org/10.3390/urbansci10060334>
- Dong, Y., Chu, T., Zhang, L., Ghaderi, H., & Yang, H. (2025). Pedestrian volume prediction using a diffusion convolutional gated recurrent unit model with dynamic time warping. *Journal of Agricultural, Biological and Environmental Statistics*.
<https://doi.org/10.1007/s13253-025-00696-4>
- Villanueva-Merino, A., Urra-Uriarte, S., Izkara, J. L., Campos-Cordobes, S., Aranguren, A., & Molina-Costa, P. (2024). Leveraging local digital twins for planning age-friendly urban environments. *Cities*, 155, Article 105458. <https://doi.org/10.1016/j.cities.2024.105458>
- Jara, A., Tamarit, P., Ortuño, J., & Faubel, E. (2025). SURPLUS: Strategic urban resource planning using LDTs and unified KQR sustainability metrics. *Street Art & Urban Creativity*, 11(7), 49–72. <https://doi.org/10.62161/sauc.v11.5997>
- Glória, A. (2025). An insight into traffic analysis with computer vision: Leveraging smart infrastructure for urban traffic flow analysis. *Street Art & Urban Creativity*, 11(7), 221–229. <https://doi.org/10.62161/sauc.v11.6019>
- Rodríguez Cánovas, B., & Guzmán Rincón, A. (2026). From Barcelona to Comillas: Gaudí's legacy in the Cantabria brand. *Street Art & Urban Creativity*, 12(2), 1–15.
<https://doi.org/10.62161/sauc.v12.6161>
- Nomeikaite, L. (2019). Street art and heritage conservation: From values to performativity. *Street Art & Urban Creativity*, 5(1), 6–19. <https://doi.org/10.25765/sauc.v5i1.181>
- Yilmaz, D., Urra-Uriarte, S., Guerriero, A., Mack, N., Usobiaga-Ferrer, E., Villanueva-Merino, A., & Kubicki, S. (2025). A path towards citizen-led participation for low-carbon and affordable neighbourhood regeneration using participatory design digital tools. In E. Petrova, M. Srećković, P. Meda, R. K. Soman, D. Hall, J. Beetz, & J. McArthur (Eds.), *Proceedings of the 2025 European Conference on Computing in Construction and 42nd International CIB W78 Conference on Information Technology in Construction*. European Council on Computing in Construction. <https://doi.org/10.35490/ec3.2025.406>
- Mejía Gómez, G. (2025). Strengthening sustainable real estate management through citizen participation. *Street Art & Urban Creativity*, 11(3), 31–45.
<https://doi.org/10.62161/sauc.v11.5736>
- Carmona-Martínez, A., Jara, A., & Molina, G. (2025). SLMs for natural language database interaction with the SENSE Citiverse of Cartagena City. *Street Art & Urban Creativity*, 11(7), 137–157. <https://doi.org/10.62161/sauc.v11.6023>
- Fernández Añez, M. V. (2025). AI cities strategy conceptual model: Application to four cities: New York, Buenos Aires, Barcelona and Madrid. *Street Art & Urban Creativity*, 11(7), 73–98. <https://doi.org/10.62161/sauc.v11.6026>
- Ondiviela García, J. A., Arvilla Gras, E. N., & Moya Tudela, P. (2025). Creativity as a driver to resilience in cities: The more creative the more resilient. *Street Art & Urban Creativity*, 11(5), 67–89. <https://doi.org/10.62161/sauc.v11.5836>
- Izkara, J. L., Usobiaga Ferrer, E., Molina-Costa, P., & Ispizua, B. (2025). Barriers to the effective implementation of AI in predicting public works. *Street Art & Urban Creativity*, 11(7), 201–219. <https://doi.org/10.62161/sauc.v11.5978>