SKET-Monitor: A Knowledge-Driven AI System for Sustainable Environmental and Territorial Monitoring

Veronica Camerada[1,†], Dario Guidotti[1,†], Salvatore Lampreu[1,†], Laura Pandolfo[1,†] and Luca Pulina[4,*,†]

[1]DUMAS, University of Sassari, Via Roma 151, 07100, Sassari, Italy

Abstract

This paper introduces SKET-Monitor, a knowledge-driven decision support initiative aimed at intelligent environmental and territorial monitoring, particularly in fragile coastal areas facing high tourism pressure. Designed to support local administrations in implementing sustainable, data-informed governance strategies, the system combines real-time sensor data, a digital replica of the monitored area, and a hybrid AI approach integrating symbolic reasoning and machine learning. A pilot deployment is currently underway on the Maria Pia beach in Alghero (Italy), where the system enables predictive analysis, policy scenario simulation, and interactive stakeholder engagement through a unified dashboard. By bridging sensor-based monitoring with AI-supported decision-making, SKET-Monitor lays the foundation for scalable, context-sensitive territorial governance tools aligned with the principles of environmental and social sustainability.

Keywords

AI for Environmental Monitoring, ASP Decision Support, Data-Driven Territorial Governance

1. Introduction

In recent years, the European Union has promoted a fundamental shift in the design and implementation of public policies—moving from a top-down, generic, and centralised paradigm to a decentralised, evidence-based, and context-sensitive approach. This transition reflects two strategic needs: to enhance policy effectiveness by tailoring interventions to local environmental and social specificities, and to empower local actors with actionable tools for sustainable governance.

Coastal areas represent a particularly complex domain for this challenge. Characterised by ecological fragility, seasonal overcrowding, and infrastructural limitations, they require policy models that are adaptable, transparent, and rooted in real-time environmental and behavioural data. Yet existing decision support systems are often limited to black-box machine learning models or siloed monitoring platforms, which provide limited interpretability and weak integration of domain-specific constraints.

The SKET-Monitor system (Smart Knowledge-based Environmental and Territorial Monitor) addresses this gap through an integrated framework for local policy modelling and simulation. Its architecture combines four key elements: (i) a real-time sensor network monitoring environmental and behavioural data; (ii) a Digital Twin representing spatial, infrastructural, and regulatory dynamics; (iii) a hybrid reasoning engine blending statistical learning with symbolic AI; and (iv) a visual and interactive interface supporting scenario exploration and decision-making. The system is developed within the context of the SKET-Monitor project, funded by the FAIR – Future AI Research programme, Spoke 9, under NRRP Mission 4, Component 2, Investment 1.3, and supported by the European Union – NextGenerationEU.

A distinctive contribution of SKET-Monitor is its use of Answer Set Programming (ASP) [1] to formalise and simulate public policy choices. ASP is a declarative programming paradigm widely

2nd Workshop on Green-Aware Artificial Intelligence, 28th European Conference on Artificial Intelligence (ECAI 2025), October 25–30, 2025, Bologna, Italy

*Corresponding author.

†These authors contributed equally.

vcamerada@uniss.it (V. Camerada); dguidotti@uniss.it (D. Guidotti); slampreu@uniss.it (S. Lampreu); lpandolfo@uniss.it (L. Pandolfo); lpulina@uniss.it (L. Pulina)

© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR-WS.org/Vol-4165/short17.pdf

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073

published 2026-02-07

used in non-monotonic reasoning and knowledge representation. Recent research has demonstrated its potential in knowledge-intensive applications such as scheduling, configuration, and explainable recommendation (see, e.g., [2, 3, 4, 5, 6]), confirming its suitability for complex decision-making domains. In particular, ASP's support for disjunctive rule heads and weak constraints enables the formalisation of competing policy alternatives and preference-based optimisation. Moreover, ASP benefits from a mature ecosystem of highly efficient solvers—such as CLINGO [7] and DLV [8, 9, 10]—which allow the rapid resolution of large and complex reasoning problems, making the approach practical for decision support scenarios.

SKET-Monitor is not merely a monitoring platform: it enables the construction, validation, and comparison of public policy alternatives under dynamic environmental conditions and institutional constraints. Through its symbolic reasoning core, it supports both tailor-made policies—anchored in local geographies—and evidence-based policies—grounded in structured data and domain knowledge. The system aims to provide local administrations with an accessible and explainable tool for exploring the consequences of their choices.

A first pilot of the system is underway in the coastal zone of Maria Pia beach in Alghero (Italy), an area under intense tourism pressure and ecological vulnerability. The deployment includes environmental sensors, mobility tracking, and a knowledge base of municipal constraints and regulatory thresholds. This real-world case serves as a demonstrator for future extensions to other fragile or high-impact regions.

The remainder of the paper is structured as follows: Section 2 presents the system architecture and its core modules. Section 3 describes the Maria Pia pilot and selected policy simulations, while Section 4 concludes the paper and discusses future work.

2. System Architecture

The architecture of SKET-Monitor is organised as a modular pipeline that transforms heterogeneous sensor data into structured, explainable, and localised policy recommendations. It supports decision-makers in environmentally and socially sensitive areas by combining real-time environmental sensing, semantic spatial modelling, and symbolic AI. Figure 1 illustrates the processing chain.

Environment and Sensors. The monitored territory includes natural and infrastructural elements such as sand dunes, pine forests, footpaths, and tourist access points. In the pilot site of Maria Pia beach in Alghero, Sardinia, a combination of IoT sensors and presence-detection devices provides continuous monitoring. Specifically, devices such as SenseCAP S2103, S1000, and S800 capture temperature, humidity, pressure, wind, solar radiation, and air quality (CO_2, PM), while the Xplore DataNode estimates visitor flows through Wi-Fi probe signals.

Unified Sensor Interface. A middleware layer handles heterogeneous data acquisition and standardisation. It normalises incoming records into a unified JSON format, enriched with temporal and spatial metadata. This abstraction enables seamless integration of additional sensors and promotes syntactic and semantic interoperability across system modules.

UDP Emulator. The UDP Emulator, implemented via a MongoDB instance, acts as a semantic-aware storage and API layer. It stores validated, time-stamped, and geo-referenced records, and makes them accessible via RESTful APIs. Inspired by Urban Data Platform paradigms, this component bridges raw sensor data and higher-level reasoning.

Digital Twin. The Digital Twin is a dynamically updated, semantically structured representation of the monitored environment. It includes models of functional zones (e.g., beach sectors, parking areas), infrastructural nodes (paths, service points), environmental states (pollution, crowding), and simulable

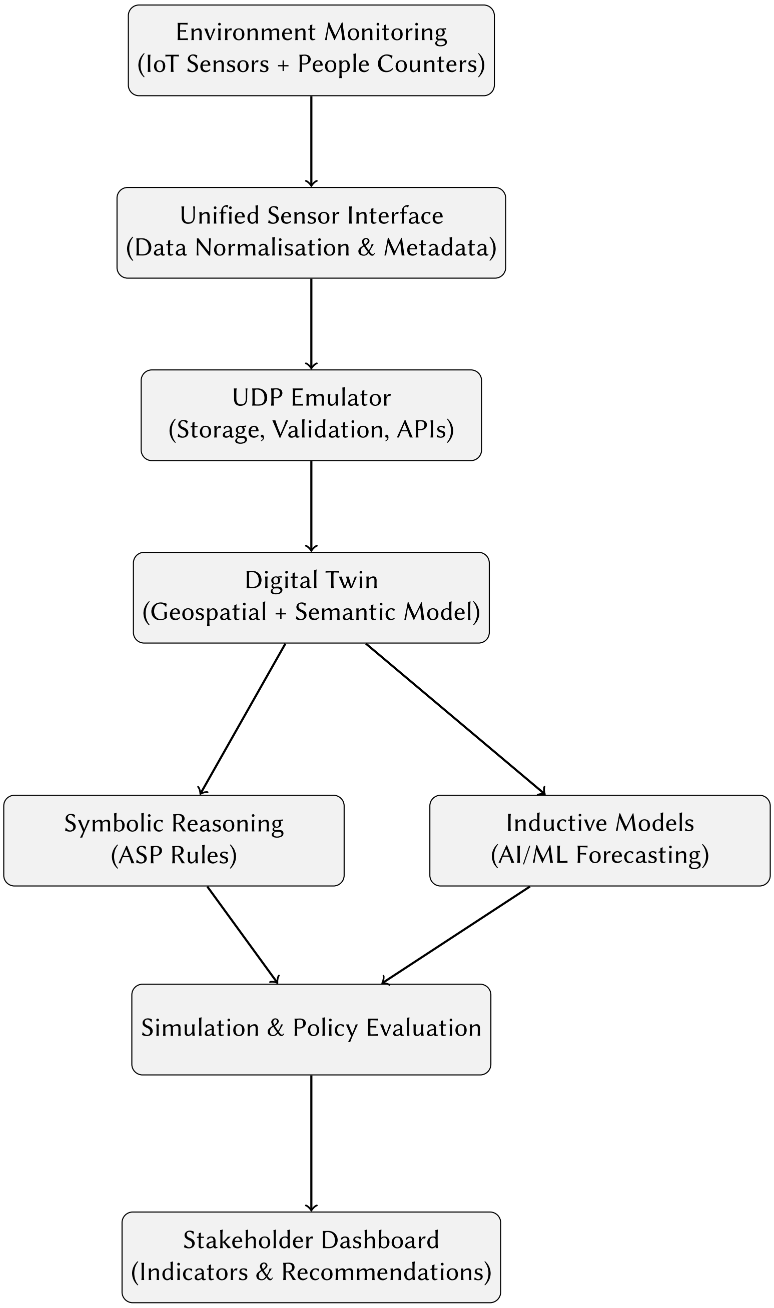

Figure 1: Modular architecture of the SKET-Monitor system.

events (e.g., overcrowding, heatwaves). The twin supports both forward simulations and retrospective queries.

AI and Reasoning Modules. From the Digital Twin, structured state descriptions are passed to two complementary reasoning modules. The *symbolic engine*, based on ASP, parses the state into logical facts and applies policy constraints, thresholds, and preferences. The *inductive engine* includes ML models trained to forecast environmental trends and crowding dynamics. Their outputs feed into a simulation module that estimates the impacts of potential interventions.

Simulation and Dashboard. The simulation module consolidates the outputs of symbolic and inductive reasoning to evaluate the feasibility and impact of policy alternatives. These results are made accessible through a dashboard for stakeholders, showing real-time indicators, justifications for actions, and ranked recommendations. This promotes transparency and supports human-in-the-loop decision-making.

3. Use Case: Coastal Monitoring in Alghero

The pilot site for SKET-Monitor is a critical segment of Maria Pia beach in Alghero, Sardinia—an ecologically sensitive coastal area subject to intense seasonal tourism. Characterised by delicate ecosystems such as sand dunes and a pinewood forest, as well as narrow access points and limited infrastructure, the site suffers from severe overcrowding, increased vehicle emissions, and mobility bottlenecks during the high season. The selected portion of the territory, highlighted in Figure 2, was chosen for its representativeness of the challenges facing coastal municipalities in managing environmental sustainability and tourist pressure. This area serves as the focus of the pilot implementation, allowing the system to capture and simulate critical dynamics through real-world data. These pressures place significant strain on both the environment and local services, underscoring the need for proactive, data-driven territorial management.

Figure 2: Excerpt of the Alghero map with the selected area highlighted for the use case.

To address these challenges, the pilot deployment involves a network of environmental and behavioural sensors installed throughout the selected area. These devices are strategically distributed to capture key indicators such as CO_2 levels, particulate matter (PM2.5 and PM10), temperature, solar radiation, and visitor flows. Their spatial arrangement is shown in Figure 3, which displays the satellite view of the monitored area. Each labelled point (`P#`) corresponds to a monitoring station equipped with a specific combination of sensing technologies—including environmental sensors, crowd detection devices, and weather stations—designed to capture real-time information under diverse conditions.

The data collected from these stations continuously feeds into a dynamically updated Digital Twin, which semantically models the current state of the territory. This structured representation is used by the symbolic reasoning engine, which leverages ASP to evaluate the system's state against defined regulatory thresholds, simulate viable policy responses, and generate explainable, ranked recommendations for decision-makers.

To demonstrate the reasoning process, we consider a hypothetical but realistic scenario derived from synthetic data emulating the summer of 2024. During this simulated period, the system detected consistently high crowding (an average of 330 people/day over 20 consecutive days), elevated CO_2 levels during peak hours, PM concentrations exceeding safe thresholds, and unfavourable meteorological conditions such as low wind speed. Based on these inputs, the reasoning module constructs a set of ASP facts and rules for policy inference, as illustrated in the following excerpt:

```
% Aggregated environmental facts
avg_presence(maria_pia, summer, 330).
days_over_presence(maria_pia, 20).
```

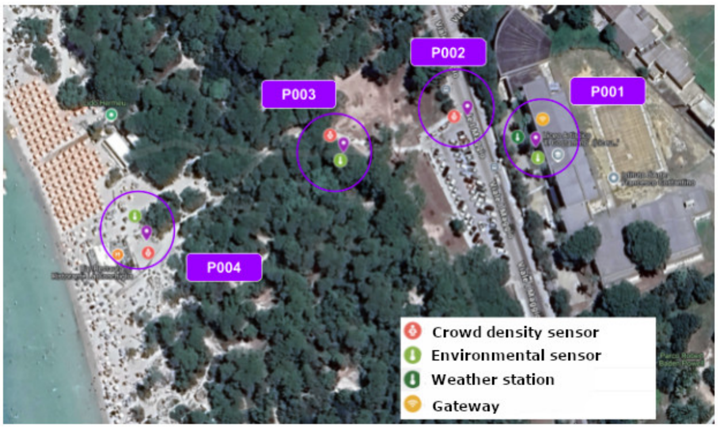

Figure 3: Distribution of monitoring stations over the satellite map of the selected area. Each point labelled as `P#` corresponds to a station, detailing the type of installed sensors (e.g., environmental sensors, crowd detection devices, weather stations).

```
avg_co2(maria_pia, summer, 620).
avg_pm25(maria_pia, summer, 38).
avg_pm10(maria_pia, summer, 58).
avg_wind_speed(maria_pia, summer, 1.2).

% Thresholds
threshold_presence(300).
threshold_days(10).
threshold_co2(600).
threshold_pm25(25).
threshold_pm10(50).
threshold_wind_min(2.0).

% Critical conditions
high_presence(Z) :- avg_presence(Z, summer, P), threshold_presence(T), P > T.
persistent_presence(Z) :- days_over_presence(Z, D), threshold_days(T), D > T.
high_co2(Z) :- avg_co2(Z, summer, C), threshold_co2(T), C > T.
high_pm25(Z) :- avg_pm25(Z, summer, V), threshold_pm25(T), V > T.
high_pm10(Z) :- avg_pm10(Z, summer, V), threshold_pm10(T), V > T.
low_wind(Z) :- avg_wind_speed(Z, summer, W), threshold_wind_min(M), W < M.

% Critical zone definition
critical_zone(Z) :-
high_presence(Z),
persistent_presence(Z),
2 { high_co2(Z); high_pm25(Z); high_pm10(Z); low_wind(Z) }.

% Policy options (exclusive)
activate_policy(Z, shuttle_system);
activate_policy(Z, traffic_restriction);
activate_policy(Z, partial_access_closure) :- critical_zone(Z).

% Preferences via weak constraints
:~ activate_policy(Z, shuttle_system).        [2@1, Z]
:~ activate_policy(Z, traffic_restriction).   [3@1, Z]
```

```
:~ activate_policy(Z, partial_access_closure). [5@1, Z]
```

This encoding formalises the environmental conditions and thresholds, defines when a zone becomes critical, and specifies mutually exclusive policy responses. The use of weak constraints enables a preference ordering: the system favours shuttle deployment, followed by traffic restrictions, and uses partial closures only as a last resort.

Based on these rules, the reasoning engine produces a ranked list of policy alternatives, which are evaluated for feasibility and impact using the Digital Twin and domain-specific constraints. Table 1 summarises the available interventions.

Table 1
Simulated policy alternatives for critical zones

Policy	Description	Rationale
`shuttle_system`	Deployment of shuttle buses from remote parking areas.	Reduces traffic near the beach and lowers CO_2 and PM concentrations.
`traffic_restriction`	Limits vehicle access to the coastal road during peak hours.	Decreases direct emissions and promotes sustainable mobility.
`partial_access_closure`	Caps access to the beach during critical days.	Extreme measure: protects ecosystem and health, but limits tourist accessibility.

The ASP solver ultimately selects the most preferred feasible action, generating justifications alongside each policy recommendation. These explanations, displayed through the system's stakeholder dashboard, help promote transparency and accountability in local decision-making. This symbolic reasoning layer complements the system's inductive components, demonstrating how explainable AI can be effectively used to support sustainable, context-aware governance in complex territorial scenarios.

4. Conclusion and Future Work

SKET-Monitor showcases the potential of combining real-time environmental sensing, semantic modelling, and symbolic reasoning to support evidence-based, explainable, and adaptive policy-making at the local scale. The system was conceived to address the growing complexity of governance in fragile and tourism-intensive territories, where sustainability imperatives meet operational constraints and social expectations. By leveraging ASP as a core component of its policy engine, the system is able to represent legal norms, evaluate alternative interventions, and express contextual trade-offs in a transparent and modular way.

The pilot case of Maria Pia beach in Alghero demonstrates how this hybrid approach—combining live sensor data and structured policy logic—can yield practical, interpretable recommendations for territorial management. The architecture's modularity and domain-agnostic interfaces also favour transferability to other ecological or urban contexts.

Looking ahead, several directions are being pursued to further develop and expand the system's scope. In the framework of the Spoke 2 of ecosystem e.INS (https://eins-spoke2.uniss.it), SKET-Monitor will be integrated with sentiment analysis data, enabling local administrations to couple environmental indicators with qualitative feedback from visitors. This will enrich the system's reasoning capabilities and foster the co-design of services attuned to users' perceptions and needs.

In synergy with the C-WISE project (RAISE ecosystem, https://www.raiseliguria.it), new use cases are being developed to model policies for inclusive urban mobility, with particular attention to the needs of vulnerable citizens. The ASP-based reasoning module is under active development to support personalised recommendations and inclusive accessibility policies, building upon shared infrastructures such as the Urban Data Platform and the Digital Twin.

Additionally, through the regional initiatives InnTerr and T3, financed by the Sardinian Regional Government, the described infrastructure is currently being deployed across over 30 coastal and inland municipalities in Sardinia. These testbeds offer a unique opportunity to assess the system's generalisability, resilience, and policy impact in diverse territorial settings—ranging from touristic shorelines to depopulated rural areas.

Acknowledgments

The contributions of V. Camerada, S. Lampreu, and L. Pulina are funded by the project FAIR – Future AI Research, Spoke 9, under NRRP Mission 4, Component 2, Investment 1.3, funded by the European Union – NextGenerationEU.

The research activities of D. Guidotti and L. Pandolfo are funded by the project e.INS – Ecosystem of Innovation for Next Generation Sardinia (code ECS00000038), financed by the Italian Ministry for Research and Education (MUR) within the framework of NRRP – Mission 4, Component 2, "From research to business," Investment 1.5, "Creation and strengthening of ecosystems of innovation" and construction of "Territorial R&D Leaders" (CUP J83C21000320007).

We also wish to express our sincere gratitude to Prof. Gavino Mariotti, scientific coordinator of the projects InnTerr and T3, for his fundamental support in facilitating the deployment of the described system across a wide range of Sardinian municipalities through direct engagement with local administrations.

Declaration on Generative AI

During the preparation of this work, the author(s) used ChatGPT in order to: Grammar and spelling check. After using this tool, the authors reviewed and edited the content as needed and takes full responsibility for the publication's content.

References

[1] M. Gelfond, V. Lifschitz, Classical negation in logic programs and disjunctive databases, New generation computing 9 (1991) 365–385.

[2] E. Erdem, M. Gelfond, N. Leone, Applications of answer set programming, Ai Magazine 37 (2016) 53–68.

[3] K. Reale, F. Calimeri, N. Leone, F. Ricca, Smart devices and large scale reasoning via asp: tools and applications, in: International Symposium on Practical Aspects of Declarative Languages, Springer, 2022, pp. 154–161.

[4] S. Caruso, G. Galatà, M. Maratea, M. Mochi, I. Porro, Scheduling pre-operative assessment clinic with answer set programming, Journal of Logic and Computation 34 (2024) 465–493.

[5] M. Cardellini, P. De Nardi, C. Dodaro, G. Galatà, A. Giardini, M. Maratea, I. Porro, Solving rehabilitation scheduling problems via a two-phase asp approach, Theory and Practice of Logic Programming 24 (2024) 344–367.

[6] A. Ielo, S. Falco, S. Iiritano, P. Piro, A. Polizzi, F. Ricca, An asp-based approach to water distribution system reconstruction, in: International Conference on Logic Programming and Nonmonotonic Reasoning, Springer, 2024, pp. 140–153.

[7] M. Gebser, R. Kaminski, B. Kaufmann, T. Schaub, Multi-shot asp solving with clingo, Theory and Practice of Logic Programming 19 (2019) 27–82.

[8] N. Leone, G. Pfeifer, W. Faber, T. Eiter, G. Gottlob, S. Perri, F. Scarcello, The dlv system for knowledge representation and reasoning, ACM Transactions on Computational Logic (TOCL) 7 (2006) 499–562.

[9] C. Dodaro, F. Ricca, The external interface for extending wasp, Theory and Practice of Logic Programming 20 (2020) 225–248.

[10] F. Calimeri, D. Fuscà, S. Perri, J. Zangari, I-dlv: the new intelligent grounder of dlv, Intelligenza Artificiale 11 (2017) 5–20.