Proc. of the European Safety and Reliability Conference (ESREL 2026)
Edited by José C. Matos, Paulo B. Lourenço, Michael Beer, Daniel Oliveira,
Edoardo Patelli, Hélder S. Sousa, Rui A. Silva and Monica Santamaria

doi: 10.3850/ESREL2026061419_esrel26-p26400-cd

Integrating MCDA Methods into the Digital Twin's Decision-Making Loop

Kimberley VIARON
Lab-STICC (UMR CNRS 6285), University of South Brittany, Lorient, France.
E-mail: kimberley.viaron@univ-ubs.fr

Philippe RAUFFET
Lab-STICC (UMR CNRS 6285), University of South Brittany, Lorient, France.
E-mail: philippe.rauffet@univ-ubs.fr

Nathalie JULIEN
Lab-STICC (UMR CNRS 6285), University of South Brittany, Lorient, France.
E-mail: nathalie.julien@univ-ubs.fr

Christine PETR
Laboratoire LEGO, University of South Brittany, Vannes, France.
E-mail: christine.petr@univ-ubs.fr

Digital twins have expanded from industrial maintenance to territorial and heritage management. Cultural and archaeological heritage sites, such as the Morbihan Megaliths recently inscribed on the UNESCO World Heritage List, face a structural challenge: dynamically managing visitor flows, logistics, and maintenance while balancing the conflicting objectives of conservation, tourism enhancement, and cost control. Converting predictive scenarios into auditable operational decisions under uncertainty remains an open problem. We propose a closed-loop decision architecture embedding Multi-Criteria Decision Analysis at the core of the simulation loop, conferring a prescriptive, human-supervised, and traceable status on the digital twin. A comparative framework positions major MCDA families with respect to four operational axes: execution latency, robustness to disturbances, configuration effort, and interpretability. A context-driven meta-module then selects and switches the appropriate method according to decision tempo and data quality, supporting operational, tactical, and strategic management horizons.

Keywords: Digital twin, territorial digital twin, multi-criteria decision analysis (MCDA), UNESCO World Heritage, territorial management.

1. Introduction

1.1. *From Industrial Digital Twins to Territorial Heritage Management*

The Digital Twin (DT) concept emerged in the early 2000s within industrial engineering and product lifecycle management, formalized as a virtual counterpart of a physical system continuously updated by sensor and operational data to support closed-loop monitoring, simulation, and decision-making (Fuller et al., 2020; Gabellone, 2022). Early deployments were driven by Industry 4.0, particularly in manufacturing and aerospace.

Building on these industrial foundations, DTs have expanded to smart cities and territorial systems, providing digital representations of city infrastructures to support planning and real-time operations (Lehtola et al., 2022; Ketzler et al., 2020). Well-known implementations include Virtual Singapore and the Angers Loire Métropole territorial twin, though city-scale deployments still face challenges including GIS-BIM interoperability and governance factors (Government Technology Agency of Singapore (GovTech), 2017; Angers Loire Métropole, 2025a,b).

Cultural and archaeological heritage sites represent a natural yet underexplored extension of this trajectory. A UNESCO site digital twin integrates 3D surveys, in-situ sensing (temperature, humidity, vibration), and heterogeneous Earth observation data to track asset condition and detect early deterioration signs (Gabellone, 2022; Ioannidis et al., 2024; Themistocleous et al., 2025). The Morbihan Megaliths, recently inscribed on the UNESCO World Heritage List, comprise more than 550 megalithic monuments across 28 communes and exemplify this challenge concretely. The Alignements de Carnac alone attract approximately 500,000 visitors per year, a figure projected to rise by 20–30% following the 2025 inscription (ADT Morbihan, 2024). Peak daily counts regularly exceed 5,000 persons in July and August, concentrating mechanical and erosion stress on the most fragile monuments precisely when operational resources are most stretched. Meanwhile, tourism across the Morbihan department generates 1.6 billion euros in annual economic consumption, making unrestricted access closure politically and economically untenable.

This creates a structural three-way conflict: structural conservation, tourism enhancement, and cost control are permanently competing objectives that no single-criterion optimisation can resolve. Moreover, the UNESCO inscription imposes explicit accountability requirements, as decisions affecting public access to a World Heritage site must remain auditable and justifiable to protection authorities. These two constraints together, conflicting objectives under uncertainty and mandatory traceability, motivate the integration of Multi-Criteria Decision Analysis (MCDA) at the core of the DT decision loop. To address this gap, this paper proposes embedding MCDA within a closed-loop DT architecture, conferring a prescriptive, adaptive, and traceable decision-making status on the twin. Section 2 surveys the relevant MCDA families, and Section 3 details the resulting integration architecture.

1.2. *MCDA Methods for Decision-Making under Uncertainty*

MCDA methods provide a formal framework to aggregate quantitative and qualitative criteria to solve complex trade-offs (Figueira et al., 2005; Mardani et al., 2015). In digital twin contexts, they are increasingly used to structure decisions over heterogeneous and noisy data streams (Tang et al., 2025). Among the many available approaches for objective and subjective weighting (Paradowski et al., 2021), three families are prevalent in engineering. **AHP** structures the decision problem as a hierarchy and derives criteria weights from pairwise comparisons (Saaty, 1980); its main strength is transparency, though it requires many comparisons and is sensitive to judgement inconsistencies, limiting its suitability for near-real-time updates. **TOPSIS** ranks alternatives by their distances to ideal and nadir solutions (Hwang and Yoon, 1981); it is computationally efficient ($\mathcal{O}(mn)$) but sensitive to extreme values and rank reversal when the alternative set changes (Triantaphyllou, 2000). **PROMETHEE** is an outranking method using preference functions and indifference thresholds (Brans and Vincke, 1985); these thresholds improve robustness to noise at the cost of higher modelling effort and $\mathcal{O}(m^2n)$ complexity. **Fuzzy variants** of AHP and TOPSIS replace crisp values with fuzzy numbers to handle vague expert judgements and incomplete information (Zadeh, 1965; Chang, 1996; Chen, 2000), improving robustness at the cost of additional defuzzification overhead (Mardani et al., 2015; Wieckowski et al., 2025).

These families form a documented MCDA core with complementary trade-offs, ranging from transparent preference elicitation (AHP) to fast operational ranking (TOPSIS) and threshold-based robustness (PROMETHEE). In the context of the Morbihan Megaliths, where decisions range from real-time crowd redirection to annual restoration planning, no single method can address all operational tempos simultaneously. Method selection must therefore be context-driven, which motivates the comparative framework presented in Section 2.

2. A Comparative Analysis of MCDA Methods for Digital Twins

To evaluate the operational integration of MCDA methods within a DT decision loop, we propose a framework bridging cyber-physical constraints and the specific stakeholder requirements of heritage site management, where auditability toward protection authorities is as critical as computational responsiveness.

2.1. *Evaluation Criteria for DT-Integrated MCDA*

The selection of an MCDA method for a DT architecture depends on four dimensions. *Execution latency* refers to the online time required to update rankings after a data refresh; in real-time CPS, response-time budgets can be as tight as a few hundred milliseconds to maintain operational synchronicity (Ortiz-Garcés et al., 2025). *Robustness to disturbances* captures ranking stability under sensor noise, missing data, and scenario evolution, formalised as bounded-disturbance tolerance in CPS (Pajic et al., 2014) and as ranking sensitivity under uncertainty in MCDA (Broekhuizen et al., 2015; Wieckowski et al., 2025). *Configuration effort* reflects the modelling workload, including parameters to tune (weights, thresholds, membership functions); complexity scales with method power, e.g., $\mathcal{O}(n^2)$ judgements for AHP (Saaty, 1980). *Interpretability* denotes the ability to produce stakeholder-facing justifications, and is particularly critical for UNESCO heritage sites where decisions affecting public access must remain defensible to protection authorities.

2.2. *Operational Disturbance Model in DT-driven Ranking*

In DT deployments, the decision matrix is repeatedly updated under imperfect and evolving information. Three recurrent disturbance families affect the decision loop: (1) *measurement noise*, where KPI fluctuations induce frequent rank swaps in sensitive methods; (2) *missingness and partial refresh*, where sensor outages or transmission delays require methods to handle incomplete matrices; and (3) *alternative-set evolution*, where scenarios are added or removed (e.g., monument emergency closure or new visitor circuit), triggering rank reversal in compensatory methods (Triantaphyllou, 2000).

2.3. *Synthesis and Mapping of Method Performance*

The interaction between DT disturbances (Sec. 2.2) and evaluation criteria (Sec. 2.1) creates a specific selection space for MCDA families.

As shown in Table 1, measurement noise favors outranking methods (e.g., PROMETHEE) that utilize thresholds to filter insignificant variations, albeit at the cost of higher modeling effort. Conversely, high-frequency updates favor compensatory methods (e.g., TOPSIS) due to their low computational complexity ($\mathcal{O}(mn)$), even if they remain more sensitive to outliers. AHP provides maximum auditability for preference structuring, while fuzzy extensions are required when uncertainty representation is a core requirement, despite the increased defuzzification overhead. Table 1 synthesizes these operational trade-offs.

Table 1. Operational trade-offs of MCDA families in a Digital Twin context.

Latency	Robustness	Config. effort	Interpretability
AHP : low latency*; high elicitation	Limited; sensitive to inconsistency	High (hierarchy and pairwise elicitation)	High (traceable weights and hierarchy)
TOPSIS: very low; $\mathcal{O}(mn)$	Moderate; outlier-sensitive	Low (weights and normalization)	Moderate (abstract geometric index)
PROMETHEE: low–med; $\mathcal{O}(m^2n)$	High (noise filtering via thresholds)	Med.–high (preference functions/thresholds)	Moderate (flow-based; GAIA maps)
Fuzzy MCDA: low–moderate	Very high (explicit uncertainty modeling)	High (membership functions/elicitation)	Low–moderate (requires defuzzification)

Note: m = alternatives, n = criteria. *For AHP, latency refers to the computational update; elicitation is accounted for in configuration effort.

3. Methodology for Integrating MCDA into the Digital Twin's Decision-Making Architecture

We propose a closed-loop decision architecture structured around four phases that directly implement the Observe–Predict–Decide–Act (OPDA) cycle. Phase 1 (Observe & Predict) continuously collects site data and generates predictive scenarios; Phase 2 (Decide) embeds MCDA to aggregate multi-criteria preferences and rank candidate scenarios; Phase 3 (Adapt) updates scenarios or preferences when the current recommendation fails to meet operational requirements; and Phase 4 (Act & Audit) implements the selected action and logs the full decision state. Critically, Phase 4 is not a terminal state: it triggers an outer feedback loop that restarts Phase 1 at the next operational cycle, making the DT genuinely self-improving

over time.
Figure 1 presents the closed-loop implementation of the Observe–Predict–Decide–Act cycle. Phase colours identify each stage: blue for Phase 1 (Observe & Predict), purple for Phase 2 (Decide via MCDA), orange for Phase 3 (Adapt), and green for Phase 4 (Output & Audit). The red node governs the Phase 2 $\rightarrow$ Phase 3 branching logic. Simulated scenarios are evaluated by MCDA under current preferences and iteratively refined until convergence or performance targets are met, yielding an auditable decision. The dashed arrow from the Recommendation output back to Phase 1 represents the outer feedback loop, restarting the cycle at the next operational horizon.

3.1. *Macroscopic Closed-Loop Workflow*

At decision time t, the DT evaluates a set of candidate scenarios $\mathcal{S}(t) = \{S_1, \ldots, S_m\}$. Each scenario S_i is described by a vector of criteria (KPIs) $k_i = [k_{i1}, \ldots, k_{in}]$, forming the decision matrix $K(t) \in \mathbb{R}^{m \times n}$.
Phase 1 (Observe & Predict). Multi-source observations (IoT sensors monitoring weather conditions, visitor flows, monument vibration and humidity; ticketing and calendar data) are assimilated into the DT state. Each S_i is simulated and KPI estimates are produced.
Phase 2 (Decide via MCDA). $K(t)$ is normalised and aggregated using current preferences (weights $w(t)$, and optionally thresholds or fuzzy parameters), yielding $\mathrm{Rank}(S_i)$. A stopping condition (scenario convergence, weight stability, or performance threshold met) determines whether to proceed to Phase 4 or trigger Phase 3. Ranking stability is assessed by monitoring top-k swap rates under small KPI perturbations; a maximum iteration bound $N_{\max}$ prevents non-convergence, triggering a conservative fallback (partial access restriction or human escalation) if stability is not reached within budget.
Phase 3 (Adapt). Phase 3 is triggered when one of three conditions holds: (i) no scenario meets the minimum conservation performance threshold; (ii) the ranking is unstable under minor KPI perturbations, signalling low decision confidence; or (iii) a contextual event (monument emergency

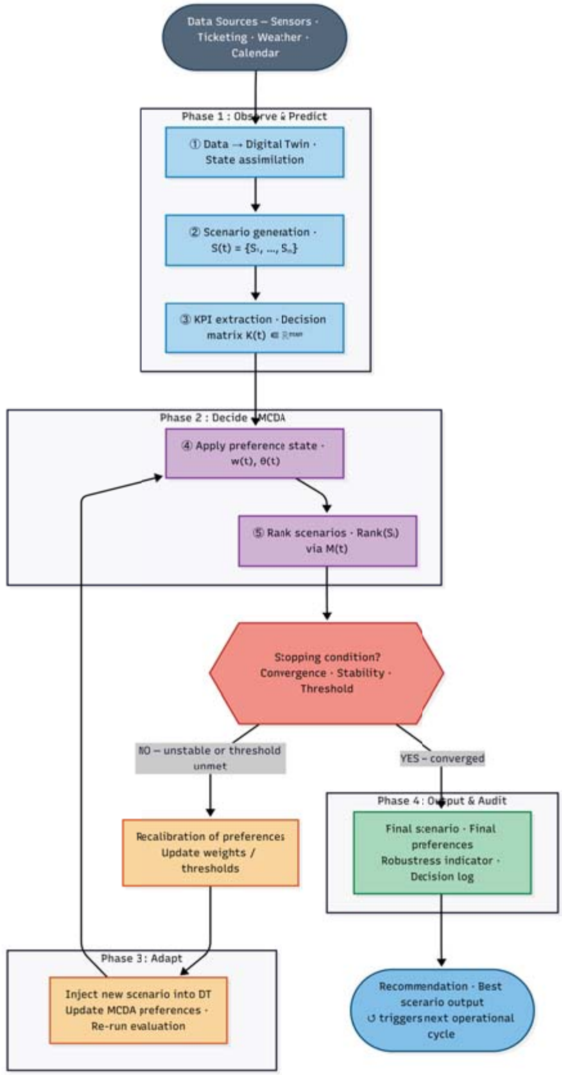

Fig. 1. Closed-loop MCDA-enabled decision pipeline structured around the Observe–Predict–Decide–Act cycle. Blue: Phase 1 (Observe & Predict); Purple: Phase 2 (Decide via MCDA); Orange: Phase 3 (Adapt); Green: Phase 4 (Output & Audit). The red node governs the Phase 2 $\rightarrow$ Phase 3 branching logic. The dashed arrow represents the outer feedback loop restarting Phase 1 at the next operational cycle.

closure or sudden weather alert) alters the feasible scenario set. The system then injects new scenarios and re-simulates (Phase 3a), or recalibrates preference weights $w(t)$ through human-validated workshops (Phase 3b), before re-entering Phase 2. **Phase 4 (Output & Audit).** Once stopping conditions are met, the DT exports the final recommended scenario, the final preference state, a robustness indicator, and a complete decision log including $M(t)$, $\theta(t)$, and $\mathrm{Rank}(S_i)$. Decisions are then implemented in the field (access control adjustments, maintenance alerts, dynamic signage). Phase 4 is not a terminal state: it triggers the outer feedback loop, restarting Phase 1 at the next operational cycle, which may be time-driven (hourly, daily, or seasonal) or event-driven.

This macroscopic loop provides a traceable and continuously adaptive decision process, from multi-source monitoring to prescriptive recommendations. In particular, Phase 2 is critical because it determines both the decision quality and the operational feasibility under time constraints. The next subsection details the supervisory meta-module used to select (and, when needed, switch) the MCDA method according to context, uncertainty, and computation budget.

3.2. *Meta-module for context-driven MCDA selection (Phase 2 zoom)*

Figure 2 details Phase 2 as a supervisory meta-module that selects (and switches) the MCDA strategy under operational constraints, operationalising the trade-offs of Table 1 across decision tempos and data-quality conditions. At time t, the DT provides $\mathcal{S}(t)$ and $K(t) \in \mathbb{R}^{m\times n}$; the meta-module then performs (2a) KPI characterisation, (2b) context-grid method selection, and (2c) MCDA ranking, while monitoring computation time and ranking stability.

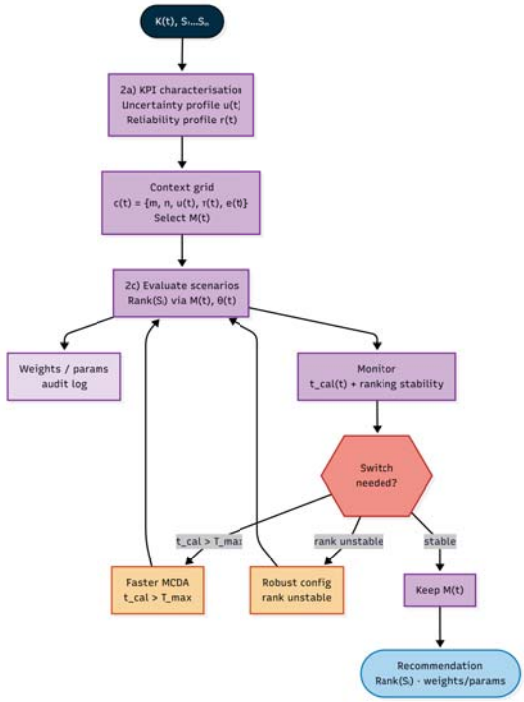

Fig. 2. Phase 2 zoom: context-driven selection and switching of the MCDA strategy under uncertainty and time-budget constraints.

Step 2a: KPI characterisation under uncertainty. Because DT inputs are heterogeneous and repeatedly refreshed, the meta-module first characterises data quality and uncertainty at the criterion level. In practice, each criterion j can be flagged as (i) crisp, (ii) interval-valued, or (iii) fuzzy, depending on measurement noise, aggregation variability, or partial refresh. This characterisation produces a compact uncertainty profile $u(t)$ and a missingness/reliability profile $r(t)$, which will influence the choice of method and parameters (e.g., indifference thresholds for outranking methods, or membership function bounds for fuzzy variants).

Step 2b: Context-grid method selection. The meta-module maps the current decision context to an MCDA family using a lightweight context grid. The context descriptor can be expressed as $c(t) = \{m, n, u(t), \tau(t), e(t)\}$, where m and n denote the number of alternatives and criteria, $\tau(t)$ the decision tempo (i.e., the allowed response time), and $e(t)$ an explainability requirement. The selection rule is operational: AHP is preferred when auditability dominates and $n \leq 10$ (e.g., annual maintenance workshops with the Association Paysages de Mégalithes du Morbihan); TOPSIS when $\tau(t) < 300\,\mathrm{ms}$ (e.g., real-time crowd redirection at the Alignements de Carnac); PROMETHEE when robustness under noisy sensors is required (e.g., vibration monitoring during high-attendance periods); and fuzzy variants when uncertainty representation is mandatory. This context-driven selection operationalises the qualitative trade-offs reported in Table 1.

Step 2c: Evaluation and ranking. Accordingly, scenario evaluation is performed with the selected method $\mathcal{M}(t)$ and its parameter set $\theta(t)$ (weights, thresholds, membership functions), ensuring that the ranking is consistent with the operational requirements identified at time t. The module computes a scenario ranking $\mathrm{Rank}(S_i)$ and returns (i) a recommendation (best scenario or top-k) and (ii) the preference state (weights/parameters) to support traceability and later auditing.

Monitoring and switching logic. A key contribution of the meta-module is that it continuously monitors computational feasibility and ranking stability. First, it measures the online computation time $t_{\mathrm{cal}}(t)$ and enforces an operational budget $T_{\max}(\tau(t))$. If $t_{\mathrm{cal}}(t) > T_{\max}$, the module triggers a switch toward a faster MCDA configuration (e.g., AHP $\rightarrow$ TOPSIS) to preserve responsiveness. Second, it evaluates ranking stability under realistic DT disturbances (measurement noise, missingness, and alternative-set evolution). If the ranking becomes unstable (e.g., frequent top-1 swaps under small KPI perturbations), the module switches to a more robust configuration (e.g., thresholds/outranking or fuzzy modelling), trading additional configuration effort for improved stability.

Traceability and integration with the closed loop. For each decision step, the DT logs $\mathcal{M}(t)$, $\theta(t)$, $t_{\mathrm{cal}}(t)$, and the resulting ranking to ensure reproducibility and stakeholder-facing justification. If the stopping condition is not met, the current best scenario and associated diagnostics are forwarded to Phase 3 (Adaptation), where new scenarios can be injected and preferences updated before re-entering Phase 2, preserving the closed-loop behaviour described in Fig. 1.

3.3. *Practical Implications for Heritage Site Management*

The proposed architecture directly supports the management of the Morbihan Megaliths across three nested decision horizons.

Real-time operational control (minutes to hours). During high-attendance periods, when visitor counts at the Alignements de Carnac exceed an operational threshold (e.g., 3,000 persons in the restricted zone), the MCDA module applies TOPSIS within the 300 ms latency budget to rank crowd management scenarios (access path closure, dynamic signage activation, guide deployment) and recommend the best option to the duty manager with full criteria weights displayed.

Tactical planning (days to weeks). At the start of each week, the DT integrates weather forecasts, booked group visits, and maintenance team availability. PROMETHEE with preference thresholds ranks tactical scenarios (preventive cleaning of engraved surfaces, staff redeployment across sites) against conservation, visitor safety, and cost criteria, filtering noise from imprecise weather predictions.

Strategic scheduling (seasonal and annual). Annual restoration priorities are determined through a full AHP elicitation workshop with the Association Paysages de Mégalithes du Morbihan. The DT provides the decision matrix (KPIs: structural degradation index, restoration cost, visitor impact, ecological sensitivity) and produces a fully auditable ranked list reviewable by heritage protection authorities.

Across all three horizons, every decision is logged with its complete preference state, ensuring that no operational choice affecting a UNESCO World

Heritage site is made without a traceable, explainable justification.

3.4. *Sensitivity and Latency Considerations*

The practical deployment of MCDA within a DT loop raises two evaluation concerns that future empirical work must address. First, *sensitivity to parameter tuning*: criteria weights $w(t)$ and preference thresholds $\theta(t)$ are set through expert elicitation or historical calibration, yet small perturbations in these parameters can alter the top-ranked scenario. The meta-module mitigates this by monitoring top-k swap rates (Sec. 3.1) and triggering recalibration when instability is detected, but systematic sensitivity analysis across representative weight ranges remains necessary to quantify decision robustness for each MCDA family. Second, *latency distribution across scenarios*: while theoretical complexity bounds ($\mathcal{O}(mn)$ for TOPSIS, $\mathcal{O}(m^2n)$ for PROMETHEE) provide worst-case guidance, actual computation times depend on matrix density, missingness patterns, and hardware. The evaluation approach envisaged for the Morbihan deployment will measure empirical latency distributions under three reference scenarios (peak-season crowd management with $m \approx 15$ alternatives, weekly tactical planning with $m \approx 8$, and annual strategic scheduling with $m \approx 5$) to validate that the 300 ms operational budget is met in practice.

4. Conclusion and perspectives

In closed-loop Digital Twins (DTs), decision-making is not a one-shot ranking problem: it results from an iterative process in which candidate scenarios $\mathcal{S}(t)$ are progressively updated, the decision matrix $K(t) \in \mathbb{R}^{m\times n}$ is refreshed as new information becomes available, and multi-criteria trade-offs are continuously arbitrated under uncertainty and explicit time constraints. In such settings, relying on a single and static MCDA method is rarely optimal, since operational requirements evolve with decision tempo, data quality, institutional accountability, and traceability expectations.

This paper made three methodological contributions. (C1) We introduced a DT-oriented comparison framework structured around four operational axes, namely execution latency, robustness to disturbances, configuration effort, and interpretability, in order to position major MCDA families, including AHP, TOPSIS, PROMETHEE, and fuzzy variants, with respect to Digital Twin constraints. (C2) Building on this framework, we proposed a complete closed-loop decision architecture embedding MCDA at the *Decide* stage of the Observe–Predict–Decide–Act cycle. This architecture includes an explicit Phase 4 $\to$ Phase 1 outer feedback loop, making the DT genuinely self-improving across operational cycles, and a formally specified Phase 2 $\to$ Phase 3 branching condition that triggers adaptation when conservation thresholds, ranking stability, or contextual requirements are not met, thereby ensuring prescriptive, reproducible, and explicitly traceable decision outputs. (C3) Finally, we detailed a context-driven meta-module which, at each decision time t, characterises KPI uncertainty using appropriate representations, selects a context-appropriate strategy $\mathcal{M}(t)$ with parameters $\theta(t)$, and monitors computational feasibility and ranking stability so as to enable adaptive method switching when operational constraints are violated.

From a practical standpoint, the proposed architecture offers heritage site managers a concrete operational benefit: every decision affecting visitor access, maintenance scheduling, or resource allocation is grounded in explicit, auditable criteria rather than informal judgment. For a UNESCO World Heritage site such as the Morbihan Megaliths, where preservation obligations coexist with significant tourism pressure, this traceability is not merely a technical feature but a governance requirement. The framework thus bridges the gap between real-time sensor data and defensible management decisions, providing protection authorities, local stakeholders, and site operators with a shared, reproducible decision record across all operational horizons.

Future work will instantiate this framework on empirical data from the Morbihan Megaliths site, with the explicit objective of supporting public authorities and local stakeholders in regulating tourism pressure under preservation, safety, and quality-of-experience constraints. The proposed MCDA–DT integration will be evaluated on representative territorial scenarios to assess its ability to structure, justify, and adapt collective decision-making across operational, tactical, and strategic horizons, without assuming a fully autonomous or fully deployed system. In this perspective, the framework is intended as a transferable decision-support layer for territorial governance, rather than as a site-specific optimisation solution.

Acknowledgement

The authors thank the "Association Paysages de Mégalithes du Morbihan" and the "CRT Bretagne" for valuable discussions and for their support through information sharing.

References

ADT Morbihan (2024). Rapport d'activité 2024. Technical report, Agence de Développement Touristique du Morbihan.

Angers Loire Métropole (2025a). Anticiper les crues grâce au jumeau numérique. Accessed 2025-12-22.

Angers Loire Métropole (2025b). Territoire intelligent, le jumeau numérique, un outil d'aide à la décision. Accessed 2025-12-22.

Brans, J.-P. and P. Vincke (1985). A preference ranking organisation method: (the promethee method for multiple criteria decision-making). *Management Science 31*(6), 647–656.

Broekhuizen, H. et al. (2015). A review and classification of approaches for dealing with uncertainty in multi-criteria decision analysis for healthcare decisions. *PharmacoEconomics 33*, 445–455.

Chang, D.-Y. (1996). Applications of the extent analysis method on fuzzy ahp. *European Journal of Operational Research 95*(3), 649–655.

Chen, C.-T. (2000). Extensions of the topsis for group decision-making under fuzzy environment. *Fuzzy Sets and Systems 114*(1), 1–9.

Figueira, J., S. Greco, and M. Ehrgott (Eds.) (2005). *Multiple Criteria Decision Analysis: State of the Art Surveys*. Springer.

Fuller, A., Z. Fan, C. Day, and C. Barlow (2020). Digital twin: Enabling technologies, challenges and open research. *IEEE Access 8*, 108952–108971.

Gabellone, F. (2022). Digital twin: A new perspective for cultural heritage management and fruition. *Acta IMEKO 11*(1).

Government Technology Agency of Singapore (GovTech) (2017, 03). 5 things to know about virtual singapore. Accessed 2025-12-22.

Hwang, C.-L. and K. Yoon (1981). *Multiple Attribute Decision Making: Methods and Applications*. Berlin, Heidelberg: Springer.

Ioannidis, C., S. Verykokou, S. Soile, D. Istrati, C. Spyrakos, A. Sarris, D. Akritidis, H. Feidas, A. K. Georgoulias, E. Tringa, et al. (2024). Safeguarding our heritage, the triquetra project approach. *Heritage 7*(2), 758–793.

Ketzler, B., V. Naserentin, F. Latino, C. Zangelidis, L. Thuvander, and A. Logg (2020). Digital twins for cities: A state of the art review. *Built Environment 46*(4), 547–573.

Lehtola, V. V., M. Koeva, S. Oude Elberink, P. Raposo, J.-P. Virtanen, F. Vahdatikhaki, and S. Borsci (2022). Digital twin of a city: Review of technology serving city needs. *International Journal of Applied Earth Observation and Geoinformation 114*, 102915.

Mardani, A., A. Jusoh, K. M. Nor, Z. Khalifah, N. Zakwan, and A. Valipour (2015). Multiple criteria decision-making techniques and their applications: a review of the literature from 2000 to 2014. *Economic Research–Ekonomska Istraživanja 28*(1), 516–571.

Ortiz-Garcés, I. et al. (2025). Autonomous cyber-physical security middleware for IoT. *Frontiers in Artificial Intelligence 8*.

Pajic, M. et al. (2014). A notion of robustness for cyber-physical systems. *IEEE Transactions on Automatic Control 59*(11), 3055–3060.

Paradowski, B., A. Shekhovtsov, A. Baczkiewicz, B. Kizielewicz, and W. Sałabun (2021). Similarity analysis of methods for objective determination of weights in multi-criteria decision support systems. *Symmetry 13*(10), 1874.

Saaty, T. L. (1980). *The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation*. New York, NY: McGraw-Hill.

Tang, Z., D. Zhuang, and J. Zhang (2025). Evaluation framework for domain-specific digital twin platforms. *Scientific Reports 15*, 10544.

Themistocleous, K., V. Evripidou, and K. Toumbas (2025). Innovative methods for monitoring cultural heritage affected by climate change by enhancing visitor experience. In *Earth Resources and Environmental Remote Sensing/GIS Applications XVI*, Volume 13671 of *Proceedings of SPIE*, pp. 136710O.

Triantaphyllou, E. (2000). *Multi-criteria Decision Making Methods: A Comparative Study*. Boston, MA: Springer.

Wieckowski, J., B. Kizielewicz, and W. Sałabun (2025). Fuzzy RANCOM: A novel approach for modeling uncertainty in decision-making processes. *Information Sciences 694*, 121716.

Zadeh, L. A. (1965). Fuzzy sets. *Information and Control 8*(3), 338–353.