AI-Empowered Urban Digital Twins: Integrating Geospatial and Legal Data for Spatial Planning and Governance

Mila Koeva[1], Dessislava Petrova-Antonova[2], Sisi Zlatanova[3], Rob Peters[4], Claudio Persello[1], Ville Lehtola[1], Pirouz Nourian[1], Johannes Flacke[1], Raoul Grouls[5], Alfred Stein[1]

[1]Faculty of Geo-Information Science and Earth Observation (ITC), University of Twente, Enschede, The Netherlands – (m.n.koeva, c.persello, v.v.lehtola, p.nourian, j.flacke, a.stein)@utwente.nl
[2]GATE Institute, Sofia University "St. Kliment Ohridski", Sofia, Bulgaria – dessislava.petrova@gate-ai.eu
[3]School of Built Environment, University of New South Wales (UNSW), Sydney, Australia – s.zlatanova@unsw.edu.au
[4]Provincie Utrecht (Province of Utrecht), Utrecht, The Netherlands – rob.peters@provincie-utrecht.nl
[5]HU University of Applied Sciences Utrecht, Utrecht, The Netherlands – raoul@pttrn.io

Keywords: Artificial Intelligence, Digital Twins, Smart Cities, Earth Observation, Geospatial Data Science

Abstract

Urban planning increasingly requires accessible ways to connect regulatory frameworks with clear spatial visualisation and machine-assisted guidance. However, the integration of legal and geospatial information into platforms that support decision-making, communication, and participation remains limited. In particular, local authorities lack an operational way to translate legal and regulatory provisions into spatially explicit permit systems, and existing Urban Digital Twin architectures rarely formalise this legal-geospatial link. This paper uses Urban Digital Twins (UDTs) and Artificial Intelligence (AI), with a focus on generative AI, for smart, participatory, and responsible spatial planning. It proposes a conceptual framework for an AI-empowered Urban Digital Twin that integrates legal and 2D/3D geospatial data in both static and dynamic forms. Its aim is to support decision-making, information provision, and participation in spatial planning processes. The framework comprises five layers: data acquisition and harmonisation, data modelling, analytics and simulation, visualisation and interaction, and governance and decision support. Each of them are supported by AI models, and is illustrated through four case studies developed with the Province of Utrecht. The framework is supported by theoretical and methodological reflections and by experimental case studies. We conclude that integrating legal and geospatial data through AI-empowered UDTs strengthens spatial planning practices, improves regulatory interpretation, supports governance processes, and enables more informed and transparent decisions, provided that current limitations in AI reliability, explainability, and human oversight are explicitly addressed.

1. Introduction

Rapid urban growth and mounting environmental pressure have made sustainable spatial planning, balancing economic, environmental, and social objectives, a global priority. More than 74% of the population in Europe and 86% in Australia now live in urban areas, straining infrastructure and complicating planning amid climate and demographic change (Krishnan et al., 2021). Short-term, sector-specific responses often ignore how policy domains interact, fragmenting interventions and weakening the integration of social, environmental, and economic priorities (de Roo, 2012). Technological innovation is therefore increasingly seen as a route to more integrated planning and governance (Batty and Yang, 2022).

Planning Support Systems (PSS) were developed to address urban planning challenges, particularly in land use and transportation (Geertman and Stillwell, 2020), though the growing complexity and interdependence of urban systems continue to demand further adaptation (Ketzler et al., 2020). Communicative planning later responded to the limits of purely prescriptive approaches by shifting attention from zoning as a technical exercise towards a more policy-oriented, deliberative form of governance (Healey, 1992). Yet the regulatory texts behind this governance remain ambiguous and hard for people and machines to interpret consistently.

Governments have also invested in spatial data infrastructures to strengthen legal and environmental coherence. The Infrastructure for Spatial Information in the European Community (INSPIRE), for example, supports EU directives such as the Habitat Directive and the Water Framework Directive by making planning-relevant data more interoperable (van Loenen, 2006), building the technical foundation on which current Digital Twin agendas now rest (European Commission, 2026).

Geographic Information Systems (GIS) have long translated selected legal norms into spatial representations, such as zoning maps expressed as polygon-based layers linked to attributes like regulations, rights, or restrictions. These remain valuable for strategic and statutory planning, but local authorities still struggle to turn policy and legal provisions into permit systems precise enough to authorise specific actions, developments, and land-use changes.

Urban planning is further complicated by the coexistence of overlapping and, at times, conflicting policy frameworks across domains such as housing, energy, transport, and climate adaptation. Conventional legal-geospatial approaches often struggle to address this degree of institutional and semantic complexity. Several countries are therefore exploring new pathways. For example, the Association of Dutch Municipalities (VNG) is spearheading a national policy architecture for federated Digital Twins. Initiatives such as netherlands3d.eu, Rotterdam's Urban Data Platform, and the Dutch Metropolitan Innovations (DMI) ecosystem illustrate efforts towards geospatial legislative interoperability. In parallel, Australian state-level Urban Digital Twin initiatives are investigating mechanisms for linking legal information using GeoSPARQL and exploring the potential of Large Language Models for more context-sensitive interaction with planning rules. We draw on these Australian initiatives and on co-author expertise from the University of New South Wales because they offer methodologically relevant precedents for legal-geospatial linkage via GeoSPARQL and LLMs; the conceptual framework and case studies presented here nonetheless remain grounded in the European, and specifically Dutch, planning and legal context.

In this work, we propose an AI-empowered Urban Digital Twin conceptual framework to integrate geospatial and legal data in support of spatial planning and governance. Here, integration is understood as a multi-layered process comprising: (i) the integration of policy goals in response to fragmented local planning approaches; (ii) the technical integration of geospatial and legal data; and (iii) the integration of public values, including transparency, explainability, and contestability, in relation to the algorithms used.

2. Background and literature review

Urban Digital Twins (UDTs) have emerged as a powerful new approach for urban planning and governance. Conceptually, UDTs function as layered systems that combine static geospatial models with dynamic real-time data (Acharya et al., 2025; Xu et al., 2024; Lehtola et al., 2022). They support simulation, optimisation, and scenario modelling across planning domains, enabling the evaluation of interventions such as green space additions or mobility changes before implementation and thereby helping to mitigate risks such as congestion or pollution (Petrova-Antonova and Ilieva, 2019; Caprari et al., 2022).

Documented UDT applications span infrastructure maintenance, disaster response and hazard analysis (Xu et al., 2025), waste management (Cárdenas-León et al., 2024), cross-sectoral indicators for drinking water, urban heat and housing (Cárdenas-León et al., 2026), energy transition planning (Romero Rodríguez et al., 2017), facility-siting optimisation (Sun et al., 2026), LiDAR-enhanced traffic simulation (Wibisana, 2024), and urban heat, PET, environmental-performance, groundwater, shadow and hazard analyses (Afzalinezhad, 2024; Cárdenas-León et al., 2023; Tripathy et al., 2024; Sukma et al., 2024; Morales Ortega, 2023; Diakite et al., 2022; Aleksandrov et al., 2021). Digital Twins are also increasingly discussed in relation to SDG-oriented urban planning (Koeva et al., 2023), though significant conceptual and implementation challenges remain in aligning data-driven systems with legal and policy frameworks.

One promising development is AI's use within UDTs to enhance data search and processing, prediction, optimisation and real-time responsiveness (Li et al., 2018); embedded generative AI (GAI) extends this by enabling synthetic content and simulation data for future scenario planning, regulatory testing and automated urban management (Shirowzhan et al., 2020).

Still, early experiences highlight limits: in the Netherlands, rule-based AI used in environmental permitting handles straightforward cases (e.g., backyard sheds) but fails on complex, multivariate assessments involving legal thresholds such as sound levels or emissions. As Peters and van Engers (2004, 2016) note, citizens ask questions like "Where can I do this activity?", which rule-based systems cannot answer; LLMs and generative AI may help by offering more contextual, interactive and explainable support for urban spatial governance (Pan et al., 2026).

3. Conceptual framework for the Novel AI-empowered Urban Digital Twin

Traditional UDTs have been schematically represented as a dual of physical and virtual environments (Batty, 2018): the physical city is sensed via optical, range, environmental, meteorological and human-based systems, while the virtual counterpart holds data models for analysis, simulation and prediction, combining this real-world data with urban plans, legal and geospatial sources to generate alternatives and prognoses supporting governance decisions. This vision, however, falls short in linking legal aspects with geospatial components. We therefore propose a dedicated AI-empowered conceptual framework covering the entire lifecycle of legal, regulatory and geospatial data (Figure 1), in which each layer can draw on AI models alongside traditional data processing and analytics.

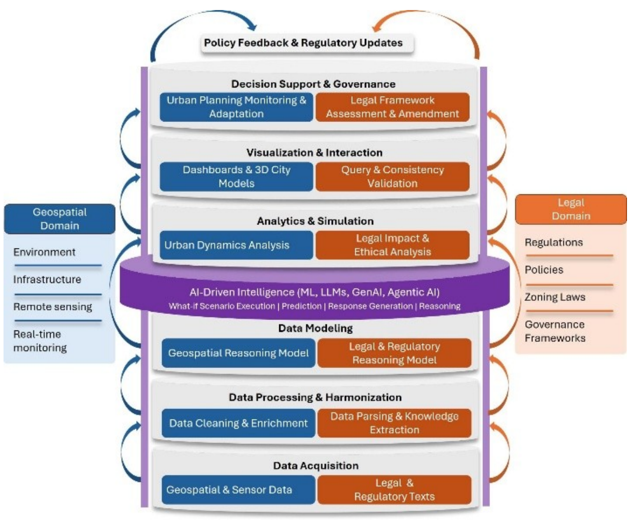

Figure 1. Conceptual view of an integrated legal, regulatory and geospatial architecture for UDT

Each layer builds on AI capabilities to prepare data for further analysis and simulations, process spatial and natural-language queries, and analyse, visualise, and explain geospatial, regulatory, and legal data. On the one hand, the UDT supports the assessment of the legal and regulatory framework and potential amendments, enabling the identification of outdated, conflicting, or inefficient regulatory elements based on empirical evidence and what-if scenario simulations. In addition, it facilitates urban planning, monitoring, and adaptation by evaluating real-time spatial and policy developments for their legal compliance, social impact, and environmental consequences. The continuous interaction between the geospatial and legal domains ensures that planning decisions remain legally grounded and that legal frameworks evolve to accommodate emerging urban challenges. AI thus acts as connective tissue for dynamic, informed, and legally consistent urban governance.

3.1 Data Acquisition, Processing and Harmonisation Layer

The Data Acquisition Layer collects data from heterogeneous sources with different spatial and temporal resolutions, including legal and regulatory documents, which the Data Preprocessing and Harmonisation Layer then prepares. Legal data may include government databases, legal document repositories, zoning maps, land-use regulations, and international conventions; legal texts must be parsed, normalised, and translated into machine-readable formats for logical querying and rule-based reasoning, while geospatial legal data, such as zoning maps, cadastral boundaries, or formal policy documents, must be digitised into standardised geospatial formats. Large Language Models (LLMs) can largely automate this process for non-legal users and also generate methods and strategies for linking it to geospatial datasets and tools. Urban datasets (e.g., imagery, sensor readings, land use data) must similarly be cleaned, georeferenced, and harmonised across coordinate systems and semantic/geometric representations, then mapped onto the discretised spatial framework, forming the spatiotemporal decision space that supports the real-time analytics, scenario testing, and automated planning that distinguish Digital Twins from conventional GIS systems.

Applicable AI Models:

- Image analysis using deep learning techniques (e.g. Convolutional Neural Networks or Vision Transformers): interprets satellite and aerial imagery to extract features such as buildings, roads, and vegetation, and to generate 2D/3D raster or vector maps.
- Data fusion/integration models: combine multiple sources, such as aerial imagery, satellite imagery, and LiDAR, into a unified urban representation.
- Noise reduction and signal processing: clean real-time sensor data and improve the quality of historical data for more reliable trend analysis.
- Natural Language Processing (NLP): extracts text from scanned legal documents and maps, and converts it into digital text for further processing.
- Named Entity Recognition (NER): identifies legal entities, key terms, and relevant provisions in legal texts.
- Geospatial data normalisation: standardises and aligns geospatial legal data to ensure compatibility with the Digital Twin's mapping system.

3.2 Data Modelling Layer

The Data Modelling Layer ensures that diverse data types are integrated into semantically rich, standards-aligned datasets for AI pipelines, maintaining quality, consistency, and compliance across geospatial, regulatory, and legal data before feeding them to the Analytics and Simulation Layer. This links structured and unstructured data, attributes, and geospatial objects, networks, and fields, reconciling differences in legal language, zoning definitions and regulations, and combining geospatial data from multiple sources with consistent formats, resolutions, and coordinate systems. The result is stored in a continuously updated, indexed spatial database that supports efficient querying and fast retrieval.

Applicable AI Models:

- Geospatial interpolation: fill spatial data gaps, e.g. using kriging, spline interpolation or machine-learning-based regression methods, to estimate unknown values across the spatial grid.
- Ontology-based data integration: supports the creation and validation of a unified legal-semantic structure that helps bridge differences between legal systems, concepts, and standards.
- Knowledge graphs: represent and link spatial, legal, and regulatory entities and their relationships in a machine-readable form.
- Semantic matching methods: align legal definitions, terms, and categories across jurisdictions and data sources to support interoperability.

3.3 Analytics and Simulation Layer

The Analytics and Simulation Layer performs advanced analysis, generates predictive insights, and simulates what-if scenarios for urban phenomena such as traffic flow, environmental impact, and legal compliance, feeding predictions to the Visualisation and Interaction Layer. It supports compliance analysis, ensuring development scenarios comply with relevant legal frameworks and flagging potential conflicts, and can predict the legal and regulatory impacts of proposed changes such as new infrastructure, zoning changes, or environmental regulations. Geospatial analysis derives insights into urban heat islands, traffic flow, and environmental risks, while predictive modelling on historical data forecasts future developments and infrastructure needs, with changes over time monitored via techniques such as image differencing or time-series analysis.

Applicable AI Models:

- State space models such as Long Short-Term Memory (LSTM) Networks or (attention-based) transformers: Time-series forecasting, e.g., predicting real-estate value change, traffic congestion or energy usage
- Generative Models such as Generative Adversarial Networks (GANs) or diffusion models: For generating synthetic data, such as potential urban growth patterns
- Large Language Models (LLMs): For legal analysis, scenario simulations, and policy impact predictions
- Legal Reasoning & Validation Models: Uses AI or, more specifically, GenAI (LLM-based methods including human-in-the-loop (HILT)) to evaluate urban scenarios against legal standards and provide compliance checks with the rules
- Segmentation Models (e.g., vision transformers, U-Nets): For semantic segmentation of multi-modal data, e.g., augmented satellite images to classify different land uses and urban features
- Decision Impact Prediction: A wide range of Artificial Neural Network configurations, mostly in the form of Graph Neural Networks (GNNs) or Probabilistic Graphical Models (PGMs) for building data-driven predictive models for ex-ante impact assessment

3.4 Visualisation and Interaction Layer

The Visualisation and Interaction Layer communicates the previous layer's insights, disaggregated or aggregated, to stakeholders, offering an informative trade-off between legal possibilities and policy or environmental objectives; ideally, it also supports interactive Multi-Criteria Decision Analysis (MCDA) or Spatial Multi-Criteria Evaluation (SMCE) for geospatial and legal decision-making. At its most basic, this layer reveals the regulatory designation of spatial zones; at its most advanced, it elicits the impact of legal zoning regulations and spatial plans on urban development against spatial, environmental, social, and economic objectives, letting users see potential legal barriers or opportunities in planning scenarios. It can also support systematic simulation of development scenarios through Design of Experiments (DoE), factoring in legal constraints and showing, for example, the impact of new infrastructure on traffic or of climate change on the urban environment (e.g., in Figures 2 and 3).

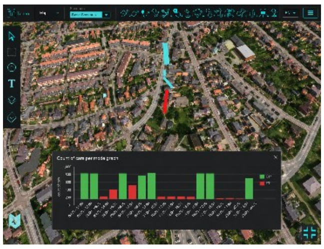

Figure 2. Policy effect calculation in Tygron. The bars indicate the effect per policy theme after a new set of housing or trees has been added to a scenario; the aim is to have AI agents calculate such scenarios automatically. This is in the testing phase, as shown below.

Figure 3. Policy values such as ‘level of required social and cultural facilities’ combined with ‘subsidised housing > 30%’ derived by AI from multiple formal documents and automatically projected directly in a twin of the city of Nieuwegein for scenario analyses.

3.5 Decision Support and Governance Layer

The Decision Support and Governance Layer provides an interface through which users interact with the urban digital twin. It allows for querying, manipulating, and interpreting data, as well as analysis and simulations. The geospatial interface lets users switch between 2D maps and 3D views, toggle data layers, and customise the visualisation to suit different analytical purposes. The legal query interface allows users to query legal data directly, enabling them to check specific regulations or zoning laws relevant to their interests or projects. Regulatory alerting can be implemented to notify users of legal risks or opportunities as they interact with the urban digital twin, providing intelligent (context-aware) feedback based on the latest data.

Feedback from domain-specific legal experts (e.g., land-use lawyers, zoning officials) and city stakeholders (e.g., planners, utility managers, residents) is systematically incorporated to refine the AI models’ rule-based components and learning parameters: inconsistencies between simulated and actual planning decisions can update zoning logic, while stakeholder input helps validate legal constraints encoded in the model, ensuring that simulations and outputs remain legally valid and grounded in real-world governance frameworks. This layer also initiates updates to legal and geospatial data as the physical city changes: newly adopted zoning amendments, permit approvals, or court rulings may alter legal constraints, while updated sensor readings or new satellite imagery can reflect new construction or shifting land use, all of which must be captured, validated, and fed back into the Digital Twin to keep future analyses current.

Applicable AI Models:

- Conversational AI (e.g., Chatbots using LLMs like GPT): For querying the urban digital twin in natural language
- Recommendation Systems: To suggest actions or interpret results based on user queries and system predictions
- AI-powered Adaptive User Interfaces: Automatically adjust and personalise the layout, functionality, and content of the user interface based on an individual's preferences, behaviours, and context
- Active Learning for Model Improvement: AI models that improve over time by incorporating feedback from users and new data inputs
- Adaptive Data Models: Models that adapt to changes in data inputs or the environment, ensuring the urban digital twin remains accurate and relevant (e.g. in Figure 4).

4. Case Studies: AI-Driven Urban Planning in Densely Populated Utrecht, the Netherlands

Utrecht, one of the densest urban centres in Western Europe, combines a high population concentration with a complex, layered urban environment typical of historic European cities. The city is mandated to expand from 350,000 to 440,000 inhabitants in line with national housing policy, while complying with the Environment and Planning Act (Omgevingswet). This regulation requires municipalities to assess and balance multiple quality indicators, such as noise levels, air quality, greenery, and water resilience. To address these challenges, Utrecht province, the regional authority of 26 municipalities, is actively investigating the use of AI. Several case studies demonstrate the potential of AI, particularly LLMs, to support spatial planning and decision-making across regulatory compliance, policy analysis, and heritage-sensitive development. The AI tools below (Cases 1 and 4) were developed in-house by the joint ITC-Province of Utrecht project team; Case 2 was built by an external Utrecht-based supplier, and Case 3 by the Province of Utrecht.

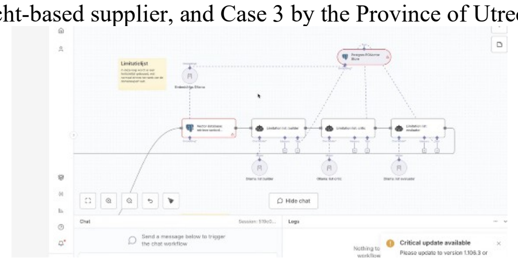

Figure 4. Part of an N8N agent architecture representing the Utrecht workflow of AI agents that support windmill identification and opportunity mapping; several LLMs contribute tasks in the workflow, directed by a skill library.

4.1 Case 1: Wind Mill Opportunity Map

This case study supports decision-making on wind turbine siting and the development of an “opportunity map.” Using an in-house, N8N-orchestrated multi-LLM workflow (Figure 4), all restrictions related to wind turbine placement are extracted from legal and regulatory documents and analysed in an integrated manner. The constraining factors are classified into three categories: complex, hard, and soft barriers. The results are visualised in a map that highlights all areas where wind turbines cannot be installed (Figure 5, GitHub -

uashogeschoolutrecht/SpatialGenAI: Multi-agent LLM system for spatial constraint analysis using RAG and Dutch geospatial data · GitHub).

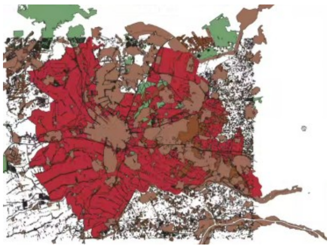

Figure 5. AI-generated opportunity map for building wind turbines (a), with the corresponding category legend (b). Colours indicate the three restriction categories defined in the text (complex, hard, and soft barriers); unshaded areas represent the resulting opportunity zones.

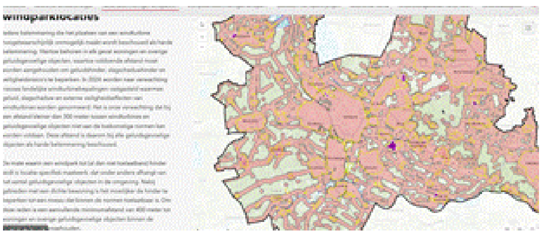

Figure 6. Opportunity map produced manually by Province of Utrecht planning staff, used to validate Figure 5.

This output has been compared with a previously produced map developed by the responsible official at the Province of Utrecht, showing a high level of agreement. Overall, this experiment clearly demonstrates the effectiveness of AI in finding opportunities (Figure 6).

4.2 Case 2: Chat with the Environmental Act

As mentioned above, the Environmental Act is a key regulatory framework in the Netherlands. The operational chatbot developed by a Utrecht supplier (Figure 7) offers a simple, accessible interface for the public to query this legislation: relevant sections are retrieved and presented in response to user questions, with associated spatial information visualised on a map where applicable. Users can pose queries such as “Can I cut trees in this location?”, and the large language model (LLM) searches regulatory documents to generate an answer, including references to the relevant clauses for that location.

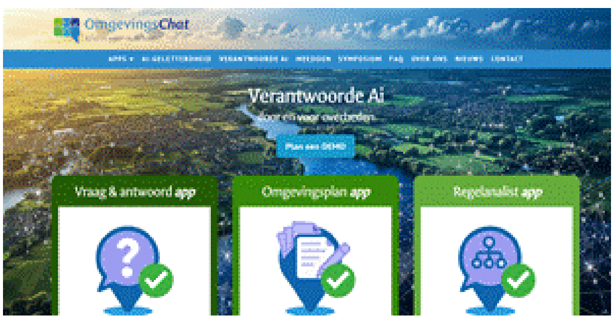

Figure 7. Chatbot for the Environmental Act: https://omgevingschat.nl/over-ons/

4.3 Case 3: Heritage Impact Assessment

This AI-based tool, built on ChatGPT and Unity by the Utrecht Province, identifies measures that permit or restrict the allocation of objects near risk-sensitive or heritage assets within a 3D environment. For example, housing or buildings near UNESCO-protected areas such as medieval fortress systems may be allowed only under specific conditions, including compensatory measures such as sound barriers, whereas wind turbines or solar panels on such structures may be prohibited near UNESCO-designated heritage assets. The AI system enables rapid geolocation and classification of relevant objects, generating outputs in formats such as Blender and JSON/Unity code to identify spatial opportunities and constraints, and supports 3D rendering and object positioning using Unity libraries (Figure 8).

Figure 8. Medieval Fortress (in red) with areas of restrictions around it, generated by AI after parsing from policy texts.

This functionality contributes to preliminary analyses required for Heritage Impact Assessments, thereby accelerating the overall permitting process. It reduces the analysis process length from months to days.

4.4 Case 4: Knowledge Graph and Decision Matrix

This case focuses on decision-making processes. The Utrecht Province’s provincial master plan is a legally binding document outlining the regional authority’s policy intentions; municipalities must adhere to its policies, ambitions, norms and rules across thematic areas such as green space, mobility and housing, complemented by a legislative website presenting the spatial implications of these themes (Figure 9). A key challenge is determining how specific actions or new developments may influence or interact with other policy domains. The scale of the policy is an additional challenge: there are about 35 documents, containing close to 400.000 words, that are all interlinked.

Figure 9. Legal Source policy-document interface (Omgevingsvisie, Provincie Utrecht, Omgevingsvisie provincie Utrecht, Actieplan geluid 2024-2029 provincie Utrecht, Omgevingsverordening provincie Utrecht (geconsolideerd) - PlanoView).

Given the impact of decision-making, any AI assistance must be fully transparent and explainable. To achieve this, we have been working to identify formal structures already present in the policies. We turned the 35 documents into nodes and edges that reflect how policy makers think about certain topics (see Figure 9), legal frameworks into decision trees, and interactions between types (activities, areas, principles, challenges, etc) into cartesian products that answered how the elements interact (for example, the cartesian product of Activity x Area answers the question “can we do activity A in location type B?”), backing up the outcomes of these interactions by linking them to existing nodes and edges in the policy. A major advantage of this approach is that humans can verify these relatively simple structures: we can create a table from a Cartesian product, link the table to specific nodes in the policy, and ask senior policymakers to verify the resulting tables as valid interpretations of the interaction. The AI can now use these human-verified structures as trustworthy and explainable sources of information.

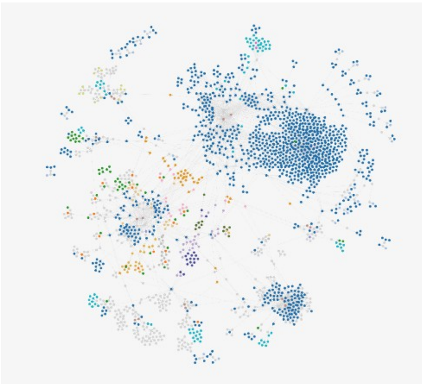

Figure 10. 35 Policy documents represented as a fully connected graph.

The policy advisor can adjust specific variables, such as introducing new areas for military practice, and assess how these changes influence other policy domains, including water management and drought resilience. Relationships between themes (e.g., military practice grounds and water or drought management) are explicitly defined by domain experts and added as new nodes and edges to the graph, and the agent in the LLM orchestration process (N8N, MCP) is instructed to use these verified structures within its analysis; when a policy is updated (eg by updating one of the 35 documents involved) it generates updated nodes and edges, triggering updates in the cartesian products that are flagged for review.

Hoofdstuk 4.1

Paragraaf	Gezondheid en veilig...	Innovatie	Klimaatadaptatie	Klimaatmitigatie	Kwaliteit en aantrekk...	Natuurinclusief	Verdienvermogen	Water en bodem stu...	Welzijn
4.1-1 - Leidende principes - Divisietekst	0	0	0	1	0	0	0	0	1
4.1-2 - Gezondheid en veiligheid - Divisietekst	2	0	1	1	1	1	0	0	1
4.1-3 - Werken aan een gezonde en veilige leefomgeving - Divisietekst	0	0	0	0	0	0	0	0	0
4.1-4 - Kwaliteit en aantrekkelijkheid - Divisietekst	1	0	0	0	2	1	1	0	1
4.1-5 - Welzijn - Divisietekst	1	1	0	0	1	0	0	0	2
4.1-6 - Klimaatmitigatie en klimaatadaptatie - Divisietekst	2	1	2	2	1	1	0	1	1
4.1-7 - Water en bodem sturend - Divisietekst	1	0	1	1	1	1	0	2	0
4.1-8 - Klimaatverkenning Utrecht 2050 - Divisietekst	-	-	-	-	0	0	0	-	0
4.1-9 - Natuurinclusief - Divisietekst	1	0	0	1	2	2	0	0	1
4.1-10 - Verdienvermogen - Divisietekst	0	2	1	1	1	0	2	0	0

Figure 11. Representing interaction between elements of the graph as a cartesian product table, verified for accuracy by a human.

5. Lessons learned

The developed tools were evaluated through stakeholder workshops with focused expert groups in the Province of Utrecht. Given the breadth of AI approaches, this study focused on LLMs as a tractable subset. In practice, this involved several layers of AI orchestration built as N8N workflows that coordinated LLaVA, Llama, Mixtral, Claude, ChatGPT, and other LLMs, with Milvus for vector storage. In particular, it was observed that LLMs:

- may produce different textual outputs for similar queries, potentially leading to variations in assessment processes and identified constraints;
- do not adequately account for the policy cycle, for example, by incorporating Environmental Impact Assessment (EIA/MER) considerations as inputs, despite these occurring at later stages in the planning process;
- may implicitly suggest or prioritise certain policy options, which is not desirable in a decision-support context.
- Semantic consistency of policy effect indicators becomes even more important when AI is involved.

These findings show the current implementation is only a first step toward layering policy goals and surfacing spatial conflicts. Future work should explicitly design features such as scenario analysis as decision-support tools for analysts and GIS experts. For now, the primary users are technical experts rather than policy-makers. Substantial development is still needed before policy professionals can rely on these tools in a legislative and politically sensitive context, and involving policy makers and planning professionals (heritage, water management, mobility, housing) at the requirements stage adds extra pressure on expectation management.

Several ethical considerations also emerged: explainability, accuracy and reliability of outputs, auditability, human accountability, and usability. As one participant noted, an LLM's self-explanatory text should always be questioned; a logically sound answer isn't necessarily accurate.

6. Discussion

AI-powered Digital Twins (DTs) are increasingly used in urban planning to support sustainable development and governance. By combining data from IoT devices, sensors, imagery, and traditional surveys, DTs create dynamic city models that enable real-time monitoring, predictive analytics, and scenario planning (Abdeen et al., 2023). Linking legal, geospatial, and historical data also enables a more objective evaluation of past policies. Some argue that embedding legal and policy-development processes directly into this architecture could streamline stakeholder consultation and better align development with sustainability goals, despite the added complexity (Batty, 2024).

6.1 Comparison with Traditional Urban Planning Models

Compared to traditional planning models and mono-disciplinary tools such as transport models or domain-specific Planning Support Systems (PSS), which often struggle to integrate diverse data and handle multi-domain complexity, the proposed AI-empowered DT framework leverages advanced data fusion and enrichment for more accurate, dynamic simulations (Jeddoub et al., 2023). It offers a generic software architecture that strengthens stakeholder engagement and SDG accountability by making decision-making trade-offs explicit (Koeva et al., 2023). This matters most in bridging legal policy development and urban planning: existing tools such as Urban Modelling or the Digital Cities Index address specific processes or digital connectivity, but not integrated legal-geospatial impact assessment. AI-empowered Digital Twins fill this gap, supporting opportunity-finding questions such as “where can we do what” or “what would this change mean for a city’s environment or economy”, while keeping planning aligned with legal standards.

6.2 Challenges and Limitations

Despite its innovations, the proposed DT framework faces several challenges. Data heterogeneity and quality remain significant issues, as integrating varied formats, structures, and standards can cause inconsistencies and inaccuracies (Latino et al., 2019). Because the case studies rely on Utrecht's comparatively rich legal and geospatial data, applying this approach in less data-mature contexts will likely require added investment in digitisation and harmonisation.

Three trade-offs constrain deployment: data richness versus privacy compliance under frameworks like GDPR (Bayat and Kawalek, 2023); model sophistication versus the computational resources realistically available in planning practice, compounded by the difficulty of balancing real-time and historical data in meaningful simulations (Jeddoub et al., 2023); and the efficiency gains of applied AI versus legislative integrity and public trust in the legal system. Integrating legal frameworks into Digital Twins also requires addressing interoperability, semantic consistency of policy indicators, and standardised data handling. Geographic information standards such as ISO 19115 (Metadata) and ISO 19157 (Data quality) need to be extended to cover AI-generated outputs to maintain quality control in spatial planning analysis.

7. Conclusion and Future Research

This paper proposes an AI-empowered Urban Digital Twin conceptual framework integrating geospatial and legal data for spatial planning and governance. Drawing on the Utrecht case, we show how such a framework supports more informed, adaptive, policy-sensitive planning by combining heterogeneous data into a dynamic, predictive representation of the urban environment. We approached integration as three interrelated dimensions: (i) aligning policy goals across fragmented planning domains, (ii) technically integrating geospatial and legal data, and (iii) embedding public values – transparency, explainability and contestability. Together, these positions view Urban Digital Twins not merely as efficiency-enhancing technologies but as socio-technical systems whose value depends on how they mediate among data, policy, and society, particularly by bridging the gap between policy intent and spatial implementation.

Future research should improve robustness, explainability, reproducibility and scalability through better data fusion, computational efficiency and usability; examine generative AI's role in model calibration and scenario generation (Xu et al., 2025); and develop richer legal-policy representations and AI-supported legal document processing (Sun et al., 2026). The decision tables, knowledge graphs, and skill libraries introduced here are only the beginning of validation instruments for AI-based digital twinning; applying the framework more broadly will help assess its transferability.

Acknowledgements

This research has been supported by the European Commission Genecity, HORIZON-MISS-2025-04-CIT-02, by Tygron and by the GIS team of the province of Utrecht.

References

Abdeen, F.N., Shirowzhan, S. and Sepasgozar, S.M.E., 2023. Citizen-centric digital twin development with machine learning and interfaces for maintaining urban infrastructure. *Telematics and Informatics*, 84, 102032.

Acharya, D., Yang, D., Wachowicz, M., Zhang, M., Shi, W. and Zlatanova, S., 2025. A conceptual framework for updating urban infrastructure in city digital twins. *ISPRS Archives*, XLVIII-G-2025, pp. 31–38. https://doi.org/10.5194/isprs-archives-XLVIII-G-2025-31-2025.

Afzalinezhad, A., 2024. Digital twin-based planning support system for urban heat island mitigation. M.Sc. thesis, University of Twente, Faculty of Geo-Information Science and Earth Observation (ITC), Enschede, the Netherlands.

Aleksandrov, M., Heslop, D.J. and Zlatanova, S., 2021. 3D indoor environment abstraction for crowd simulations in complex buildings. *Buildings*, 11(10), 445.

Batty, M. and Yang, W., 2022. A digital future for planning: spatial planning reimagined. Digital Task Force for Planning, London, UK. Available at: www.digital4planning.com.

Batty, M., 2018. Digital twins. *Environment and Planning B: Urban Analytics and City Science*, 45(5), pp. 817–820. https://doi.org/10.1177/2399808318796416.

Batty, M., 2024. Digital twins in city planning. *Nature Computational Science*, 4(3), pp. 192–199.

Bayat, A. and Kawalek, P., 2023. Digitization and urban governance: the city as a reflection of its data infrastructure. *International Review of Administrative Sciences*, 89(1), pp. 21–38.

Caprari, G., Castelli, G., Montuori, M., Camardelli, M. and Malvezzi, R., 2022. Digital twin for urban planning in the green deal era: a state of the art and future perspectives. *Sustainability*, 14, 6263.

Cárdenas-León, I., Koeva, M., Nourian, P. and Davey, C., 2024. Urban digital twin-based solution using geospatial information for solid waste management. *Sustainable Cities and Society*, 115, 105798.

Cárdenas-León, I., Nourian, P., Koeva, M.N. and Pfeffer, K., 2026. Cross-sectoral digital planning Drinking Water-Urban Heat-Housing Provision indicators dataset. *Data in Brief*, 64, 112310.

Cárdenas-León, I.L., Morales, R., Koeva, M., Atun, F. and Pfeffer, K., 2023. Digital twins for Physiological Equivalent Temperature calculation guide. Zenodo. https://doi.org/10.5281/zenodo.8306456.

de Roo, G., 2012. Spatial planning, complexity and a world 'out of equilibrium': outline of a non-linear approach to planning. In: de Roo, G., Hillier, J. and van Wezemael, J. (eds), *Complexity and Spatial Planning: Systems, Assemblages and Simulations*, Ashgate Publishing, Farnham, UK, pp. 129–165.

Diakite, A.A., Ng, L., Barton, J., Rigby, M., Williams, K., Barr, S. and Zlatanova, S., 2022. Liveable city digital twin: a pilot project for the city of Liverpool (NSW, Australia). *ISPRS Annals*, X-4/W2-2022, pp. 45–52.

European Commission, 2026. European Digital Infrastructure Consortium – EDIC. Shaping Europe's digital future. Available at: https://digital-strategy.ec.europa.eu/en/policies/edic

Geertman, S. and Stillwell, J., 2020. Planning support science: developments and challenges. *Environment and Planning B: Urban Analytics and City Science*, 47(8), pp. 1326–1342.

Healey, P., 1992. Planning through debate: the communicative turn in planning theory. *Town Planning Review*, 63(2), pp. 143–162.

Jeddoub, I., Nys, G.-A., Hajji, R. and Billen, R., 2023. Digital twins for cities: analyzing the gap between concepts and current implementations with a specific focus on data integration. *International Journal of Applied Earth Observation and Geoinformation*, 122, 103440. https://doi.org/10.1016/j.jag.2023.103440.

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

Koeva, M.N., Cárdenas, I.L. and Morales, L.R., 2023. Developments, challenges and applications of city digital twins in the context of the SDGs. *GIM International*, 37(7), pp. 26–29.

Krishnan, S., Aydin, N. and Comes, T., 2021. Planning support systems for long-term climate resilience: a critical review. In: Geertman, S.C.M., Pettit, C., Goodspeed, R. and Staffans, A. (eds), *Urban Informatics and Future Cities*, Springer, Cham, pp. 465–498.

Latino, F., Naserentin, V., Öhrn, E., Zhao, S., Fjeld, M., Thuvander, L. and Logg, A., 2019. Virtual City@Chalmers: Creating a prototype for a collaborative early stage urban planning AR application. *Proceedings of the 7th Regional International Symposium on Education and Research in Computer Aided Architectural Design in Europe*, pp. 15–24.

Lehtola, V. V., Koeva, M., Elberink, S. O., Raposo, P., Virtanen, J.-P., Vahdatikhaki, F., Borsci, S., 2022. Digital twin of a city: Review of technology serving city needs. *Int. J. Appl. Earth Obs. Geoinf.*, 114, 102915. https://doi.org/10.1016/j.jag.2022.102915.

Li, X., Cai, B.Y. and Ratti, C., 2018. Using street-level images and deep learning for urban landscape studies. *Landscape Architecture Frontiers*, 6(2), pp. 20–29. https://doi.org/10.15302/J-LAF-20180203.

Morales Ortega, L.R., 2023. A digital twin for ground water table monitoring. M.Sc. thesis, University of Twente, Faculty of Geo-Information Science and Earth Observation (ITC), Enschede, the Netherlands.

Pan, Y., Wang, M., Lu, L., Lamsal, R., Pärn, E., Zlatanova, S. and Brilakis, I., 2026. LLM-enabled multi-agent framework for natural language interaction with graph-based digital twins. *Automation in Construction*, 183, 106791.

Peters, R. and van Engers, T., 2004. The Legal Atlas ©: map-based navigation and accessibility of legal knowledge sources. In: Wimmer, M.A. (ed.), *Knowledge Management in Electronic Government*, Lecture Notes in Computer Science, 3035, Springer, Berlin, Heidelberg, pp. 228–236.

Peters, R.M., 2016. The Law, the Map and the Citizen: Designing a legal service infrastructure where rules make sense again. PhD thesis, University of Amsterdam, Amsterdam.

Petrova-Antonova, D. and Ilieva, S., 2019. Methodological framework for digital transition and performance assessment of smart cities. In: *2019 4th International Conference on Smart and Sustainable Technologies (SpliTech)*, pp. 1–6. https://doi.org/10.23919/SpliTech.2019.8783170.

Romero Rodríguez, L., Duminil, E., Sánchez Ramos, J. and Eicker, U., 2017. Assessment of the photovoltaic potential at urban level based on 3D city models: a case study and new methodological approach. *Solar Energy*, 146, pp. 264–275.

Shirowzhan, S., Tan, W. and Sepasgozar, S.M.E., 2020. Digital twin and CyberGIS for improving connectivity and measuring the impact of infrastructure construction planning in smart cities. *ISPRS International Journal of Geo-Information*, 9(4), 240.

Sukma, A.I., Koeva, M.N., Reckien, D., Bočkarjova, M., Da Silva Mano, A., Canilli, G., Vicentini, G. and Kerle, N., 2024. 3D city digital twin simulation to mitigate heat risk of urban heat islands. *ISPRS Archives*, XLVIII-4/W11-2024, pp. 129–136. https://doi.org/10.5194/isprs-archives-XLVIII-4-W11-2024-129-2024.

Sun, Y., Shen, X., Zlatanova, S., Barati, K. and Linke, J., 2026. Text-based automatic knowledge graph construction for road infrastructure operations management. *Automation in Construction*, 182, 106733.

Tripathy, A., Koeva, M.N., Cárdenas-León, I.L. and Nourian, P., 2024. Planning support system development for analysing environmental performance of form-based design decisions: case study of Enschede. *ISPRS Archives*, XLVIII-4/W11-2024, pp.137–144. https://doi.org/10.5194/isprs-archives-XLVIII-4-W11-2024-137-2024.

van Loenen, B., 2006. Developing geographic information infrastructures: the role of information policies. PhD thesis, Delft University of Technology, Delft. DUP Science.

Wibisana, M.I., Koeva, M., Nourian, P., Petrova-Antonova, D. and Karamitov, K., 2024. A LiDAR-based digital twinning workflow for traffic monitoring and simulation. *ISPRS Annals*, X-4-2024, pp. 411–418. https://doi.org/10.5194/isprs-annals-X-4-2024-411-2024.

Xu, H., Omitaomu, F., Sabri, S., Zlatanova, S., Li, X. and Song, Y., 2024. Leveraging generative AI for urban digital twins: a scoping review on the autonomous generation of urban data, scenarios, designs, and 3D city models for smart city advancement. *Urban Informatics*, 3(1), 29.

Xu, H., Zlatanova, S., Liang, R. and Canbulat, I., 2025. Generative AI as a pillar for predicting 2D and 3D wildfire spread: beyond physics-based models and traditional deep learning. *Fire*, 8(8), 293.