

Article

Generative AI-Driven Digital Twin Architecture for Urban Mobility Simulation and Decision Support

Pablo Vicente-Martínez ^{1,*}, Emilio Soria-Olivas ², Adrián Chust-Ros ¹, María Ángeles García-Escrivà ^{3,*}, Edu William-Secin ⁴ and Manuel Sánchez-Montañés ⁵

¹ SPV Scala, Gran Canaria, 35100 San Bartolomé de Tirajana, Spain; c.datos34@salascalea.com

² Intelligent Data Analysis Laboratory (IDAL), Department of Electronic Engineering, Universitat de València, 46100 Valencia, Spain; emilio.soria@uv.es

³ Fundación Canaria Living Lab, 35017 Las Palmas de Gran Canaria, Spain

⁴ Department of Economics and Business Management, Institute of Tourism and Sustainable Development (TIDES), Universidad de Las Palmas de Gran Canaria, 35017 Las Palmas de Gran Canaria, Spain; edu.william@ulpgc.es

⁵ GNB, Department of Computer Science, Universidad Autónoma de Madrid, 28049 Madrid, Spain; manuel.smontanes@uam.es

* Correspondence: c.datos12@salascalea.com (P.V.-M.); coordinacionit@canariaslivinglab.org (M.Á.G.-E.)

Highlights

What are the main findings?

- A generative AI-driven digital twin architecture was developed to translate natural language requests into validated urban mobility simulation configurations through a conversational interface.
- The proof-of-concept demonstrates a 96% success rate in intent recognition and parameter mapping, while completing 19 h multimodal mobility simulations in less than 20 min.

What are the implications of the main findings?

- The integration of large language models with digital twins can significantly reduce technical barriers, enabling non-specialist stakeholders to participate in urban mobility planning and scenario analysis.
- A decoupled validation layer based on schema enforcement improves the reliability and safety of AI-generated configurations, supporting scalable and trustworthy decision-support systems for smart city applications.

Abstract

Urban mobility planning in smart cities requires sophisticated simulation tools, yet their complexity often creates a technical barrier for non-expert stakeholders. This paper presents a novel architecture that integrates generative artificial intelligence with digital twin technology to create an accessible and decision-support prototype. The framework employs a conversational AI agent based on Gemini 2.5 Flash Lite to interpret natural language intentions and translate them into validated simulation parameters. A critical safety layer, built using Pydantic, ensures that the agent's stochastic outputs adhere to strict technical schemas and predefined logical bounds before execution. The underlying digital twin, developed with SimPy, NetworkX, and OSMnx, features a multi-source data integration strategy that includes demographic density (INE), tourism activity (ISTAC), and high-resolution traffic statistics (TomTom) to calibrate vehicle behavior. The architecture was technically demonstrated through a Technology Readiness Level (TRL) 4 proof-of-concept in Las Palmas de Gran Canaria, simulating multimodal scenarios including buses, the future *MetroGuagua* (BRT), and pedestrian flows. Results demonstrate a 96% success rate



Academic Editors: Pierluigi Siano and Jian Sun

Received: 2 May 2026

Revised: 16 June 2026

Accepted: 21 June 2026

Published: 30 June 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

in intent recognition and configuration mapping, with end-to-end execution times under 20 min for a 19 h simulated day. This study demonstrates that LLM-driven orchestration, coupled with automated data pipelines and a decoupled microservice architecture, can lower technical barriers to urban simulation, which could support broader participation in future smart city deployments.

Keywords: smart cities; digital twin; generative AI; urban mobility; decision support system; conversational AI; sustainable mobility; urban planning

1. Introduction

1.1. Context and Motivation

Urban mobility represents one of the most critical challenges faces by smart cities. As urban populations continue to grow, with projections indicating that 68% of the world's population will reside in urban areas by 2050 [1], cities face unprecedented pressures on their transportation infrastructure. Traffic congestion, air pollution and energy consumption issues constitute persistent challenges that demand innovative solutions [2]. The concept of smart cities emerges as a comprehensive framework that leverages digital technologies, data analytics, and intelligent systems to address these urban challenges while promoting sustainability, efficiency, and quality of life [3].

Consequently, modern urban mobility planning has evolved from traditional static approaches to dynamic, data-driven methodologies capable of accommodating the complexity and variability inherent in modern transportation networks [4]. Decision-makers require tools that not only simulate current conditions but also enable the exploration of alternative scenarios, predict the impact of interventions, and support evidence-based policy formulation [5]. However, traditional simulation and planning tools often present significant barriers: they require specialized technical expertise, involve lengthy configuration processes, and lack the flexibility to rapidly iterate through multiple scenarios [6].

1.2. State of the Art

Addressing these limitations requires the integration of three distinct technological domains: digital twin technology, generative artificial intelligence, and conversational AI systems.

1.2.1. Digital Twins for Urban Planning

Digital twins (virtual replicas of physical systems that enable real-time monitoring, simulation, and optimization) have gained substantial traction in urban planning and smart city applications [7,8]. Urban digital twins integrate heterogeneous data sources including geospatial information, sensor networks, and historical records to create dynamic models of city infrastructure and services [9]. Applications span traffic management [10,11], energy systems [12,13], and comprehensive city-scale planning [14].

In the mobility domain specifically, digital twins enable the simulation of multimodal transportation networks, incorporating vehicles, public transit, pedestrians, and emerging mobility services [15]. These systems typically employ agent-based modeling (ABM) frameworks where individual entities (vehicles, passengers) follow defined behavioral rules within a spatial network representation [16]. Tools such as SUMO (Simulation of Urban MObility) [17], MATSim [18], and custom simulation engines built with discrete-event simulation libraries like SimPy have been successfully applied to urban mobility analysis [19].

Despite their sophistication, most digital twin implementations require significant technical expertise for configuration and scenario design, limiting their accessibility to specialized analysts rather than broader stakeholder groups including urban planners, policymakers, and community representatives [20].

1.2.2. Generative AI for Simulation and Urban Design

Large language models (LLMs) and generative AI offer new paradigms for automating complex tasks through natural language interaction [21,22]. Recent work has explored the application of generative AI to urban design [23,24] and scenario generation [25].

Generative models excel at translating high-level user intentions expressed in natural language into structured system configurations [26]. This capability addresses a critical bottleneck in simulation systems: the translation of domain expertise and planning objectives into technical parameters. Google's Gemini models, including the Gemini 2.5 Flash variants, represent state-of-the-art multimodal LLMs with demonstrated capabilities in complex reasoning, structured output generation, and tool use [27]. This structural translation aligns with recent paradigms in intelligent transportation systems that bridge natural language processing with complex spatial databases. For instance, frameworks like MAGAE [28] leverage multi-level alignment over semantic graphs with attribute enhancement to optimize vehicle retrieval from text queries. Similarly, our architecture maps conversational user intent to localized simulation graphs, ensuring that high-level planning objectives are dynamically embedded into the digital twin's technical parameters.

However, the integration of generative AI with domain-specific simulation engines for urban mobility remains largely unexplored. Existing applications primarily focus either on standalone AI assistants for information retrieval or on simulation systems with conventional interfaces, lacking a deep integration where AI actively configures and controls simulation parameters [29].

1.2.3. Conversational AI for Decision Support

Conversational AI systems, employing natural language processing (NLP) for human-computer interaction, are currently deployed across numerous domains including healthcare [30,31], customer service [32], and business analytics [33]. In the context of decision support systems, conversational interfaces reduce cognitive load, accelerate task completion, and democratize access to complex analytical tools [34].

Several studies have investigated conversational interfaces for urban data exploration and smart city applications [35]. However, these implementations typically provide read-only access to data or simple parameter adjustments, rather than comprehensive configuration management for complex simulation systems. Furthermore, most conversational AI systems for urban applications operate independently from digital twin infrastructures, limiting their capacity to provide actionable insights grounded in rigorous simulation [36].

1.3. Research Gap and Motivation

While substantial progress has been made in each of these three domains independently, a critical gap exists at their intersection. Specifically, there is a lack of integrated architectures that combine:

1. Digital twin technology for realistic, data-driven simulation of urban mobility systems;
2. Generative AI for intelligent interpretation of user requirements and automated configuration;
3. Conversational interfaces for accessible, natural language interaction with complex simulation systems.

This gap is particularly significant given the pressing need for decision support tools that can be effectively utilized by diverse stakeholders in smart city planning—not only technical specialists but also urban planners, policymakers, and community representatives who may lack programming or simulation expertise [6].

To clearly position our contribution within the existing landscape, Table 1 contrasts the proposed generative AI-driven architecture with established simulation frameworks and urban dashboards. While traditional tools excel in simulation fidelity, they frequently lack accessible natural language interfaces and automated configuration mechanisms. Our approach bridges this gap by prioritizing accessibility, safety validation, and rapid scenario iteration.

Table 1. Comparison of the proposed architecture with existing urban simulation and digital twin approaches.

Feature	Traditional DES (e.g., SUMO, MATSim)	Urban Dashboards/Virtual Singapore	Proposed Generative AI-Driven Digital Twin
Primary Interface	Code/Complex GUIs	Visual/Interactive Maps	Natural Language (Conversational AI)
Configuration	Manual XML/Code editing	Pre-configured scenarios	Automated YAML generation via LLM
Technical Barrier	High (Programming expertise)	Low (Read-only)	Low (Conversational with underlying complex execution)
Safety Validation	Manual/Post-execution errors	N/A	Automated schema validation (Pydantic)

Moreover, existing digital twin frameworks for urban mobility rarely incorporate mechanisms for dynamic reconfiguration based on stakeholder input, limiting their utility in iterative, participatory planning processes essential for equitable smart city development [37].

To address these systemic gaps, this research explores the following overarching Research Questions (RQs):

- RQ1: How can generative AI and large language models be architecturally integrated with discrete-event simulation engines to automate the configuration of urban mobility digital twins?
- RQ2: To what extent can a decoupled validation layer ensure the technical safety and consistency of stochastic LLM outputs in smart city planning scenarios?
- RQ3: Does a conversational interface significantly reduce the operational latency and technical barriers for setting up and executing complex, multimodal ‘what-if’ urban simulations?

By answering these questions, this study demonstrates the broader applicability of using standardized open data formats combined with generative AI to create scalable decision-support systems.

1.4. Proposed Approach

This paper addresses the identified gap by presenting a novel integrated architecture that synergistically links generative AI with digital-twin-oriented simulation prototypes for urban mobility simulation and decision support. Our approach employs a conversational AI agent built on the Gemini 2.5 Flash Lite model to enable natural language configuration of simulation parameters, executed on a digital twin constructed using established simula-

tion tools (SimPy for discrete-event simulation, NetworkX and OSMnx for spatial network representation, and GeoPandas for geospatial data processing).

The architecture covers the complete decision support cycle: from data acquisition and processing, through scenario configuration and simulation execution, to visualization and interpretation of results. Here, the conversational AI agent acts not merely as an interface layer but as an intelligent intermediary that interprets user intentions, validates configuration consistency, and maintains context across multiple interactions.

1.5. Contributions

This work makes the following key contributions to the field of smart cities and urban mobility planning. It is important to note that the core contribution lies in systems integration and architectural design rather than the development of novel AI algorithms or simulation methods:

1. **Integrated Architecture:** We present a novel architecture that seamlessly integrates generative AI with digital twin technology, demonstrating how these complementary technologies can be combined to create more accessible and flexible decision support systems for urban mobility.
2. **Conversational Configuration:** We introduce and evaluate a conversational AI agent capable of interpreting natural language requests, identifying relevant simulation parameters (through entity recognition), and dynamically modifying configuration files, reducing the technical barrier to entry for non-specialist users.
3. **Multimodal Mobility Simulation:** We demonstrate the application of our architecture to a realistic multimodal urban mobility scenario incorporating buses, metro systems (MetroGuagua), and pedestrian flows. While the current TRL-4 evaluation focuses primarily on technical feasibility, the architecture provides a robust foundation for future explicit consideration of sustainability indicators (environmental, economic, and social dimensions).
4. **Proof-of-Concept Evaluation:** We provide comprehensive documentation and evaluation of a functional prototype at Technology Readiness Level (TRL) 4, demonstrating technical feasibility in a controlled laboratory environment and establishing a foundation for future development toward operational deployment.
5. **Open and Extensible Design:** The modular architecture facilitates integration with diverse data sources, simulation engines, and visualization tools, supporting adaptation to different urban contexts and planning requirements.

1.6. Paper Structure

The remainder of this paper is organized as follows. Section 2 describes the materials and methods, detailing the system architecture, data acquisition processes, digital twin construction, conversational AI implementation, and visualization approach. Section 3 presents the results of our proof-of-concept evaluation, including specific simulation scenarios and performance metrics. Section 4 discusses the implications of our findings for smart city planning, compares our approach with prior work, and acknowledges current limitations. Section 5 concludes with a summary of key contributions and directions for future research.

2. Materials and Methods

This section presents the comprehensive methodology employed to develop and validate the proposed design. We describe the system architecture, data acquisition and processing pipelines, digital twin construction approach, the implementation of the conversational AI agent, simulation generation mechanisms, and visualization strategies. The

implementation follows established software engineering practices for modular, scalable system design appropriate for a Technology Readiness Level (TRL) 4 proof-of-concept.

The selection of this specific technological stack was explicitly driven by the need to balance computational efficiency, deterministic execution, and accessibility. Specifically, Gemini 2.5 Flash Lite was selected not only for its conversational capabilities but for its robust support of Structured Outputs, which is critical for translating stochastic human intent into rigid simulation parameters. SimPy was chosen over continuous simulation frameworks because discrete-event modeling is optimally suited to resolve resource contention (e.g., vehicle capacity bottlenecks) without prohibitive computational overhead. Finally, Pydantic acts as an essential architectural firewall, driving innovation in safe AI deployment by ensuring that all LLM-generated configurations are strictly validated before execution.

2.1. System Architecture Overview

The proposed architecture integrates three primary components: (1) a conversational AI agent for natural language configuration, (2) a digital twin simulation engine for multi-modal urban mobility modeling, and (3) a visualization layer for interactive exploration of simulation results. Figure 1 presents the high-level system architecture.




Figure 1. Detailed system architecture highlighting four primary layers: (1) the Conversational AI interface incorporating Gemini 2.5 and Pydantic for validated YAML generation and user feedback loops; (2) Data Ingestion pipelines unifying open and proprietary sources into a PostGIS database; (3) the Asynchronous Simulation Engine powered by SimPy and orchestrated via AWS Fargate; and (4) the output visualization layer using Kepler.gl.

The architecture follows a modular design pattern, in which each component operates independently while maintaining well-defined interfaces for data exchange. This modularity facilitates system maintenance, component upgrades, and adaptation to different urban contexts. The system operates primarily in Python 3.11 and is containerized using Docker to ensure reproducibility and deployment consistency across different computing environments.

Cloud infrastructure services (AWS ECS for container orchestration, AWS RDS for PostgreSQL database hosting, and AWS S3 for data storage) provide scalable computational resources and persistent storage, although the current TRL 4 implementation operates in a controlled laboratory environment rather than production deployment.

2.2. Digital-Twin-Oriented Simulation Construction

The reliability of such a data-driven simulation platform is intrinsically linked to the quality, diversity, and resolution of its underlying data. In this work, we implemented a multi-source data acquisition strategy that combines open-access institutional repositories with high-resolution private traffic data. This approach ensures that the simulation is not

merely a theoretical model but a data-driven representation of the urban dynamics of Las Palmas de Gran Canaria.

2.2.1. Data Sources

The digital twin integrates heterogeneous geospatial and demographic data to describe the urban transportation network and mobility patterns:

- **Road Network Data:** OpenStreetMap (OSM) provides the street network topology, including road classifications and connectivity. Data is accessed programmatically via the OSMnx Python library [38], which retrieves and processes network data for specified geographic bounding boxes to construct a directed graph representation.
 - **Public Transit Data (GTFS and Custom):** Transit stop locations, line geometries, and official schedules were processed from General Transit Feed Specification (GTFS) datasets provided by the local authority, *Guaguas Municipales*, through the Transitland API [39]. Additionally, a custom GeoJSON dataset was developed to represent the future *MetroGuagua* (BRT) infrastructure, incorporating planned stops and corridors based on official project specifications.
 - **Demographic and Resident Dynamics (INE):** Population density data was obtained from the Spanish National Institute of Statistics (*Instituto Nacional de Estadística*, INE) at the census tract level (*secciones censales*). This allows for a high-granularity distribution of resident agents based on residential density.
 - **Tourism and Temporary Population (ISTAC and OSM):** To model tourist flows, we integrated occupancy rates by accommodation type from the INE, complemented by precise geographic locations of hotels from OSM. Furthermore, data on licensed holiday rentals and apartments were retrieved from the Canary Islands Statistics Institute (ISTAC) to ensure a comprehensive mapping of tourist ‘spawn points’.
 - **Points of Interest (POIs):** Over 200 geographically referenced locations of tourist attractions, commercial centers, and activity generators were extracted from OSM using custom importers that query specific amenity types [40]. These served as primary destinations in the agents’ activity agendas, categorized by their profile (resident vs. tourist).
 - **High-Resolution Traffic Statistics (TomTom):** To overcome the limitations of static schedules, we incorporated segment-level traffic data from TomTom. This dataset provides percentiles for vehicle speeds for every street segment in the city, sampled at 30 min intervals from 06:00 to 18:00. While the provider also offers volume percentiles, the current TRL-4 simulation engine relies solely on historical speed data to calibrate segment traversal times; integrating volume data for dynamic congestion modeling is reserved for future work.
- This enables the simulation to account for historical congestion patterns and provides a first-order approximation of transit behavior, particularly in areas where transit delays are frequent.
- **Synthetic Pedestrian Profiles:** We generate synthetic pedestrian populations with differentiated behavioral profiles representing tourists and local residents [41]. Because this TRL-4 prototype is not yet calibrated against observed pedestrian counts or mobile phone data, these profiles are designed to generate plausible synthetic flows for system testing, rather than a validated representation of actual pedestrian demand.

2.2.2. Data Processing and Graph Unification

Raw data undergoes a systematic pipeline to ensure quality and compatibility with the simulation engine. A significant contribution of this research lies in the automated and

manual processing required to unify these heterogeneous sources into a single, simulation-ready graph.

Automated Importers and Scalability: For highly dynamic sources such as INE, ISTAC, GTFS, and OSM, we developed modular Python-based importers. These scripts facilitate the periodic update of the digital twin's baseline, ensuring that any changes in census data or transit schedules can be reflected with minimal manual intervention. This design also supports the framework's replicability in other Spanish cities where these standardized open data sources are available, although the highest-fidelity traffic calibration still requires complementary proprietary data.

Spatial and Network Processing: The digital twin represents the urban environment as a geospatially embedded network graph $G = (N, E)$, where N denotes the set of nodes (intersections, stops) and E represents the set of directed edges (road segments). Each edge $e \in E$ is associated with the following attributes:

$$e = \{\text{geometry}, \text{length}, \text{speed_limit}\} \quad (1)$$

GeoPandas [42] and Shapely [43] libraries handle geospatial manipulation, including coordinate system transformations and geometric validation. The road network is converted into a directed graph using NetworkX [44]. For each bus and *MetroGuagua* route, a dedicated subgraph is generated, ensuring that transit agents are constrained to realistic paths.

Transit Route Calibration and Map-Matching: GTFS data is parsed to extract route shapes and sequences. Routes are aligned with the road network graph through map-matching algorithms that project trajectories onto the nearest network edges [45]. One of the most complex tasks involved the spatial alignment of proprietary TomTom traffic segments with the OSMnx network. This unification allows for the assignment of historical speed profiles to each edge. Given that transit vehicles typically exhibit lower average speeds than private cars, we implemented a calibration factor. In traffic-based mode, bus speeds are estimated as 70% of the median speed (V_{median}) reported by TomTom:

$$V_{bus} = 0.7 \times V_{median}(\text{segment}, \text{time}) \quad (2)$$

The static coefficient of 0.7 operates as a macroscopic heuristic penalty. It was statistically derived by comparing historical GTFS scheduled traversal times against private vehicle free-flow speeds during peak operational hours in Las Palmas. This factor accounts for kinematic constraints inherent to heavy transit vehicles, specifically, lower acceleration rates, deceleration profiles when approaching stops, and the negotiation of tight urban geometries; which are not reflected in the aggregate TomTom speed percentiles. While a dynamic, segment-specific kinematic model would yield higher fidelity, this static adjustment provides a computationally inexpensive baseline suitable for a TRL-4 strategic planning tool.

For simulation hours outside the TomTom data coverage window (i.e., 05:00–06:00 and 18:00–24:00), the system defaults to the segment's legal speed limit (free-flow conditions). This serves as a reasonable and standard approximation, given that urban congestion is minimal during these early morning and night-time periods.

This heuristic approach provides a more plausible baseline for transit travel times compared to models that rely solely on GTFS theoretical times, which are frequently unmet in Las Palmas due to urban congestion. The scaling factor was also added to the configuration file to allow users to simulate different scenarios and analyze the impact of changing bus speeds.

2.2.3. Dynamic Entity Modeling

The digital twin incorporates three primary classes of dynamic entities, each governed by specific behavioral rules:

1. **Bus Entities:** Buses follow predefined routes extracted from GTFS data, adhering to scheduled departure times and stop sequences. The simulation models boarding and alighting processes at each stop, tracking vehicle occupancy dynamically. Bus movement between stops incorporates:
 - **Schedule Adherence:** Buses attempt to maintain scheduled arrival times but may experience delays due to traffic conditions (modeled as stochastic edge traversal times).
 - **Capacity Constraints:** Each bus has a maximum passenger capacity. If a bus reaches capacity, waiting passengers must wait for the next service.
 - **Dwell Time:** Time spent at stops depends on the number of passengers boarding and alighting, following empirical distributions derived from transit operations literature [46].
2. **MetroGuagua Entities:** The MetroGuagua (a bus rapid transit system) operates similarly to conventional buses but with higher frequencies, dedicated lanes (modeled as edges with reduced travel times), and larger capacities. MetroGuagua routes are treated as high-priority corridors in the network graph.
3. **Pedestrian Agents:** Pedestrians are modeled as autonomous agents with individualized origin-destination pairs, behavioral profiles (tourist vs. local resident), and activity schedules. Pedestrian movement incorporates:
 - **Route Choice:** Pedestrians select routes based on shortest path algorithms (Dijkstra's algorithm implemented in the NetworkX library) considering walking distance.
 - **Mode Choice:** When pedestrians decide to visit a Point of Interest (POI), they decide whether to board available services based on factors including destination proximity and distance to transit stops.
 - **Activity Scheduling:** Pedestrian agents have time-dependent activity patterns (visiting POIs, commuting) that influence their spatial-temporal trajectories.

2.2.4. Discrete-Event Simulation Engine

The digital twin employs SimPy [19], a process-based discrete-event simulation (DES) framework in Python, to orchestrate entity behaviors and interactions. In DES, the simulation advances through discrete events (e.g., bus departure, passenger boarding, agent arrival at destination) rather than continuous time steps, improving computational efficiency for large-scale scenarios.

Mathematically, the simulation models edge traversal as a time-dependent resource consumption. The base travel time $T_e(t)$ for a private vehicle traversing an edge e of length L_e at a given time t is determined by the historical speed profile:

$$T_e(t) = \frac{L_e}{V_{\text{TomTom}}(e, t)} \quad (3)$$

For transit entities (e.g., buses), the total traversal time between nodes incorporates both the transit speed penalty and the stochastic dwell times D_{stop} at designated stops. The effective traversal time $T_{e,\text{bus}}(t)$ is formulated as follows:

$$T_{e,\text{bus}}(t) = \frac{L_e}{0.7 \times V_{\text{TomTom}}(e, t)} + D_{\text{stop}} \quad (4)$$

In this discrete-event paradigm, traffic propagation is implicitly handled by these time-varying speed percentiles rather than explicit microscopic car-following models, thereby ensuring fast execution times for city-scale scenarios.

SimPy's environment manages the simulation clock, schedules events, and resolves resource contention (e.g., multiple agents competing for limited bus capacity). Each entity is implemented as a SimPy process that yields control to the simulation engine during waiting periods or resource acquisition attempts.

The simulation operates over a configurable time horizon (typically representing a full day of urban activity, e.g., 05:00 to 24:00) with a temporal resolution sufficient to capture meaningful mobility dynamics while maintaining computational tractability. All entity states (positions, occupancies, delays) are logged at configurable intervals for subsequent analysis.

2.2.5. Optimization and Pre-Computation

To achieve the execution times required for an interactive decision-support system, several parameters are pre-calculated using the baseline configuration. This includes the generation of arrival/departure matrices for every stop under both schedule-based and traffic-based scenarios. Pre-computing these temporal variables reduced the computational overhead during simulation execution, allowing 19 h simulations to be processed in minutes.

2.2.6. Quality Assurance and Architectural Readiness

The pipeline includes automated routines to detect errors in the matching process between different graph sources (e.g., disconnected transit stops or misaligned segments). For data consistency, Pydantic [47] schemas are employed to enforce strict type checking and validation, particularly for the configuration files generated by the conversational AI agent. This ensures that parameters such as frequencies (integers), capacities, and probability distributions (which must sum to 1.0) are valid before simulation execution.

While a mature digital twin implies continuous, real-time synchronization with the physical system, the current TRL 4 implementation operates as a digital-twin-oriented prototype. It handles high-volume TomTom traffic data through specialized manual pre-processing rather than live ingestion. However, the architecture is designed to support real-time data streams in future operational deployments. The internal data structures and API endpoints are prepared to transition toward live sensor integration (e.g., traffic cameras or GPS feeds), representing a primary direction for future operational scaling and real-time analytics.

2.2.7. Data Management and Infrastructure

All processed data and simulation results are stored in a PostgreSQL database with PostGIS extensions for spatial queries, hosted on AWS RDS. This centralized repository facilitates efficient data access and supports version control for different simulation scenarios, providing a scalable foundation for the digital twin infrastructure.

2.3. Conversational AI Agent for Configuration

2.3.1. Natural Language Processing Architecture

The conversational AI agent serves as the primary user interface, enabling urban planners and decision-makers to configure simulation scenarios through natural language dialogue rather than manual editing of configuration files or programming. The system utilizes Gemini 2.5 Flash Lite [27] as the core reasoning engine, specifically deployed via the Google Cloud Vertex AI API (model version: `gemini-2.5-flash-lite`). To ensure deterministic behavior and minimize hallucinations during parameter extraction, the model's

inference settings were strictly configured with a temperature of 0.0. Furthermore, we utilized the model's native Structured Outputs capability (requesting application/json response MIME type) to guarantee that the raw output matched our predefined schemas before reaching the Pydantic validation layer [48].

While frontier models introduced in 2026 exhibit superior generalized reasoning, Gemini 2.5 Flash Lite was deliberately selected to optimize the cost-latency trade-off essential for a scalable urban architecture. In our framework, the LLM functions primarily as an orchestration and extraction engine rather than an open-ended consultant. The Flash Lite variant provides exceptional inference speed and proven reliability in generating rigid, schema-compliant JSON outputs (Structured Outputs). Deploying heavier computational models for these deterministic mapping tasks would introduce unnecessary latency and significantly increase the operational costs for municipal authorities without a proportional gain in configuration accuracy.

The agent makes use of four primary capabilities:

1. **Intent Recognition:** The agent analyzes user utterances to identify the underlying intent (e.g., modifying bus frequencies, adjusting MetroGuagua capacity, offering results). Intent recognition leverages Gemini's in-context learning capabilities, where the system prompt provides examples of intent-query mappings.
2. **Named Entity Recognition (NER):** For configuration modification requests, the agent extracts relevant entities including:
 - Parameter names (frequency, capacity, etc.)
 - Parameter values (numeric quantities, temporal expressions)
 - Scope identifiers (specific points of interest, vehicle types, or system-wide changes)
 NER is performed through structured output generation, where Gemini is prompted to return JSON-formatted entity extractions rather than free-form text responses [26].
3. **Configuration Management:** Extracted entities are validated against predefined schemas (using Pydantic) and applied to the simulation configuration file (YAML format shown in Appendix A). The agent maintains a conversational context allowing users to iteratively refine configurations through multi-turn dialogues.
4. **Response Generation:** The agent generates natural language confirmations of applied changes, explanations of system capabilities, and guidance for next steps. Response generation is contextualized by the current configuration state and conversation history.

2.3.2. Conversational Interface Implementation

The conversational interface is implemented using Chainlit [49], a Python framework for building production-ready conversational AI applications. Chainlit provides:

- **Session Management:** Persistent conversation histories that enable context-aware interactions across multiple user queries.
- **Streaming Responses:** Progressive display of agent responses as they are generated, improving perceived responsiveness.
- **Rich Media Support:** Display of structured data (configuration previews), images (results visualizations), and file downloads (simulation outputs).

The frontend interface, developed with React and TypeScript, communicates with the Python backend via RESTful API endpoints. This separation of concerns follows modern web application architecture patterns, facilitating independent development and testing of frontend and backend components.

2.3.3. Prompt Engineering and System Instructions

The effectiveness of the conversational agent depends critically on carefully designed system prompts that define the agent's role, capabilities, and behavioral guidelines. The system prompt includes:

1. **Role Definition:** Clear specification that the agent is a specialized assistant for urban mobility simulation configuration, not a general-purpose chatbot.
2. **Domain Knowledge:** Contextual information about the simulation system's capabilities and parameter constraints.
3. **Output Format Instructions:** Explicit templates for Structured Outputs (JSON schemas for entity extraction, YAML formatting for configuration updates).
4. **Safety and Validation Guidelines:** Instructions to validate user requests against system constraints and to alert users to potentially problematic configurations (e.g., the sum of the probabilities of visiting the different POIs for the simulated agents is above one and the configuration needs to be revised).

The system prompt (exposed in Appendix B) is version-controlled and subject to iterative refinement based on the evaluation of the agent performance across various user queries.

2.4. Simulation Generation

2.4.1. Simulation Execution Workflow

Once the user confirms a configuration via the conversational agent, the simulation execution workflow proceeds as follows:

1. **Configuration Loading:** The validated YAML configuration file is parsed and converted into Python data structures representing simulation parameters (vehicle schedules, capacities, pedestrian population sizes, temporal boundaries).
2. **Entity Initialization:** Based on configuration parameters, the simulation engine instantiates entity objects:
 - Bus and MetroGuagua vehicles are created according to scheduled service patterns.
 - Pedestrian agents are generated with stochastic origin–destination pairs sampled from distributions defined by POI locations and population density estimates.
3. **Event Scheduling:** Initial events (vehicle departures, pedestrian activity starts) are scheduled in the SimPy environment's event queue.
4. **Simulation Loop:** The SimPy environment processes events in chronological order, advancing the simulation clock. At each event:
 - Entity states are updated (positions, occupancies).
 - New events are scheduled (next stop arrivals, passenger arrivals).
 - Interactions are resolved (boarding decisions, capacity constraints).
5. **Data Logging:** Throughout the simulation, state variables are logged to capture:
 - Vehicle trajectories (sequence of edge traversals with timestamps).
 - Occupancy levels (passenger counts at each stop and along route segments).
 - Pedestrian paths (origin, destination, mode choices, travel times).
 - System-level metrics (total vehicle-kilometers traveled, kilometers traveled, average wait times).
6. **Output Generation:** Upon simulation completion, logged data is post-processed and exported in formats suitable for visualization and analysis.

2.4.2. Asynchronous Simulation Management

To accommodate potentially long-running simulations without blocking user interactions, the system implements asynchronous task execution using AWS Fargate (ECS). When a simulation is initiated:

1. A unique simulation ID is generated and returned immediately to the user.
2. The simulation task is dispatched to a Fargate container where it executes independently.
3. The user can query simulation status via the conversational agent using the simulation ID.
4. Upon completion, results are packaged and made available for download as compressed archives (ZIP format containing GeoJSON and Excel files).

This architecture prevents interface freezing during computation and supports concurrent execution of multiple simulation scenarios.

2.5. Visualization and Decision Support

Geospatial Visualization with Kepler.gl

Simulation outputs are visualized using Kepler.gl (see Figure 2), an open-source geospatial analysis tool developed by Uber [50]. Kepler.gl provides:

- **Multi-layer Rendering:** Simultaneous display of multiple geospatial datasets (road networks, bus routes, pedestrian flows) as distinct layers with customizable styling.
- **Temporal Animation:** Time-series data (vehicle trajectories over the simulation period) can be animated to visualize traffic dynamics and congestion patterns.
- **Interactive Filtering:** Users can filter data by attributes (e.g., vehicle type, occupancy level, time of day) to focus analysis on specific subsets.
- **Data Aggregation:** Spatial aggregation functions (e.g., hexagonal binning) summarize point-level data into regional statistics, revealing spatial patterns in mobility demand.




Figure 2. Interactive map made using Kepler.gl and the results of a simulation. The yellow dots correspond to metroguaguas' stops, while red dots are bus stops. The coloured lines represented are the different bus routes and the planned MetroGuagua itinerary (in yellow). Each element can be clicked to read its properties. The user can interact with the layers and watch the simulated elements move across the city using controls similar to most video players.

Simulation outputs are exported as GeoJSON files, a widely supported standard for encoding geographic features with associated properties. Each GeoJSON file contains:

- **Geometry:** LineString geometries for vehicle/pedestrian trajectories, point geometries for stops and POIs.

- **Properties:** Attributes including entity IDs, timestamps, occupancy levels, route identifiers, and sustainability indicators.

It is important to note that the architecture does not currently feature an automated cloud API integration with geospatial platform providers (such as Foursquare Studio). Instead, users manually upload the exported GeoJSON files to the standard Kepler.gl web interface, where they can interactively explore results and generate custom visualizations tailored to specific decision-making questions. The visualization layer is intentionally decoupled from the core simulation pipeline. This architectural boundary was established to computationally isolate the asynchronous simulation engine (AWS Fargate) from the high-memory demands of browser-based 3D geospatial rendering, ensuring system stability during complex multi-scenario testing. While this manual interaction step slightly limits the fully automated experience, it provides the necessary robust backend validation for TRL-4. Subsequent TRL-5 iterations will resolve this by natively embedding the Kepler.gl React component directly within the Chainlit frontend, achieving seamless end-to-end automation.

2.6. Implementation Technologies and Infrastructure

Table 2 summarizes the key technologies and infrastructure components employed in the system implementation.

Table 2. Summary of implementation technologies and infrastructure.

Component	Technology/Tool	Version/Provider
Programming Language	Python	3.11
LLM for Conversational AI	Gemini 2.5 Flash Lite	Google Vertex AI
Conversational Framework	Chainlit	2.6.2
Discrete-Event Simulation	SimPy	4.1.1
Network Analysis	NetworkX, OSMnx	3.4.2, 2.0.5
Geospatial Processing	GeoPandas, Shapely	1.1.1, 2.1.1
Data Validation	Pydantic	2.11.7
Database	PostgreSQL with PostGIS	17.5
Containerization	Docker	29.0.2
Cloud Infrastructure	ECS Fargate, RDS, S3	AWS
Frontend Framework	React, TypeScript	19.1.0, 5.8

2.7. Evaluation Approach

Given the TRL 4 status of this proof-of-concept, evaluation focuses on demonstrating technical feasibility: specifically the system’s functional robustness, architecture orchestration, and computational performance. A rigorous empirical assessment of the simulation’s fidelity, including data-driven calibration against physical sensors, benchmarking against established platforms like SUMO, and comprehensive user-centered evaluations through participatory planning workshops, falls outside the scope of this initial prototype. As detailed in the limitations section, these empirical validations are reserved for subsequent TRL 5-7 deployment phases.

3. Results

This section details the experimental evaluation of the proposed generative AI-driven digital twin architecture. The evaluation focuses on the technical feasibility of the system, its functional robustness as a decision-support tool, and the efficiency of the conversational interface in managing complex urban mobility simulations. Rather than focusing solely on mobility outcomes, the presented results emphasize seamless integration of large language models (LLMs) with discrete-event simulation (DES) engines. The evaluation is structured

to assess the accuracy of the conversational agent in parameter mapping, the reliability of the evaluation layer, and the overall performance metrics of the decoupled architecture.

3.1. Case Study: Las Palmas de Gran Canaria

The system was validated using a case study focused on Las Palmas de Gran Canaria, a major European mid-sized and the capital city of Gran Canaria in the Canary Islands, Spain. Las Palmas presents an ideal testbed for urban mobility analysis due to several challenges and structural characteristics:

- **Complex Morphological and Linear Constraints:** The city exhibits a pronounced linear morphology and includes a prominent northern peninsula that hosts the industrial harbor, residential areas, and major tourist attractions. This geographic configuration concentrates most north–south transit flows into a limited number of critical corridors, particularly the isthmus connecting the peninsula to the mainland. As a result, simulating new transit systems such as the *MetroGuagua* is especially relevant for congestion management.
- **Tourist–Resident Dynamics:** The city manages two overlapping mobility layers: a stable resident population of approximately 380,000 people and a high tourist influx with distinct travel patterns and preferences.
- **Multimodal Public Transit System:** The current infrastructure relies on an extensive bus network. The ongoing implementation of the *MetroGuagua* (a high-capacity system with dedicated lanes and metro-like frequencies) provides a perfect ‘what-if’ scenario for our architecture, as it requires planners to reconfigure existing bus lines and evaluate multimodal connectivity.
- **Standardized Data Availability:** The choice of this city is also strategic for its adherence to Spanish and European data standards. By mainly using open institutional data, complemented by proprietary traffic statistics, the case study serves as a foundational blueprint for replicability across other Spanish and European municipalities.

The study area encompasses the central urban core, covering approximately 100 km². The digital twin represents a network of more than 40 active bus routes and the planned 12 km *MetroGuagua* line, comprising 20 dedicated stops [51]. For the purposes of this evaluation, we simulated a standard 19 h operational day (05:00 to 24:00), allowing for the observation of peak-hour dynamics and transition periods. This scenario provides enough complexity to test whether the conversational agent can interpret high-level planning instructions and translate them into a valid, executable simulation configuration.

3.2. Evaluation of Conversational Configuration

A primary objective of this evaluation was to demonstrate the conversational AI agent’s capability to interpret natural language requests and correctly modify simulation parameters. We conducted structured tests that encompassed not only common configuration tasks that urban planners might perform, but also typical error scenarios. During the development of the system, we held formal consultation meetings with the Mobility Department of the Municipal Public Company Sagulpa, responsible for the Mobility and Street Cleaning of the City of Las Palmas de Gran Canaria, ensuring alignment with real operational needs and public service requirements.

3.2.1. Example Interaction 1: Viewing Current Configuration

User: “Show me the current configuration settings.”

Agent Response: The agent successfully retrieved and displayed the YAML configuration file content, presenting key parameters including:

- Bus service headway: 10 min (peak), 15 min (off-peak)

- MetroGuagua service headway: 5 min (peak), 8 min (off-peak)
- Bus capacity: 80 passengers
- MetroGuagua capacity: 150 passengers
- Simulation temporal bounds: 05:00–24:00

The agent formatted the configuration in a human-readable structure with clear labeling of parameter categories (POI opening hours, specifications by vehicle type, temporal settings, etc.).

3.2.2. Example Interaction 2: Modifying Service Headway

User: “Change the MetroGuagua headway to every 10 min.”

Agent Processing:

1. Intent Recognition: Identified modification intent.
2. Entity Extraction:
 - Target parameter: `metroguagua.headway`.
 - New value: 10 min.
 - Scope: MetroGuagua headway only.
3. Configuration Update: Modified YAML file to set headway inside the `metroguagua` section to 10.
4. Confirmation: “I have updated the MetroGuagua headway to 10 min. The new setting will be applied in the next simulation run.”

Evaluation confirmed that the YAML file was correctly modified, and subsequent simulation execution used the updated parameter.

3.2.3. Example Interaction 3: Complex Multi-Parameter Adjustment

User: “I want to test a scenario where buses run less frequently but have higher capacity. Set bus headway to 20 min and capacity to 100 passengers.”

Agent Processing:

1. Recognized compound modification intent affecting two parameters.
2. Extracted entities:
 - `buses.headway: 20`
 - `buses.capacity: 100`
3. Applied both modifications atomically (ensuring configuration consistency).
4. Confirmed: “Configuration updated: Bus headway set to 20 min, capacity increased to 100 passengers.”

Table 3 summarizes the conversational agent’s performance across 50 structured test queries covering diverse interactions with the configuration and simulation engine of the digital twin.

Table 3. Conversational AI agent performance on configuration modification tasks.

Task Category	Test Cases	Success Rate (%)
View current configuration	10	100
Modify single parameter	15	93.3
Modify multiple parameters	10	90.0
Launch the simulation using the modified configuration	10	100
Detect invalid or ambiguous requests	5	100
Overall	50	96

The agent demonstrated strong performance in interpreting user intentions and correctly extracting entities (some examples are shown in Table 4). Errors primarily occurred in ambiguous queries lacking explicit parameter names (e.g., ‘make buses run faster’ without specifying whether to adjust frequency or speed limits). In such cases, the agent requested clarification rather than making incorrect assumptions, which is the desired behavior for safety-critical configuration management. It also handled out of scope requests, like adding new vehicle types or trying to modify the bus itineraries.

Table 4. Examples of Conversational Agent Processing and Configuration Mapping.

User Prompt	Recognized Intent	Extracted Entities/Parameters	Validation Result & Final Action
“Show me the current configuration settings.”	View_Configuration	None	Success: Retrieved current YAML state and displayed formatted parameters to the user.
“Change the MetroGuagua headway to 10 min.”	Modify_Parameter	Target: metroguaagua.headway Value: 10	Success (Pydantic): Value is within valid integer bounds. YAML updated and user notified.
“Set bus headway to 20 min and capacity to 100.”	Modify_Multiple_Parameters	Target 1: buses.headway (Value: 20) Target 2: buses.max_capacity (Value: 100)	Success (Pydantic): Both values verified. Atomic update applied to YAML file.
“Make buses run faster.”	Ambiguous_Request	Target: Unknown	Clarification Requested: Agent prompts user to specify whether to modify speed limits or headways.

3.3. System Performance and Computational Efficiency

A critical component of the architecture is the safety/validation layer implemented via Pydantic. This layer ensures that the stochastic nature of the LLM does not result in broken simulations. During the testing phase, the system monitored the types of errors caught by the validator before the simulation was triggered.

The most frequent validation triggers were:

1. Type Consistency: Attempting to pass string values to integer fields (e.g., ‘high’ as maximum bus capacity).
2. Probability Constraints: Ensuring that the sum of probabilities for POI visits does not exceed 1.0. This was the most usual case, since most times the user does not know that, in order to increase the probability of agents visiting a group of POIs, it is needed to reduce the probability of other groups.
3. Domain Bounds: Restricting parameters like the speed ratio of buses vs other vehicles to a logical range of [0, 1].

This decoupled approach allows the simulation API to remain agnostic of the AI’s internal state, receiving only verified and sanitized YAML configurations. However, it is crucial to distinguish between technical and empirical validity. While the Pydantic layer guarantees that a configuration is syntactically correct, type-safe, and within defined numerical bounds (technical validity), it cannot guarantee that the user’s requested scenario makes sense from a transportation-planning perspective (empirical validity). A configuration can be technically flawless while representing an entirely unrealistic mobility scenario; assessing empirical realism remains the responsibility of the human planner.

To assess the practical utility for urban planners, we measured the end-to-end time of the ‘what-if’ cycle to assess the computational performance of the system and evaluate its scalability potential. The metrics were averaged over five full simulation runs (simulating a standard 19 h operational day).

- **Simulation Execution Time:** Average execution time for a full 19 h simulation: 14.2 min (SD = 2.1 min) on AWS Fargate containers with 2 vCPUs and 4 GB RAM.
- **Conversational Agent Latency:** Average response time for configuration queries: 3.8 s. Parameter modification confirmation: 6.3 s (including YAML file I/O).
- **Visualization Loading:** Kepler.gl visualization load times: 10–20 s for datasets with more than 50,000 trajectory points using standard client-side rendering within the open-source Kepler.gl web application.

These performance characteristics indicate that the system operates within acceptable latency bounds for interactive use by urban planners, supporting iterative scenario exploration without prohibitive wait times.

It is important to note that while the AWS Fargate deployment is architecturally designed for high scalability, evaluating the system under massive concurrent workloads (stress testing) falls outside the scope of this proof-of-concept. The current results validate the reliability of the asynchronous dispatch and retrieval mechanism for single and sequential user requests. Comprehensive load testing, queue behavior analysis, and auto-scaling performance are reserved for future operational deployment phases.

Ablation Study and Technical Calibration

To empirically evaluate the architectural necessity of the decoupled validation layer, we conducted an ablation study utilizing a benchmark of 50 standardized urban planning queries. We compared the full architecture against an ablated configuration where the Structured Output of the Gemini 2.5 Flash Lite model is fed directly into the SimPy simulation engine without Pydantic schema constraints.

As summarized in Table 5, the ablated system exhibited a 12% failure rate in basic syntax parsing (e.g., malformed JSON structures or broken lists) and an additional 12% failure rate due to semantic boundary violations (e.g., negative transit headways or invalid time windows). These formatting anomalies caused immediate execution crashes in the underlying discrete-event simulation engine, yielding a final simulation execution rate of only 76%. Conversely, the full architecture successfully intercepted and corrected these structural hallucinations via Pydantic type-coercion and schema validation, achieving a 100% simulation execution rate at the expense of a negligible latency overhead (≈ 0.30 s).

Table 5. Ablation study metrics over 50 evaluation planning queries.

Architecture Configuration	Syntactic Success	Semantic Validation	Simulation Execution	Avg. Pipeline Latency
Ablated System (No Pydantic)	88% (44/50)	76% (38/50)	76% (38/50)	≈ 6.00 s
Full Architecture (With Pydantic)	100% (50/50)	100% (50/50)	100% (50/50)	≈ 6.33 s

4. Discussion

This section interprets the results presented in Section 3, analyzes the implications of the integrated AI-digital twin architecture for smart city planning, compares the proposed approach with prior work, and acknowledges the limitations inherent to this TRL 4 proof-of-concept.

4.1. Interpretation of Results

The evaluation results demonstrate that the proposed generative AI-driven digital-twin-oriented architecture successfully addresses key challenges in urban mobility planning through three primary mechanisms: accessibility enhancement, flexibility in scenario exploration, and evidence-based decision support.

4.1.1. Accessibility Through Natural Language Interaction

The conversational AI agent achieved a 96% success rate in interpreting and executing configuration modification tasks, demonstrating that natural language interfaces can effectively lower technical barriers to sophisticated simulation tools. This finding is particularly significant given that traditional digital twin platforms require specialized technical expertise in programming, geospatial analysis, and simulation modeling [20]. By abstracting technical complexity behind a conversational interface, the system enables urban planners, policymakers, and domain experts without computational backgrounds to directly engage with simulation-based analysis.

A key factor in this success is the integration of Pydantic as a validation layer. Although LLMs are inherently stochastic, the system ensures that the agent's output is strictly mapped to the required technical schemas before any tool is executed. The two observed failure cases (4.0% of queries) involved minor entity extraction errors during complex parameter modifications, which were safely caught by the validation layer. Conversely, when users submitted genuinely ambiguous requests, the agent successfully detected the ambiguity and requested clarification rather than hallucinating values (achieving a 100% success rate in these specific test cases). This behavior, requesting clarification aligns with responsible AI principles for safety-critical applications [52]. This prevents potentially harmful misconfigurations that could lead to invalid simulation results.

4.1.2. Flexibility in Scenario Exploration

The ability to rapidly configure and execute diverse scenarios represents a substantial improvement over conventional planning workflows. Traditional approaches often involve manual editing of configuration files and lengthy data pre-processing, processes that can take hours and are prone to human error [6]. In contrast, the proposed architecture enabled scenario modifications to be specified, validated, and executed within minutes, supporting iterative exploration of the urban mobility design space.

This speed allows for an iterative exploration of trade-offs between environmental sustainability, economic efficiency, and social equity. For instance, testing the impact of *MetroGuagua* frequencies against conventional bus service quality becomes a matter of minutes, facilitating a future use of multi-objective decision support tools that make these relationships transparent to stakeholders [53].

4.1.3. Evidence-Based Decision Support

The integration of diverse data sources, ranging from INE census data and ISTAC tourism statistics to OSM infrastructure, provides a robust baseline for evidence-based decisions. By using standardized open data formats, the architecture proves to be highly replicable across other Spanish and European cities.

The simulation results, such as identifying accessibility gaps (e.g., limit the distance pedestrians are willing to walk to transit stops), transform raw open data into actionable insights. This capability allows planners to assess whether proposed interventions, such as the introduction of the *MetroGuagua* corridor, contribute meaningfully to city-wide sustainability and equity targets [54]. It can also help identify specific geographic areas where additional transit stops, demand-responsive services, or pedestrian infrastructure

improvements could enhance equity. Without simulation-based analysis, such service gaps might remain undetected until after infrastructure investments are made, when remediation is more costly [55].

4.2. Value of the Integrated Architecture

The integration of generative AI with digital twin technology creates a complementary system in which each component enhances the other's capabilities. The conversational AI layer transforms the digital twin from a passive simulation tool into an interactive environment by interpreting high-level planning objectives and translating them into executable parameters. This process is constrained through schema validation mechanisms, ensuring logical consistency and reducing the risk of erroneous outputs while preserving the flexibility of natural language interaction [26].

At the same time, the digital twin grounds AI-generated configurations in physical and behavioral realities, enforcing constraints such as network topology, vehicle capacity, and traffic dynamics. This reduces the risk of technically invalid configurations being executed, although ensuring the empirical realism of the final scenario ultimately requires human domain expertise [56]. Together, this integration enables a closed-loop decision support process in which planners iteratively define goals, simulate outcomes, and refine strategies. By significantly accelerating this feedback cycle compared to traditional methods, the system aligns with participatory planning approaches while improving efficiency and responsiveness [57].

4.3. Comparison with Prior Work

The proposed architecture extends prior research in digital twins, urban AI, and conversational systems by integrating these domains into a unified, accessible framework. Rather than introducing a methodological breakthrough in base AI models or simulation logic, the primary scientific value lies in the engineering and orchestration of existing state-of-the-art tools (e.g., Gemini, Pydantic, SimPy, NetworkX, and AWS). Traditional urban mobility digital twins, such as those based on SUMO [17] or MATSim [18], offer strong simulation capabilities but typically require technical expertise and manual configuration. Even more advanced implementations, including Virtual Singapore [58], rely on conventional interfaces that can limit accessibility. In contrast, this approach leverages a conversational interface powered by large language models, combined with automated pipelines based on open data sources (e.g., INE, OSM, GTFS), enabling more flexible interaction while improving replicability across municipalities.

Within the broader landscape of AI in urban planning, existing approaches often focus on predictive analytics or generative design [4,23], frequently operating as opaque 'black-box' systems. Here, AI is instead positioned as an orchestration layer that translates user intent into validated simulation configurations. The incorporation of schema validation mechanisms ensures that generated outputs remain technically consistent, preserving human oversight while reducing barriers to use [59]. This paradigm also advances the role of conversational AI, which has seen limited application in urban planning beyond data querying tools such as ODIN [60], by enabling the management of complex, stateful simulation workflows rather than simple query-response interactions.

These contributions have several implications for smart city development. By lowering technical barriers, the system could support broader participation in future deployments, aligning with principles of collaborative governance [37]. Its ability to rapidly generate and evaluate scenarios facilitates more adaptive and iterative policymaking in dynamic urban contexts [5]. At the same time, reliance on standardized open data enhances transparency and reproducibility, supporting evidence-based decision-making and cross-city comparabil-

ity [57]. Finally, the system's modular, API-driven design provides a foundation for future real-time digital twin implementations, where integration with live urban data streams could enable responsive, operational decision support in evolving city environments [15].

4.4. Principal Findings and Implications for Smart City Development

The evaluation study yielded several key findings with implications for smart city planning:

- **Lowering Technical Barriers through Accessibility:** Natural language interfaces effectively lower the technical barriers to sophisticated urban planning tools. The conversational agent enables non-technical users to configure and trigger complex simulations, a task that traditionally required specialized programming expertise in simulation frameworks and geospatial analysis.
- **Replicability through Core Open Data:** A major strength of the architecture is its reliance on standardized data sources such as OSM, GTFS, and the Spanish INE. While the high-resolution traffic calibration relies on proprietary data (e.g., TomTom), this open data foundation ensures that the core methodology is not only robust for the case study of Las Palmas and partly replicable across other Spanish and European cities, facilitating a standardized approach to data-driven urban management.
- **Quantifiable Trade-offs for Decision Support:** While urban planning involves inherent conflicts between objectives (e.g., service frequency vs. operational costs), our system provides the evidence needed to make these trade-offs explicit. By allowing planners to compare scenarios, such as the impact of the *MetroGuagua* corridor on existing bus lines, the data-driven simulation platform supports more informed and transparent decision-making processes.
- **Identification of Social Equity Gaps:** The high spatial granularity of the simulation allows for the detection of accessibility gaps. By analyzing pedestrian flows and transit proximity, the system can identify neighborhoods falling below critical service thresholds (e.g., imposing a walking limit). This capability is vital for ensuring that smart city interventions promote social equity before infrastructure is physically deployed.
- **Agile Iteration and Exploration:** The efficiency of the system—capable of processing a 19 h urban day in less than 20 min—supports an agile policymaking cycle. This rapid iteration allows planners to explore a vast design space of 'what-if' scenarios, identifying robust strategies that can adapt to changing urban conditions.

The validated architecture demonstrates that generative AI and digital twin technology can be integrated to create decision support tools that are simultaneously sophisticated, accessible, and actionable. This combination addresses a critical barrier to smart city development: the gap between data-driven insights and practical implementation.

By potentially enabling broader participation in future planning processes, the system aligns with principles of inclusive governance and could lead to mobility policies that better reflect diverse community needs [37]. Furthermore, by providing evidence-based analysis of sustainability impacts, the system supports cities in achieving climate, economic, and equity goals.

4.5. Limitations and Future Work

As a TRL 4 proof-of-concept, this work demonstrates technical feasibility in a controlled environment but possesses specific limitations that outline the roadmap for future operational deployment in the Smart City ecosystem.

The proposed system presents several limitations and opportunities for improvement related to data quality, simulation fidelity, scalability, user evaluation, and governance. The simulation relies on publicly available data sources such as OpenStreetMap (OSM)

and GTFS, which, although comprehensive, may include inaccuracies or omissions. OSM coverage varies across regions, and GTFS data represents scheduled rather than actual service conditions [38]. For operational deployment, integration with higher-quality and continuously updated data sources, such as real-time vehicle tracking, automated passenger counting systems, and validated geospatial datasets maintained by transportation authorities, would be essential. The current decoupled architecture is designed to accommodate such real-time data streams from urban sensors. Additionally, the use of synthetic pedestrian populations simplifies real-world mobility demand. While agent-based modeling can capture heterogeneous travel behaviors [41], accurate calibration requires extensive empirical data, including travel surveys, mobile phone records, and transit smart card data, which were not available for this proof-of-concept. Future efforts should therefore prioritize data-driven calibration and validation against observed mobility patterns.

In terms of simulation fidelity, the digital twin incorporates simplifying assumptions to ensure computational tractability. For example, the V_{bus} calibration factor ($0.7 \times V_{median}$) serves as a practical heuristic at TRL 4 but does not capture micro-scale influences such as weather conditions, driver behavior, or intersection-specific delays. Similarly, pedestrian routing is based on shortest-path algorithms that do not account for qualitative factors such as sidewalk width, lighting conditions, or terrain characteristics like the steep topography of volcanic cities [61]. Enhancing fidelity through the integration of microscopic traffic models, more nuanced pedestrian behavior, and dynamic multimodal interactions would improve realism, though at the cost of increased computational demands. Balancing accuracy and responsiveness remains a key design trade-off.

Scalability also represents a critical challenge. While the system effectively handles a study area of approximately 100 km² with thousands of agents, extending it to metropolitan-scale environments with millions of entities will require further optimization. This includes adopting distributed computing approaches, hierarchical simulation techniques, and more efficient data structures to support large-scale analyses [16].

From a user-centered perspective, although the initial evaluation of the system was based on structured testing and was informed by multiple consultations with the municipal mobility company (Sagulpa) during the planning and development phases, comprehensive evaluation with end users has not yet been conducted. Future work should include usability studies, participatory planning workshops, and real-world deployment scenarios, as well as improved interoperability with existing municipal tools such as GIS platforms and project management systems. These steps are necessary to assess and enhance the system's practical impact on urban decision-making processes.

Furthermore, while the digital twin architecture is designed to support comprehensive sustainability assessments, evaluating specific environmental and economic indicators (such as emission proxies or detailed operational costs) remains a future application. Subsequent phases of development will integrate explicit sustainability models to fully exploit the platform's decision-support capabilities.

Finally, the integration of generative AI into urban planning workflows raises important ethical and governance considerations, particularly regarding transparency and potential algorithmic bias [52]. The current architecture adopts a human-in-the-loop approach, whereby the AI agent proposes configurations that must be explicitly validated by the user. However, there remains a need for greater transparency in documenting the agent's internal prompts and decision-making logic. Addressing these concerns will require the development of robust AI governance frameworks, including transparent documentation practices, algorithmic auditing, and safeguards to ensure that AI-assisted planning supports equitable outcomes and does not disproportionately disadvantage specific communities [59].

Finally, the integration of generative AI into urban planning workflows raises important ethical and governance considerations, particularly regarding transparency and potential algorithmic bias [52]. Furthermore, while the Pydantic validation layer successfully mitigates syntactic hallucinations (e.g., structural JSON errors or invalid data types), it cannot prevent semantic hallucinations. An LLM might generate a technically valid configuration that is practically disastrous (for instance, recommending a highly inefficient MetroGuagua route or a 1 s service headway that creates gridlock). If deployed without strict oversight, such AI-induced errors could lead to costly municipal infrastructure mistakes and severe service disruptions. Consequently, the architecture strictly mandates a “human-in-the-loop” governance model. The AI functions exclusively as an advisory orchestration tool; human mobility experts must critically evaluate the empirical plausibility of the generated configurations before any policy decisions or capital investments are made. However, there remains a need for greater transparency in documenting the agent’s internal prompts and decision-making logic. Addressing these concerns will require the development of robust AI governance frameworks, including transparent documentation practices, algorithmic auditing, and safeguards to ensure that AI-assisted planning supports equitable outcomes and does not disproportionately disadvantage specific communities [59].

5. Conclusions

This paper has presented a novel architecture integrating generative artificial intelligence with digital twin technology for urban mobility simulation and decision support in smart cities. The system addresses a critical gap in existing planning tools: the lack of accessible, flexible interfaces that enable diverse stakeholders to engage with sophisticated simulation-based analysis without requiring specialized technical expertise.

5.1. Summary of Contributions

The key contributions of this work are:

1. **Integrated Architecture (Answering RQ1):** We developed and demonstrated an architecture that synergistically combines a conversational AI agent (based on Google Gemini 2.5 Flash Lite) with a digital twin simulation engine to enable natural language configuration of multimodal urban mobility simulations. This integration answers RQ1 by demonstrating how complementary technologies can be composed to create capabilities exceeding what either component could achieve independently.
2. **Conversational Configuration Interface (Answering RQ3):** We implemented and tested a conversational AI agent capable of interpreting natural language planning objectives, extracting relevant simulation parameters, and autonomously modifying configuration files. Addressing RQ3, the agent achieved 96% accuracy across 50 diverse configuration tasks, proving that a natural language interface significantly lowers the technical barriers and operational latency required to execute complex urban simulations.
3. **Data-Rich Multimodal Simulation:** We developed a digital-twin-oriented prototype that integrates heterogeneous data sources (including INE census tracts for resident density, ISTAC and OSM for tourist activity, and TomTom traffic statistics) to create a realistic simulation environment. The system models buses, the future *MetroGuagua* (BRT), and pedestrian agents with differentiated behavioral profiles operating on a realistic road network extracted from OpenStreetMap. This setup provides a robust foundation for generating comprehensive outputs, such as vehicle trajectories and occupancy dynamics, which are essential for assessing sustainability indicators across environmental, economic, and social dimensions.

4. **Proof-of-Concept Demonstration (Answering RQ2) (TRL 4):** We conducted a technical evaluation of the framework in Las Palmas de Gran Canaria, demonstrating the feasibility of LLM-driven orchestration of discrete-event simulations. By executing multiple scenarios, we confirmed the technical feasibility of seamlessly integrating conversational AI with a discrete-event urban simulation engine. The evaluation directly addresses RQ2 by highlighting the effectiveness of the Pydantic-based safety layer, which achieved a high success rate in mapping natural language to executable configurations, ensuring system stability despite the stochastic nature of the generative agent.

5.2. Roadmap

Several directions for future research and development are evident:

1. **Transition to TRL 5-7:** Advance the system toward operational deployment through validation in real-world settings with actual planning agencies. This transition requires integration with authoritative data sources, comprehensive user evaluation, and refinement based on stakeholder feedback.
2. **Enhanced Simulation Fidelity:** Incorporate more detailed models of traffic dynamics, pedestrian behavior, and multimodal interactions. Integrate real-time data streams (GPS tracking, sensor networks) to enable operational decision support beyond scenario planning.
3. **Multi-City Generalization:** Validate the architecture across diverse urban contexts (varying sizes, geographic settings, transportation systems) to assess generalizability and identify necessary adaptations for different planning environments.
4. **Integration with Broader Planning Ecosystems:** Develop APIs and data exchange standards to integrate the system with existing GIS platforms, transportation demand models, and project management tools used by planning agencies.
5. **Ethical AI Frameworks:** Establish governance frameworks addressing algorithmic transparency, bias mitigation, data privacy, and accountability for AI-assisted planning decisions [52].
6. **Economic and Social Impact Assessment:** Conduct longitudinal studies to measure the actual impact of AI-digital twin decision support on planning outcomes, policy quality, and community well-being.

5.3. Closing Remarks

As urban populations continue to grow and cities face mounting pressures related to climate change, resource constraints, and social equity, the need for intelligent, evidence-based planning tools becomes increasingly urgent. This work demonstrates that the integration of generative artificial intelligence with digital twin technology offers a promising pathway toward more accessible, flexible, and effective decision support for smart city mobility planning.

By transforming sophisticated simulation engines from specialist tools into systems that respond to natural language, we can lower technical barriers to data-driven insights and potentially support broader participation in shaping the future of urban mobility. By grounding AI-generated recommendations in rigorous simulation, we ensure that decision support is based on realistic assessments of physical and behavioral dynamics rather than abstract speculation.

The validated proof-of-concept represents a significant step toward this vision, but substantial work remains to transition from laboratory demonstration to operational deployment. We hope that this research inspires continued innovation at the intersection of AI, digital twins, and smart cities, ultimately contributing to more sustainable, equitable, and livable urban environments for all.

Author Contributions: Conceptualization, E.S.-O., M.S.-M. and E.W.-S.; methodology, A.C.-R., M.Á.G.-E. and P.V.-M.; software, A.C.-R.; validation, M.Á.G.-E., P.V.-M. and M.S.-M.; writing—original draft preparation, E.S.-O., A.C.-R. and P.V.-M.; writing—review and editing M.S.-M. and P.V.-M. All authors have read and agreed to the published version of the manuscript.

Funding: This work has been carried out within the framework of the Spain Living Lab project (Grant Reference 1/1/2024-0412093852—SLLC16-01), funded by the Canarian Agency for Research, Innovation and the Information Society (ACIISI), Department of Universities, Science, Innovation and Culture of the Government of the Canary Islands, under the RETECH Programme, contributing to milestones 251, 252 and 253 of Component 16 of the Recovery, Transformation and Resilience Plan (PRTR), and co-funded by the European Union—Next Generation EU.

Data Availability Statement: The data and code supporting the findings of this study are available from the corresponding authors upon reasonable request.

Conflicts of Interest: Authors Pablo Vicente-Martínez and Adrián Chust-Ros were employed by the company SPV Scala. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Appendix A. Default Configuration File

The following segment illustrates the default values of the parameters to configure the simulations. This specification establishes the number of agents simulated, the opening hours of the points of interest and other relevant parameters.

Listing A1. YAML Configuration Structure.

```
metroguagua:
  start_time: '05:00'
  end_time: '24:00'
  max_capacity: 150
  headway_min:
    peak: 5
    off_peak: 8
  stop_duration_min: 4

buses:
  simulation_type: 'schedule'
  sampling_interval: 5
  stop_time_seconds: 20
  max_capacity: 80
  headway_min:
    peak: 10
    off_peak: 15
  speed_ratio_vs_car: 0.7

pedestrians:
  location: 'Las Palmas de Gran Canaria, Spain'
  minutes_per_day: 1140
  start_date: '2024-01-25T00:00:00'

n_tourists: 1000
n_locals: 1000

max_visits:
  tourist: 8
  canarian: 3
```

```
query_radius: 10000
zone_radius: 800

mode_weights:
  metroguagua: 0.7
  bus: 0.3

street_capacity: 200

walk_speed:
  mean: 80
  std: 10
  min: 40
  max: 120

time_windows:
  lunch: ['13:00', '15:00']
  dinner: ['21:00', '23:00']
  daylight: ['06:00', '18:00']
  night: ['8:00 PM', '12:00 AM']
  school: ['8:00 AM', '6:00 PM']
  clinic: ['9:00 AM', '5:00 PM']
  pharmacy: ['8:00 AM', '8:00 PM']

local_restaurant_rate: 0.25

poi_avg_times:
  school: '05:00'
  hospital: '01:00'
  clinic: '01:00'
  pharmacy: '00:15'
  post_office: '00:15'
  library: '02:00'
  cafe: '00:45'
  bar: '02:00'
  restaurant: '01:30'

# Tourist-only amenities
casino: '02:00'
theater: '03:00'
movie theater: '02:00'
nightclub: '04:00'

# Shops
supermarket: '00:30'
bakery: '00:20'
gift shop: '00:10'

# Nature / leisure
beach: '03:00'
park: '02:00'
water_park: '02:00'
stadium: '03:00'
```

```

# Tourism
museum:      '01:30'
artwork:     '00:30'
attraction:  '01:00'
viewpoint:   '00:15'
aquarium:    '01:00'

daylight_types:
- park
- viewpoint
- artwork
- attraction
- aquarium
- beach
- museum
- water_park

night_types:
- nightclub
- bar
- casino
- theater
- cinema

capacities:
hotel:      100
apartment:  5
guest_house: 3
hostel:     60
motel:      40

```

Appendix B. Configuration Agent System Prompt

The subsequent appendix details the operational instructions employed to direct the functional conduct of the dialogue system. These directives delineate the agent's persona, permissible operations, and procedural limitations, guaranteeing coherent and systematic performance throughout the integrated framework. The prompt was originally in Spanish, but has been translated to facilitate its understanding.

Listing A2. Agent System Prompt.

```

You are an assistant specializing in managing configuration files for simulation
models.

You can help users:
- Check or display the initial or current configuration (with or without changes)
  (tool: view_actual_config)
- Update a configuration parameter (tool: change_config_value)
- Update multiple configuration parameters (tool: change_config_value, useful for
  some validation rules)
- View the structure of the configuration file (below)
- Create a simulation using the virtual redesign API (tool:
  virtual_redesign_simulation)
- Download the zip file containing the results of a previously run simulation (
  report of statistics and files for visualization in Kepler.gl) (tool:
  virtual_redesign_report_delivery)

```

Instructions:

- Use the tools as needed.
- Display the current configuration to the user when requested. Report any changes made and ask the user if they want to save the changes before finishing.
- Save the changes only if the user requests it.
- Report any errors received by the tools.
- To create a simulation, use the "virtual_redesign_simulation" tool and provide the user with the simulation ID and status. If the status is still in progress, you can tell the user that they can check back later to verify if the process has completed. Do not tell them that they will be notified when the simulation is complete.
- To obtain a zip file containing the report and files for viewing a completed simulation, use the "virtual_redesign_report_delivery" tool with the ID provided by the user. Inform the user of the simulation's status or whether the file downloaded successfully. Do not disclose the file path to the user.

Configuration structure:

```
{config_structure}
```

Use the available tools as needed and always explain the changes you make.

References

1. United Nations, Department of Economic and Social Affairs, Population Division. *World Urbanization Prospects: The 2018 Revision*; United Nations: New York, NY, USA, 2018.
2. Bibri, S.E.; Krogstie, J. Smart sustainable cities of the future: An extensive interdisciplinary literature review. *Sustain. Cities Soc.* **2017**, *31*, 183–212.
3. Yigitcanlar, T.; Kamruzzaman, L.; Foth, M.; Sabatini-Marques, J.; da Costa, E.; Ioppolo, G. Can cities become smart without being sustainable? A systematic review of the literature. *Sustain. Cities Soc.* **2019**, *45*, 348–365.
4. Liao, F.; Arentze, T.; Timmermans, H. Incorporating space-time constraints and activity-travel time profiles in a multi-state supernetwork approach to individual activity-travel scheduling. *Transp. Res. Part B Methodol.* **2013**, *55*, 41–58.
5. Batty, M. Digital twins. *Environ. Plan. B Urban Anal. City Sci.* **2018**, *45*, 817–820.
6. Pettit, C.; Bakelmun, A.; Lieske, S.N.; Glackin, S.; Hargroves, K.; Thomson, G.; Shearer, H.; Dia, H.; Newman, P. Planning support systems for smart cities. *City Cult. Soc.* **2018**, *12*, 13–24.
7. Deng, T.; Zhang, K.; Shen, Z.J. A systematic review of a digital twin city: A new pattern of urban governance toward smart cities. *J. Manag. Sci. Eng.* **2021**, *6*, 125–134.
8. Huzzat, A.; Anpalagan, A.; Khwaja, A.S.; Woungang, I.; Alnoman, A.A.; Pillai, A.S. A comprehensive review of Digital Twin technologies in smart cities. *Digit. Eng.* **2025**, *4*, 100040. <https://doi.org/10.1016/j.dte.2025.100040>.
9. Shahat, E.; Hyun, C.T.; Yeom, C. City digital twin potentials: A review and research agenda. *Sustainability* **2021**, *13*, 3386.
10. Lv, Z.; Chen, D.; Lou, R.; Song, H. Industrial security solution for virtual reality. *IEEE Internet Things J.* **2021**, *8*, 6273–6281.
11. Nag, D.; Brandel-Tanis, F.; Pramestri, Z.A.; Pitera, K.; Frøyen, Y.K. Exploring digital twins for transport planning: A review. *Eur. Transp. Res. Rev.* **2025**, *17*, 15. <https://doi.org/10.1186/s12544-025-00713-0>.
12. Qi, Q.; Tao, F.; Hu, T.; Anwer, N.; Liu, A.; Wei, Y.; Wang, L.; Nee, A.Y.C. Enabling technologies and tools for digital twin. *J. Manuf. Syst.* **2021**, *58*, 3–21.
13. Mbasso, W.F.; Harrison, A.; Dagal, I.; Jangir, P.; Khishe, M.; Kotb, H.; Shaikh, M.S.; Smerat, A.; Donfack, E.F.; Kumar, R. Digital twins in renewable energy systems: A comprehensive review of concepts, applications, and future directions. *Energy Strategy Rev.* **2025**, *61*, 101814. <https://doi.org/10.1016/j.esr.2025.101814>.
14. Ferré-Bigorra, J.; Casals, M.; Gangoellells, M. The adoption of urban digital twins. *Cities* **2022**, *131*, 103905.
15. Mohammadi, N.; Taylor, J.E. Smart city digital twins. In Proceedings of the 2017 IEEE Symposium Series on Computational Intelligence (SSCI), Honolulu, HI, USA, 27 November–1 December 2017; pp. 1–5.
16. Heinrichs, D.; Cyganski, R. Automated driving: How it could enter our cities and how this might affect our mobility decisions. *disP-Plan. Rev.* **2015**, *51*, 74–79.

17. Lopez, P.A.; Behrisch, M.; Bieker-Walz, L.; Erdmann, J.; Flötteröd, Y.P.; Hilbrich, R.; Lücken, L.; Rummel, J.; Wagner, P.; Wiessner, E. Microscopic traffic simulation using SUMO. In Proceedings of the 2018 21st International Conference on Intelligent Transportation Systems (ITSC), Maui, HI, USA, 4–7 November 2018; pp. 2575–2582.
18. Horni, A.; Nagel, K.; Axhausen, K.W. *The Multi-Agent Transport Simulation MATSim*; Ubiquity Press: London, UK, 2016.
19. SimPy Team. SimPy: Discrete-Event Simulation for Python. Available online: <https://simpy.readthedocs.io/> (accessed on 15 December 2024).
20. Francisco, A.; Mohammadi, N.; Taylor, J.E. Smart city digital twin-enabled energy management: Toward real-time urban building energy benchmarking. *J. Manag. Eng.* **2020**, *36*, 04019045.
21. Brown, T.B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language models are few-shot learners. In Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS 2020), Virtual, 6–12 December 2020; Volume 33, pp. 1877–1901.
22. OpenAI. GPT-4 Technical Report. *arXiv* **2023**, arXiv:2303.08774.
23. Cugurullo, F.; Acheampong, R.A.; Gueriau, M.; Dusparic, I. The transition to autonomous cars, the redesign of cities and the future of urban sustainability. *Urban Geogr.* **2021**, *42*, 833–859.
24. Liu, C.; Li, G.; Huang, C.; Herzog, O.; Ritter, H.; Ma, S.; Ye, Y. Generative AI for Urban Planning and Design: Progress Review and Future Perspectives. *Engineering* **2026**. <https://doi.org/10.1016/j.eng.2026.03.001>.
25. Nishant, R.; Kennedy, M.; Corbett, J. Artificial intelligence for sustainability: Challenges, opportunities, and a research agenda. *Int. J. Inf. Manag.* **2020**, *53*, 102104.
26. Liu, P.; Yuan, W.; Fu, J.; Jiang, Z.; Hayashi, H.; Neubig, G. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Comput. Surv.* **2023**, *55*, 1–35.
27. Google DeepMind. Gemini: A Family of Highly Capable Multimodal Models. *arXiv* **2023**, arXiv:2312.11805.
28. Zhang, R.; Yang, Z.; Dai, T.; Sun, D.J.; Ma, S.; Zhao, X. MAGAE: Multi-level Alignment over Aggregation Semantic Graph with Attribute Enhancement for Text-based Vehicle Retrieval. *IEEE Trans. Intell. Transp. Syst.* **2025**, *26*, 13704–13720. <https://doi.org/10.1109/TITS.2025.3566502>.
29. Chen, X.; Lin, M.; Schärli, N.; Zhou, D. Teaching large language models to self-debug. In Proceedings of the 12th International Conference on Learning Representations (ICLR 2024), Vienna, Austria, 7–11 May 2024.
30. Palanica, A.; Flaschner, P.; Thommandram, A.; Li, M.; Fossat, Y. Physicians’ perceptions of chatbots in health care: Cross-sectional web-based survey. *J. Med. Internet Res.* **2019**, *21*, e12887.
31. Huynh, A.L.; Roy, T.J.; Jackson, K.N.; Lee, A.G.; Liaw, W.; Hossain, M.M. Applications of artificial intelligence-based conversational agents in healthcare: A systematic umbrella review. *Int. J. Med. Inform.* **2026**, *207*, 106204. <https://doi.org/10.1016/j.ijmedinf.2025.106204>.
32. Chung, M.; Ko, E.; Joung, H.; Kim, S.J. Chatbot e-service and customer satisfaction regarding luxury brands. *J. Bus. Res.* **2020**, *117*, 587–595.
33. Kar, A.K.; Haldar, N. Applying chatbots to the Internet of Things: Opportunities and architectural elements. *Int. J. Inf. Manag.* **2016**, *36*, 1233–1238.
34. Shum, H.Y.; He, X.D.; Li, D. From Eliza to XiaoIce: Challenges and opportunities with social chatbots. *Front. Inf. Technol. Electron. Eng.* **2018**, *19*, 10–26.
35. Wang, W.; Siau, K. Artificial intelligence, machine learning, automation, robotics, future of work and future of humanity: A review and research agenda. *J. Database Manag.* **2019**, *30*, 61–79.
36. White, G.; Zink, A.; Codecà, L.; Clarke, S. A digital twin smart city for citizen feedback. *Cities* **2021**, *110*, 103064.
37. Cardullo, P.; Kitchin, R. Being a ‘citizen’ in the smart city: Up and down the scaffold of smart citizen participation in Dublin, Ireland. *GeoJournal* **2019**, *84*, 1–13.
38. Boeing, G. OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Comput. Environ. Urban Syst.* **2017**, *65*, 126–139.
39. Transitland. Bus Feed for Las Palmas de Gran Canaria Urban Transport. Available online: <https://www.transit.land/feeds/f-autobús-urbano-de-las-palmas-de-gran-canaria> (accessed on 9 June 2026).
40. Ballatore, A.; Bertolotto, M.; Wilson, D.C. An evaluative baseline for geo-semantic relatedness and similarity. *GeoInformatica* **2014**, *18*, 747–767.
41. Bonabeau, E. Agent-based modeling: Methods and techniques for simulating human systems. *Proc. Natl. Acad. Sci. USA* **2002**, *99*, 7280–7287.
42. GeoPandas Development Team. GeoPandas: Python Tools for Geographic Data. Available online: <https://geopandas.org/> (accessed on 15 December 2024).
43. Gillies, S. Shapely: Manipulation and Analysis of Geometric Objects. Available online: <https://shapely.readthedocs.io/> (accessed on 15 December 2024).

44. Hagberg, A.A.; Schult, D.A.; Swart, P.J. Exploring network structure, dynamics, and function using NetworkX. In Proceedings of the 7th Python in Science Conference (SciPy2008), Pasadena, CA, USA, 19–24 August 2008; pp. 11–15.
45. Newson, P.; Krumm, J. Hidden Markov map matching through noise and sparseness. In Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, Seattle, WA, USA, 4–6 November 2009; pp. 336–343.
46. Tirachini, A. Bus dwell time: The effect of different fare collection systems, bus floor level and age of passengers. *Transp. A Transp. Sci.* **2013**, *9*, 28–49.
47. Pydantic Development Team. Pydantic: Data Validation Using Python Type Hints. Available online: <https://docs.pydantic.dev/> (accessed on 15 December 2024).
48. Google Cloud. Vertex AI Documentation: Gemini API Reference. Available online: <https://cloud.google.com/vertex-ai/docs/generative-ai/model-reference/gemini> (accessed on 25 May 2026).
49. Chainlit Development Team. Chainlit: Build Production-Ready Conversational AI Applications. Available online: <https://docs.chainlit.io/> (accessed on 15 December 2024).
50. Uber Technologies. Kepler.gl: A Powerful Web-Based Geospatial Data Visualization Tool. Available online: <https://kepler.gl/> (accessed on 15 December 2024).
51. Guaguas Municipales. MetroGuagua: Recorrido y Sistema de Transporte Masivo. Available online: <https://www.guaguas.com/lineas/metroguagua/recorrido> (accessed on 22 April 2026).
52. Jobin, A.; Ienca, M.; Vayena, E. The global landscape of AI ethics guidelines. *Nat. Mach. Intell.* **2019**, *1*, 389–399.
53. Macharis, C.; Bernardini, A. Reviewing the use of Multi-Criteria Decision Analysis for the evaluation of transport projects: Time for a multi-actor approach. *Transp. Policy* **2015**, *37*, 177–186.
54. Creutzig, F.; Jochem, P.; Edelenbosch, O.Y.; Mattauch, L.; van Vuuren, D.P.; McCollum, D.; Minx, J. Transport: A roadblock to climate change mitigation? *Science* **2015**, *350*, 911–912.
55. Foth, N.; Manaugh, K.; El-Geneidy, A.M. Towards equitable transit: Examining transit accessibility and social need in Toronto, Canada, 1996–2006. *J. Transp. Geogr.* **2013**, *29*, 1–10.
56. Ji, Z.; Lee, N.; Frieske, R.; Yu, T.; Su, D.; Xu, Y.; Ishii, E.; Bang, Y.J.; Madotto, A.; Fung, P. Survey of hallucination in natural language generation. *ACM Comput. Surv.* **2023**, *55*, 1–38.
57. Innes, J.E.; Booher, D.E. Reframing public participation: Strategies for the 21st century. *Plan. Theory Pract.* **2004**, *5*, 419–436.
58. OECD Observatory of Public Sector Innovation (OPSI). Virtual Twin Singapore. Available online: <https://oecd-opsi.org/innovations/virtual-twin-singapore/> (accessed on 29 April 2026).
59. Green, B.; Chen, Y. The principles and limits of algorithm-in-the-loop decision making. *Proc. ACM Hum. Comput. Interact.* **2019**, *3*, 1–24.
60. Wondyifraw, D.; van Hillegersberg, J. Open and Extendable Urban Digital Twins: Real-Time Data Pipelines and Conversational Interfaces for Sustainable Cities. *Procedia Comput. Sci.* **2026**, *278*, 791–799. <https://doi.org/10.1016/j.procs.2026.03.052>.
61. Ewing, R.; Handy, S. Measuring the unmeasurable: Urban design qualities related to walkability. *J. Urban Des.* **2009**, *14*, 65–84.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.