Urban Digital Twins for Megalopolises: Requirements, Challenges and Opportunities

Daniel Cordeiro[1], Jiefu Song[2], Kelly R. Braghetto[1], Fabio Kon[1], Frédéric Amblard[2]

[1]Universidade de São Paulo – São Paulo – SP – Brazil

[2]Institut de Recherche en Informatique de Toulouse (IRIT) – Toulouse – France

{daniel.cordeiro, kellyrb, fabio.kon}@usp.br
{jiefu.song, frederic.amblard}@irit.fr

Abstract. *Urban digital twins (UDTs) are emerging as critical tools for integrating heterogeneous data and models to support urban decision-making in areas such as mobility and energy management. However, broader adoption of these systems in large cities is constrained by scientific challenges in their architecture related to three interconnected dimensions: (1) scalability, through multi-modeling and surrogate modeling strategies that balance accuracy and resource efficiency; (2) interoperability, via adaptive and opportunistic workflows that dynamically integrate models and datasets based on context and granularity of decision-making; (3) frugality, by optimizing energy consumption across model and workflow executions. This paper details innovative data science and Urban Digital Twin approaches for collecting and analyzing urban data to simulate complex urban phenomena. By proposing scalable, interoperable, and energy-efficient architectures, this study seeks to advance systems supporting evidence-based public policy, promoting broader sustainable development.*

1. Introduction

Urban Digital Twins (UDTs) are rapidly emerging as critical tools for integrating heterogeneous data sources and computational models, aiming to enhance decision-making processes across multiple urban domains. As epicenters of cultural, human, and economic capital, cities play a fundamental role in driving global economic growth and socio-economic advancement. Simultaneously, they are major consumers of resources—particularly energy and materials—and significant contributors to waste and greenhouse gas (GHG) emissions.

As of 2022, more than half of the global population resided in urban areas, with projections indicating an increase to 70% by 2050. The case of Latin America is particularly illustrative, where approximately 80% of the population was urbanized by 2015, and by 2000, nearly one-quarter of the region's population was concentrated in just four megacities. This rapid and often unplanned urbanization has precipitated unsustainable development trajectories: cities consume up to 80% of the global energy supply and account for approximately 75% of global carbon emissions, even as they generate about 75% of global GDP [United Nations Human (UN) Settlements Programme 2012].

Smart city technologies have enabled the application of advanced computer science methodologies to address these complex urban challenges. However, effectively managing the scale and intricacy of contemporary megalopolises requires a paradigm shift toward more decentralized, network-oriented infrastructure planning. This shift requires moving beyond limited local-scale interventions and toward reimagining urban centers as components of broader "smart regions." This regionalized approach aligns with global sustainability frameworks such as the COP21 Paris Agreement and the 2030 Agenda of Sustainable Development of the United Nations, both of which emphasize the importance of sustainable cities and access to clean, efficient, and affordable energy.

Urban Digital Twins (UDTs) function as dynamic, continuously evolving virtual representations of urban environments. By integrating real-time data streams, Artificial Intelligence (AI), and the Internet of Things (IoT), UDTs offer actionable insights to inform evidence-based policymaking. These systems shall be underpinned by data management architectures that support high-volume, multimodal data ingestion and serve as the backbone for simulation and analytical models.

Despite their transformative potential, the deployment of UDT systems at the scale required by megacities remains constrained by several key scientific and technical challenges. These challenges can be categorized into three interrelated dimensions:

Scalability. Scaling UDT systems from metropolitan to megalopolitan scales presents significant computational and architectural difficulties. High-fidelity paradigms such as agent-based models (ABMs), while offering granular expressiveness, are computationally intensive and require extensive datasets, often limiting their feasibility to large urban populations.

Interoperability. Current UDT architectures often suffer from fragmented data schemas, limited interoperability, and inflexible workflows, which inhibit the holistic modeling of urban systems and the anticipation of emergent, system-wide effects. To overcome these limitations, UDTs must evolve into interoperable ecosystems capable of real-time integration across diverse datasets and model types. This entails robust data integration pipelines that harmonize heterogeneous sources, including IoT sensors, satellite imagery, transportation systems, and energy networks.

Frugality. The substantial computational demands of large-scale UDTs risk introducing significant energy consumption and associated environmental impacts, thereby undermining their role in promoting sustainability. Recognizing that digital technologies can both mitigate and exacerbate climate change, the project advocates for a carbon-responsible approach to UDT execution. This involves the design of frugal runtimes that optimize scheduling decisions based on a triad of constraints: performance, cost, and carbon emissions.

This paper presents innovative approaches that leverage data science and UDTs to collect and analyze urban data, enabling the simulation of complex urban phenomena. By proposing scalable, interoperable, and energy-efficient (frugal) UDT architectures, the study aims to advance the development of systems capable of supporting evidence-based public policy design, with a particular focus on reducing greenhouse gas (GHG) emissions and promoting broader goals of sustainable urban development.

2. Urban Digital Twins for Sustainable Megalopolises Development

2.1. Scaling Up UDT Systems

The scaling up of simulation models is a critical scientific challenge for the broader adoption of Urban Digital Twin systems, particularly in megacities. To effectively address the complex problems prevalent in these large urban areas, models must possess the capability to scale up from small regions of neighborhoods to the megalopolis level.

One significant challenge to achieving scalability is the computational cost and inherent limitations of existing models. Agent-based models (ABMs), despite their high expressiveness and flexibility for simulating individual and collective behaviors within urban systems, are computationally intensive and demand substantial data resources. Their complexity and computational burden frequently grow exponentially as the number of agents and interactions increase, limiting their feasibility for real-time decision-making and comprehensive scenario analysis [Michel et al. 2018, Gaudou et al. 2014].

The computational power required to simulate high-fidelity scenarios for megalopolis-scale Digital Twins can be massive, potentially leading to significant energy consumption and associated greenhouse gas emissions, which in itself becomes a sustainability problem. Furthermore, traditional urban models often concentrate on isolated subsystems, lacking comprehensive integration, and are not always developed with scalability and real-time adaptability as requirements. It is known that the existing tools and technologies employed by smaller cities for greenhouse gas (GHG) emission neutralization and other sustainability objectives may not be adequate for larger megalopolises.

To overcome these challenges, several strategies and methodologies for scaling can be explored. Multi-modeling and surrogate modeling strategies can effectively balance accuracy and resource efficiency. Surrogate models approximate system behavior by learning from a relatively sparse set of high-fidelity simulation data [De Leeuw et al. 2022, Llacay and Peffer 2025]. By employing various machine learning techniques, these models can generate fast, low-cost predictions while maintaining an acceptable level of accuracy. Such approaches are particularly beneficial for extensive parameter exploration, optimization, and uncertainty quantification, which would often be computationally prohibitive with full agent-based models (ABM) simulations.

Scaling microscopic model simulations also necessitates parallelizing execution across multiple computing nodes. This requires algorithms capable of dynamically distributing computational loads and adapting to real-time variations in execution times and available nodes. Dynamic load balancing continuously adapts to evolving agent interactions and workload imbalances, minimizing inter-process communication and optimizing resource utilization. Solutions explored to enhance load-balancing efficiency include task replication, redundancy transitioning, and nature-inspired optimization algorithms. New topology-aware partitioning methods have also been introduced to optimize large-scale spiking neural network (SNN) simulations, ensuring efficient distribution of neurons and connections across nodes [Zeng et al. 2024].

Adaptive simulation execution engines can also be used to adjust to resource availability dynamically. These engines allow detailed ABM simulations to be executed when sufficient computational power is available, while machine learning models can serve as surrogates in low-power or resource-constrained conditions, providing a lightweight alternative for fast, approximate predictions and significantly reducing computational demands [De Leeuw et al. 2022]. This adaptive strategy facilitates dynamic load balancing and ensures the system remains robust and efficient even amid fluctuating resources.

2.2. Interoperability of Data and Models

Cities are inherently complex, and to effectively address their intricate problems, UDTs must integrate diverse data and models in a dynamic and flexible manner. Achieving this interoperability is challenging due to the fragmented data models and limited integration capabilities within existing UDT architectures [Malleson et al. 2022]. These limitations restrict the system's capacity to fully capture city-wide interactions and anticipate the systemic effects of public policies.

The heterogeneity of urban data sources presents another major challenge. Data originates from numerous sources, including IoT sensors, satellite imagery, transportation networks, and energy grids [Mehmood et al. 2017]. Differences in data formats, spatiotemporal resolutions, and semantic meanings complicate efficient integration [Al-Yadumi et al. 2021]. Furthermore, scaling up UDTs is difficult because adding new data and models often requires rebuilding existing workflows, thereby limiting adaptability and slowing down updates.

A persistent challenge in integrated urban system modeling is ensuring consistency across different models, especially when confronted with data heterogeneity. Achieving seamless interoperability between various heterogeneous models remains an open research question. Concerns also persist regarding data reliability and minimizing biases within integrated datasets.

To overcome these obstacles, UDTs must evolve into interoperable ecosystems capable of dynamically integrating diverse datasets and models while ensuring scalability, adaptability, and efficiency. A cornerstone of this evolution is the implementation of robust data integration pipelines [Agarwal 2024], which harmonize heterogeneous data sources through efficient Extract, Transform, Load (ETL) processes. These pipelines should support (near-)real-time workflows for continuous data updates from distributed sources such as IoT devices, mobile sensors, and external databases [Weil et al. 2023]. Additionally, they must address data quality by detecting and correcting inconsistencies, missing values, and anomalies to ensure reliability and usability [Ilyas and Chu 2019]. Approaches such as knowledge graphs and semantic data models further enhance interoperability between diverse urban datasets [Voelz et al. 2023, Rocha et al. 2019].

Beyond syntactic harmonization, semantic interoperability is essential for meaningful integration of urban data. Semantic approaches leverage knowledge graphs to encode the contextual meaning of data elements, relationships, and constraints, enabling systems to interpret and interlink data across domains [Hogan et al. 2021]. Knowledge graphs represent data as interconnected entities and relationships, enabling rich querying capabilities and inference over heterogeneous urban datasets [Wang et al. 2024]. For example, Consoli et al. (2015) demonstrate how semantic data models can be applied to integrate transportation, environmental, and demographic data, providing a unified knowledge base that enhances the reasoning capacity of UDTs. Similarly, Consoli et al. (2017) illustrate the use of linked data to bridge different urban information systems, in order to facilitate cross-domain interoperability and support comprehensive urban analytics. Such semantic enrichment enables UDTs not only to integrate raw data but also to understand and model complex urban phenomena more effectively.

Finally, creating a repository of reusable datasets and models is crucial, ensuring proper versioning and accessibility across different Digital Twin applications, and requiring the definition of standardized interfaces for model and data integration to facilitate cross-domain simulations and multi-model interactions. By centralizing validated datasets and well-documented models, repositories facilitate cross-domain simulations and multi-model interactions, which are crucial for capturing the systemic dynamics of complex urban environments [Jeddoub et al. 2024]. Ensuring interoperability at this level demands the adoption of standardized interfaces and APIs that abstract model functionalities and data schemas, enabling different Digital Twin components to communicate and integrate seamlessly.

The integration of data with digital twins, facilitated by modern data integration architectures such as data meshes [Goedegebuure et al. 2024] and data lakehouses [Harby and Zulkernine 2025], that merge the scalability of data lakes with the benefits of traditional data warehouses, can yield new knowledge to support the development of public policies for more carbon-efficient megalopolises. This hybrid model supports the large-scale ingestion of raw urban data alongside structured, curated datasets, enabling complex analytical queries and machine learning tasks essential for Digital Twin operations [Schneider et al. 2024].

2.3. Frugal Urban Digital Twins

A major challenge for applying Urban Digital Twins to large cities is the high computational cost of large-scale simulations and their environmental impact, particularly those relying on paradigms like agent-based modeling.

In the context of UDTs, frugality refers to a foundational design principle aimed at minimizing computational, data transfer, and energy costs without compromising the functional adequacy of the system. It includes a form of computational sustainability, where resources are used efficiently, workflows are dynamically adapted to contextual constraints, and the overall environmental footprint of large-scale simulations is reduced [Kim et al. 2023, Violos et al. 2025, Knebel et al. 2020]. Examples of frugal UDT approaches include implementing strategies to optimize energy consumption across several dimensions: by learning surrogate models, evaluating the energy cost associated with specific workflows, and optimizing the execution of those workflows.

A key methodology involves proposing adaptive and opportunistic workflow solutions that combine data and models intelligently. These solutions are designed to generate suitable workflows based on analytical needs, considering factors like simulation granularity and frugality. Recent studies emphasize the importance of adaptive hybrid workflows in Urban Digital Twins, which dynamically combine high-fidelity models (e.g., ABMs) with lightweight surrogate models, based on contextual needs and resource constraints. This approach enables systems to adapt model granularity in real time, optimizing both computational cost and analytical accuracy. Ullrich et al. (2024) present a hybrid workflow connecting network-based and agent-based models for predictive pedestrian movement. The system dynamically selects the appropriate model type based on data availability and performance requirements, illustrating a concrete instantiation of context-aware model switching in urban simulation tasks. When computational power is ample, detailed agent-based model (ABM) simulations can be executed to capture intricate dynamics [Cheng et al. 2025, Shin et al. 2025]. Conversely, under low-power or resource-constrained conditions, surrogate models—which are faster, low-cost approximations learned from high-fidelity data—provide a lightweight alternative, significantly reducing computational demands while maintaining acceptable accuracy [De Leeuw et al. 2022, Llacay and Peffer 2025].

This adaptive strategy also allows for dynamic load balancing, ensuring the simulation system remains robust and efficient even with fluctuating resource availability [Ahmadzadeh and Sarbazi-Azad 2024, Chippagiri et al. 2024]. Efficient load balancing in distributed simulation environments is crucial for optimizing resource utilization, particularly in dynamic environments like traffic simulations, where agent interactions can vary unpredictably [Mastio et al. 2017].

A computational platform for running megalopolis-scale digital twins should adopt a hybrid Computing Continuum approach, integrating the high-performance, scalable infrastructure of cloud computing with the low-latency, privacy-preserving features of edge computing. A frugal Digital Twin runtime within this continuum must be Carbon Responsible, requiring scheduling algorithms that balance performance, cost, and emissions while adapting to resource and power constraints [Nafus et al. 2021]. GreenScale [Kim et al. 2023] is a state-of-the-art framework designed for carbon-aware scheduling across edge cloud infrastructures. It dynamically allocates tasks based on time and location-specific carbon intensity, reducing emissions by up to 29% compared to performance-focused schedulers. Ecovisor [Souza et al. 2023] introduces a concept of virtual energy systems, exposing power usage and renewable generation controls to applications. This enables UDT runtimes to optimize operations based on grid carbon intensity and energy availability. EASE [Perin et al. 2022] illustrates resource allocation in edge environments with renewable energy. Using model-predictive control and consensus protocols, EASE achieves near carbon-neutral operations while maintaining QoS.

A promising strategy involves integrating the moldable task scheduling model with *follow-the-renewables* approaches, such as *follow-the-sun* (assigning tasks to locations with greater solar energy availability) or *follow-the-moon* (prioritizing locations with lower energy demand), to enable more efficient utilization of renewable energy sources [Liu et al. 2011]. This also includes adjusting the fidelity of the digital twin dynamically based on environmental performance metrics, using surrogate models.

3. System Architecture for UDTs

Figure 1 presents the system architecture for our vision of Urban Digital Twins. In this architecture, a *Smart City Platform* is required to collect data from sensors deployed around the city and invoke actuators, sensors' dynamic counterparts. State-of-the-art examples of such platforms are the FIWARE project[1] and the InterSCity platform [Del Esposte et al. 2019].

All data collected from the smart city platform is fed to the *Data Broker*. This component is responsible for filtering the data relevant to the deployed UDTs, thus helping

[1] `https://www.fiware.org/`

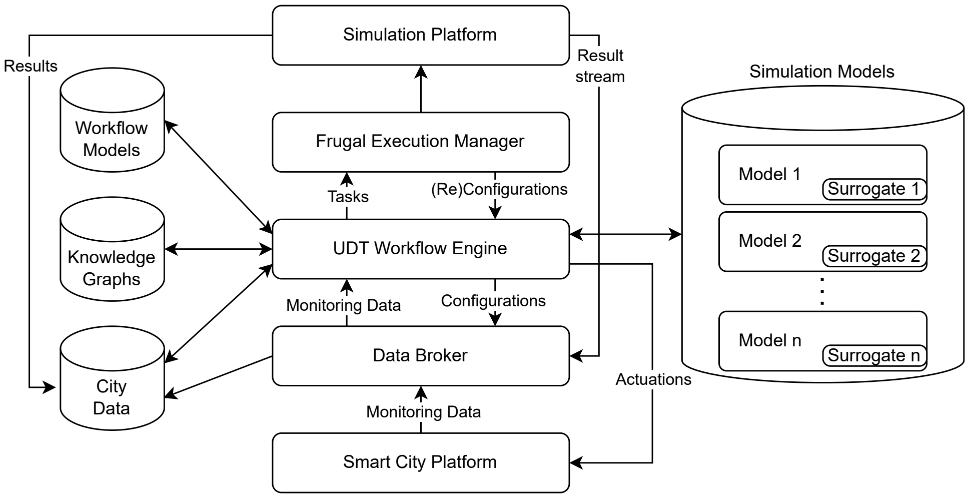

Figure 1. Proposed system architecture for megalopolises-scale UDTs.

to ensure frugality and to integrate the relevant data into the data mesh/data lakehouse systems as presented in Section 2.2 and indicated as the *City Data* in the figure.

The *UDT Workflow Engine* is a core component of the architecture. It will orchestrate the different simulations executed by UDTs, integrate the data being monitored using robust integration pipelines, manage the repository of reusable datasets and models (both *Workflow Models* and *Simulation Models*), while ensuring the interoperability between digital twins with the use of *Knowledge graphs* and semantic data models.

Tasks from the workflow will be executed by the *Frugal Execution Manager*. This component is responsible for managing all computational resources needed to run, explore, diagnose, and optimize the simulated models. Tasks will be scheduled on resources in the Cloud–Edge Continuum in a carbon-responsible way [Nafus et al. 2021], i.e., tasks will be assigned to the resources that minimize the GHG emissions associated with the execution. The runtime can also ask the workflow engine to reconfigure the execution based on the measured impact of the execution. For instance, if the GHG emissions are increasing city-wide, the runtime can reconfigure the parameters of the execution of the workflow engine, that, in turn, can decide to freeze the execution of the workflow or to replace the simulated model by the task with a *Surrogate model* that will provide a simple (but less accurate) model to be executed. This component could be based on existing projects such as the OpenMOLE model exploration platform [Reuillon et al. 2013].

The execution manager will choose the resource(s) to execute the simulation. The *Simulation Platform* will then launch the execution of the simulation model using the appropriate framework. Depending on the model being simulated, the execution manager could apply different simulation strategies. Taking urban mobility simulations, for instance, the simulation could be carried out by an agent-based model platform such as GAMA [Gaudou et al. 2014], giving detailed information for all elements of the simulation, or using a mesoscopic simulation strategy that uses fewer computational resources, but provides less accurate results, such as the InterSCSimulator [Rocha et al. 2021]. The data stream resulting from the simulations will be stored in the City Data data mesh/data lakehouse and immediately made available to the system, which will allow continuous ingestion, transformation, and real-time analysis of streaming data.

4. Conclusions

The technological solutions discussed in this paper represent a significant step toward integrating environmental efficiency into managing megalopolises through advanced computing technologies. By leveraging interdisciplinary tools, it establishes a robust framework for sustainable urban development. The development and deployment of digital twins for large cities will not only enhance our capacity to assess environmental impacts and test potential solutions efficiently but also reduce the computational and economic costs associated with large-scale simulations. Furthermore, the solutions align with the principles of digital sufficiency by advocating for a more intentional and environmentally conscious use of technology.

Acknowledgments

This project was partially funded by FAPESP (procs. 23/00811-0, 23/18026-8, and 24/01115-0) and by the CNRS International Research Center (IRC) 'Transitions'.

References

Agarwal, G. (2024). Robust data pipelines for ai workloads: Architectures, challenges, and future directions. *International Journal of Advanced Research in Science, Communication and Technology*, 5(2):622–632.

Ahmadzadeh, A. and Sarbazi-Azad, H. (2024). Performance analysis and modeling for quantum computing simulation on distributed gpu platforms. *Quantum Information Processing*, 23(11):1–40.

Al-Yadumi, S., Xion, T. E., Wei, S. G. W., and Boursier, P. (2021). Review on integrating geospatial big datasets and open research issues. *IEEE Access*, 9:10604–10620.

Cheng, K., Puyang, H., Li, X., Lee, P. P., Hu, Y., Li, J., and Wu, T.-Y. (2025). Toward load-balanced redundancy transitioning for erasure-coded storage. *IEEE Transactions on Parallel and Distributed Systems*.

Chippagiri, S., Ravula, P., and Gangwani, D. (2024). Optimizing load balancing and task scheduling in cloud computing based on nature-inspired optimization algorithms. *European Journal of Theoretical and Applied Sciences*, 2(6).

Consoli, S., Mongiovic, M., Nuzzolese, A. G., Peroni, S., Presutti, V., Reforgiato Recupero, D., and Spampinato, D. (2015). A smart city data model based on semantics best practice and principles. In *Proceedings of the 24th International Conference on World Wide Web*, pages 1395–1400.

Consoli, S., Presutti, V., Recupero, D. R., Nuzzolese, A. G., Peroni, S., Gangemi, A., et al. (2017). Producing linked data for smart cities: The case of catania. *Big Data Research*, 7:1–15.

De Leeuw, B., Mohammadi Ziabari, S. S., and Sharpanskykh, A. (2022). Surrogate modeling of agent-based airport terminal operations. In *International Workshop on Multi-Agent Systems and Agent-Based Simulation*, pages 82–94. Springer.

Del Esposte, A. d. M., Santana, E. F., Kanashiro, L., Costa, F. M., Braghetto, K. R., Lago, N., and Kon, F. (2019). Design and evaluation of a scalable smart city software platform with large-scale simulations. *Future Generation Computer Systems*, 93:427–441.

Gaudou, B. et al. (2014). The MAELIA multi-agent platform for integrated analysis of interactions between agricultural land-use and low-water management strategies. *Multi-Agent-Based Simulation XIV. MABS 2013. LNCS*, 8235.

Goedegebuure, A., Kumara, I., Driessen, S., Van Den Heuvel, W.-J., Monsieur, G., Tamburri, D. A., and Nucci, D. D. (2024). Data mesh: a systematic gray literature review. *ACM Computing Surveys*, 57(1):1–36.

Harby, A. A. and Zulkernine, F. (2025). Data lakehouse: A survey and experimental study. *Information Systems*, 127:102460.

Hogan, A., Blomqvist, E., Cochez, M., d'Amato, C., Melo, G. D., Gutierrez, C., Kirrane, S., Gayo, J. E. L., Navigli, R., Neumaier, S., et al. (2021). Knowledge graphs. *ACM Computing Surveys (Csur)*, 54(4):1–37.

Ilyas, I. F. and Chu, X. (2019). *Data cleaning*. Morgan & Claypool.

Jeddoub, I., Nys, G.-A., Hajji, R., and Billen, R. (2024). Data integration across urban digital twin lifecycle: a comprehensive review of current initiatives. *Annals of GIS*, pages 1–20.

Kim, Y. G., Gupta, U., McCrabb, A., Son, Y., Bertacco, V., Brooks, D., and Wu, C.-J. (2023). Greenscale: Carbon-aware systems for edge computing. *arXiv preprint arXiv:2304.00404*.

Knebel, F. P., Wickboldt, J. A., and de Freitas, E. P. (2020). A cloud-fog computing architecture for real-time digital twins. *arXiv preprint arXiv:2012.06118*.

Liu, Z., Lin, M., Wierman, A., Low, S. H., and Andrew, L. L. (2011). Geographical load balancing with renewables. *ACM SIGMETRICS Performance Evaluation Review*, 39(3):62–66.

Llacay, B. and Peffer, G. (2025). Categorical surrogation of agent-based models: A comparative study of machine learning classifiers. *Expert Systems*, 42(1):e13342.

Malleson, N., Birkin, M., Birks, D., Ge, J., Heppenstall, A., Manley, E., McCulloch, J., and Ternes, P. (2022). Agent-based modelling for urban analytics: State of the art and challenges. *AI Communications*, 35(4):393–406.

Mastio, M., Zargayouna, M., Scemama, G., and Rana, O. (2017). Distributed agent-based traffic simulations. *IEEE Intelligent Transportation Systems Magazine*.

Mehmood, Y., Ahmad, F., Yaqoob, I., Adnane, A., Imran, M., and Guizani, S. (2017). Internet-of-things-based smart cities: Recent advances and challenges. *IEEE Communications Magazine*, 55(9):16–24.

Michel, F., Ferber, J., and Drogoul, A. (2018). Multi-agent systems and simulation: A survey from the agent community's perspective. In *Multi-Agent Systems*, pages 17–66. CRC Press.

Nafus, D., Schooler, E. M., and Burch, K. A. (2021). Carbon-responsive computing: Changing the nexus between energy and computing. *Energies*, 14(21):6917.

Perin, G., Meneghello, F., Carli, R., Schenato, L., and Rossi, M. (2022). Ease: Energy-aware job scheduling for vehicular edge networks with renewable energy resources. *IEEE Trans. on Green Communications and Networking*, 7(1):339–353.

Reuillon, R., Leclaire, M., and Rey-Coyrehourcq, S. (2013). Openmole, a workflow engine specifically tailored for the distributed exploration of simulation models. *Future Generation Computer Systems*, 29(8):1981–1990.

Rocha, B., Cavalcante, E., Batista, T., and Silva, J. (2019). A linked data-based semantic information model for smart cities. In *2019 IX Brazilian Symposium on Computing Systems Engineering (SBESC)*. IEEE.

Rocha, F. W., Fukuda, J. C., Francesquini, E., and Cordeiro, D. (2021). Accelerating smart city simulations. In *Latin American High Performance Computing Conference*, pages 148–162. Springer.

Schneider, J., Gröger, C., Lutsch, A., Schwarz, H., and Mitschang, B. (2024). The lakehouse: State of the art on concepts and technologies. *SN Computer Science*, 5(5):449.

Shin, Y., Moul, V., Kang, K., and Lee, B. (2025). Polypal: A parallel microscale virtual specimen generator. *Computer Physics Communications*, 308:109458.

Souza, A., Bashir, N., Murillo, J., Hanafy, W., Liang, Q., Irwin, D., and Shenoy, P. (2023). Ecovisor: A virtual energy system for carbon-efficient applications. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, pages 252–265.

Ullrich, A., Hunger, F., Stavroulaki, I., Bilock, A., Jareteg, K., Tarakanov, Y., Gösta, A., Quist, J., Berghauser Pont, M., and Edelvik, F. (2024). A hybrid workflow connecting a network and an agent-based model for predictive pedestrian movement modelling. *Frontiers in Built Environment*, 10:1447377.

United Nations Human (UN) Settlements Programme (2012). Sustainable Urban Energy. Technical report, United Nations Human (UN) Settlements Programme.

Violos, J., Diamanti, K.-C., Kompatsiaris, I., and Papadopoulos, S. (2025). Frugal machine learning for energy-efficient, and resource-aware artificial intelligence. *arXiv preprint arXiv:2506.01869*.

Voelz, A., Amlashi, D. M., and Lee, M. (2023). Semantic matching through knowledge graphs: a smart city case. In *International Conference on Advanced Information Systems Engineering*, pages 92–104. Springer.

Wang, Z., Han, F., and Zhao, S. (2024). A survey on knowledge graph related research in smart city domain. *ACM Transactions on Knowledge Discovery from Data*, 18(9):1–31.

Weil, C., Bibri, S. E., Longchamp, R., Golay, F., and Alahi, A. (2023). Urban digital twin challenges: A systematic review and perspectives for sustainable smart cities. *Sustainable Cities and Society*, 99:104862.

Zeng, Y., Feng, G., Chen, Z., Lu, Y., and Xiao, N. (2024). Atm: Area-based partition and topology-aware mapping for large-scale snn simulation. In *2024 IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA)*, pages 1841–1848. IEEE.