



## Mobility utopias for realists: Generative design of topologically consistent transport system scenarios in existing cities

Lukas Ballo

Raphael Eder

Ayda Grisiute

Nina Wiedemann

Kay W. Axhausen

Martin Raubal

**Conference Paper**

**April 22, 2025**

**STRC** | **25th Swiss Transport Research Conference**  
Monte Verità / Ascona, May 14-16, 2025

# Mobility utopias for realists: Generative design of topologically consistent transport system scenarios in existing cities

Lukas Ballo  
IKG  
ETH Zurich  
lballo@ethz.ch

Ayda Grisiute  
IKG  
ETH Zurich  
agrisiute@ethz.ch

Kay W. Axhausen  
IVT  
ETH Zurich  
axhausen@ethz.ch

Raphael Eder  
Sociology Department  
UC San Diego  
reder@ucsd.edu

Nina Wiedemann  
IKG  
ETH Zurich  
nwiedemann@ethz.ch

Martin Raubal  
IKG  
ETH Zurich  
mraubal@ethz.ch

April 22, 2025

## Abstract

Urban planning faces systemic challenges such as climate change, urbanization, and demographic changes that require bold, integrated responses. However, traditional land use and transport planning processes are fragmented and not suitable for the transformations needed. This paper suggests to overcome this blockade by adopting a “deep planning” approach that spans across the traditional silos. It introduces an approach that combines large language models (LLMs) with a network-level urban design tool SNMan, to enable rapid generation of disruptive yet realistic urban scenarios—a process we term “structured ideation”. A proof-of-concept in Lucerne, Switzerland, demonstrates how the LLM converts natural language inputs into fundamentally redesigned yet consistent traffic networks using SNMan. We suggest reinforcement learning to improve the scenario generation process. The “deep planning” approach unifies creative ideation with real-world consistency, enabling cities to move beyond incrementalism and develop visionary yet actionable plans. While centered on urban design, this method introduced in this paper is also applicable to many other complex systems with or without spatial components.

## Keywords

Large Language Model, E-Bike City, SNMan, Transport Planning, Sustainability

**Preferred citation**

Ballo, L., R. Eder, N. Wiedemann, A. Grisiute, K. W. Axhausen and M. Raubal (2025) Mobility utopias for realists: Generative design of topologically consistent transport system scenarios in existing cities, paper presented at the *25th Swiss Transport Research Conference (STRC 2025)*, Ascona, May 2025.

## Contents

List of Figures . . . . .	1
1 Introduction . . . . .	2
2 Previous Work . . . . .	3
3 Vision . . . . .	4
4 Proof of Concept . . . . .	5
5 Future Development . . . . .	9
6 Discussion and Conclusion . . . . .	10
7 References . . . . .	11

## List of Figures

1 Converting unstructured inputs to instructions for SNMan in GPT-4o . . . . .	7
2 Automated creation of consistent scenarios . . . . .	8

# 1 Introduction

Societies around the world face systemic challenges such as climate change (IPCC, 2022), urbanization (UN, 2019), and demographic changes. However, existing land use and transport planning processes do not allow for an appropriate treatment of these pressing issues. Abstractly, planning processes usually follow three key steps: (1) a spatial master plan, (2) a transport plan, and (3) an implementation plan. Despite some overlaps, each step is based primarily on separate experts, data sources, and software tools. Constrained by these silos, the incorporation of additional complexity caused by systemic challenges remains incremental, fragmented, and slow. This stifles experimentation and prevents the exploration of visions beyond mere extrapolation of the status quo.

We argue that traditional planning processes are limited by two inherent “conflicts-by-design”: The “predict conflict” is due to the system’s reliance on extrapolating the status quo, limiting the opportunities for experimentation. The “provide-conflict” arises from the narrow range of possible solutions in the implementation plan, preventing the exploration of bold innovations. Addressing the systemic challenges requires a combination of an integrated planning process with structured assessments of possible future scenarios.

The “E-Bike City” project (Axhausen, 2022; Ballo *et al.*, 2023) made important steps toward breaking through this first impasse by testing a hypothetical policy of dedicating 50% of road space to cycling. Integrating macroscopic transport planning with the physical allocation of road space on every street, the team has developed a toolkit SNMan (Street Network Manipulator) (Wiedemann *et al.*, 2025; Ballo *et al.*, 2024) that spans across silos of the traditional planning process by allowing an automated planning of road space and the resulting transportation networks within the same model. In its original form, it generates alternative multimodal transportation systems that can be implemented solely by repurposing the existing road space. For example, it can be used to design networks of one-way streets, while reallocating the remaining road space to cycling paths, bus lanes, and green spaces. The results are controlled by a structured set of technical design parameters that need to be provided for each area of study. This approach serves as a first step toward a holistic planning process that we call “deep planning.”

In this paper, we take the next step toward enabling such a comprehensive planning approach. We propose to integrate SNMan with a large language model (LLM) that facilitates a translation between unstructured planning ideas and the technical design parameters. We term this process “structured ideation”. It allows users to quickly generate disruptive yet realistic urban futures within real-life constraints, such as existing

infrastructure.

The rest of this paper is structured as follows: Section 2 shows previous work on AI-based approaches in urban planning, Section 3 summarizes the vision of this paper in terms of technical interaction, Section 4 presents a proof of concept, Section 5 elaborates its possible future development, and Section 6 concludes the paper.

## 2 Previous Work

Scenario generation and exploration have long been essential elements of urban planning, supporting processes of conversion, visioning, and evaluation. However, existing approaches often remain limited by a lack of integration across domains and scales, leading to siloed workflows and excessive complexity.

Recent developments in generative AI have opened new pathways for producing spatial plans that respond more flexibly to complex demands. For instance, Zheng *et al.* (2023) propose a collaborative human–AI workflow to automate tedious aspects of urban design, enabling the creation of multiple spatial layout options aligned with different visions. Yet, the complexity and technical demands of such approaches have so far limited their practical uptake among planning practitioners.

Another branch of generative tools focuses on the rapid creation of visualizations and the facilitation of idea exchange among stakeholders. Applications such as UrbanistAI<sup>1</sup>, Hectar<sup>2</sup>, and PlacemakingAI<sup>3</sup> can suggest disruptive changes to street segments or public spaces, often responding to contextual inputs such as pedestrian activity (Valença *et al.*, 2025). However, these tools typically operate on isolated street segments or locations, without considering street connectivity or functional interdependencies within broader street networks.

Advanced applications, combining large language models with deterministic tools, offer a promising way to overcome this limitation by acting as an interface between users and specialized planning tools, facilitating structured workflows while maintaining accessibility. To date, many LLM applications in urban contexts have focused on individual-level tasks

---

<sup>1</sup><https://site.urbanistai.com/>

<sup>2</sup><https://www.parametric.se/>

<sup>3</sup><https://www.placemaking.ai/>

such as itinerary planning. Examples include UrbanLLM (Jiang *et al.*, 2024), which decomposes complex mobility queries into API calls, and ITINERA (Tang *et al.*, 2024), which supports custom route generation based on points of interest and optimization tools.

Only recently have researchers begun exploring the role of LLMs in more systemic, infrastructure-level planning tasks. A promising step was taken by Jin and Ma (2024), who developed an LLM-based assistant for planning parking infrastructure. Their model supports users in selecting appropriate data sources, weighting key factors, and preparing configuration files for existing simulation tools. While this work marks an important shift toward more integrated LLM-driven planning, it remains focused on a narrow domain (parking) and depends heavily on specific tools and datasets.

Beyond technical assistance, LLMs also hold potential for making urban planning more inclusive. For example, Zhou *et al.* (2024) simulate participatory design processes by generating LLM agents with varied personas to assess and comment on urban proposals. Similarly, LLMs have been used for automating urban audits, such as by combining street view imagery with text-based evaluation (Jang and Kim, 2025), thereby streamlining the assessment of existing environments.

Other projects aim to blend advanced capabilities with user-friendly design interfaces. For instance, Digital Blue Foam (DBF)<sup>4</sup> integrates GPT-3, a particular LLM model, into an agent-based tool for designing 15-minute cities. The system supports qualitative scenario evaluation through SWOT analyses and textual explanations, bridging technical depth and policy communication.

### 3 Vision

In contrast to previous work, we envision the use of LLMs for a wide range of planning applications in combination with databases and specialized tools, such as SNMan. We imagine three levels of autonomy for the LLM:

1. *Restricted execution*: The user provides specific instructions on where to place what kind of infrastructure, using natural language. LLMs translate these instructions

---

<sup>4</sup><https://www.digitalbluefoam.com/>



into street network configurations and provide visualizations.

2. *Single-mode decision making*: The user sets a goal with respect to one type of infrastructure, such as “reduce on-street parking by 50% and double the amount of green space and cycling infrastructure”. The LLM must break this problem down into subtasks and use query optimization tools and simulators to conduct a suitable analysis.
3. *Multi-modal high-level planning*: The user provides a high-level goal, such as “reducing the CO2 emissions by 20%”. This requires scenario development and running multiple tools, e.g., placing new bike lanes *and* new EV charging stations. The LLM could come up with suitable scenarios that provide a broad landscape of possible modifications and translate them into inputs to the available tools.

## 4 Proof of Concept

In this section, we present a proof of concept applied to a planning case study. We cover a small area of four street blocks in downtown Luzern, Switzerland. The area consists of urban streets with adjacent mixed-use buildings, as well as a thoroughfare with a high traffic volume and multiple major bus routes. See the “Status Quo” network in Figure 2. We will test a hypothetical reorganization of road space to reach the city’s climate change mitigation and adaptation targets.

The network contains information about streets, their width, and their allocation of road space among travel lanes, on-street parking, cycling lanes, and bicycle parking (represented by different colors). The underlying data model also contains information about the (approximate) number of residents and jobs in each building. The network is based on open geodata from OpenStreetMap, enriched with location-specific publicly available data sources such as the location of on-street parking spots, public transport routes, and aggregated jobs and population datasets. See Ballo *et al.* (2024) for more details on the process and the data sources used.

To integrate the process in an LLM, we use prompt engineering in GPT-4o, instructing it to convert the provided prompts to a structured JSON format with predefined attributes. Figure 1 provides an example of the prompts and corresponding answers. Approximately twenty manual prompts were necessary to train the LLM toward this behavior. They included overall instructions like “convert every provided input into a JSON and return it

without any further information”, a template for the intended JSON structure, and several corrective prompts to avoid any errors that occurred in the later responses. In a future development, this process can be automated by adding a further step that validates the responses and provides corrective feedback to the LLM if needed. Performing experiments on output structures with different complexity levels and with different LLMs, we can gain quantitative insights into the computational effort needed and the limits of this approach.

At this stage, the LLM can reliably convert the unstructured inputs into exact instructions for SNMan, which then generates the corresponding scenarios. However, note that at this stage, even the unstructured instructions still need to contain rather specific information about the individual elements used, like one-way streets, or the desired provision of parking. So far, it cannot create useful scenarios based on high-level inputs like “adapt the city to climate change”. In the future, adding further context and learning from impact assessments of the results can help overcome this challenge. Section 5 elaborates on such future development.

Next, we use the JSON output to guide the scenario generation in SNMan. The results for two scenarios are shown in Figure 2. Scenario 1 enforces one-way traffic on all streets and provides 0.5 parking spaces and 0.3 m<sup>2</sup> of green space per person (within a walking distance of 300 meters). SNMan uses a gravity model to distribute resources like parking and green spaces such that they correspond to the approximate demand, in this case, the number of residents and full-time job equivalents in each building. Further, Scenario 1 also transforms the thoroughfare Hirschmattstrasse into a bus-only street and adds cycling infrastructure where possible. Scenario 2 reduces the parking supply to 0.1 and increases the green space supply to 1 m<sup>2</sup> per person, which also results in more cycling infrastructure. Note that the resulting scenarios are always compliant with the existing road infrastructure, parking demand, and connectivity constraints (e.g., the one-way streets form a strongly connected graph, and the parking is distributed based on expected demand), without the user having to deal with this complexity.

The resulting networks can be evaluated for their impacts on key factors, such as accessibility or heat island effects. A simplified version of accessibility calculations without congestion effects can be conducted directly in SNMan.

Figure 1: Converting unstructured inputs to instructions for SNMan in GPT-4o

One-way streets everywhere, per person 0.5 parking spaces and 0.3 m2 green space. Make Hirschmattstrasse into a bus-only thoroughfare in both directions.

json

Copy

Edit

```
{
  "car_parking_supply_factor": 0.5,
  "green_space_supply_factor": 0.3,
  "motorized_traffic_in_both_directions": false,
  "constraints": [
    {
      "street_name": "Hirschmattstrasse",
      "lanes": ["T<*", "T>*"]
    }
  ]
}
```

Now reduce parking by 80% and provide 1m2 of green space per person. Make the middle section of Habsburgerstrasse where it crosses with Hirschmattstrasse into a pedestrian/cycling space. It should be about 15 meters long in each direction.

json

Copy

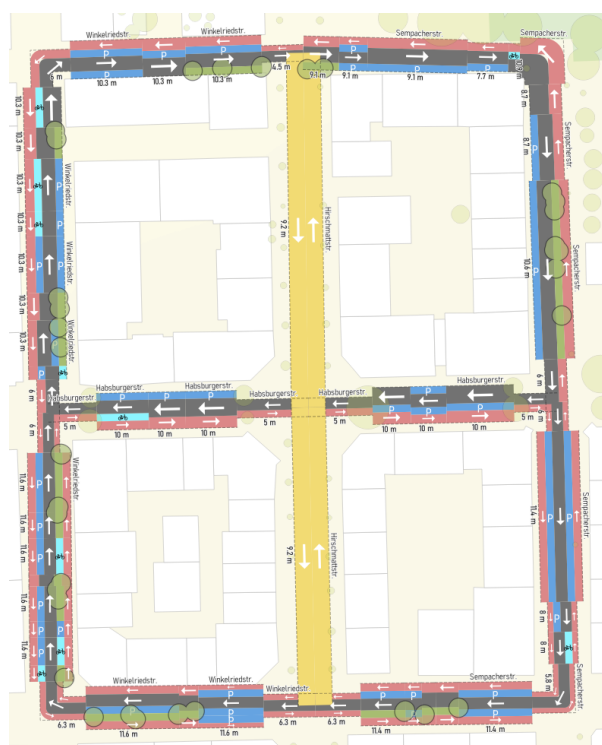
Edit

```
{
  "car_parking_supply_factor": 0.1,
  "green_space_supply_factor": 1.0,
  "motorized_traffic_in_both_directions": false,
  "constraints": [
    {
      "street_name": "Hirschmattstrasse",
      "lanes": ["T<*", "T>*"]
    },
    {
      "street_name": "Habsburgerstrasse",
      "lanes": ["X-*"],
      "intersection_with": "Hirschmattstrasse",
      "length_before": 15,
      "length_after": 15
    }
  ]
}
```

Figure 2: Automated creation of consistent scenarios



Status Quo



Scenario 1



Scenario 2

## 5 Future Development

In future developments, the scenario generation in SNMan could incorporate further dimensions, for example, zoning changes, population growth, energy infrastructure, fiscal restrictions, or tax codes, all while ensuring consistency. The LLM can be conditioned with domain-specific knowledge and common design patterns, such as the 15-minute city or superblocks, and major trade-offs embedded in these concepts. Generating multiple scenarios in SNMan simultaneously allows users to explore the resulting trade-offs and select the most desirable scenario.

However, achieving acceptable solutions to complex issues will likely require generating a large number of scenarios each time. To reduce the computational cost, we need to guide the LLM to instruct SNMan to generate only the most promising ones, rather than an extensive range of possible scenarios. This can be solved by adding a loop of reinforcement learning that feeds the results from each scenario’s impact assessment back to the LLM. With an adaptation of the “Dreamer” algorithm (Hafner *et al.*, 2025), the LLM could autonomously explore different scenarios and learn to anticipate the consequences of future configurations without having to iteratively try them out upon each user’s request. With such a configuration, the LLM becomes an AI agent that autonomously searches for the best solutions, while keeping the human planner in the loop for providing the problems, constraints, and a judgment of the feasibility of resulting scenarios.

In a final stage, even SNMan’s native design constraints, which ensure permanent consistency (e.g., road network connectivity) within the resulting scenarios, could be extracted by learning from real cities, rather than being hard-coded in the design framework. Like images or written text, cities exhibit many regularities that must be met to be considered “consistent”. Future studies may reveal the extent to which such consistency rules can be inferred from open geodata.

With the above future developments, the resulting method could offer suggestions for addressing complex urban challenges while respecting the constraints of existing urban environments. The solutions will span both horizontally, across domains, and vertically, across levels of detail. For example, it will be able to integrate population growth with changes to the fiscal situation of a city, future accessibility changes, provision of public infrastructure, and changes to the transport system. If needed, it can model each scenario down to the detail of road space allocation to unwind the tradeoffs that typically only emerge at a detailed phase, where it is often impossible to change the initial vision.

While we have so far focused on a use case in a civic, urban-planning context, the proposed methodology is generalizable to any complex system, whether or not it includes spatial components.

## 6 Discussion and Conclusion

Current planning processes are fragmented, time-consuming, and ill-equipped to explore the full range of solutions required to tackle today’s systemic challenges. The “E-Bike City” project made a crucial contribution by expanding the solution space through the integration of transport planning with the physical allocation and design of streets. “Deep planning” takes this a step further, generalizing this approach to enable the integration of any relevant aspect needed to tackle complex, systemic challenges through a process of “structured ideation”. This approach combines the creative potential of generative AI (ideation) with the rigor of logical consistency (structured). As a result, visions for the future of cities move beyond (unrealistic) utopias and instead are designed in ways that align with real-world constraints. These visions can be quantitatively assessed within the existing context of land use and transport planning, with the assessment results feeding back into guiding the generation of future proposals.

Still, “deep planning” does not resolve all challenges facing current planning practices. Transport planning, in particular, relies on a range of assumptions, such as the valuation of costs and benefits (see Zani *et al.* (2023) for a discussion of the difficulties in quantifying these factors). Future assessments of urban utopia must critically reflect on these input parameters to avoid the unintentional reproduction of existing paradigms.

The proof of concept presented in this paper illustrates the impressive capabilities of LLMs to deal with unstructured inputs and use them to control precise, specialized tools. Future work will explore the possibilities of widening the solution space beyond road space allocation, leveraging reinforcement learning and “structured ideation” to support planners in developing truly visionary solutions for existing cities facing systemic challenges.

## Author Contributions

Lukas Ballo: Conceptualization, Software, Writing - Original Draft, Visualization

Raphael Eder: Conceptualization, Writing - Original Draft

Nina Wiedemann: Conceptualization, Writing - Original Draft

Ayda Grisiute: Conceptualization, Writing - Original Draft

Kay W. Axhausen: Writing - Review & Editing

Martin Raubal: Writing - Review & Editing

## Note on Generative AI

Generative AI was used to improve writing and wording.

## 7 References

- Axhausen, K. W. (2022) The dilemma of transport policy making and the COVID-19 accelerator, in M. Attard and C. Mulley (eds.) *Transport and Sustainability, Transport and Pandemic Experiences, Volume 17*, 39–51, Emerald Publishing Limited, Bingley.
- Ballo, L., L. Meyer de Freitas, A. Meister and K. W. Axhausen (2023) The E-Bike City as a radical shift toward zero-emission transport: Sustainable? Equitable? Desirable?, *Journal of Transport Geography*, **111**, 103663.
- Ballo, L., M. Raubal and K. W. Axhausen (2024) Designing an E-Bike City: An automated process for network-wide multimodal road space reallocation, *Journal of Cycling and Micromobility Research*, **2**, 100048, ISSN 2950-1059.
- Hafner, D., J. Pasukonis, J. Ba and T. Lillicrap (2025) Mastering diverse control tasks through world models, *Nature*, ISSN 0028-0836, 1476-4687.
- IPCC (2022) Climate change 2022, mitigation of climate change, summary for policymakers, *Technical Report*, Intergovernmental Panel on Climate Change, Geneva.

- Jang, K. M. and J. Kim (2025) Multimodal Large Language Models as Built Environment Auditing Tools, *The Professional Geographer*, **77** (1) 84–90, ISSN 0033-0124, 1467-9272.
- Jiang, Y., Q. Chao, Y. Chen, X. Li, S. Liu and G. Cong (2024) Urbanllm: Autonomous urban activity planning and management with large language models, *arXiv preprint arXiv:2406.12360*.
- Jin, Y. and J. Ma (2024) Large language model as parking planning agent in the context of mixed period of autonomous vehicles and human-driven vehicles, *Sustainable Cities and Society*, **117**, 105940.
- Tang, Y., Z. Wang, A. Qu, Y. Yan, Z. Wu, D. Zhuang, J. Kai, K. Hou, X. Guo, H. Zheng *et al.* (2024) Itinera: Integrating spatial optimization with large language models for open-domain urban itinerary planning, *arXiv preprint arXiv:2402.07204*.
- UN (2019) World urbanization prospects: The 2018 revision, *Technical Report*, United Nations, Department of Economic and Social Affairs, Population Division, New York.
- Valença, G., C. Azevedo, F. Moura and A. Morais De Sá (2025) Creating visualizations using generative AI to guide decision-making in street designs: A viewpoint, *Journal of Urban Mobility*, **7**, 100104, ISSN 26670917.
- Wiedemann, N., C. Nöbel, L. Ballo, H. Martin and M. Raubal (2025) Bike network planning in limited urban space, *Transportation Research Part B: Methodological*, **192**, 103135, ISSN 0191-2615.
- Zani, D., C. Kielhauser and B. T. Adey (2023) Overcoming challenges in cost-benefit analysis of urban cycling infrastructure, paper presented at the *23rd Swiss Transport Research Conference*, Monte Verità. May 10-12.
- Zheng, Y., Y. Lin, L. Zhao, T. Wu, D. Jin and Y. Li (2023) Spatial planning of urban communities via deep reinforcement learning, *Nature Computational Science*, **3** (9) 748–762, ISSN 2662-8457.
- Zhou, Z., Y. Lin, D. Jin and Y. Li (2024) Large Language Model for Participatory Urban Planning, <http://arxiv.org/abs/2402.17161>. ArXiv:2402.17161 [cs].