Computational urban science

https://doi.org/10.1038/s43588-023-00515-1

AI improves the design of urban communities

Paolo Santi

Check for updates

A reinforcement-learning-based framework is proposed for assisting urban planners in the complex task of optimizing the spatial design of urban communities.

Urban planning and design can be better described as a process than a task: design and planning decisions need to be integrated with political guidelines, societal goals and different types of constraints, and discussed with community stakeholders, typically across several iterations. Computational tasks, related for instance to spatial allocation of land use and facilities, are only a small part of the process – yet one of the most time-consuming. In urban planning practice, a solution that might look optimal from a quantitative and optimization viewpoint might be less preferable than another one that is maybe quantitatively suboptimal but more acceptable for the local community. In fact, the outcome of a planning task is not merely an allocation of space to buildings, parks and functions, but the design of a place where urban communities will live, work, interact and, hopefully, thrive for a very long time. Writing in *Nature Computational Science*, Yu Zheng and colleagues[1] propose tackling this problem with the help of artificial intelligence (AI).

How can AI help and assist urban planners and local communities in this process? Even if we restrict the focus of AI to the specific task of spatial land use and road layout, the application of AI methods needs to deal with the irregularity in the land and road layout, which makes the problem more challenging than well-known applications of AI to other spatial decision-making tasks such as the game of Go[2] and chip design[3].

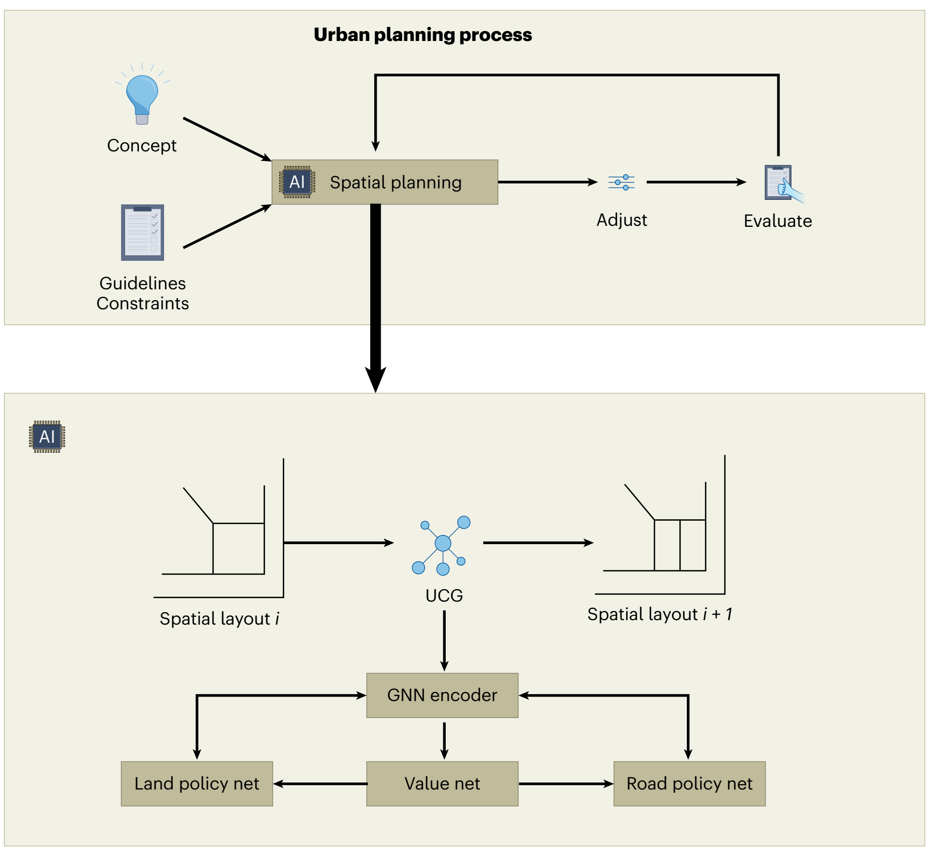

Fig. 1 | AI-assisted urban planning process. Top: the urban planning process is an iterative process where the quantitative optimization step of spatial planning is integrated into steps that require ideation, interaction with policy and local stakeholders, public engagement, and thorough evaluation of the obtained plan. Bottom: the AI-assisted spatial planning framework proposed by Zheng et al.[1]. The spatial planning problem is translated into a Markov decision process (MDP) on the dynamic urban contiguity graph (UCG). Through a graph neural network (GNN) state encoder, the information in the UCG is exchanged with neural networks that are used to constrain the action space of the MDP (the land policy and road policy network) and to evaluate the current layout situation (value network).

To address the spatial irregularity challenge, Zheng et al. introduced a theoretical framework for translating the computational task of spatial layout planning into a sequential Markov decision process and solving it using reinforcement learning. One of the key elements of the framework is the notion of the urban contiguity graph (UCG), which expresses the topological relationships between the elements of a spatial layout, namely land plots, road segments and junctions. Hence, potential irregularities in the spatial elements are abstracted into a model that accounts only for their topological and proximity relationships. As planning decisions are made (for instance, introducing a new road), the topological relationships between the urban elements modeled in the UCG evolve to reflect such decisions. The spatial layout planning problem can then be formulated as a sequential decision-making process on a dynamic graph. As decisions are taken at each step, the current layout is evaluated through a reward function that accounts for three metrics related to access to basic services (service), access to parks (ecology) and efficiency of the road network (traffic).

With the challenge related to spatial irregularity addressed, the spatial layout design problem is still far from being solved. To steer the sequential decision-making process, the AI approach needs to explore a vast search space, through an even wider action space that comprises about $4{,}000^{100}$ actions even for a moderate-sized community of a few square kilometers. Without a strategy for shrinking such exploration spaces, the number of iterations needed to converge to a good solution would be prohibitive and impair the feasibility of the proposed framework. To tackle this computational challenge, the authors use different types of neural networks to constraint the action space for land and road layout decisions (the policy networks), and a value network to judge current planning situations and predict planning performance. The policy and value networks exchange and use the information embedded in the UCG through a graph neural network (GNN) state encoder.

Thanks to this deep reinforcement-learning-based architecture (Fig. 1), the authors show that AI-assisted spatial land and road layout outperforms layouts produced by human experts with respect to all considered metrics (services, ecology and traffic) by nearly 50%.

Despite these results, the authors acknowledge the nuisances of the urban planning process and by no means advocate for replacing human planners with AI. Rather, they suggest an integrated human–AI model for spatial planning, where human experts generate conceptual plans using centers and axes (task 1), which are then fed to AI spatial layout models for quantitative optimization (task 2), and finally are reviewed and adjusted by human experts (task 3). The authors claim that such an integrated workflow might reduce — by as much as a factor of 3,000 — the time needed for task 2 as compared with a fully human workflow. With AI, human planners can be liberated from computationally intensive, multi-objective quantitative tasks and focus entirely on the more nuanced and inherently human tasks related to conceptualization and comprehensive evaluation. The authors corroborate such an important claim with carefully designed experiments involving human planners, which show that the integrated human–AI framework can achieve Pareto optimal solutions in service and ecology, improving their efficiency by 12% and 5% with respect to a fully human workflow, while abating the time needed to build such solutions.

While addressing important conceptual and computational challenges and demonstrating the feasibility of an integrated human–AI workflow for spatial layout planning, the work of Zheng et al. leaves many avenues open for future research. Although the authors demonstrate that their proposed framework displays transfer learning and some scaling features, they argue that extending the use of AI to the spatial design of an entire city is still an open challenge. The challenge is not only related to computational issues, but also to the fact that a city can be described as a complex system formed of intertwined subsystems, where changes affecting one subsystem might have unexpected consequences in another part of the city. For instance, it has been observed that the recent 15-minute city concept in urban planning[4], which is motivated by the idea of promoting short-range active mobility and vibrant local neighborhoods, might have the unexpected consequence of exacerbating social segregation[5]. Whether AI can assist urban planners also in identifying such unexpected consequences of urban planning is another open question.

Paolo Santi [1,2] ✉
[1]MIT Senseable City Lab, Cambridge, MA, USA. [2]IIT-CNR, Pisa, Italy.
✉e-mail: psanti@mit.edu

Published online: 11 September 2023

References

1. Zheng, Y. et al. *Nat. Comput. Sci.* https://doi.org/10.1038/s43588-023-00503-5 (2023).
2. Silver, D. et al. *Nature* **529**, 484–489 (2016).
3. Mirhoseini, A. et al. *Nature* **594**, 207–212 (2021).
4. Moreno, C., Allam, Z., Chabaud, D., Gall, C. & Pratolong, F. *Smart Cities* **4**, 93–111 (2021).
5. Glaeser, E. The 15-minute city is a dead end — cities must be places of opportunity for everyone. *LSE* (28 May 2021); https://blogs.lse.ac.uk/covid19/2021/05/28/the-15-minute-city-is-a-dead-end-cities-must-be-places-of-opportunity-for-everyone/

Competing interests

The author declares no competing interests.