Sebastiao *et al. Applied Network Science* (2026) 11:34
https://doi.org/10.1007/s41109-026-00792-5

Applied Network Science

RESEARCH Open Access

The trade-off between directness and coverage in transport network growth

Clément Sebastiao[1*], Anastassia Vybornova[1,2], Ane Rahbek Vierø[3], Luca Maria Aiello[1,4] and Michael Szell[1,4,5]

*Correspondence:
Clément Sebastiao
clse@itu.dk
[1]Networks, Data, and Society (NERDS), IT University of Copenhagen, Rued Langgaards Vej 7, 2300 Copenhagen, Denmark
[2]Copenhagen Center for Social Data Science (SODAS), Copenhagen University, Øster Farimagsgade 5, 1353 Copenhagen, Denmark
[3]Mobility, Space, Place and Urban Studies (MOSPUS), Roskilde University, Universitetsvej 1, 4000 Roskilde, Denmark
[4]Pioneer Centre for AI, Øster Voldgade 3, 1350 Copenhagen, Denmark
[5]Complexity Science Hub, Metternichgasse 8, 1030 Vienna, Austria

Abstract

Designing spatial networks, such as transport networks, commonly deals with the problem of how to best connect a set of locations through a set of links. In practice, it can be crucial to order the implementation of the links in a way that facilitates early functioning of the network during growth, like in bicycle networks. However, it is unclear how this early functional structure can be achieved by different growth processes. Here, we systematically study the growth of connected planar networks, quantifying functionality of the growing network structure. We compare random growth with various greedy and human-designed, manual growth strategies. We evaluate our results via the fundamental performance metrics of directness and coverage, finding non-trivial trade-offs between them. Manual strategies fare better than greedy strategies on both metrics, while random strategies perform worst and are unlikely to be Pareto efficient. Centrality-based greedy strategies tend to perform best for directness but are worse than random strategies for coverage, while coverage-based greedy strategies can achieve maximum global coverage as fast as possible but perform as poorly for directness as random strategies. Directness-based greedy strategies get stuck in local optimum traps. These results hold for a number of stylized urban transport network topologies. Our insights are crucial for applications where the order in which links are added to a spatial network is important, such as in urban or regional transport network design problems.

Keywords Spatial network, Greedy algorithm, Bicycle network, Network evolution, Transport planning

Introduction

The design of transport networks like street, rail, or bicycle networks comes with many constraints, such as planarity (Barthelemy 2022) or the provision of direct connections, serving a given population, avoiding bottlenecks (Kirkley et al. 2018), or multimodality (Alessandretti et al. 2023). When designing by optimization, such constraints naturally come with trade-offs. In general network design, trade-offs include average shortest path versus link density (Rfi and Solé 2003) or total length (Brede 2010), costs versus benefits (Louf et al. 2013), or comparing measures of efficiency (Latora and Marchiori 2002; Cardillo et al. 2006; Morer et al. 2020). In more applied transport settings, trade-offs include directness versus coverage in bicycle networks (Szell et al. 2022), densification

© The Author(s) 2026. **Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by-nc-nd/4.0/.

versus exploration in road networks (Strano et al. 2012; El Gouj et al. 2022b), centralization in air transport networks (Wang et al. 2014), or growing versus shrinking in rail networks (Bottinelli et al. 2019). Bio-inspired and historical development processes can inform modeling efforts to help deal with such trade-offs (Douady et al. 2020; El Gouj et al. 2022b).

However, a specific formulation of a transport network design problem and its trade-offs is underexplored so far, emerging from recent studies of bicycle network growth (Szell et al. 2022; Steinacker et al. 2022; Vybornova et al. 2023; Paulsen and Rich 2023; Mauttone et al. 2017; Ospina et al. 2022; Steinacker et al. 2025), and called out as a research gap by Diogo Pinto et al. (2026). Unlike the slow growth of streets from historical processes (El Gouj et al. 2022a), these networks are developed as subnetworks on a larger, already existing street network. Therefore, the growth of such subnetworks can proceed relatively quickly under the right political conditions (but still take years or decades), such as the growth of the bicycle networks of Seville (Marqués et al. 2015) or Paris (Moran 2022).

Nevertheless, this development can take many years or even decades, since high political friction and extremely limited funding due to car-centrism imply long construction times (Carstensen et al. 2015) that can leave network gaps open for many decades (Natera Orozco et al. 2020; Vybornova et al. 2023). If the developing network does not provide uninterrupted or direct routes, it could severely obstruct the uptake of cycling (Schoner and Levinson 2014; Ferenchak and Marshall 2025). Therefore, a good growth strategy should implement a *functional* network as soon as possible, i.e., *a network that is useful for its users as early as possible in the construction process*.

The starting point in this context is either to consider all links of the street network for potential future protected bicycle infrastructure, or a subset of all links that has been identified by planners to connect a number of locations (Szell et al. 2022). While this problem is motivated by the bicycle network context, it is a special case of a research question that can be formulated more generally, where "functionality" depends on the particular use case:

> *Given a set of locations (nodes) and a set of possible connections (links) between them, in which order should these connections be added to ensure a functional network structure already during growth?*

Motivated by Szell et al. (2022), here we operationalize "functionality" via the two performance metrics of directness and coverage, which quantify how direct routes are in the network and how much of the area the network covers, respectively – more formally defined in Sect. 2.

The posed network design question has been addressed from both transport planning using multi-parameter optimization (Mauttone et al. 2017; Ospina et al. 2022; Paulsen and Rich 2023; Ballo et al. 2024; Wiedemann et al. 2025) and from network science using single or multi-parameter optimization (Steinacker et al. 2022, 2025; Szell et al. 2022). In several such previous studies the growth algorithm is greedy, adding only the best possible next link without considering the long-term future. Further, instead of such an additive approach, Steinacker et al. (2022, 2025), Ballo et al. (2024), Wiedemann et al. (2025) have explored a *subtractive* design process where starting backwards from the final network the least important link is removed in each step.

While all these approaches provide useful applications, so far a *systematic* study of spatial network growth processes and the evolution of their performance metrics is missing, as also identified by Diogo Pinto et al. (2026). Such a theoretical groundwork is the contribution of our paper. In particular, we study the growth of connected planar networks and explore the trade-off between the two most important metrics, directness and coverage. We develop and study our framework on synthetic networks, then demonstrate its applicability with two empirical networks.

The rest of the paper is structured as follows. Section 2 introduces the methodology, describing the network metrics used as cost functions during the growth, and the overall framework for optimal growth of planar networks. Section 3 presents the results of applying these methods to one of the simplest planar networks—a regular grid lattice—and then extends the analysis to other stylized urban street patterns. Section 4 discusses the results and puts them into context. The method developed here is available as reproducible open-source code at Sebastiao (2024a) which can be applied to any network structure or metric.

Methods

Strategies and orders

In this article we study the concept of link ordering *strategy*, which we define as an algorithm S that takes as input the spatial network $\mathcal{G} = (\mathcal{N}, \mathcal{L})$ made up of the set of nodes $\mathcal{N} = \{n_1, \ldots, n_N\}$, and links $\mathcal{L} = \{e_1, \ldots, e_L\}$. We call this network $\mathcal{G}$ the *final network*. It can represent a future network that planners are deliberately targeting to build, which can be a subnetwork on top of a pre-existing network with a large number of links—for example a bicycle network on top of a street network. The strategy outputs a sequence (ordered list) of links encoding in which order they should be implemented: $S(\mathcal{G}) = (\sigma(e_i))_{i=1}^{L}$, where σ is a permutation determined by S. We call $(\sigma(e_i))_{i=1}^{L}$ a (link) *order* of strategy S for the network $\mathcal{G}$. For example, a strategy might determine that the best order for a final network $\mathcal{G}' = (\{n_1, n_2, n_3\}, \{e_1, e_2, e_3\})$ is the sequence (e_3, e_1, e_2), in which case link e_3 should be implemented first, followed by link e_1, followed by link e_2. We call a subnetwork $\mathcal{H} = (\mathcal{N}, S(\mathcal{G})_{i=1}^{k}) \subseteq \mathcal{G}$ that has built the first k links according to strategy S the *kth step* of network $\mathcal{G}$ for strategy S.

If a strategy's output is always the same for the same input, it is fully deterministic. However, in our study, each considered strategy has multiple possible orders for the same set of nodes and links because randomness is always part of the process—either by construction in a random strategy, or when a new added link has to be selected randomly among multiple equivalent link choices. So, for the example strategy above, calling the algorithm another time on the same final network could yield a different order (e_1, e_3, e_2). To observe the variation of strategies, we create 50 orders for each strategy, except for the random strategy where we create 1000 orders as it generates wider variations.

Network metrics

We use several network metrics for two purposes: (1) to evaluate the quality of a given link order, and (2) for defining greedy growth strategies. Our framework works with any metric, but here we pick specific ones that are most relevant for the bicycle network link ordering problem motivated in the introduction. The formalism is general though and

could be applied to a wide range of other use cases, such as designing road, rail or other spatial networks. Further, the networks we consider are planar by construction because urban street networks are mostly planar (Boeing 2020), but strict planarity is not a necessary prerequisite. It is important though that the networks are spatial.

The Dutch CROW Manual (CROW 2016), a key reference for bicycle infrastructure design, reports four main requirements for a good bicycle network: cohesion, directness, safety, and attractiveness. While safety and attractiveness are mostly about the quality of the physical infrastructure and its surroundings, cohesion and directness relate to the topology of the bicycle network. Cohesion was operationalized by Szell et al. (2022) as being able to reach all points of interest (accessibility, or coverage) in an interconnected manner (connectedness). Directness is the ratio between shortest path distances and beeline distances.

As we do not deal with spatial attributes other than distances between nodes, we are dropping the safety and attractiveness aspects in the following. We focus on the topological/geometrical aspects of coverage, directness, and connectedness for measuring performance of our growth process, defined in detail in the following subsections. Figure 1 illustrates the specific multi-criteria optimization that we address. We aim to design a growing network where both coverage and directness are high, but we also expect a trade-off between the two. For example, previous work has shown that growing disconnected networks via a random strategy leads to high coverage at the expense of other network metrics (Szell et al. 2022).

At the same time, we impose the condition of connectedness through all growth stages (i.e., to consist of one single connected component), because disconnected components are an undesirable trait, especially in transport networks (Natera Orozco et al. 2020). Further, greedy network growth would not be meaningful when isolated links or components were allowed, as it would be trivial and not desirable, see leftmost panel of Fig. 1: A growth process in which isolated pairs of nodes connected by single, straight links are continuously added to the network would optimize both directness and coverage. In this case, the optimization of directness happens because the only existing origin–destination pairs are neighboring nodes; the optimization of coverage happens because the buffers around scattered links do not overlap, thus contributing their full area while the

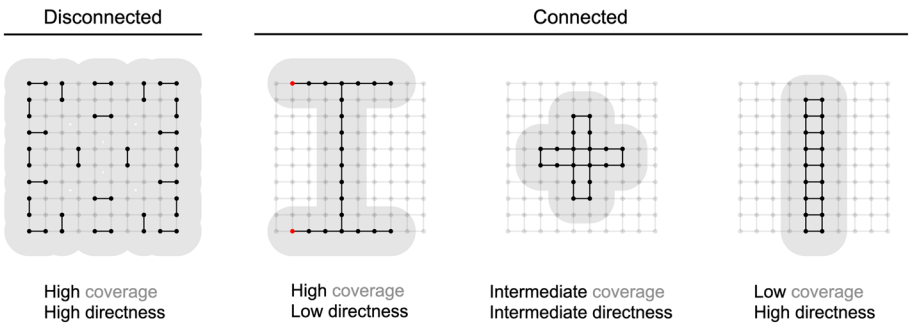


Fig. 1 The directness-coverage trade-off. Four examples of growing networks at step 24 are illustrated for the final network of the 10×10 grid (gray network). Gray areas denote buffered network elements which determine coverage. From left to right: Growth snapshot from (1) a disconnected strategy with high coverage and directness. Due to being disconnected, this strategy is not useful for most applications; (2) a connected strategy with high coverage and low directness—the red nodes highlight one of the large detours when directness is not prioritized; (3) a connected strategy with intermediate coverage and directness, exemplifying a trade-off between the two metrics; (4) a connected strategy with low coverage and high directness

network is sparse. If this was a bicycle network, everybody would be close to a bicycle lane and coverage would be high, but nobody who relies on the infrastructure's uninterrupted safety could use the network because it is not connected. It is exactly the situation that must be avoided (Szell et al. 2022).

The connectedness condition implies that we will never create a new connected component when growing the network. While this condition is unrealistic in uncoordinated or in random network growth, Szell et al. (2022) have shown that it is useful for *planned* network growth over the spatial scale of a city, where on average at most four components exist over only few initial growth steps before consolidating into one component. Indeed, connectivity was a high priority behind Seville's planned network which practically grew as one component (Marqués et al. 2015). Over larger spatial scales, our framework could be applied to each single component independently, or it would need to be extended to account for multiple components. We deem it adequate here that the fundamental performance metric of connectedness is fulfilled by construction and will therefore not be considered further.

Besides the two relevant performance metrics of coverage and directness, we consider several other optimization metrics in a greedy process where we grow by optimizing the optimization metric step by step, i.e., link by link. As we explore the growth of *connected* planar networks from scratch, potential performance metrics for disconnected networks such as the number of connected components or the length of the largest connected component are not relevant (Szell et al. 2022). The two selected performance metrics of coverage and directness coincide with two of our optimization metrics: For example, we can grow by optimizing directness step by step and observe the effect on directness itself, expecting high performance that is locally optimal by construction but may or may not be globally optimal. In addition, we consider optimization metrics that are often used for spatial networks, like closeness and betweenness centrality (Barthelemy 2022). Some of these metrics allow a direct comparison with previously studied growth strategies of planar networks (Szell et al. 2022). We define all the performance and optimization metrics in the following subsections.

Coverage

Coverage is the area covered by the spatially buffered links of the network, see gray areas in Fig. 1. For buffering we use rounded line endings, which gives

$$\begin{aligned} A(l,d) &= 2ld + 2\frac{1}{2}\pi d^2 \\ &= \left(\frac{2l}{d} + \pi\right) d^2 \end{aligned} \tag{1}$$

for the area of a buffered link of length l with buffer size d. Buffer size heavily influences the results and must be chosen reasonably: For a buffer size too small, $d \ll l$, all links account for roughly the same amount of added coverage; for a buffer size too large, $d \gg l$, there is too much buffer overlap and links that do not lie on the periphery of the network will not account for coverage. In both these cases, little to no information would be gained, which makes the choice of an intermediate buffer size crucial. Following this argument, for a grid a reasonable buffer size is 1.5 times the link length l: for three consecutive links, the link in the middle is completely covered by the buffered outer links.

If coverage in this simple, purely areal definition was interpreted as a proxy for coverage of certain spatial elements like a population or locations of interest, it would assume a homogeneous distribution of such elements. Thus, additional refinements could make the coverage concept more realistic in specific applications, such as accounting for an empirical, non-homogeneous distribution of the spatial elements that are aimed to be covered.

As our chosen lengths l are arbitrary, in the following we report coverage in a normalized form where a normalized coverage value of 1 means that the final network's covered area is reached.

Adaptive coverage

As explained above, the buffer size should be intermediate. However, choosing a single constant buffer size cannot account for different network scales. For example, road networks are typically hierarchical, having different meaningful scales (Batty 2006), which also holds for other networks such as the cardiovascular system or the branching structure of a tree (West and Brown 2005).

This issue can be addressed with adaptive coverage, where the buffer changes throughout the growth process. Starting with a large buffer size that is decreased over time, different scales can be covered iteratively. In the context of human transport networks, this idea allows building first a "network of main roads", i.e. a sparse network covering the whole area crudely, and then iteratively adding refined "side roads" that cover more and more nooks.

To implement adaptive coverage, we start from a large buffer size and reduce it whenever an added link contributes less than a given threshold to the area of the buffered network. As mentioned previously, the area of a buffered link of length l with a buffer size d is $A(l,d) = (\frac{2l}{d} + \pi)d^2$. As our growth process imposes connectedness, with each new link there will be necessarily a minimum share of the new link's buffered area with the existing coverage, which is the area of the circle πd^2 around at least one existing connecting node. We therefore consider the ratio $r = \frac{\Delta A}{A(l,d)-\pi d^2}$ between the area ΔA that is effectively added and the area $A(l,d) - \pi d^2$ of the buffered link minus the necessary overlap πd^2. If at any step the ratio r falls below a specified threshold (which we set to 0.1, but also perform robustness checks for other values), the buffer size is halved for the next step.

This procedure is adequate if the network intuitively resembles a simplified road network, i.e., if nodes are not too clustered but roughly homogeneously distributed in space. For this reason, in practical applications to transport networks, geometric or topological preprocessing would be required to simplify areas with high auxiliary node densities such as roundabouts or large intersections (Boeing 2017, 2025; Fleischmann et al. 2025).

Directness

We compute directness D as the average of the ratio between Euclidean distance $d_E(i,j)$ and shortest network path length $d_G(i,j)$ over all node pairs i and j:

$$D = \left\langle \frac{d_E(i,j)}{d_G(i,j)} \right\rangle_{i \neq j} \tag{2}$$

If $D = 1$, any path between any pair of nodes is straight, i.e. identical with the beeline. In other words, the more direct a network is, the shorter the path from any node to any other node. In the literature, this value (or the inverse of this value) can be found under many other names: detour factor, detour index, detour ratio, circuity, route factor, and straightness (CROW 2016; Reggiani et al. 2021; Barthelemy 2022). From a geometric perspective, *D* is also equivalent to the average sinuosity of the shortest paths. Directness is typically used to determine how efficient a spatial network is, as the closer *D* is to 1, all else being equal, the faster travelers will reach their destinations. The choice of the origins and destinations impacts the final estimation. Some studies take random origins to sample the directness (Giacomin and Levinson 2015), others use trajectories from real trips to determine the typical directness experienced (Yang et al. 2018). For computations of accessibility, another way of finding meaningful origins and destinations is selecting households and amenities, respectively (Lowry et al. 2016). Since here we work with synthetic networks, and do not want to depend on an arbitrary choice of selected origin–destination pairs, we consider the all-to-all network between all pairs of nodes in line with the literature (Vybornova et al. 2023). For a grid of size $N \times N$, we thus compute directness *D* as:

$$D = \frac{\sum_{i=1}^{N}\sum_{j=1}^{N}\sum_{k=1}^{N}\sum_{l=1}^{N} d(i,j,k,l)}{N^2(N^2-1)} \tag{3}$$

where

$$d(i,j,k,l) = \begin{cases} 0 & \text{if } (i,j) = (k,l) \\ \frac{\sqrt{|k-i|^2+|l-j|^2}}{|k-i|+|l-j|} & \text{otherwise} \end{cases} \tag{4}$$

For a grid, the directness decreases and converges with increasing size to $D \approx 0.79$. The CROW Manual (CROW 2016) uses this argument, together with the imperfection of empirical street networks which allow for higher directness, to argue for exactly this recommended minimum value for directness (the manual actually argues with a maximal detour factor of 1.26, which is equivalent to the inverse of the directness $D = 1/1.26 \approx 0.79$).

Directness is highly non-linear as it depends on the relative position of each node to all other nodes. The only ways in which a network can have the maximum directness $D = 1$ is either by being fully connected or by having only straight paths. For planar networks, which transport networks are to a large extent (Boeing 2020), these are Fig. 2, the fully connected network with four nodes K_4; the family of path networks P_N of arbitrary length *N*, which includes K_2; and the family of one/two-sided fan networks $F_{1/2,N}$ which are the unions of one/two nodes and a path network P_N (Brandstädt et al. 1999). One-sided fan networks $F_{1,N}$ include K_3. All other K_N for $N > 4$ are not planar (Kuratowski 1930). Additional non-intersecting, curved links could be added to P_N or to $F_{1/2,N}$ until reaching maximal planarity, but would be irrelevant as *D* only considers lengths of shortest paths.

While coverage is monotonously non-decreasing with growth, this behavior is different for directness. A larger network will usually have a lower directness than a smaller one, as it is harder to keep a planar network close to being fully connected (and thus

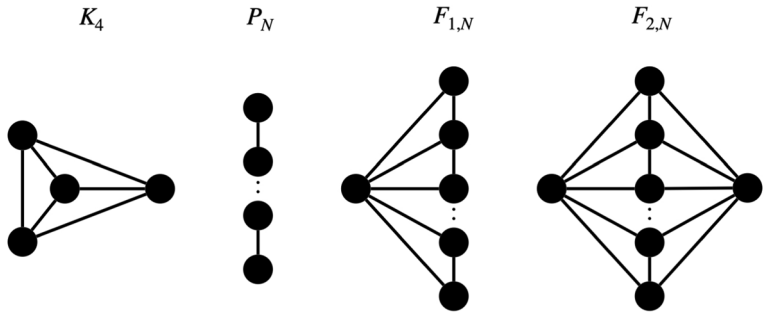


Fig. 2 The set of all planar networks with directness $D = 1$. K_4: The fully connected network with four nodes; P_N: The family of path networks of arbitrary length *N*, which includes K_2; $F_{1,N}$: The family of one-sided fan networks, which includes K_3; and $F_{2,N}$: Two-sided fan networks

direct) than a smaller one: a complete network K_N has $\frac{N(N-1)}{2} \sim N^2$ links, while a maximally planar network has only $3N - 6 \sim N$ links.

Relative directness

We want to evaluate the different orders in which a fixed set of links can be added to a fixed set of nodes. Thus, the directness for each node pair is capped to the directness for the node pair in the final network. From this observation, we compute the relative directness D_R, where instead of comparing the Euclidean distance to the shortest network path length, it compares the actual shortest network path length $d_G^A(i,j)$ to the final shortest network path length $d_G^F(i,j)$:

$$D_R = \left\langle \frac{d_G^A(i,j)}{d_G^F(i,j)} \right\rangle_{i \neq j} \tag{5}$$

If $D_R = 1$, the path between any randomly selected pair of nodes is as short as it is on the final network.

Out of all considered metrics in this study, D and D_R are the most computationally demanding ones, because the shortest path computations, already computationally expensive in themselves, have to be conducted for every connected pair of nodes, for every tested link, at every growth step, yielding a computational complexity of $O(L^5)$ where L is the number of links.

Closeness

Closeness centrality (Sabidussi 1966) measures how close (in terms of shortest path length) a node *u* is to any other node of the network,

$$C_C(u) = \frac{N-1}{\sum_{u \neq v}^{N-1} d(u,v)} \tag{6}$$

where *N* is the total number of nodes, and *d*(*u*, *v*) is the shortest path length between nodes *u* and *v*. Since the closeness centrality is only defined on nodes and we need to optimize by links, we compute the link closeness centrality $C_C^E(e)$ as the mean of the closeness of the nodes (e_1, e_2) connected by the link *e*:

$$C_C^E(e) = \frac{C_C(e_1) + C_C(e_2)}{2} \tag{7}$$

Betweenness

Link betweenness centrality (Girvan and Newman 2002) measures how frequently a link (*i*, *j*) is included in the shortest paths between any pair of nodes of the network,

$$C_B^E(e) = \sum_{s \neq t} \frac{\sigma_{st}(e)}{\sigma_{st}} \frac{2}{N(N-1)} \tag{8}$$

where σ_{st} is the number of shortest paths between node *s* and node *t*, and $\sigma_{st}(e)$ is the number of shortest paths between node *s* and node *t* that include the link *e*. The normalization factor $\frac{2}{N(N-1)}$ is the number of node pairs $\binom{N}{2}$.

Assuming an equal number of travelers coming from and going to every node, the link betweenness centrality is a proxy for traffic flow and the reason why it is widely used in street network analysis (Barthelemy 2022).

Require: Planar network G, network metric $f(G)$
  Initialize G_1 with the link of highest closeness
  for (# of links in G) - 1 **do**
    Find the list of links $l \in G$ connected to G_i.
    for **e** in l **do**
      Add **e** to G_i, producing $G_i^{\mathbf{e}}$
      Compute $f(G_i^{\mathbf{e}})$
    end for
    Add **e** to G_i for which $f(G_i^{\mathbf{e}})$ is maximal
  end for
  return Order of the links added

Algorithm 1 Compute the optimal order of growth of the network *G* for the metric *f(G)*

Growth strategies

Random growth

The simplest growth strategy is a random ordering of links. On the one hand, this strategy serves as a null model, or a lower bound, for other strategies. On the other hand, random growth can be a reasonable model for the observed growth of certain networks such as bicycle networks (Szell et al. 2022).

For all random growth strategies we start growing at the link with highest closeness centrality. In symmetric networks where this link is not unique, we choose a specific seed link out of the possible candidates, for all growth strategies—without loss of generality due to symmetry. Then, we consider adding a link randomly from all possible links from the final network that would preserve connectedness (due to the connectedness condition). We follow this procedure at each step until all links from the final network have been added.

Greedy optimization

Greedy optimization considers at each step the optimal choice for that particular step only. It means always choosing the local optimum, not necessarily leading to the global optimum. See Algorithm 1 for how we implement our greedy optimization. When multiple links provide the same maximal value, we choose one among them randomly.

One point of our study is to show the pros and cons of greedy strategies relative to random and manual strategies, and whether it is possible to "get away with" a greedy

strategy instead of trying to solve the network growth problem with a more advanced heuristic or a global approach (Wiedemann et al. 2025).

The advantages of greedy optimization are its ease of use and implementation, and its explainability. The latter is crucial for applications of the method, such as support in urban planning. The disadvantages of greedy optimization are well known (Cormen et al. 2022). As it is a simple method taking very limited knowledge into account, it can easily give suboptimal results due to falling into local optimum traps, especially for highly non-linear cost functions.

We show here the results for what we call an "additive" strategy, also called "forwards percolation", "forward" or "bottom-up" strategy (Ballo et al. 2024; Steinacker et al. 2025; Wiedemann et al. 2025), where we start from an initial network and grow it by adding the link which yields the network with the maximal chosen optimization metric at each step until reaching the final network. Apart from that, we study a subtractive strategy, also called "backwards percolation", "reverse" or "top-down" strategy (Ballo et al. 2024; Steinacker et al. 2025; Wiedemann et al. 2025), where we instead start from the final network and remove the link which yields the network with the maximal chosen optimization metric at each step until reaching a single link. In Sect. 3 we show results for the additive strategy as it usually outperforms the subtractive strategy (see Sect. 3.6).

Manual growth orders

We want to compare our results from the greedy growth with other growth methods. Overall, these methods fall into three classes: random strategies as a lower bound; greedy strategies analyzed for each metric described in Sect. 2.2; and manual orders, introduced below.

The general idea behind manual ordering is the incorporation of design strategies that imagine a human planner who aims to implement a global optimum for specific metrics, or certain macroscopic structures first such as a "ring road" or other hierarchical structures (Barthelemy et al. 2013; Walker 2024), being able to circumvent local optimum traps, see Fig. 3. Thus, the manual orders can serve as an upper bound for single metrics, or as an implementation of a globally optimal or "human-designed" strategy.

Since here we focus on directness and coverage, we manually design globally optimal growth orders for each of these two performance metrics, providing the best growth we can expect. This way we expect to discover how far away our greedy strategies are from the global optimum. We also design an intermediate "hierarchical order", explained

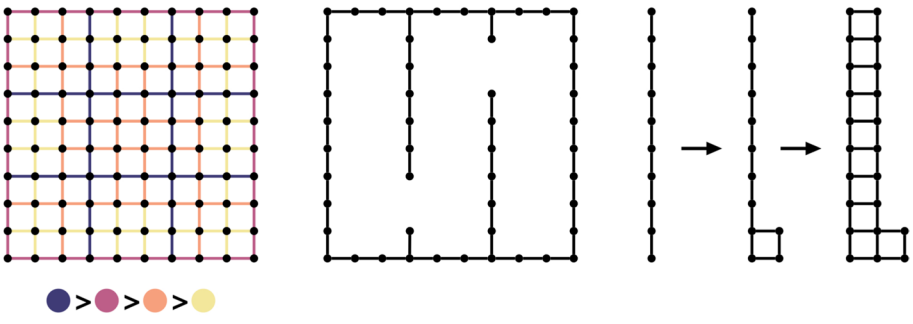


Fig. 3 The three manual growth orders: Left: Manual order H following an implicit hierarchy of the network, from blue to pink to orange to yellow. Center: Manual order C optimizing for coverage. Shown here is step 48 out of 180, when reaching the full coverage. Right: Manual order D optimizing for directness

below. We are able to design globally optimal manual orders for the regular grid, but not for other test networks, as their added complexity and inflated solution space make it infeasible to argue for global optimality. For comparability with the other growth strategies, in each of our manual orders we keep the constraint of connectedness, and we start with a link similarly central as for the random and greedy strategies—we have verified that growth processes remain comparable.

To explain the link order in manual growth orders, we define rows and columns for horizontal and vertical links respectively, starting from the top left of the grid. For instance, the horizontal link connected to the top left corner of the grid will be on the first row and between the first and the second column. The complete ordered list of links for each manual order is available in the code (Sebastiao 2024a).

Manual hierarchical order H. We design a manual hierarchical order "H" to represent what a planner could consider a good hierarchical growth strategy. Without any additional knowledge of links, this order is a good compromise between different hard-to-define factors a planner could have in mind, i.e., balancing coverage and directness through iteratively refined macroscopic structures (like "ring roads", etc.). This order does not follow any realistic growth plans from the transport planning literature, it only exemplifies how general design guidelines could be interpreted. We create different levels of priority from the implicit hierarchy of the network (see left panel of Fig. 3) and start by building from the highest hierarchy to the lowest, keeping the growth symmetrical at all stages. The complete ordered list of links is reported in Fig. 16.

We first grow the blue links, filling out the fourth and seventh rows and columns. We then grow the pink links, forming the border of the grid. The finished border together with the blue links now forms a coarse grid of 3 by 3 links on the network. Next, we grow the orange links, filling out the third and eighth columns and rows, then adding every missing link in the center between them. Finally, we add the remaining yellow links.

Manual coverage order C. The manual coverage order "C" is providing the upper bound for coverage, for the fixed buffer size of 1.5 times the link length. Until reaching full coverage, this order follows closely greedy optimization for coverage. We start by growing the border from one corner, until we miss only three links to close the square. We then grow 6 links of long lines along the fourth and the seventh columns on opposite sides of the square, and add a single link on their complementary side. We reach the full coverage by adding the missing three links from the border, see the central panel of Fig. 3. Next we fill up the fourth and seventh columns, then the fourth and seventh rows. Once a coarser grid of 3 by 3 squares is formed, the strategy continues in the same way as "H" from the orange and yellow links.

Manual directness order D. The manual directness order "D" is providing the upper bound for directness. We first grow a straight line on the first column, as it is a network with maximum directness (see Fig. 2). Once the line cannot be expanded further, we add a link on the first row, on the second row, and connect both by adding a link on the second column, creating a square. We repeat the process square by square until reaching the end of the second line, and continue the same way until completing the grid (see right panel of Fig. 3).

Results

We first describe our results on a square grid. Without loss of generality, we select a square grid with 10×10 nodes. The chosen buffer size for both the Coverage strategy and for all subsequent analysis is 1.5 times the length of a link, as argued above. For the Adaptive coverage strategy, we initiate with a doubled buffer size of 3 times the link length, halving it to 1.5 and then again to 0.75 times the link length (see Sect. 3.5 for more details). We set the threshold for the adaptive coverage at 0.1, later checking also other values.

As the network has multiple axes of symmetry and all links have the same length, there can be steps in the greedy optimization process when several links are equally desirable, and the next link is chosen randomly among these equally good options. To observe the impact of these random choices, we run each greedy strategy 50 times. For the random strategy, we perform 1000 runs, as this strategy can theoretically give results from the entire solution space. We report the average performance of each strategy on coverage and directness in the left and right columns of Fig. 4, respectively. The outcomes of all single runs are reported via gray curves. We also visualized the same data in Fig. 17a and b, respectively, to allow a more direct comparison of the curves.

Each of our two performance metrics (coverage and directness) provides for each strategy a value for each growth step, yielding a curve over all growth steps. As we aim to quantify the performance across the entire growth process, the Area Under the Curve (AUC) is an appropriate tool. Since the final network is pre-defined, the initial and final values of all curves are known, while the interesting variations happen during intermediate stages of growth. While it is possible to visually compare curves, as in Fig. 4, the AUC summarizes a whole curve with a single value, and thus allows comparison between multiple strategies in coverage-directness space, see Fig. 5. Additionally, the semantics of AUC align naturally with our objective of summarizing strategy performance across all possible levels of built length, encompassing a range of scenario-dependent levels of infrastructure investment.

From a planning perspective, the interpretation of AUC differs slightly between coverage and directness. For coverage, the interpretation is straightforward: higher AUC values correspond to lower investment requirements to achieve high levels of coverage. For directness, the interpretation is more nuanced due to non-monotonicity. A high AUC may obscure intermediate development stages characterized by very low directness, where partial investments create networks with poor average travel efficiency. For this reason, practical planning decisions should consider the full directness trajectory alongside AUC, rather than relying on the summary statistic alone.

AUC performance metrics are normalized in the range [0, 1]. To do so, we first normalize the built length and coverage values by min-max normalization, i.e., we compute the normalized values x' from the original values x using:

$$x' = \frac{x - \min(x)}{\max(x) - \min(x)} \tag{9}$$

We do not normalize directness, as it is by definition bounded within [0, 1]. As the initial and final values of our performance metrics are fixed, we normalize the computed AUC for our normalized variables by the AUC obtained with the optimal curve for our

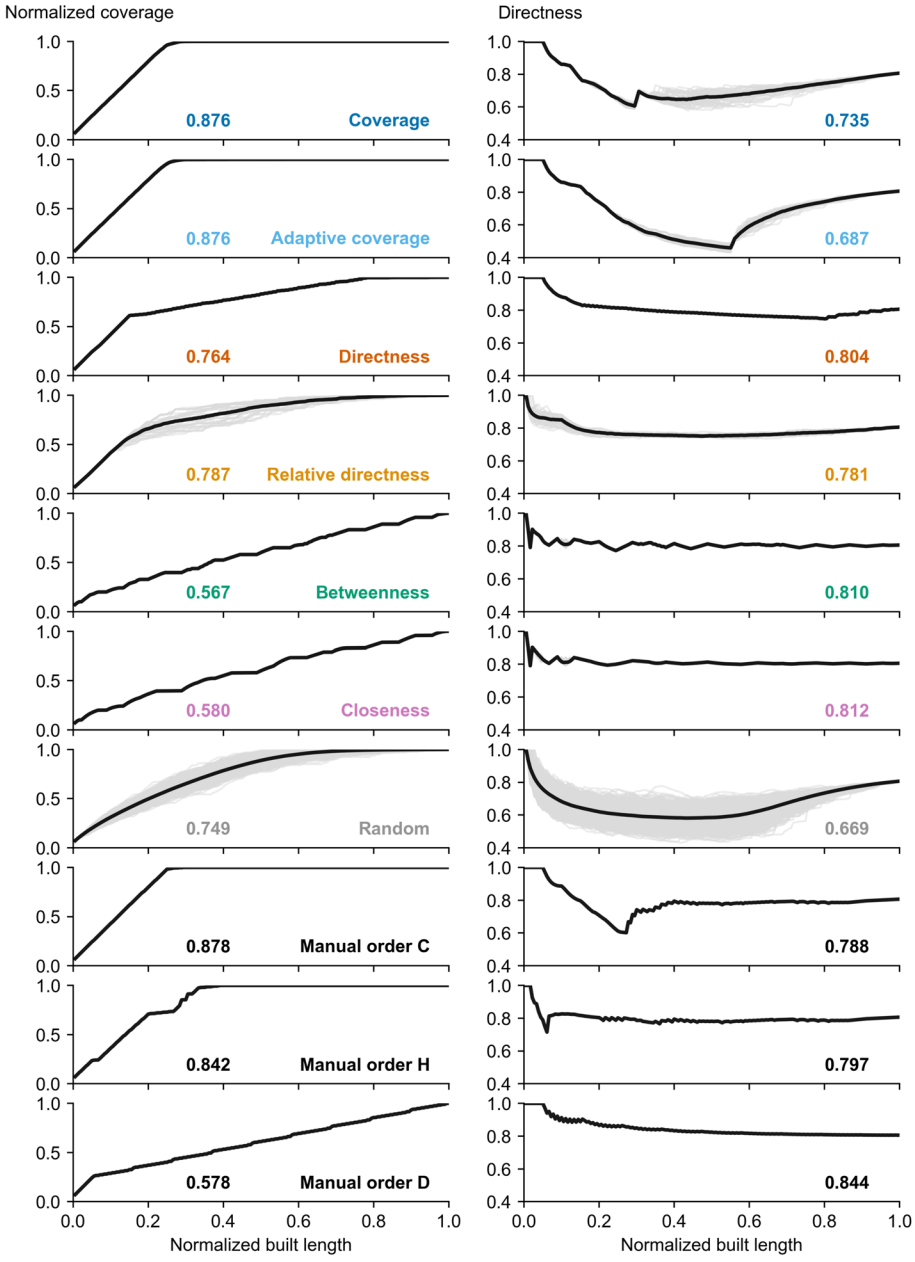


Fig. 4 Coverage (left column) and directness (right column) versus built length for each growth strategy (rows). All 1000 runs are shown in gray and average curves are highlighted in black. Inset values report the areas under the curve (AUC)

performance metrics, i.e., a curve starting at the given initial value, reaching the maximal value for all intermediate stages, and ending at the final value.

The AUC tool, as applied above, weights all growth steps equally. Since it could be argued that earlier growth steps are more important, we explore in Appendix E also a variant of the AUC with exponential discounting to give more value to earlier growth stages. This variant leads to similar results.

We report our main observations in the following subsections.

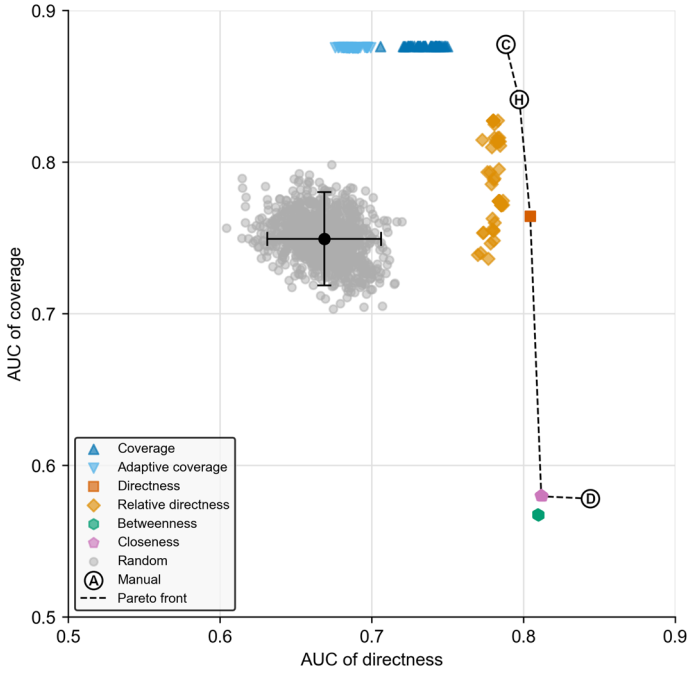


Fig. 5 AUC of coverage against AUC of directness for all additive growth strategies. The black dot is the mean Random strategy over 1000 runs, each error bar depicts two standard deviations. The random orders are over six standard deviations away from the Pareto front, showing that brute-forcing is infeasible. Coverage and Adaptive coverage yield maximum AUC of coverage, but so is the manual order C which also outperforms both in AUC of directness. The Pareto front is made up of all the manual strategies and Directness and Closeness. Betweenness and Relative directness are slightly sub-optimal

Greedy strategies outperform random strategy

We find that typical random orders are highly sub-optimal as their directness curves stay below most other growth strategies and their coverage curves are also suboptimal, Fig. 4. This suboptimality becomes more clear with the AUC values reported in Fig. 5 (cloud of 1000 gray markers) which are far and well separated from the Pareto front (computed for all considered orders, dashed line). In fact, quantifying the distance of the AUC values from the Pareto front, we find that the average random order (black dot) is at least six standard deviations away from the front for each axis (error bars in the black cross denote two standard deviations). Assuming normality, which is justified at least for coverage (Shapiro-Wilk test does not reject at $p = 0.89$), it means that on average the Random strategy would need to be run at least 500 million times before finding one order close to the Pareto front. This result shows that, although in theory a random ordering could recreate any other order high performing solutions are unlikely to be reached by brute force. Therefore, at least for the grid, if the aim is to reach a good level of directness or coverage (or both), it is clear that smarter strategies are needed.

The AUC of directness spreads slightly more than the AUC of coverage for random orders. Further, as visible in Fig. 4, while coverage increases monotonously and by a maximum possible value, Eq. (1), directness is non-linear and can feature long jumps triggered by the addition of one single link. These observations give a first hint that optimizing for directness is harder (see Sect. 3.4).

Concerning the spread of the greedy strategies, three of them are deterministic (Directness, Closeness, and Betweenness), leading just one marker being visible for each in Fig. 5 because the 50 runs produce the same AUC values. The three others (Coverage,

Adaptive coverage, and Relative directness) involve a strong random element leading to scattered AUC values from their 50 runs. The Coverage and Adaptive coverage orders (blue triangle markers) reach the same AUC of coverage but different AUC of directness, with Coverage always outperforming Adaptive coverage (all dark blue markers are further right than the light blue markers). Relative directness orders (orange diamond markers) are spread along both dimensions but more vertically, along AUC of coverage. The Directness strategy (red square marker) always reaches higher AUC of directness but most Relative directness orders reach higher AUC of coverage. Interestingly, both the Closeness and Betweenness strategies reach slightly higher AUC of directness than the Directness strategy (albeit with much lower AUC of coverage), implying that the Directness strategy is not globally optimal. We give a more detailed explanation of the scattering of AUC values in Sect. 3.4.

Manual orders outperform greedy strategies

When accounting for manual orders, we first find that the only greedy orders on the Pareto Front are Directness and Closeness (red square and purple pentagram markers in Fig. 5, respectively). The Directness strategy has slightly more AUC of directness than manual order H, while the Closeness strategy has slightly more AUC of coverage than the manual order D.

The manual orders C and D lie on the Pareto Front, which was expected since we designed them to be optimal strategies. Manual order H, despite only following a hierarchical heuristic, is also on the Pareto Front, showing the performance of a simple intuitive heuristic.

The manual order C is on the same maximal level of AUC of coverage as the Coverage or Adaptive coverage orders, but it beats all of them in terms of AUC of directness. This difference in directness is caused by the randomness of choices made by the greedy strategies when full coverage is reached (see Sect. 3.3).

The manual order D has much worse AUC of coverage than all the Directness and Relative directness orders, but outperforms them in terms of AUC of directness. This gap can be explained by the local optimal traps a greedy algorithm can fall into (see Sect. 3.4).

Coverage strategies reach full coverage, then become random

When growing a planar network with straight links, eventually a point is reached where no additional link adds more coverage. We call this point "reaching full coverage". The number of links needed to reach full coverage depends on the buffer size. Let us consider our 10×10 grid with 180 links.

If the buffer size is at most 0.5 times the link length, all 180 links (100%) are needed to reach full coverage, since each new link will add coverage, no matter how many previous links there are. In the other extreme, if the buffer size is at least 4.5 times the link length, only the 36 links (20%) on the periphery are needed to reach full coverage, since links on the inside would not add any extra coverage. If the buffer size is intermediate, the network reaches full coverage with an intermediate number of links.

In the case of the selected intermediate buffer size of 1.5 times the link length, with a single connected component the full coverage can be reached with 50 links (28%), which is achieved by the manual order C, see central panel of Fig. 3. Starting from the link with

largest average closeness, 58 links (32%) are necessary to reach full coverage, which is achieved by the additive Coverage strategy, see left panel of Fig. 6.

After reaching full coverage, the remaining links are equivalent when optimizing for coverage. Thus, every link added after the first 59 is chosen at random. This random choice results in additive Coverage orders sharing the same curve for coverage, but varying in directness after reaching the full coverage at 33% of total length built (see Fig. 4). In contrast, manual order C follows the hierarchical order H (see Sects. 2.3 and 2.3) after having reached full coverage optimally. By switching the objective here, manual order C manages to be well-performing on both directness and coverage.

Let us now consider what happens with a subtractive strategy. As explained in Sect. 2.3, subtractive strategies start from the final network, removing the link of the least importance for the chosen metric. When optimizing for coverage however, any link on the inside of the network removes the same amount of coverage and is thus equivalently "least important". Thus, the first links removed are random links from the inside. The remaining reverse growth process, while optimizing for coverage, follows highly constrained, typically meandering shapes (see right panel of Fig. 6). These shapes have worse directness than in the additive Coverage strategy (see left panel of Fig. 6), leading to the subtractive Coverage strategy underperforming on both coverage and directness, compared to the additive Coverage strategy (see Sect. 3.6).

Directness falls into a local optimum trap

As directness is heavily non-linear, locally optimal choices can easily trap the greedy optimization process. Cycles are a good example to understand why, as shown in Fig. 7. With a fixed link length, the longer the cycle (i.e., the more links it contains), the lower is its directness. Therefore, when the Directness strategy is additive, cycles are only created whenever they are one step away to be completed—a state that is increasingly harder to reach the longer the cycle is, see left panel of Fig. 7. However, when the Directness strategy is subtractive, the longer the cycle, the bigger the impact of removing one single link, which can result in holding on to cycles for too long. In other words, the subtractive algorithm grows large cycles first. Both additive and subtractive issues are undesirable behaviors stemming from local optimum traps for directness.

Local optimum traps also explain why strategies based on centrality (i.e., Betweenness and Closeness) perform in directness better than the Directness strategy itself. As the centrality value of each link is fixed by the final network, centrality-based strategies

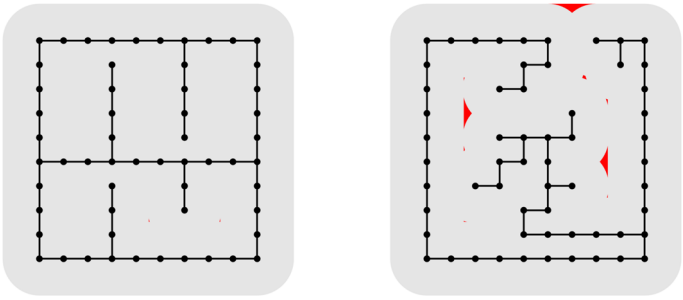


Fig. 6 Coverage strategy snapshot illustrating the effect of randomness. Left: additive, Right: subtractive. Shown are two typical networks at step 58. The additive strategy has almost reached full coverage. At the same stage, the subtractive strategy is non-deterministic and has reached a meandering form that is underperforming in both coverage and directness. Highlighted in red is the difference to the full coverage

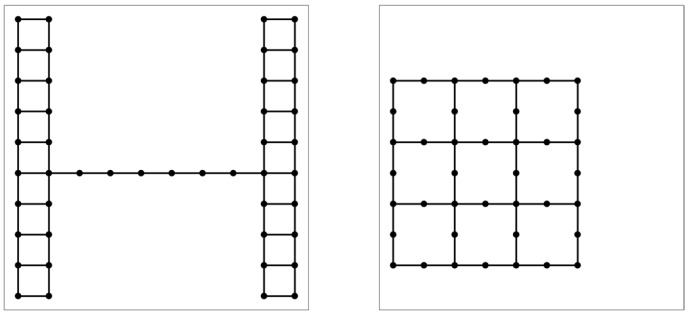


Fig. 7 Directness strategy snapshots illustrating local optimum traps. Left: additive, step 63. Right: subtractive, step 48. The additive strategy avoids building cycles, but if it has to, they will be as short as possible. On the contrary, the subtractive strategy avoids breaking up cycles, leading to large cycles constructed first

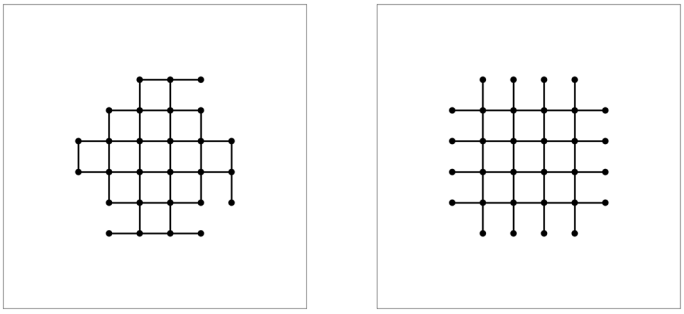


Fig. 8 Centrality-based strategy snapshots illustrating "well-behaved" growth. Network at step 40 with the Betweenness (left) and Closeness (right) strategies. Both growth processes create small cycles from the center, which keeps directness relatively high and stable (at the cost of low coverage)

follow a simple heuristic, creating small cycles from the center, see Fig. 8. They keep a value close to the final directness throughout the growth process, with some oscillations, see the Betweenness and Closeness curves in the right column of Fig. 4. On the other hand, the Directness strategy locks itself into a local optimum trap after constructing a straight line of maximal directness, see left panel of Fig. 7. This trapping leads to a slow but steady decrease of the directness for most of the growth process, see the Directness curve in the right column of Fig. 4. Putting higher value on early growth stages via exponential discounting (see Appendix E) makes the Directness strategy perform better, but does not change the relative position of most strategies overall.

While we can easily optimize to get the best or worst AUC of coverage possible using greedy growth, we cannot do the same for AUC of directness due to its non-linear nature. However, a greedy process could account for the 4-cycles in a grid lattice by increasing the time horizon from one to three future steps (see Sect. 3.7 for other examples).

Coverage vs. adaptive coverage, directness vs. relative directness

In this section, we explore the variants of coverage (adaptive coverage) and directness (relative directness).

Adaptive coverage orders are worse than Coverage orders for both AUC of coverage and directness, compare light blue with dark blue markers in Fig. 5. To measure coverage, we choose a single buffer size. The area measured is dependent on both the network structure, and the selected buffer size. Since the Adaptive coverage strategy optimizes

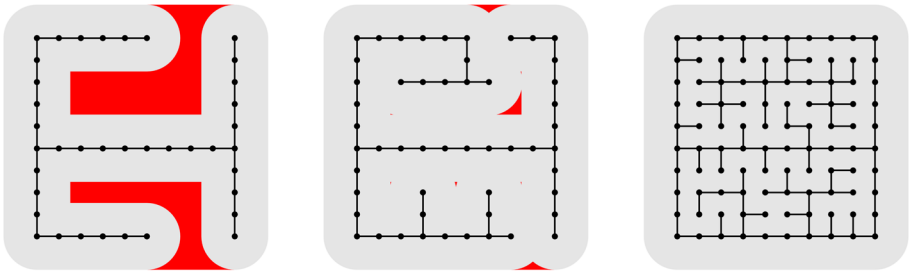


Fig. 9 Adaptive coverage strategy snapshots. From left to right: network at step 37, 51, and 101, respectively. Full coverage is reached only at step 101. Highlighted in red is the difference to the full coverage

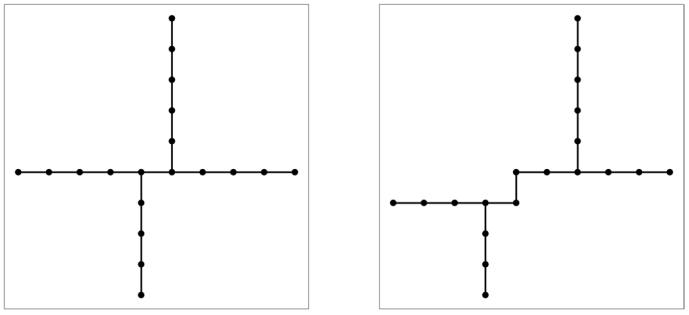


Fig. 10 Relative directness strategy snapshots illustrating random variations. Two examples of a network at step 18 having maximal relative directness

across different scales through growth via a buffer size change, we expect its AUC of coverage—which is measured for one single buffer size – to be worse than the AUC of coverage for the Coverage strategy which only optimizes for that single buffer size (see Fig. 4). Both strategies optimize by covering the greatest area possible, which for a lattice results in a sparse network of high coverage but low directness. While the Coverage strategy reaches full coverage with 59 links, see left panel of Fig. 6, the Adaptive coverage strategy does so only when it reaches 101 links, with the buffer size decreasing at 37 and 51 links, see Fig. 9. After reaching full coverage, 70 links are randomly added in the Adaptive coverage strategy versus 121 links in the Coverage strategy. Consequently, the decrease of directness lasts longer for the Adaptive coverage orders, see Fig. 4.

Results are shown for a threshold of 0.1 (see Sect. 2.2). We also tested a threshold of 0.05 and 0.2, resulting in similar values, thus confirming robustness: the average AUC of Coverage between thresholds is almost constant (varying by 0.0003), and there is more variation in the AUC of Directness between trials for a single threshold (0.005) than between thresholds (0.003).

We defined previously relative directness as the directness where the Euclidean distance is replaced by the final shortest network path length. Given that we work on a grid, the final shortest network path between any node pair will be a sequence of horizontal and vertical links, implying that the path's length is then equivalent to the Manhattan distance. Consequently, maximal relative directness is achieved by different link orderings with varying AUC of coverage, see Figs. 10 and 4.

Additive strategies outperform subtractive strategies

As the subtractive strategies start with more information on the network, one could expect these strategies to outperform the additive ones. However, when taking together additive and subtractive greedy strategies and comparing them, we find that no subtractive strategy comes to lie on the Pareto front for the grid, except centrality-based strategies that are equivalent in both orders, see Fig. 18.

The lower performance of subtractive strategies can be explained by already noted behaviors: As optimizing for coverage after having reached maximal coverage implies random link choice, subtractive Coverage and Adaptive coverage strategies start randomly. The optimization phase is then heavily constrained by the random configuration of maximal coverage obtained. On the other hand, additive Coverage and Adaptive coverage strategies start by efficiently covering the area, creating a more optimal growth, see Fig. 6.

The additive Directness strategy starts with a straight line along the first link, then creates perpendicular straight lines at the extremities, from where small cycles are created. Subtractive Directness strategies, on the other hand, can remove different links that are all producing subnetworks of equal directness, creating the variability in AUC of coverage. From those multiple configurations typically emerge large cycles that are harder to break, see Fig. 7. While both behaviors are suboptimal, additive strategies outperform subtractive ones.

Additive strategies outperform subtractive ones on all tested networks, except for the subtractive Adaptive coverage and Coverage strategies on the diagonal grid and the radio-concentric networks. The difference can be small: On the grid, the highest AUC of coverage for the additive and subtractive Coverage strategies is 0.87647 and 0.87633, respectively, and the highest AUC of directness for Directness strategies is 0.80438 versus 0.80157. The difference between additive and subtractive strategies is much smaller than the difference between random and greedy strategies. Subtractive strategies can perform better for specific conditions, depending on the chosen performance metric and the network structure.

Results for other stylized networks are similar

Expanding our analysis, we run the random and greedy strategies on three other stylized networks, see Fig. 11. All networks have similar size, around 100 nodes and 180 links. The selection of stylized networks follows typical urban street patterns, explained in Appendix C.

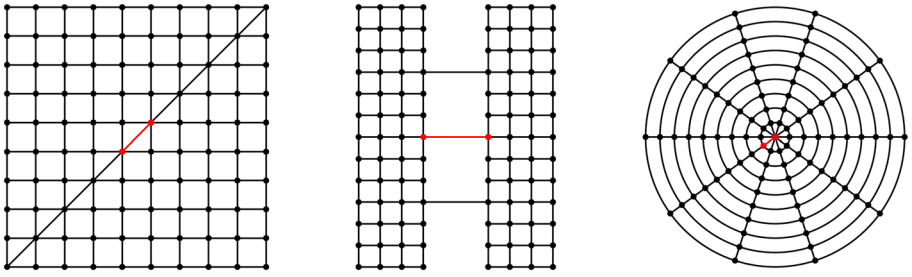


Fig. 11 Additional tested stylized networks. Left: Grid with a diagonal, Center: Three bridges network, Right: Radio-concentric network. Highlighted in red are the initial link and nodes for the additive strategies

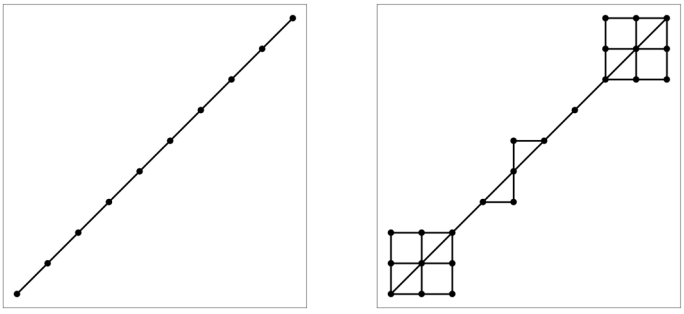


Fig. 12 Grid with a diagonal illustrating the Directness strategy closing 3-cycles early. Left: step 9, Right: step 37

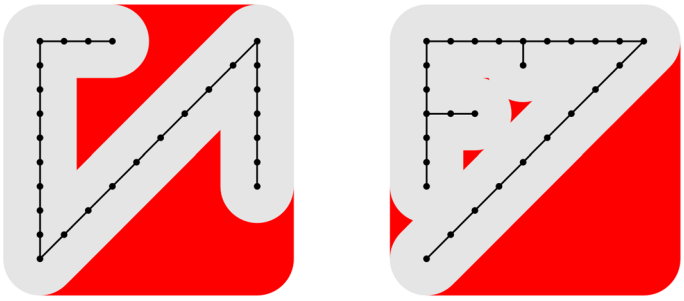


Fig. 13 Grid with a diagonal illustrating the Adaptive coverage slightly outperforming the non-adaptive Coverage strategy in terms of coverage. Left: Adaptive coverage strategy, covering 68.2% of the full coverage. Right: non-adaptive Coverage strategy, covering 65.9% of the full coverage. Both at step 27. Highlighted in red is the difference to the full coverage

Grid with a diagonal

The "Grid with a diagonal" network is a 10×10 grid with links along one diagonal, resulting in 189 links. Links along the diagonal are $\sqrt{2}$ times longer than the other links. This network is the most similar to the regular grid, with the diagonal acting both as added asymmetry and a "shortcut" through the network. The initial link is at the center of the diagonal.

In contrast to to the regular grid, the AUC of directness of all growth strategies is higher for the Grid with a diagonal for almost all growth strategies (except Coverage and Adaptive coverage), see Fig. 18c. This can be explained by the higher directness of the final network ($D \approx 0.81$ for the grid versus $D \approx 0.85$ for the grid with a diagonal), as the Grid with a diagonal has nine additional links compared to the grid. The Directness strategy outperforms all other strategies on directness. After growing the diagonal, one adjacent link is added. Since cycles adjacent to the diagonal cycles contain three (rather than four) links, the Directness strategy can find these 3-cycles and create them throughout the growth process, in this manner avoiding some local optimum traps, see Fig. 12.

The Adaptive coverage strategy slightly outperforms the non-adaptive Coverage strategy on coverage. Since the starting link is on the diagonal, both these growth strategies first construct all links along the diagonal. Thereafter, the Adaptive coverage strategy starts building the periphery, because for a larger buffer size these are the only links that improve the coverage, see left panel of Fig. 13. However, for the non-adaptive Coverage strategy, the links on the periphery adjacent to the diagonal act as a local optimum trap due to the small buffer size. Building a new side for the periphery to grow is more costly

than continuing growth on a single side of the periphery. This behavior leads the non-adaptive Coverage strategy to delay constructing the periphery, see right panel of Fig. 13.

Three bridges

The "Three bridges" network consists of two 4×13 nodes grids connected by three links 3 times longer than the other links, resulting in 104 nodes and 177 links. While most of the network is still grid-like, the three bridges are critical links of high value for directness, coverage, and the shortest path affecting the centrality-based metrics. The initial link is the central bridge.

Obtained results are similar to the regular grid, see Fig. 18e. The gap in AUC of coverage between the Coverage and Adaptive coverage orders and some other orders is larger. Since the three bridge links contribute a significant share to the total coverage, random orders have lower AUC of coverage. While centralities are limited by the growth diffusing from the center, the bridge links have high betweenness, thus they are more quickly built by the Betweenness strategy than by the Closeness strategy, see Fig. 14. As the Directness strategy grows in a similar manner as the regular grid, see Fig. 7, bridges are added at the latest stages. The consequences are compensated for by the reduced amount of connections between the extremities of the network making shortest paths longer, and thus lowering the overall directness ($D \approx 0.8064$ for the regular grid, while $D \approx 0.7923$ for the Three bridges network).

Radio-concentric

The "Radio-concentric" network is made up of 10 radial lines intersecting with 9 circles, resulting in 91 nodes and 180 links. Links along the radial lines are all as long, while links along the circles are of length $2\pi k/10$ radial link length, with k numbering the circles from the center ranging from 1 for the smallest to 9 for the largest, from around 0.63 radial link length to 5.65 radial link length. The initial link is inside the first circle, along a radial line.

Obtained results are similar to the regular grid, see Fig. 18g. The range of values of AUC of directness is larger though, reaching less than 0.6 for random strategies and over 0.9 for some non-random strategies. This larger range is caused by multiple factors. First, the higher ratio of links per nodes increases the directness of the final network ($D \approx 0.8064$ for the regular grid, while $D \approx 0.8930$ for the Radio-concentric network). Second, shortest paths are less evenly distributed than on the grid. Third, the first circle

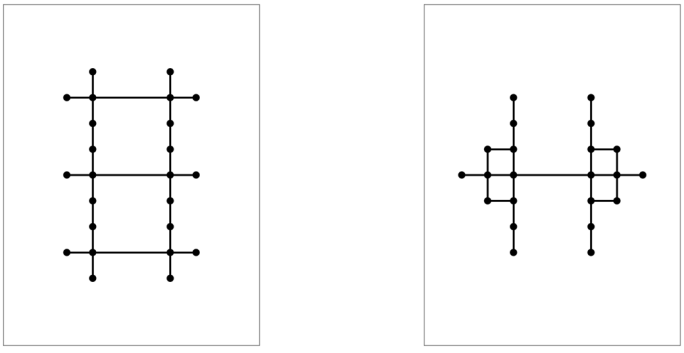


Fig. 14 All three bridges are built early with the Betweenness but not with the Closeness strategy. Left: Betweenness strategy, Right: Closeness strategy. Both at step 25

is made up of three-link cycles that are easier to grow for the Directness strategy, as explained in Sect. 3.4 (see Fig. 15). Finally, angles between links from different radial lines are smaller than the $90°$ angle of the grid, reducing the gap between Euclidean distances and shortest network paths.

Robustness to exponential discounting and distortions

Benefits of transport infrastructure is cumulative: The earlier it is built, the larger are its total benefits for a fixed time horizon. Therefore, we check results for a growth variant where the early stages are prioritized, using exponential discounting. For operational details see Appendix E. We report results in Fig. 26. The results are similar to no discounting, compare with Fig. 18. All results are also robust to distortions of the grid, except for Centrality strategies, see Appendix B.

This robustness shows that effects from discounting or spatial noise are negligible or less dominant than the choice of optimization metric.

Discussion and conclusion

Our study explored systematically the growth of spatial networks and their performance in directness-coverage space, to provide fundamental insights into the trade-offs between these two key metrics in network design applications. We focused on transport applications, but our study could be relevant for other applications that involve spatial networks. Motivated by the transport perspective, we focused on the case of a given set of nodes with a final set of links between them that could be built over time, and on the order in which these links are added. We pursued and compared the peculiarities of several strategies down to specific network configurations trapped in local optima, involving random growth, greedy optimization of several metrics, and manual designs.

Overall, we confirm the intuitive expectation that greedy growth is generally outperforming random growth, and that well-designed manual growth strategies can surpass greedy strategies, even those that locally optimize for the global target metric. The main added benefits of our study is (1) *a quantification* of how much worse such strategies are, and (2) *an ordering* of several metrics in the directness-coverage space, showing which of them are worth pursuing through greedy optimization for which goal. For example, on a grid-like city plan, Directness always outperforms Relative directness, and Coverage always outperforms Adaptive coverage, in terms of being closer to the Pareto front.

In a practical design process, our study shows that, for example, it can make sense to first generate a link ordering based on greedy coverage optimization towards global coverage optimization, to be then taken as a starting point for a manual tuning to achieve

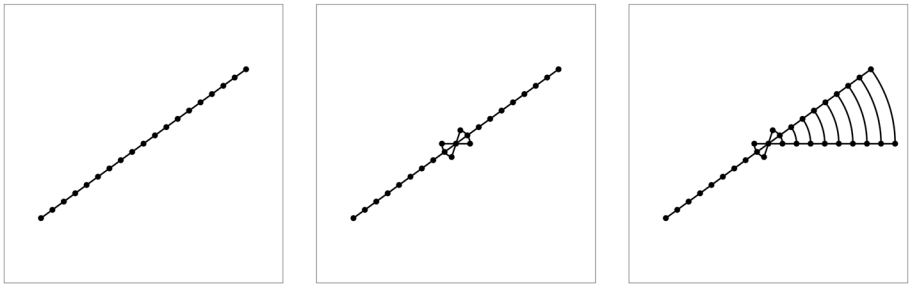


Fig. 15 The Radio-concentric network builds up radial links first with the Directness strategy. Left: step 18, Center: step 26, Right: step 42

higher directness. Holding for a number of stylized urban transport network topologies, we found that centrality-based greedy strategies tend to perform best for directness but are worse than random strategies for coverage, while coverage-based greedy strategies can achieve maximum global coverage as fast as possible but perform similarly bad for directness as random strategies.

Since the solution space is exponentially large and thus intractable, we quantify how greedy optimization—despite usually not providing the globally optimal solution—is giving much better results than random trials. In fact we quantify the large extent to which brute-forcing is infeasible already for small networks, where over 500 million random trials would be needed to find a solution close to the Pareto front. Techniques such as simulated annealing or evolutionary algorithms (Plumed et al. 2016) could yield better solutions than our greedy algorithms and could be a future algorithmic improvement. A further improvement could be the creation of a step-dependent optimization function, for example switching from a greedy coverage strategy to a greedy directness strategy. Such a switch could be justified if there are different objectives at different stages of the growth process.

Supervised learning frameworks for graph generation represent also a promising alternative approach to model network growth (Guo and Zhao 2022). Autoregressive (Kong et al. 2023), adversarial (You et al. 2018), and reinforcement-learning approaches (Zhou et al. 2019) have been proposed and tested in multiple domains and are potentially applicable to our scenario. Recent approaches emphasizing the temporal dimension (Touwen et al. 2024) or structural constraints (Madeira et al. 2024) of graph growth are particularly promising in relation to the task of transport network planning.

The intuitive insight of "manual beats greedy beats random" being confirmed in our context has practical consequences. In practice, the growth of transport networks can indeed be close to random or greedy. This feature is most prevalent in networks where the cost or friction of adding links is high and where single links are perceived as negligible for network overall function, such as in bicycle networks (Szell et al. 2022). Here, the addition of many single bicycle lanes can be seen as a politically contested competition for car parking space (Oldenziel and de la Bruhèze 2011; Gössling 2020; Henderson and Gulsrud 2019). Further, political competences can be fragmented: for example, a city might have an all-encompassing bicycle network plan in place, but individual districts may ignore it. We are aware of anecdotal evidence for this effect in cities like Paris, Vienna, or Copenhagen and Frederiksberg (Carstensen et al. 2015; Paulsen and Rich 2023; Henderson and Gulsrud 2019). Due to such effects, the addition of links proceeds by lowest political friction, which in network terms is indistinguishable from a random growth strategy (Szell et al. 2022), since these local decisions are not accounting for the wider network structure, thus do not follow a greater vision.

Short-term horizons further appear naturally in the context of bicycle networks, where urban or regional planners might aim to implement a time-limited plan for a specific neighborhood or region (Rich et al. 2021; Paulsen and Rich 2023) (i.e., locally optimizing coverage and directness), or to close certain gaps (Natera Orozco et al. 2020; Vybornova et al. 2023) (i.e., locally optimizing for connectedness and directness). Even when city-wide "master plans" exist, there can be long delays in their realization (Flyvbjerg et al. 2005; Carstensen et al. 2015). Our results show that in general it is better to implement a global plan than to follow locally or randomly motivated growth. At the same time, we

must be aware of the long history of failed "central plans" in urban planning (Batty 2013), thus practical feasibility and a natural bottom-up component of city development must be accounted for.

The finding that centrality-based growth strategies outperform greedy directness-based strategies is particularly interesting from a broader network science perspective. One possible interpretation is that centrality measures implicitly capture more global or long-term structural information, effectively introducing a form of "memory" that helps avoid local optima associated with purely myopic optimization. This interpretation calls for a more systematic study and comparison of other possible indicators that are either local, for example degree-based, or global like eigenvector, flow-based, or accessibility-related measures. Several of these would also be relevant in the transport network context. We expect that results would also crucially depend on the "complexity" of the studied network. The transport networks studied here are planar with narrow degree distributions, but non-planar networks (Boeing 2020) or networks with broader degree distributions could yield different insights. Also, a scale-dependence of directness, coverage, or the buffer size needs to be accounted for in different applications.

One further notable insight of our study is that additive strategies generally outperform subtractive strategies. This insight gives an unexpected contrasting result to four recent studies on bicycle network design (Steinacker et al. 2022, 2025; Ballo et al. 2024; Wiedemann et al. 2025) that argue in favor of the subtractive "backwards percolation" method. Although our greedy single metric optimization is not directly comparable to the greedy optimization of the bikeability metrics used in those studies, our result calls for care when deciding between additive and subtractive methods: both should always be tested before settling for one of them.

Our last discussion point is on realism. Our study purposefully analyzes network growth on stylized topologies first, in order to provide a *systematic* insight into the trade-offs between directness and coverage. We also use metrics like betweenness as a crude proxy for flow, while in reality there is an underlying population distribution and origin–destination matrix, where travelers, especially cyclists, might not take shortest paths but could be strongly influenced by the presence of slopes, green and blue spaces, or vehicular traffic. As such, this research is highly reductionist and fundamental, and thus in itself not directly applicable to practical network design problems. Nevertheless, by having applied our method to two empirical networks, we took the first step towards applicability. Our results have profound implications for the general design of spatial networks, especially planar transport networks, and provide different starting points for planners who aim to achieve higher performance in specific network metrics. As Steinacker et al. (2025) have recently argued, a network science based approach like ours is justifiable despite being more approximative compared to dynamic cost-benefit approaches from transport planning, due to its qualitatively similar results and ease of implementation.

Although here we focused on abstract synthetic networks, we find the CROW Manual's normative guidelines to be adequate to use, as we see our research as groundwork for future applications in this context, like how to best order implementation of links in a bicycle network plan. Of course future research should explore the potential limitations or suboptimality of the CROW Manual's guidelines. Nevertheless, following our research and its preliminary application to relatively small empirical networks, the next

step should be its application to larger empirical street networks or real-life (bicycle network) plans. While this is a straightforward call to make, it is also a challenging task as we expect difficulties due to computational complexity. Directness calculations could scale as $O(L^5)$ with L being the number of links, due to the need to recalculate all-to-all shortest paths in each step. Furthermore, directness and coverage calculations require expensive spatial operations. The running time on an ordinary laptop (2.4GHz quad-core Intel Core i5 with 16GB RAM) to compute the directness growth strategy varies on average between 11 s for the grid and 13 s for the radio-concentric network, which both have 180 links. For the Paris metro system with 386 links (see Appendix D), the running time to compute the directness growth strategy already blows up to 110 s, corresponding roughly to a ten-fold increase in running time for a doubling of links.

There is also a cognitive complexity: for larger networks, manual orders would be much harder to design, especially given the complex, organic, and layered nature of a real transport network.

Nonetheless, we anticipate such computational issues to be eventually resolved, breaking grounds for the application of our framework to real transport networks, and potentially also to unexpected applications to other network design problems such as telecommunication networks or electric grids. Out of necessity, we anticipate climate change to eventually increase political pressure towards downsizing infrastructure for unsustainable modes of transport like road or highway networks in favor of much more effective modes like cycling and mass transit (Brand et al. 2021). Therefore, our research could also become an important groundwork for the development or restoration of once dismantled (Bottinelli et al. 2019) rail, tram, or bus networks.

Appendix A: Additional plots

See Figs. 16, 17 and 18.

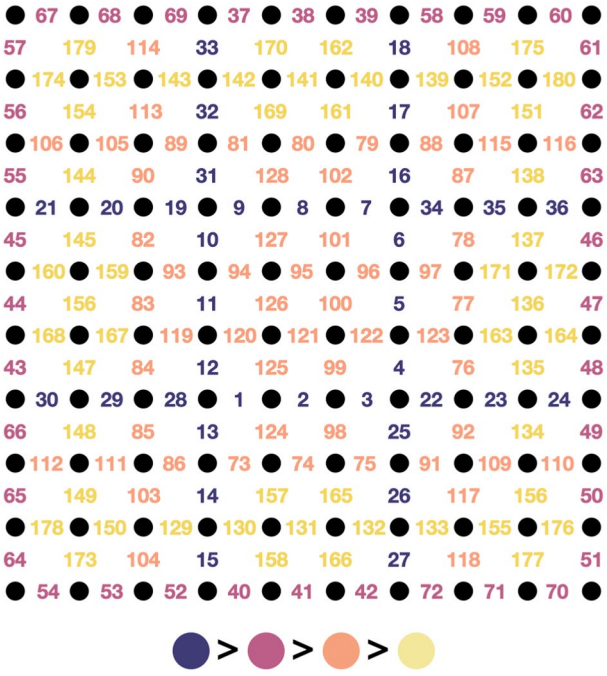


Fig. 16 Manual order H, with each link numbered by order of growth and colored by macroscopic structure

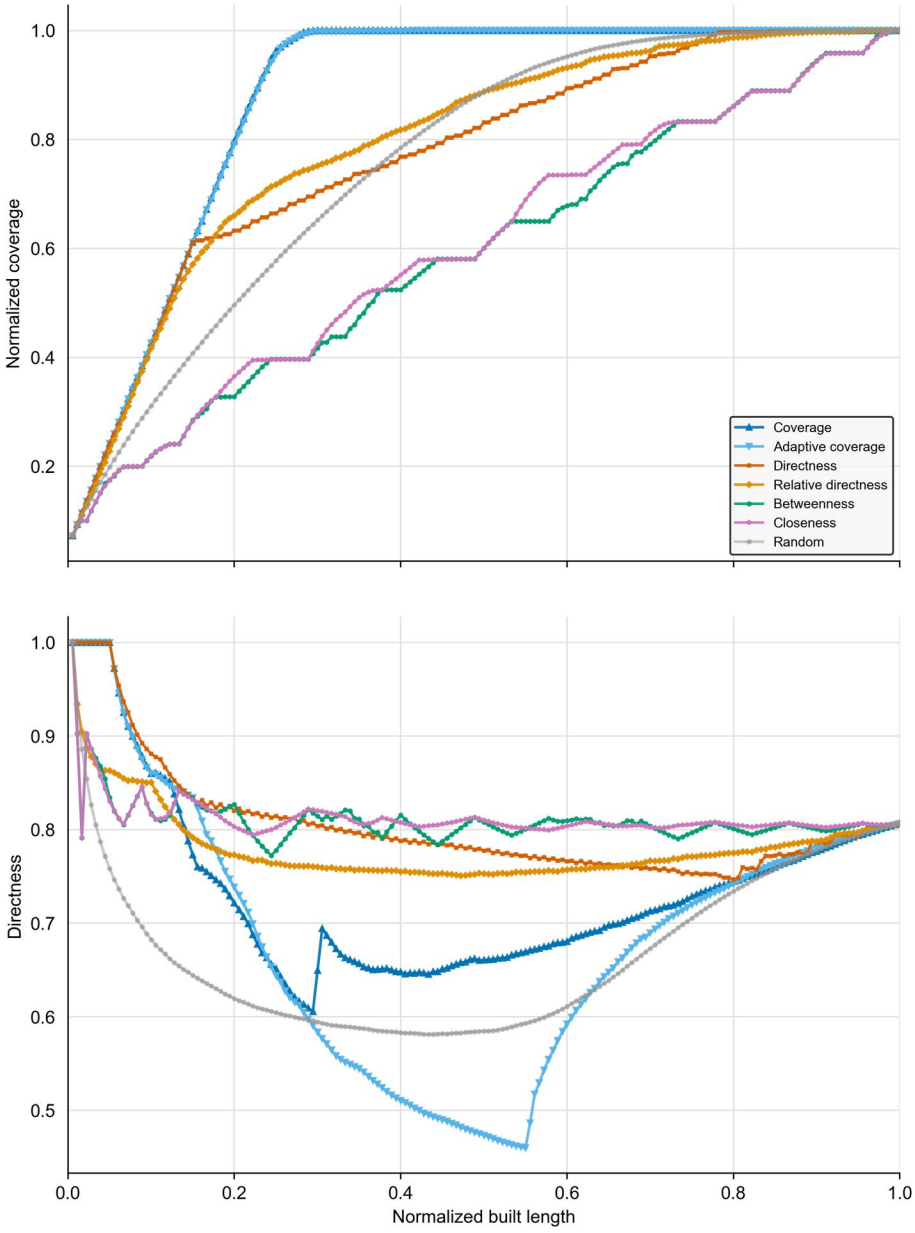


Fig. 17 Performance metrics per built length. Average coverage (top) and directness (bottom) per built length for all additive growth strategies, except manual strategies. Coverage is monotonically non-decreasing, while directness can fluctuate heavily due to its non-linear nature. Both coverage and built length are normalized by their final value

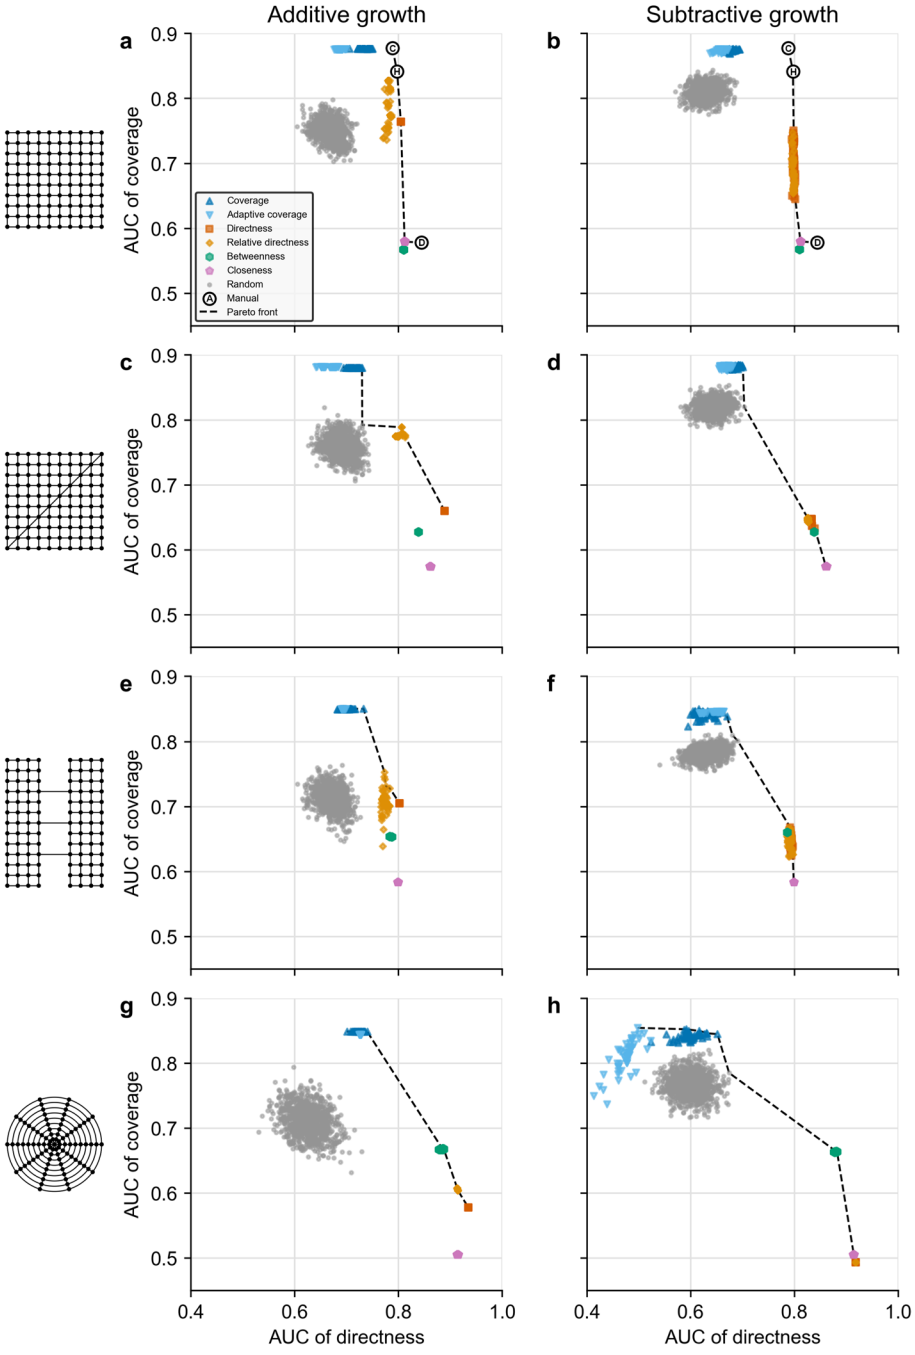


Fig. 18 Comparison of the AUC of coverage and directness for **a** all additive growth strategies on a square grid, **b** all subtractive growth strategies on a square grid, **c** all additive growth strategies on a square grid with diagonal, **d** all subtractive growth strategies on a square grid with diagonal, **e** all additive growth strategies on a square grid with bridges, **f** all subtractive growth strategies on a square grid with bridges, **g** all additive growth strategies on a radio-concentric network, **h** all subtractive growth strategies on a radio-concentric network

Appendix B: Results are robust to grid distortion

We tested the robustness of our results to grid distortion by re-running the analysis on randomly distorted grids (see Fig. 19). Grids are distorted by randomly shifting the position of each node by up to ± 0.45 link length in a uniform distribution. This method is chosen as it ensures that the network stays planar.

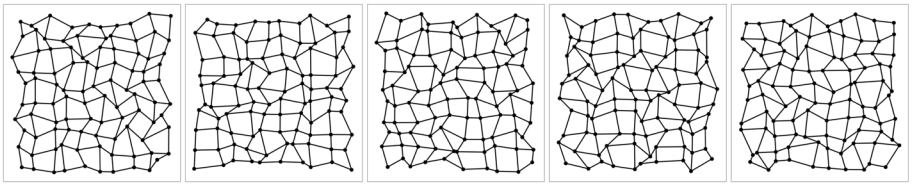


Fig. 19 Distorted grids tested

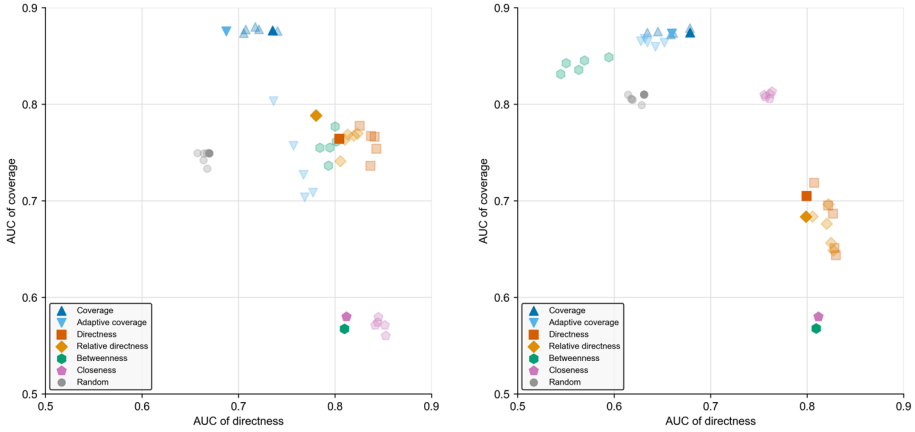


Fig. 20 Comparison of the median AUC of coverage and directness for different growth strategies on the grid and distorted grids for the additive (left) and subtractive (right) order

As is shown in Fig. 20, overall, the relative position of all strategies is preserved, except for Centralities and additive Adaptive coverage. We also observe that the directness of all tested distorted grids ($D \in [0.846, 0.858]$) is higher than on the Regular grid lattice ($D \approx 0.806$).

Appendix C: Stylized urban street networks

Since we are interested in the temporal evolution of network growth efficiency, relevant use cases concern network usage during growth in a long enough timescale. The most typical example is the case of an urban transportation network. The completion of bicycle network plans, for example, takes years or even decades to complete in practice, but they are used throughout the entire development period.

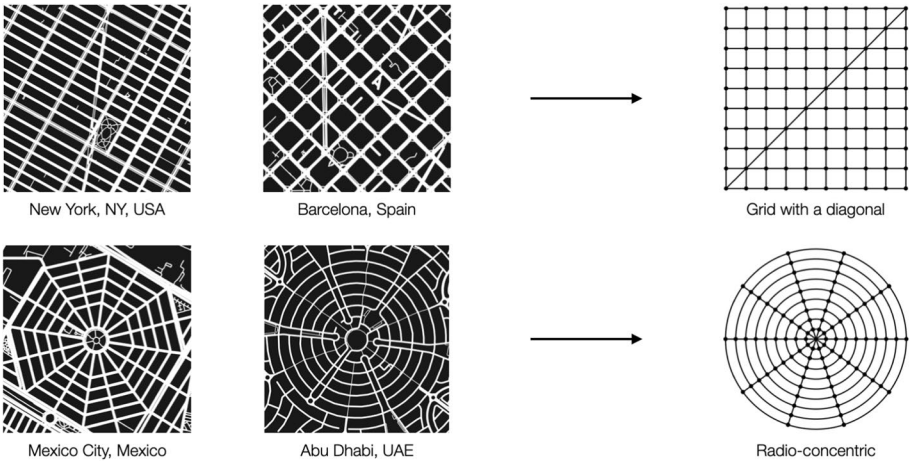


Fig. 21 Examples of empirical street networks and their stylized topologies

Since this is the most typical use case, we use stylized street networks that are representing typical urban patterns, in line with the urban planning literature. With a connectedness requirement in place, the growth of a tree network would be highly constrained; therefore, we are limiting ourselves to urban patterns that are not trees. We explore four different structures (see Fig. 11) and use them as ideal proxies for real street networks, see Fig. 21. We find that most results are qualitatively similar for non-gridlike structures, so for reasons of simplicity in the main part of the paper we focus on the regular grid.

For access to the repository producing the stylized networks used here, see Sebastiao (2024b).

Square grid

The grid, also called rectilinear, gridiron, or Hippodamian plan (Larice and Macdonald 2012), is one of the most ubiquitous shapes found in street networks. From the Centuriation to the contemporary layout of most US cities, it allows for a regular, highly-connected, simple plan. After human mobility systems became increasingly car-centered, the square grid was also celebrated by modernist urban planning as it allows for fast, uninterrupted flow of motorized traffic. As Le Corbusier puts it: "The circulation of traffic demands the straight line; it is the proper thing for the heart of a city. The curve is ruinous, difficult and dangerous; it is a paralyzing thing. [...] The winding road is the Pack-Donkey's Way, the straight road is man's way." (Corbusier 1929).

We motivate the grid size used in the main part of this study by yielding a balanced number of 180 links, allowing to observe growth behavior without the need for long computational times. It also allows to create larger blocks within the grid of 9 blocks of 3×3 faces, as in Barcelona's Superblocks. Having an even number of nodes further avoids link formation along the lines of symmetry.

Grid with a diagonal

A variant of the grid is to add a diagonal street along the grid, to allow for better connectivity and straightness through a major street crossing through the grid. Famous examples are the Avinguda Diagonal in Barcelona's neighborhood l'Eixemple, or Broadway in New York City's neighborhood Manhattan. Our Grid with a diagonal has 100 nodes and 189 links.

Three bridges

Many cities grew around rivers because of the benefits of having direct access to water: Cairo, Paris, Hangzhou, New York, Budapest, etc. Bridges play an important role as being the sparse links between parts of the cities. They have been studied since what is usually referred as the first example of applied network science in the "Seven bridges of Königsberg" problem of Euler (Hopkins and Wilson 2004). Bridges are important because they represent a small number of links that connect entire parts of the city, constituting a bottleneck for traffic and being critical for the network's resilience. Our Three bridges network has 104 nodes and 177 links.

Radio-concentric

Another typical urban pattern is the radio-concentric pattern, in which streets emerge radially outwards from a central point of interest (historically a stronghold, a central

plaza, or a church). Those radial streets are then connected to one another by multiple concentric streets. Archetypal examples can be found in French "circulade" (Watteaux 2003) like Bram, or any city physically embodying the "concentric zone model" (Park et al. 1925). Our Radio-concentric network has 91 nodes and 180 links.

Appendix D: Empirical networks

While this paper focuses on a fundamental study of synthetic network growth, we demonstrate as a proof of concept an application of our methodology on two empirical transport networks.

First we study a simplified sample of the road network of Rome, Italy. We extract the network from OpenStreetMap using `OSMnx` (Boeing 2017), and simplify its topology using `neatnet` (Fleischmann et al. 2025). We selected this set of roads as it displays in a small area an organic combination of several urban patterns, notably a grid and a radio-concentric pattern. Once topologically simplified (see Fig. 22), the network consists of 134 nodes and 222 links.

Second, we study a rail network. We selected a network of the Paris metro system, using a publicly available dataset from a previous study (Asgari et al. 2016). It was simplified by Asgari et al. (2016) to represent each station as a single node, and we simplify it by representing parallel metro lines between the same stations as a single link. The network consists of 303 nodes and 367 links, see Fig. 23. Compared to all the other network studied, this network is bigger, has dead-ends, and is mostly made of nodes of degree 2. Châtelet-Les Halles station, one of the largest underground stations in the world and also a major underground train station, is the node of highest closeness.

Because the networks are of different scales, we estimate a typical distance between nearby links that we use to set a buffer size to compute coverage. We follow our previous argumentation for a perfect grid where all link lengths are a constant l, where we chose a buffer size of 1.5 times l as a reasonable intermediate buffer size: for three consecutive links, the link in the middle is completely covered by the buffered outer links. Here, however, our empirical networks have varying link lengths, so we take the *median* link length $\tilde{l}$ as a rough approximation for a typical distance and choose for the buffer size 1.5

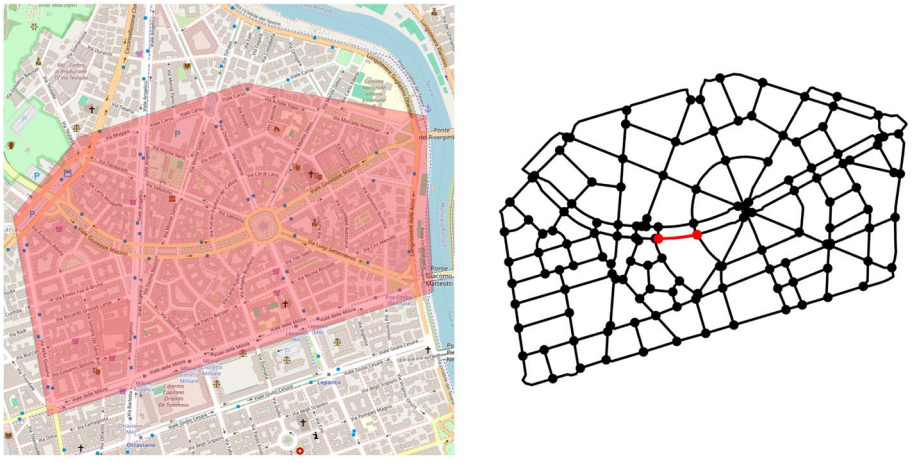


Fig. 22 Empirical road network studied. Left: The location in central Rome from which the street network is extracted, within the red polygon. Map data by OpenStreetMap. Right: The extracted and topologically simplified network. Highlighted in red is the initial, most central link and its adjacent nodes for the additive strategies

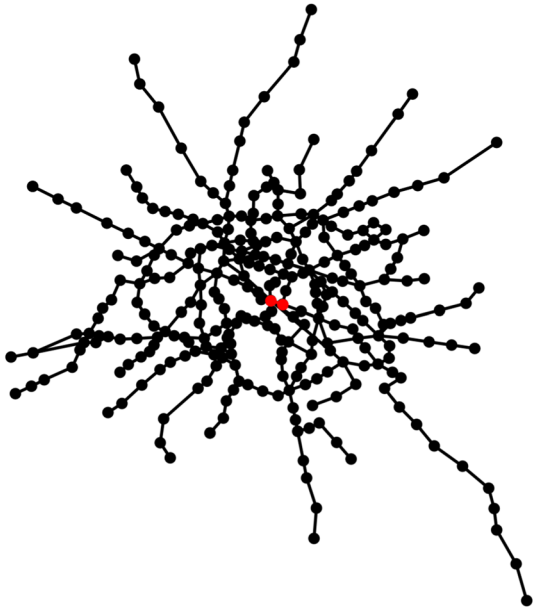


Fig. 23 Empirical rail network studied. Simplified metro network of Paris. Highlighted in red is the initial, most central link and its adjacent nodes for the additive strategies

times $\tilde{l}$. The resulting buffer sizes for the road and the rail networks are approximately $142\,\mathrm{m}$ and $770\,\mathrm{m}$, respectively. This approximation is acceptable as we are reaching full coverage with 36% for the road network, a similar order of magnitude as the grid. The rail network, of a different topology with its long dead-ends, reaches full coverage only with 62% of all links, a higher value where coverage is still a relevant metric. Before applying our framework to further empirical networks in future works, the method of setting buffer sizes should be investigated in more detail, for example by accounting for the fractal dimension of the network.

As we can see in both Figs. 24 and 25, the results are following a similar structure as for our synthetic networks, with differences arising from the organic and irregular topology of the networks. The tradeoff between coverage and directness is even more pronounced than for synthetic networks, as local optimal traps are more prevalent. Ties between links are highly unlikely for greedy strategies (except for coverage related metrics). Finally, link betweenness centrality is concentrated on a smaller but more dispersed set of links. Links that are spatially centered have naturally a higher link betweenness centrality, but irregularities in the network create shortcuts and a hierarchy, reducing this effect, in line with Kirkley et al. (2018).

The results slightly differ for the empirical rail network (see in Fig. 25), with a Pareto front further away from random strategies. This additional distance is likely caused by the larger size of the network and the prevalence of interstitial nodes of degree 2, both drastically increasing the solution space with non-optimal configurations.

As shown in this appendix, the methodology described in this article is applicable to real transport networks. However, real use cases are embedded within a multimodal transport system, and more broadly in a complex urban substrate. Studying a transport network would benefit from local knowledge and additional datasets (traffic flow, population distribution, accessibility to amenities, planning constraints, etc.) to design more adequate performance metrics.

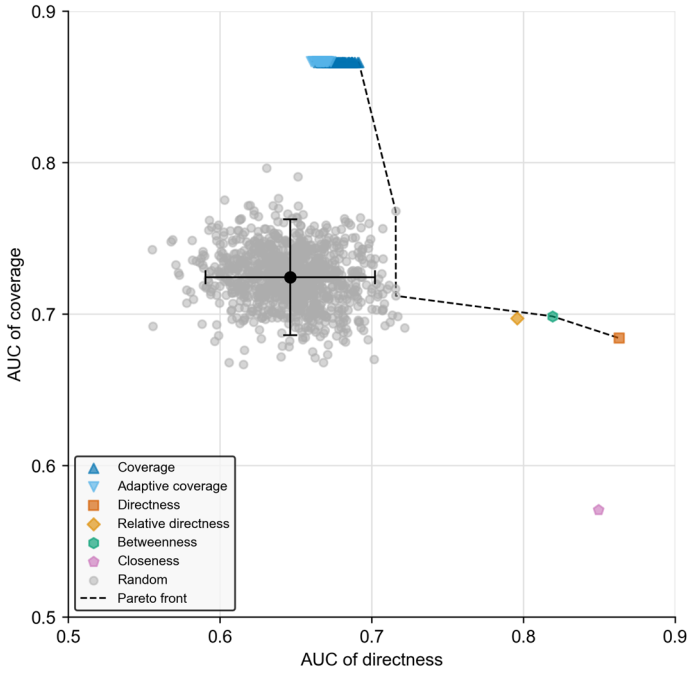


Fig. 24 AUC of coverage against AUC of directness for all additive growth strategies for the empirical road network

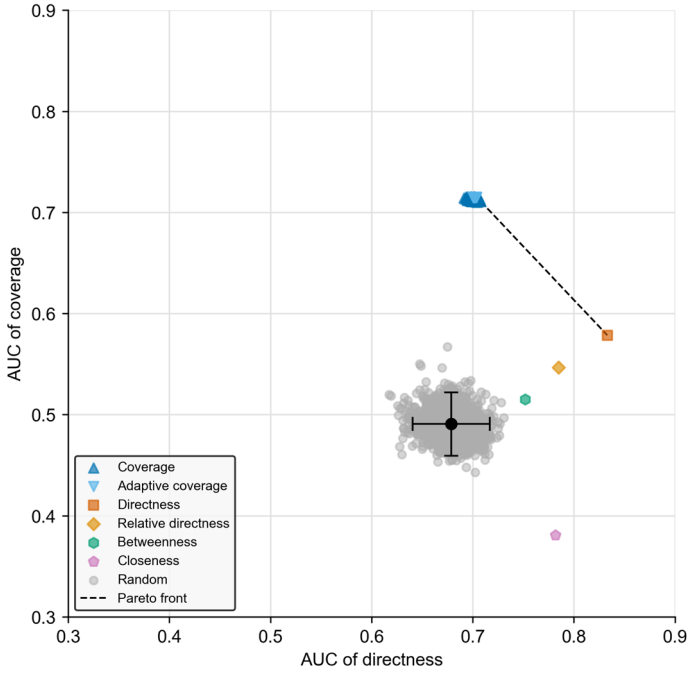


Fig. 25 AUC of coverage against AUC of directness for all additive growth strategies for the empirical rail network

Appendix E: Area under curve with exponential discounting

Benefits of transport infrastructure is cumulative: The earlier it is built, the larger are its total benefits for a fixed time horizon. Therefore, we check results for a growth variant where the early stages are prioritized, using exponential discounting. We define the exponentially discounted performance metric y_g as:

$$y_g(x, y) = ye^{-x \ln(g)},$$

where y is the performance metric, x is the normalized built length and g is the discounting gap, where $y_g(0, y)/y_g(1, y) = g$, meaning that the last link is valued g times less than the first link. We report results in Fig. 26 for $g = 10$. As focused growth at an early

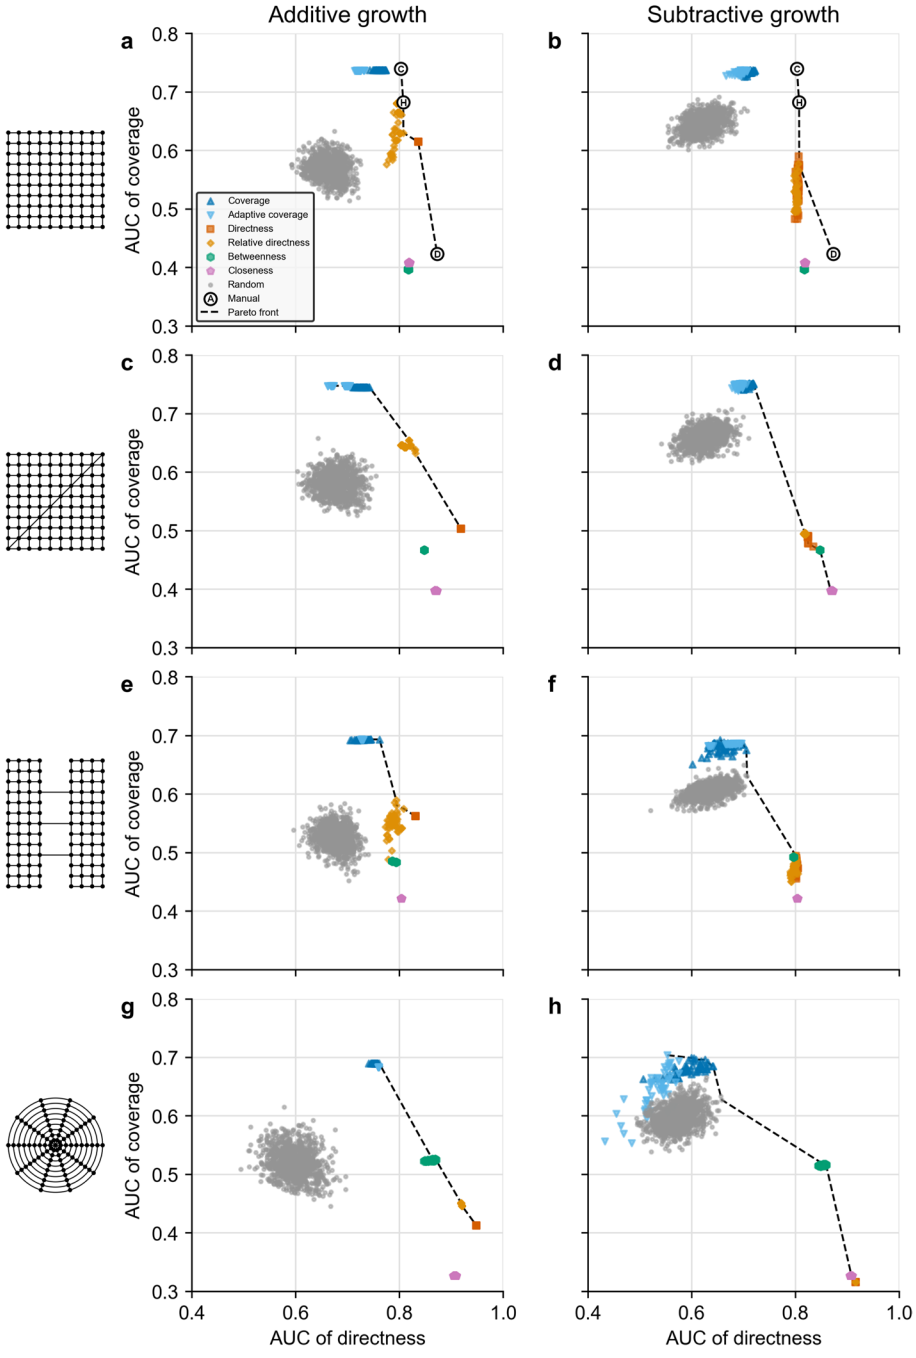


Fig. 26 Comparison of the AUC of coverage and directness with exponential discounting for **a** all additive growth strategies on a square grid, **b** all subtractive growth strategies on a square grid, **c** all additive growth strategies on a square grid with diagonal, **d** all subtractive growth strategies on a square grid with diagonal, **e** all additive growth strategies on a square grid with bridges, **f** all subtractive growth strategies on a square grid with bridges, **g** all additive growth strategies on a radio-concentric network, **h** all subtractive growth strategies on a radio-concentric network

stage is rewarded, we observe a larger gap between random and greedy strategies, exacerbating previous observations. Qualitatively, results remain similar to no discounting, compare with Fig. 18.

Acknowledgements
Not applicable.

Author Contributions
Clément Sebastiao: Conceptualization, Formal analysis, Investigation, Methodology, Software, Visualization, Writing – original draft, Writing – review and editing. Anastassia Vybornova: Conceptualization, Writing - review and editing. Ane Rahbek Vierø: Conceptualization, Writing – review and editing. Luca Maria Aiello: Conceptualization, Supervision, Writing – review and editing. Michael Szell: Conceptualization, Funding acquisition, Methodology, Project administration, Supervision, Visualization, Writing – original draft, Writing – review and editing.

Funding
C.S. and M.S. acknowledge support by the European Union through the Horizon Europe grant JUST STREETS (Grant agreement ID: 101104240). A.V. acknowledges support from Villum Fonden through the Villum Young Investigator programme (project number: 00037394).

Data availability
Road network data is available in OpenStreetMap. Rail network data is available online, see Asgari et al. (2016). All materials are available online for reproducing figures, results, and generated data, see Sebastiao (2024a).

Declarations

Conflict of interest
The authors declare no conflict of interest.

Received: 16 July 2025 / Accepted: 11 March 2026
Published online: 27 March 2026

References

Alessandretti L, Natera Orozco LG, Saberi M et al (2023) Multimodal urban mobility and multilayer transport networks. Environ Plann B Urban Anal City Sci 50(8):2038–2070
Asgari F, Sultan A, Xiong H et al (2016) Ct-mapper: mapping sparse multimodal cellular trajectories using a multilayer transportation network. Comput Commun 95:69–81
Ballo L, Raubal M, Axhausen KW (2024) Designing an e-bike city: an automated process for network-wide multimodal road space reallocation. J Cycl Micromobility Res 2:100048
Barthelemy M (2022) Spatial networks: a complete introduction:from graph theory and statistical physics to real-world applications. Springer, Cham. https://doi.org/10.1007/978-3-030-94106-2
Barthelemy M, Bordin P, Berestycki H et al (2013) Self-organization versus top-down planning in the evolution of a city. Sci Rep 3(1):215. https://doi.org/10.1038/srep02153
Batty M (2006) Hierarchy in cities and city systems. Hierarchy Nat Soc Sci 3:143–168
Batty M (2013) The new science of cities. MIT Press
Boeing G (2017) Osmnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. Comput Environ Urban Syst 65:126–139
Boeing G (2020) Planarity and street network representation in urban form analysis. Environ Plann B Urban Anal City Sci 47(5):855–869
Boeing G (2025) Topological graph simplification solutions to the street intersection miscount problem. Trans GIS 29(3):e70037
Bottinelli A, Gherardi M, Barthelemy M (2019) Efficiency and shrinking in evolving networks. J R Soc Interface 16(154):20190101. https://doi.org/10.1098/rsif.2019.0101
Brand C, Dons E, Anaya-Boig E et al (2021) The climate change mitigation effects of daily active travel in cities. Transp Res Part D Transp Environ 93:102764. https://doi.org/10.1016/j.trd.2021.102764
Brandstädt A, Le VB, Spinrad JP (1999) Graph classes: a survey. SIAM
Brede M (2010) Coordinated and uncoordinated optimization of networks. Phys Rev E-Stat Nonlinear Soft Matter Phys 81(6):066104
Cancho RFi, Solé RV (2003) Optimization in complex networks. In: Statistical mechanics of complex networks. Springer, pp 114–126
Cardillo A, Scellato S, Latora V et al (2006) Structural properties of planar graphs of urban street patterns. Phys Rev E 73(6):066107. https://doi.org/10.1103/PhysRevE.73.066107
Carstensen TA, Olafsson AS, Bech NM et al (2015) The spatio-temporal development of Copenhagen's bicycle infrastructure 1912–2013. Geografisk Tidsskrift-Danish J Geogr 115(2):142–156. https://doi.org/10.1080/00167223.2015.1034151
Corbusier L (1929) The City of to-morrow and its planning
Cormen TH, Leiserson CE, Rivest RL et al (2022) Introduction to algorithms. MIT Press, Cambridge
CROW (2016) Design manual for bicycle traffic
Diogo Pinto J, Pinheiro FL, Bação F (2026) Frameworks for prioritizing cycling network investments: a systematic literature review. Transp Res Interdiscipl Perspect 35:101782. https://doi.org/10.1016/j.trip.2025.101782
Douady S, Lagesse C, Atashinbar M et al (2020) A work on reticulated patterns. Comptes Rendus Mécanique 348(6–7):659–678

El Gouj H, Rincón-Acosta C, Lagesse C (2022a) Urban morphogenesis analysis based on geohistorical road data. Appl Netw Sci 7(1):6. https://doi.org/10.1007/s41109-021-00440-0
El Gouj H, Rincón-Acosta C, Lagesse C (2022b) Urban morphogenesis analysis based on geohistorical road data. Appl Netw Sci 7(1):6. https://doi.org/10.1007/s41109-021-00440-0
Ferenchak NN, Marshall WE (2025) The link between low-stress bicycle facilities and bicycle commuting. Nat Cities 2(6):555–559
Fleischmann M, Vybornova A, Gaboardi JD, et al (2025) Adaptive continuity-preserving simplification of street networks. arXiv:2504.16198
Flyvbjerg B, Skamris Holm MK, Buhl SL (2005) How (in) accurate are demand forecasts in public works projects? The case of transportation. J Am Plann Assoc 71(2):131–146
Giacomin DJ, Levinson DM (2015) Road network circuity in metropolitan areas. Environ Plann B Plann Des 42(6):1040–1053. https://doi.org/10.1068/b130131p
Girvan M, Newman MEJ (2002) Community structure in social and biological networks. Proc Natl Acad Sci 99(12):7821–7826. https://doi.org/10.1073/pnas.122653799
Gössling S (2020) Why cities need to take road space from cars-and how this could be done. J Urban Des 25(4):443–448
Guo X, Zhao L (2022) A systematic survey on deep generative models for graph generation. IEEE Trans Pattern Anal Mach Intell 45(5):5370–5390
Henderson J, Gulsrud NM (2019) Street fights in Copenhagen: bicycle and car politics in a green mobility city. Routledge
Hopkins B, Wilson RJ (2004) The Truth about Königsberg. Coll Math J 35(3):198–207. https://doi.org/10.1080/07468342.2004.11922073
Kirkley A, Barbosa H, Barthelemy M et al (2018) From the betweenness centrality in street networks to structural invariants in random planar graphs. Nat Commun 9(1):2501
Kong L, Cui J, Sun H, et al (2023) Autoregressive diffusion model for graph generation. In: International conference on machine learning, PMLR, pp 17391–17408
Kuratowski C (1930) Sur le probleme des courbes gauches en topologie. Fundam Math 15(1):271–283
Larice M, Macdonald E (eds) (2012) The Urban Design Reader, 2nd edn. Routledge, Londo. https://doi.org/10.4324/9780203094235
Latora V, Marchiori M (2002) Is the Boston subway a small-world network? Phys A Stat Mech Appl 314(1–4):109–113
Louf R, Jensen P, Barthelemy M (2013) Emergence of hierarchy in cost-driven growth of spatial networks. Proc Natl Acad Sci 110(22):8824–8829. https://doi.org/10.1073/pnas.1222441110
Lowry MB, Furth P, Hadden-Loh T (2016) Prioritizing new bicycle facilities to improve low-stress network connectivity. Transp Res Part A Policy Practice 86:124–140. https://doi.org/10.1016/j.tra.2016.02.003
Madeira M, Vignac C, Thanou D et al (2024) Generative modelling of structurally constrained graphs. Adv Neural Inf Process Syst 37:137218–137262
Marqués R, Hernández-Herrador V, Calvo-Salazar M et al (2015) How infrastructure can promote cycling in cities: lessons from Seville. Res Transp Econ 53:31–44. https://doi.org/10.1016/j.retrec.2015.10.017
Martínez Plumed F, Ferri Ramírez C, Contreras Ochando L (2016) Cycling network projects: a decision-making aid approach
Mauttone A, Mercadante G, Rabaza M et al (2017) Bicycle network design: model and solution algorithm. Transp Res Procedia 27:969–976. https://doi.org/10.1016/j.trpro.2017.12.119
Moran ME (2022) Treating COVID with bike lanes: design, spatial, and network analysis of 'pop-up' bike lanes in Paris. Findin. https://doi.org/10.32866/001c.33765, https://findingspress.org/article/33765-treating-covid-with-bike-lanes-design-spatial-and-network-analysis-of-pop-up-bike-lanes-in-paris
Morer I, Cardillo A, Díaz-Guilera A et al (2020) Comparing spatial networks: a one-size-fits-all efficiency-driven approach. Phys Rev E 101(4):042301. https://doi.org/10.1103/PhysRevE.101.042301
Natera Orozco LG, Battiston F, Iñiguez G et al (2020) Data-driven strategies for optimal bicycle network growth. R Soc Open Sci 7(12):201130. https://doi.org/10.1098/rsos.201130
Oldenziel R, de la Bruhèze AA (2011) Contested spaces: bicycle lanes in Urban Europe, 1900–1995. Transfers 1(2):29–49
Ospina JP, Duque JC, Botero-Fernández V et al (2022) The maximal covering bicycle network design problem. Transp Res Part A Policy Practice 159:222–236. https://doi.org/10.1016/j.tra.2022.02.004
Park RE, Burgess EW, McKenzie RD (1925) The city chicago. University of Chicago Press, Chicago
Paulsen M, Rich J (2023) Societally optimal expansion of bicycle networks. Transp Res Part B Methodol 174:102778. https://doi.org/10.1016/j.trb.2023.06.002
Reggiani G, van Oijen T, Hamedmoghadam H et al (2021) Understanding bikeability: a methodology to assess urban networks. Transportation. https://doi.org/10.1007/s11116-021-10198-0
Rich J, Jensen AF, Pilegaard N et al (2021) Cost-benefit of bicycle infrastructure with e-bikes and cycle superhighways. Case Stud Transp Policy 9(2):608–615
Sabidussi G (1966) The centrality index of a graph. Psychometrika 31(4):581–603. https://doi.org/10.1007/BF02289527
Schoner JE, Levinson DM (2014) The missing link: bicycle infrastructure networks and ridership in 74 US cities. Transportation 41:1187–1204
Sebastiao C (2024a) https://github.com/csebastiao/orderbike
Sebastiao C (2024b) https://github.com/csebastiao/utg
Steinacker C, Storch DM, Timme M et al (2022) Demand-driven design of bicycle infrastructure networks for improved urban bikeability. Nat Comput Sci 2(10):655–664. https://doi.org/10.1038/s43588-022-00318-w
Steinacker C, Paulsen M, Schröder M et al (2025) Robust design of bicycle infrastructure networks. Sci Rep 15(1):15471
Strano E, Nicosia V, Latora V et al (2012) Elementary processes governing the evolution of road networks. Sci Rep 2(1):296. https://doi.org/10.1038/srep00296
Szell M, Mimar S, Perlman T et al (2022) Growing urban bicycle networks. Sci Rep 12(1):6765. https://doi.org/10.1038/s41598-022-10783-y
Touwen L, Bucur D, van der Hofstad R et al (2024) Learning the mechanisms of network growth. Sci Rep 14(1):11866
Vybornova A, Cunha T, Gühnemann A et al (2023) Automated detection of missing links in bicycle networks. Geogr Anal 55(2):239–267. https://doi.org/10.1111/gean.12324
Walker J (2024) Human transit, revised edition: how clearer thinking about public transit can enrich our communities and our lives. Island Press, Washington, DC

Wang J, Mo H, Wang F (2014) Evolution of air transport network of China 1930–2012. J Transp Geogr 40:145–158. https://doi.org/10.1016/j.jtrangeo.2014.02.002
Watteaux M (2003) Le plan radio-quadrillé des terroirs non planifiés. Études rurales 167–168:187–214. https://doi.org/10.4000/etudesrurales.8024
West GB, Brown JH (2005) The origin of allometric scaling laws in biology from genomes to ecosystems: towards a quantitative unifying theory of biological structure and organization. J Exp Biol 208(9):1575–1592
Wiedemann N, Nöbel C, Ballo L et al (2025) Bike network planning in limited urban space. Transp Res Part B Methodol 192:103135. https://doi.org/10.1016/j.trb.2024.103135
Yang H, Ke J, Ye J (2018) A universal distribution law of network detour ratios. Transp Res Part C Emerg Technol 96:22–37. https://doi.org/10.1016/j.trc.2018.09.012
You J, Ying R, Ren X, et al (2018) Graphrnn: generating realistic graphs with deep auto-regressive models. In: International conference on machine learning, PMLR, pp 5708–5717
Zhou Z, Kearnes S, Li L et al (2019) Optimization of molecules via deep reinforcement learning. Sci Rep 9(1):10752

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.