Random networks embedded in the plane

David Aldous

7 June 2019

Work of Jean-Francois Le Gall, Gregory Miermont and many others has created a theory of random planar maps that is elegant, mathematically deep and surprisingly well-connected with other areas of mathematics and physics. In that theory, a **map** is a topological isomorphism class of networks. But this setting is not so relevant to study of macroscopic real-world networks in the plane, where vertices have given positions.

Nowadays we are aware of many kinds of network, and the quantitative study of networks has attracted a huge literature in many areas of the mathematical sciences over the last 20 years. The phrase *spatial networks* has come into use to indicate contexts such as

(a) transportation (roads, railways) and distribution (electricity, gas, water) networks where the edges represent connections physically situated in two-dimensional space,

or more generally, contexts such as
(b) airline routes, wireless networks, where the vertices are situated in two-dimensional space and the physical distance between vertices is important.

Spatial networks have been studied implicitly in fields outside Probability:

- Operations Research
- Computational Geometry
- Theory of Algorithms

One might assume that the 2-dimensional setting made study easier than the abstract setting of a network as an (edge-weighted) graph, but not necessarily so, and one can ask rather different questions. Consider the question:

what is the shortest network linking n given points?

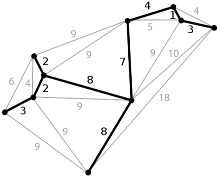

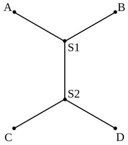

On an abstract network this is the minimum spanning tree (MST), obtained by a simple greedy algorithm. In 2-dimensional space it is the minimal Steiner tree – algorithmically NP.

Random spatial networks in mathematical probability

Two relevant well-studied areas with authoritative monographs

- stochastic geometry, e.g. the Delaunay Triangulation (left)
- continuum percolation/random geometric graph (right)

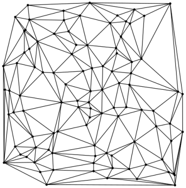

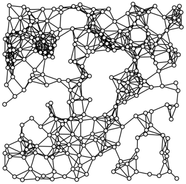

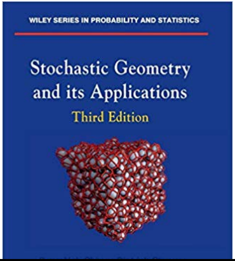

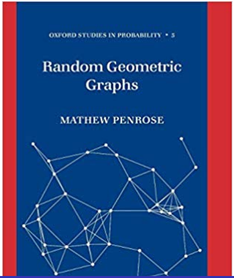

Within modeling, random geometric graphs are used in theoretical study of ad hoc wireless networks (e.g. Gupta - Kumar (2000) *The capacity of wireless networks*, cited by 10030).

Much academic study of general networks has been done within a statistical physics style. A Google Scholar search on "spatial networks" gets a 2011 survey paper by Marc Barthelemy (IPhT, Paris Saclay) and also his 2018 monograph "Morphogenesis of Spatial Networks" in statistical physics style. In contrast, toy models for spatial networks with physical links have been little studied within theorem-proof mathematics. I will talk about work of mine and others over the last 10 years seeking to bridge the gap.

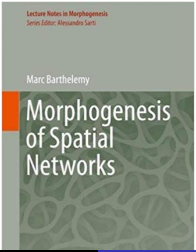

Seeking because I don't claim much success – no remarkable theorems. Perhaps because

- Real world networks are finite and heterogeneous – our familiar models tend to be spatially homogeneous and we tend to study asymptotics.
- Orders of magnitude often obvious, anything more often very difficult.

So we often rely on simulation. This actually has a positive aspect

- One can invent many easy-to-define models not studied before, and so one can involve undergraduates in simulation projects.

I will describe models for transportation networks in 3 different settings.

- Inter-city road or rail networks. As in classical models earlier, we first define a network for arbitrary point-sets, then study the random networks obtained from a PPP. But instead of a simple rule for constructing the network we want to consider *optimal* networks (in some sense). [Aldous-Shun 2010]
- A model of scale-invariant networks in the continuum (cf. random walk and Brownian motion, finite trees and continuum random tree). The only sophisticated math underlying this talk – many open theory problems. [Aldous-Ganesan 2013]
- Subway network – for a given population density in a city, where do you build subway lines (using some optimality criterion)? This addresses the non-homogeneous setting. [Aldous-Barthelemy 2019]

Setting 1: Optimal inter-city networks, model cities as rate-1 PPP on the infinite plane.

In considering optimal networks, a natural simplification is to regard the **cost** of a network as its total length. And to measure the **benefit** as having short routes.

Formalizing such benefits turns out to be rather subtle, as shown by the following very elementary observation from Aldous - Kendall (2008). Given n cities in square of area n, first construct the (minimum connected length) Steiner tree, then overlay a sparse Poisson line process.

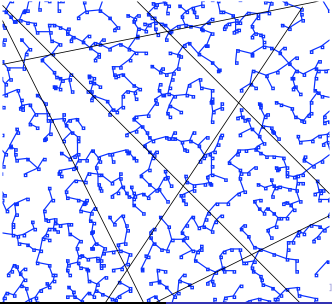

So we can construct networks $\mathcal{N}_n$ that are first-order optimal for both total length and average route-lengths, that is attain the optimal constants for

$$\lim_n \ n^{-1}\mathrm{length}(\mathcal{N}_n) \text{ and}$$

$$(*) \quad \lim_n \ n^{-1/2}\mathrm{ave}_{v_1,v_2}[\mathrm{route\_length}(v_1,v_2)].$$

So in the $n \to \infty$ asymptotic regime there is no trade-off between cost and benefit as defined by (*) – we can optimize both at once.

This answer is clearly unrealistic, and a warning against naive asymptotics.
Let's instead require short route-lengths on all scales. That is, consider

$$\rho(d) = \frac{\mathbb{E}(\text{network distance between cities at distance } d)}{d} - 1$$

Observation: In both simple math models and real-world data we see a characteristic shape for the function $\rho(d)$.

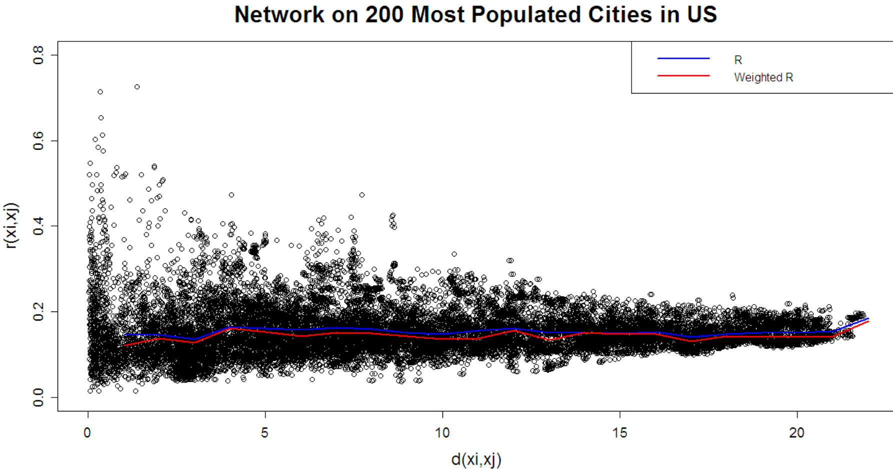
Network on 200 Most Populated Cities in US
R
Weighted R
r(xi,xj)
0.0
0.2
0.4
0.6
0.8
d(xi,xj)
0
5
10
15
20

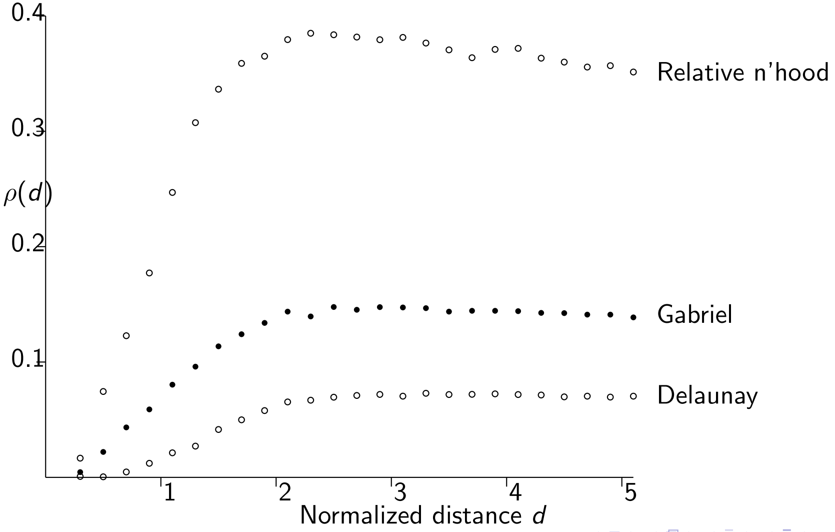
0.4
0.3
0.2
0.1
$\rho(d)$
1
2
3
4
5
Normalized distance d
Relative n'hood
Gabriel
Delaunay

This prompts us to use the statistic

$$R := \max_{0<d<\infty} \rho(d).$$

In words, $R = 0.2$ means that on every scale of distance, route-lengths are on average at most 20% longer than straight line distance.

Next figure compares values of R and L for different networks over a Poisson point process.

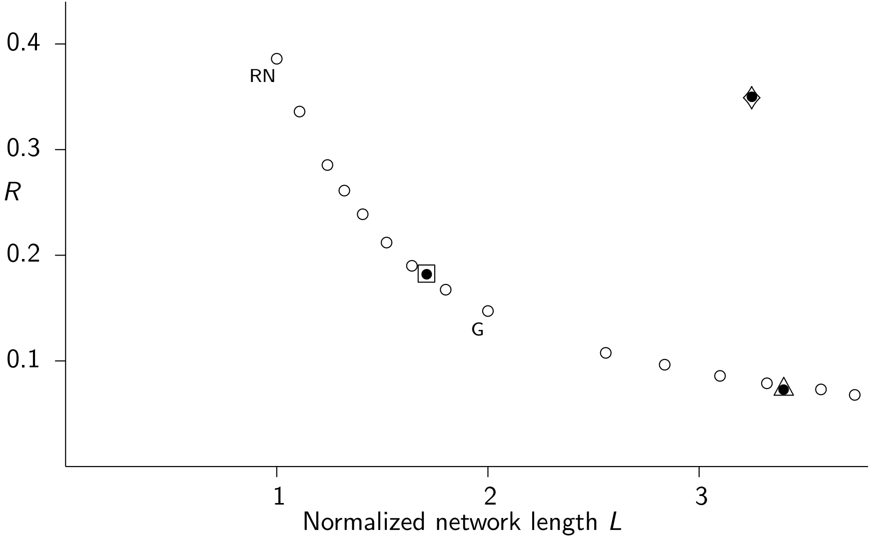
0.4
0.3
0.2
0.1
R
RN
G
1
2
3
Normalized network length L

The smooth curve is from the β-skeleton family of proximity graphs, which are defined by a simple "local" rule and look like

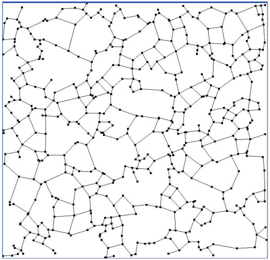

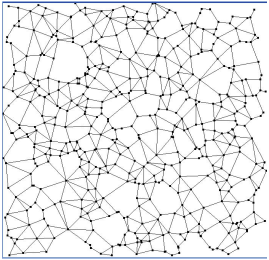

We conjecture that this family is almost optimal in our sense, but have no theorems.

Conclusions from Setting 1: inter-city networks.

1. Graphic is plausible as qualitative trade-off between length and "short routes" in such networks. Indeed leads to a prediction that in an efficient network designed to link N cities in a region of area A, the total length should be around $2\sqrt{NA}$.
2. But visually wrong. Real world networks have long motorways (or main railway lines) running roughly straight; cannot reproduce this by local rules, even if we vary city sizes.
3. Real-world road networks have a hierarchy of "size" or "importance" of road, measured by number of lanes or traffic volume or numbering system. We do not know how to obtain this as an emergent property of optimal networks.

Our next, completely different, model paradigm does lead to such a spectrum of "sizes" of road.

* * * start simulation * * *

Setting 2: Scale-invariant random spatial networks

A map on **paper** has a definite scale: maybe every street in a small town, or maybe the major road network of an entire country. These have intrinsically different real-world patterns, but we can attempt a (not so realistic) a modeling framework which is consistent over different spatial scales.

Conceptual starting point:

Online road maps differ from paper maps in 2 [obvious] ways that will motivate our modeling.

- Can zoom in – see greater detail in window covering less area.
- Can get routes between any two specified addresses.

Idea behind our set-up: start with **routes** between **addresses** instead of **roads**.

We abstract *Google maps* as an "oracle" that for any start/desination pair (z_1, z_2) in the plane gives us a route $\mathrm{r}(z_1, z_2)$.

Analogous to ergodic theory regarding the *Don Quixote* text as one realization from a stationary source, we regard *Google maps* as containing one realization of a "continuum random spatial network" with some distribution. We will define a class of such random networks by axiomatizing properties of random routes $\mathcal{R}(z_1, z_2)$.

The key assumption is **scale-invariance**, described intuitively as follows.

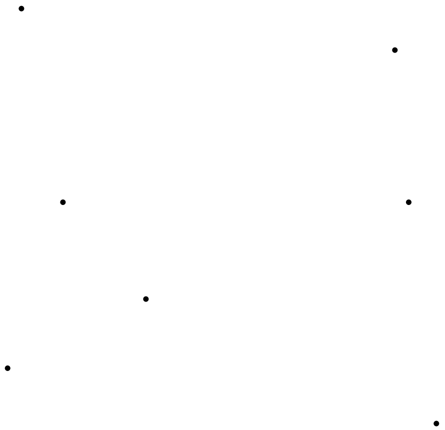

7 points in a window.

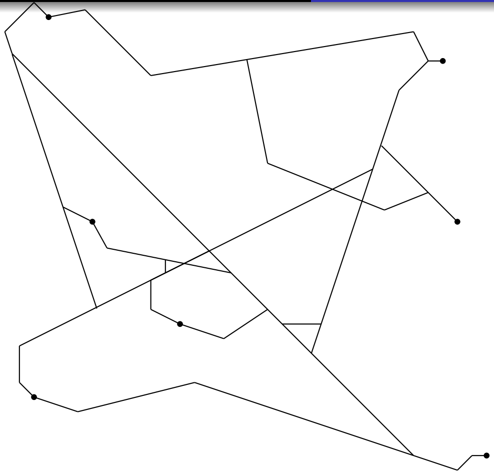

Scale-invariance means: doing this within a randomly positioned window, the statistics of the subnetwork observed don't depend on the scale, i.e. don't depend on whether the side length is 10 km or 100 km.

Comments re scale-invariance:

- To have a network model which is exactly scale-invariant, we need to work in the continuum (cf. random walk and Brownian motion).
- Naive Euclidean scaling, not "scaling exponent".

Axiomatic setup: 1

Details are pretty technical, but ......

Process is presented via FDDs of random routes $\mathcal{R}(z_1, z_2)$; in other words we are given a distribution for the random subnetwork spanning each finite set $\{z_1, \ldots, z_k\}$, Kolmogorov-consistent.

Assume

- Translation and rotation-invariant
- Scale-invariant

So route-length D_r between points at (Euclidean) distance r apart must scale as $D_r \stackrel{d}{=} rD_1$.

Assume $ED_1 < \infty$ so **not fractal.**

Axiomatic setup: 2

Envisage the route $\mathcal{R}(z_1, z_2)$ as the path that optimizes *something* (e.g. travel time) but do not formalize that idea; instead

Assume a route-compatability property.

Convenient to study the process via the subnetwork $\mathcal{S}(\lambda)$ spanning a Poisson point process (rate λ per unit area).

Define a statistic

$$\ell = \text{length-per-unit-area of } \mathcal{S}(1).$$

Assume $\ell < \infty$.

Axiomatic setup: 3

Sample points of a rate-λ PPP; draw only routes between points in A and points in B.

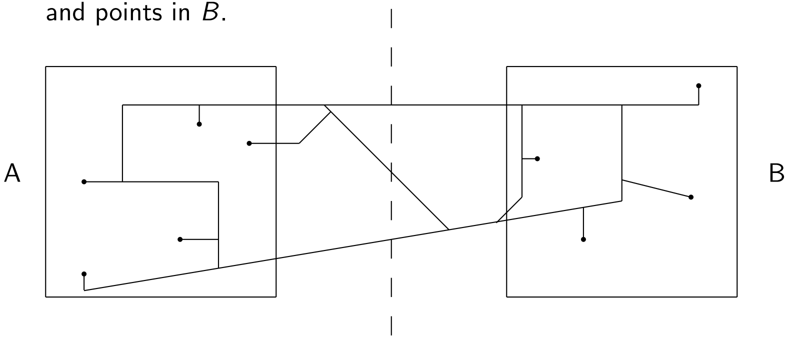

A real-world road network would have the property: as $\lambda \to \infty$ the number of places where one of these routes crosses an intervening line stays finite. We can define a numerical constant ρ (details omitted) indicating the rate of such crossings, and require $\rho < \infty$.

We have defined a class of processes we'll call

SIRSN: Scale-invariant random spatial networks.

for which we have very many questions but very few answers.

- Do a broad variety of SIRSNs actually exist?
- Can we specific particular canonical or optimal ones?
- What are their math properties?
- Any realistic aspects?

Elegant construction by Kendall (2014) *From Random Lines to Metric Spaces* uses a Poisson line process, different lines representing roads with different speeds, and routes are minimum-time path.

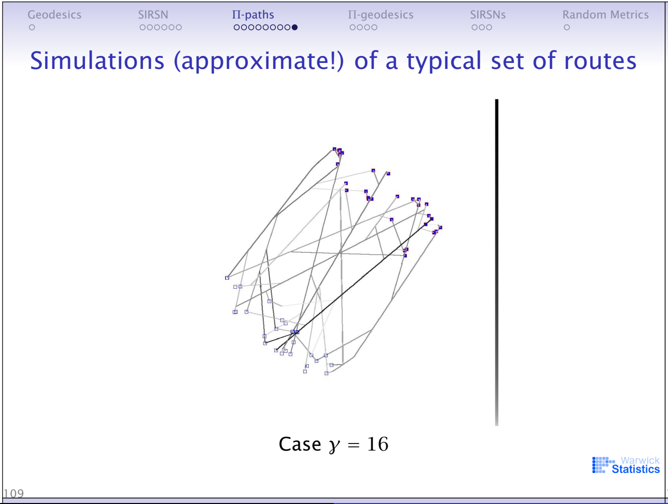

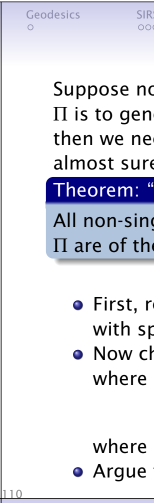

Heuristically, we can construct via “dynamic” variants of the “static” proximity graphs like

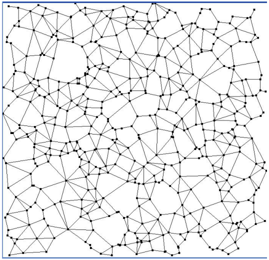

Throw points one-by-one; for a new point ξ, put an edge to an existing point ξ' iff the disc with diameter (ξ, ξ') contains no third existing point. * * * show simulations * * *

Want to define routes as shortest path; technically difficulty to verify a.s. uniqueness. But given the limit process exists, it must be scale-invariant.

- Do a broad variety of SIRSNs actually exist?
- **Can we specific particular canonical or optimal ones?**

Canonical analogous to Brownian motion or the continuum random tree?
Optimal in terms of cost-benefit parameters ℓ and $\mathbb{E}D_1$.

I don't know.

- What are their math properties?
- Any realistic aspects?

First-passage percolation (FPP) on the plane lattice is a well-studied model where routes appear: questions studied there can be asked of SIRSNs (but remember this is a class of processes, and no bottom level of pure randomness).

For a SIRSN:

- At least one semi-infinite geodesic from $\mathbf{0}$.
- By self-similarity, behavior near $\mathbf{0}$ relates to behavior near infinity.
- No doubly-infinite geodesics.
- hierarchy of regularity conditions, starting with "unique semi-infinite geodesic from $\mathbf{0}$" and increasing strength to ...
- $\forall \varepsilon \quad \exists$ random $\delta > 0$ s.t. all routes $\mathrm{disc}(\mathbf{0}, \delta)$ to $\mathrm{disc}(\mathbf{1}, \delta)$ coincide outside $\mathrm{disc}(\mathbf{0}, \varepsilon) \cup \mathrm{disc}(\mathbf{1}, \varepsilon)$. This formalizes idea that map $(z_1, z_2) \to \mathcal{R}(z_1, z_2)$ is a.e. continuous.
- $\exists$ random $\delta > 0$ s.t. all routes $\mathrm{disc}(\mathbf{0}, \delta)$ to $\mathrm{disc}(\mathbf{0}, 1/\delta)$ cross $\mathrm{circle}(\mathbf{0}, 1)$ at the same point.

In proving uniqueness of routes in the two models constructed rigorously, we actually prove this last property.

Define $\mathcal{E}(r)$ to be the "skeleton" of the network consisting of points ξ in the network which are in some route $\mathcal{R}(\xi_1, \xi_2)$ for which both endpoints are at distance $\geq r$ from ξ.

This gives a notion of "major road", and then the "importance" of a point ξ in the network can be defined as

$$m(\xi) = \sup\{r : \ \xi \in \mathcal{E}(r)\}.$$

Two conjectures is this context:

1. Heuristically if the volume of traffic flow between (ξ_1, ξ_2) is proportional to $||\xi_1 - \xi_2||^{-\alpha}$ then the volume of flow through a typical point of "importance" m is proportional to $m^{\beta}(\alpha)$.
2. Starting from the intersection of two major roads, there is a boundary between the region reached by starting along the first road and the region reached by starting along the first road; we conjecture this boundary is fractal.

As undergraduate project we have looked at real-world subnetwork topologies:

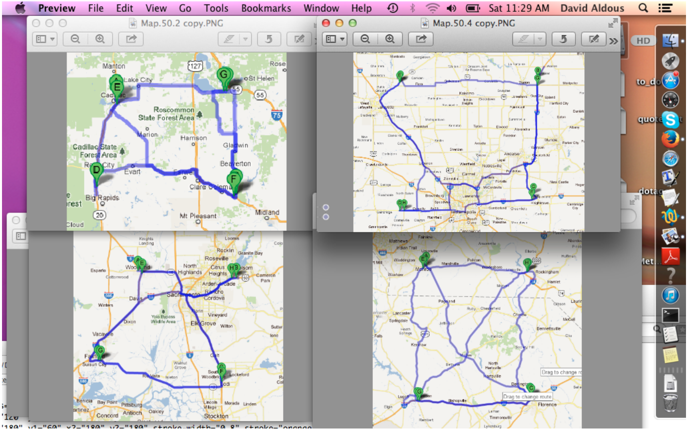

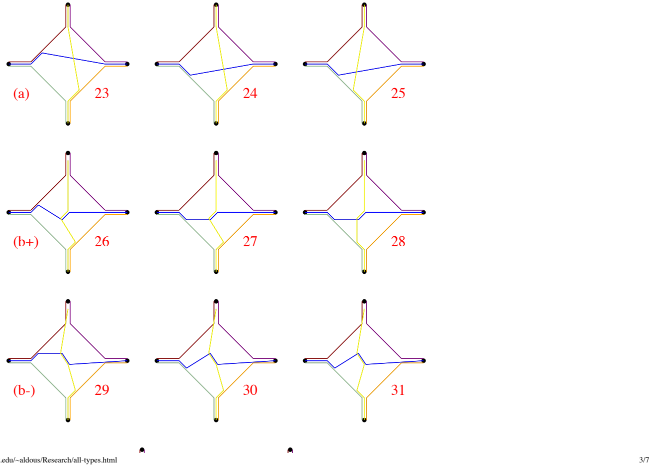

and listed all topologies on 4 addresses – different conventions from usual planar graph theory. Could compare distributions over these topologies in real-world and models.

Conclusions from Setting 2: Scale-invariant random spatial networks.

Mostly open problems. Someone please think about them!

- Do a broad variety of SIRSNs actually exist?
- Can we specific particular canonical or optimal ones?
- What are their math properties?
- Any realistic aspects?

Setting 3: Optimal subway networks.

In my "real world" context I describe typical math probability models (like the SIRSN) as "fantasy" – unconnected to any real data. But now I'll tell you an even more extreme fantasy – which will lead to an elementary-to-state math problem.

Imagine that somewhere there's an eccentric multi-billionaire with a taste for dramatic projects

and imagine there's a large spread-out metropolitan region without good public transport but with bad road traffic.

And then the billionaire has an idea ......

Where should one put a hypothetical such network?

Background. *Wikipedia – rapid transit* shows typical topologies (shapes) for subway-type networks.

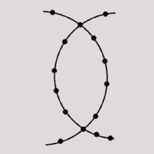
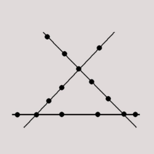


An interesting problem – see Aldous-Barthelemy (2019) but not discussed today – is can we reproduce these as optimal under some slightly-realistic toy model?

Musk's hypothetical tunnel network suggests an extreme model: "infinite speed, no wait time, no discrete stations".

Problem (a) *Find the connected network of length L that minimizes the expected distance from a random start to the closest point on the network.*

This depends on the density ρ of starting point; as default take 2-dimensional standard Normal.

This model implies constant speed outside the network, infinite speed within the network, but one is forced to use the network. Slightly more realistic to allow a direct route outside the network:

Problem (b) Find the connected network of length L that minimizes the expected time $t(\xi_1, \xi_2)$ between independent(ρ) points,

$$t(\xi_1, \xi_2) = \min(||\xi_1 - \xi_2||, s(\xi_1) + s(\xi_2))$$

$$s(\xi) = \text{ distance from } \xi \text{ to the closest point on the network.}$$

(For large L the two problems are essentially the same).

- We seek the actual optimal network for each L – how does the shape evolve as L grows?
- We work numerically. Mostly we consider specific parameterized shapes and optimize over parameters
- Alternatively try simulated annealing to optimize over all networks.

Lemma: *An optimal network must be a tree (or single path).*

Because: If there is a circuit, removing a length ε segment costs order ε^2 but reattaching it elsewhere benefits order ε.

We do have a theorem concerning the $L \to \infty$ behavior. This result is not so interesting, so [recall name of Musk's tunneling company]

Take a starting density ρ. Write $d(L)$ for the expected distance-to-network in the optimal network of length L.

Theorem (The Boring Theorem)

$$d(L) \sim \frac{1}{4L} \left(\int_{\mathbb{R}^2} \rho^{1/2}(z)\ dz \right)^2 \quad as\ L \to \infty.$$

What the argument actually shows is that a sequence of networks is asymptotically optimal as $L \to \infty$ if and only if the rescaled local pattern around a typical position z consists of asymptotically parallel lines with spacing proportional to $\rho^{-1/2}(z)$, but the orientations can depend arbitrarily on z. Visualize a fingerprint.

Near-optimal network for uniform density on square.

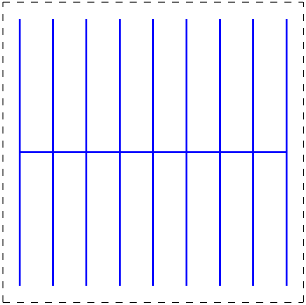

A simple topology is the star network, with $n \geq 2$ branches of lengths L/n from the center, with optimal choice of $n = n_L$. Comparing with the other shapes we have examined leads us to the (rather unexciting)

Observation. For the Gaussian density, the star networks are optimal or near-optimal over the range $0 < L \leq 16$.

[We guess this is quite robust – true for other densities]

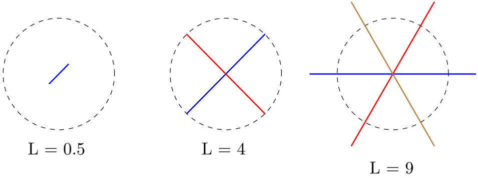

Figure 3: The optimal star networks for L = 0.5 and L = 4 and L = 9 (Gaussian density: the dashed circle indicates 1 s.d., so contains about 40% of the population).

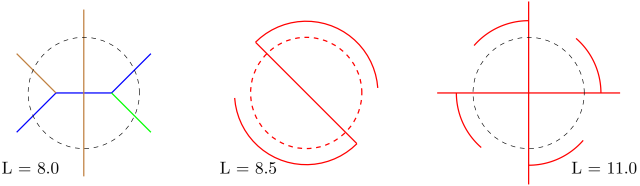

Figure 4: The "horse" network (left), the 2-arc network (center) and the 4-arc network (right).

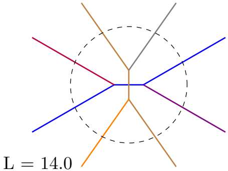

Figure 6: The spider network.

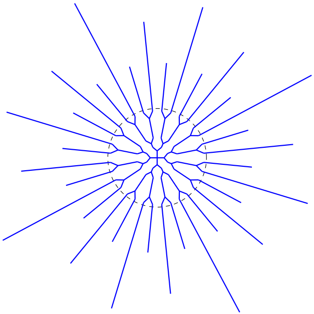

Figure 9: A finite approximation to an asymptotically optimal tree network.

As L grows an asymptotically optimal network becomes a branching tree.

Also one can construct **spirals** as asymptotically optimal. But contrary to our intuition, numerics say the tree is better (at second order).

Recall

Observation. For the Gaussian density, the star networks are optimal or near-optimal over the range $0 < L \leq 16$.

This was originally rather surprising.

By "reverse engineering" the Boring Theorem we see that star networks are asymptotically optimal for the non-Gaussian density of the rotationally invariant distribution on the radius-r_0 disc with R uniform on $[0, r_0]$.

Suggests robustness to density.

Conclusions from Setting 3: Optimal subway networks.

- Model is too unrealistic.
- Our intuition was poor.
- Don't hold your breath for the global sensation.