Some topics in spatial networks

David Aldous

4 October 2023

As you know . . . . . .

- Since 2000 there has been a huge literature on quantitative aspects of **networks** in general.
- In the more specific setting of **spatial networks**, the definitive reference has been the 2011 survey by Marc Barthélemy – cited by 2773 on Google Scholar. Expanded to a 2022 book *Spatial Networks: A complete introduction* – 400 pages of mathematics without theorems. (Like this talk!)
- But only 61 citations (with *Spatial Networks* in title) within MathSciNet, which covers most of theorem-proof mathematics.
- In fact, theorem-proof literature contains deep study of a few specific models/topics related to spatial networks.

Relevant well-studied topics in theorem-proof literature:

- Random planar graphs (up to isomorphism) = random mappings.
- The *random geometric graph* model.
- First passage percolation on a lattice.
- Stochastic geometry (random triangulations etc).

Also

- KPZ relation (statistical physics).
- Geometric spanner networks (CS computational geometry).

I take road networks as prototype examples, for two reasons

- readily accessible data
- none of above theory is very relevant to road networks (except geometric spanner networks)

which makes "road networks" appropriate for undergraduate projects.

3 reasons why there is little theory relevant to road networks.

- A key issue that no-one discusses: saying a network is **spatial** is saying that position matters: but population density varies hugely over a typical country. If we want a toy model representing infrastructure in a typical country then we should use some generic model of density variation. But no-one does. In the context of road networks, people use different models for urban or rural or inter-city networks.
- Within our toy models, orders of magnitude are usually rather obvious, but anything sharper seems impossibly difficult. For instance, for the length of TSP on n uniform random points in an area-n square, mean length $\sim cn$ but no way to calculate c. Also, unlike most of combinatorial optimization, average-case is same order of magnitude as worst-case.

- Freeways may (and do) intersect outside cities. So to model inter-city road networks we need *Steiner* networks, for which there is essentially no quantitative theory.

So what can we hope to say about road networks?

- Instead of deep mathematics, I am seeking to study different toy models for different properties. So this is *exploratory* math.
- **Toy models:** What would be an optimal network of a given type under some toy model? How similar are actual networks to such optima?
- I will mostly just talk about some of my older work and "neglected" topics that interest me: not intended to be comprehensive.

Topic 1: Elementary scaling and optimality issues.

Radio Shack had a peak of 7,300 stores in the United States in 1999. Were they located as theory suggests?

Toy model: minimize average distance to nearest store.

If we used a constant store density ρ:

$$\text{mean distance to closest store} \propto \rho^{-1/2} .$$

With variable population density $\mu(z)$ we should choose store density of the form $\rho(z) = g(\mu(z))$ for some $g(\cdot)$.

Math problem:
Minimize $\int \mu(z)\rho^{-1/2}(z)\ dz$ subject to $\int \rho(z) dz = 7300$.

Solution: Take $\rho(z) \propto (\mu(z))^{2/3}$.

This idea is classical – see e.g. Gastner-Newman (2006) for variants. And roughly fits data.

Is there anything similar for networks?

Suppose we use the same "minimize average distance to network' criterion to define an optimal network in a country.

Take population density $\mu(z)$. Write $d(L)$ for the mean distance-to-network in the optimal connected network of length L.

Theorem (The right answer to the wrong question.)

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

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.

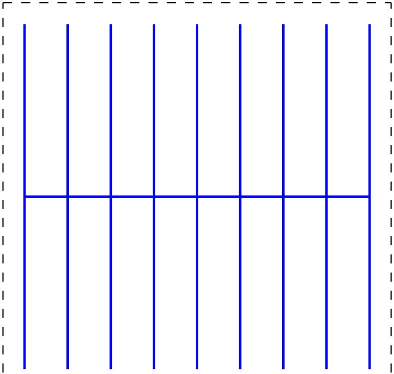

Topic 2: A slightly less silly model for inter-city networks.

Consider an inter-city road network: specifically, the network linking the N largest cities in a country of area A.

Simple heuristics say the total length will be $\propto \sqrt{NA}$.

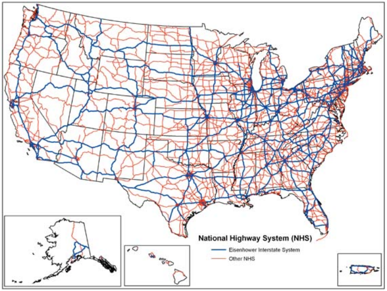

To make an analogy with the Radio Shack setting, for a toy model we take cities as N uniform random points in area A and ask for the optimal (?) network of given length $c\sqrt{NA}$, which will depend on c.

We want short routes!

Scale distance so that cities have density 1: that is, N cities in area N.

Let's consider route-lengths on all scales. That is, consider

$$\rho(d) = \frac{\mathbb{E}(\text{ route-length between city-pairs at distance} \approx d)}{d} - 1$$

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

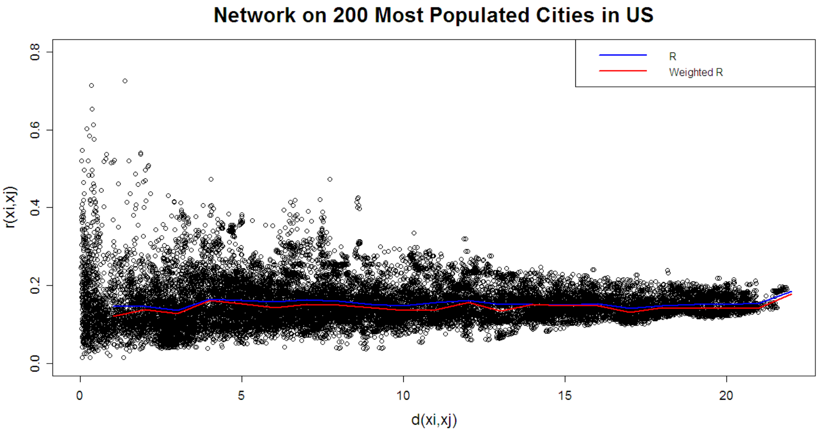
Network on 200 Most Populated Cities in US
R
Weighted R
r(xi,xj)
d(xi,xj)
0.0
0.2
0.4
0.6
0.8
0
5
10
15
20

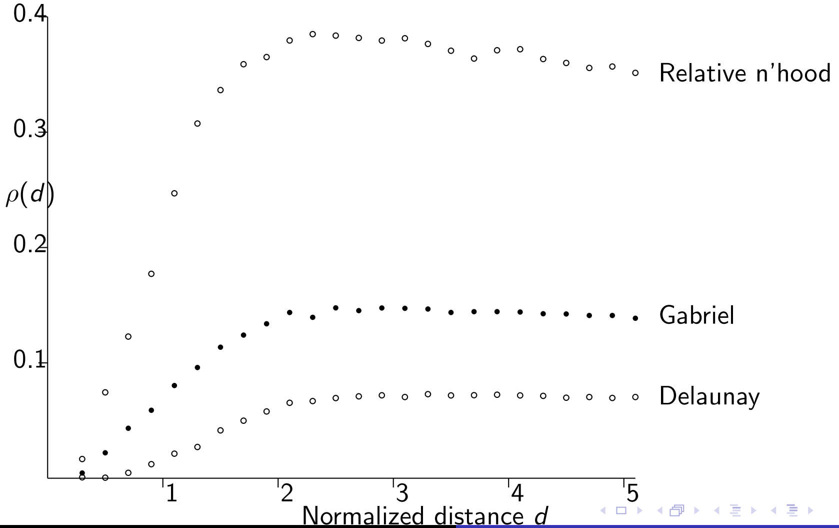
0.4
0.3
0.2
0.1
$\rho(d)$
1
2
3
4
5
Normalized distance d
Relative n'hood
Gabriel
Delaunay

This prompts us to use (for optimization criterion) 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.

Relate R to normalized length $L :=$ average road length per unit area.

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

This material is from Aldous - Shun (2010). Julian Shun, then an undergrad, now Associate Prof in EECS at MIT.

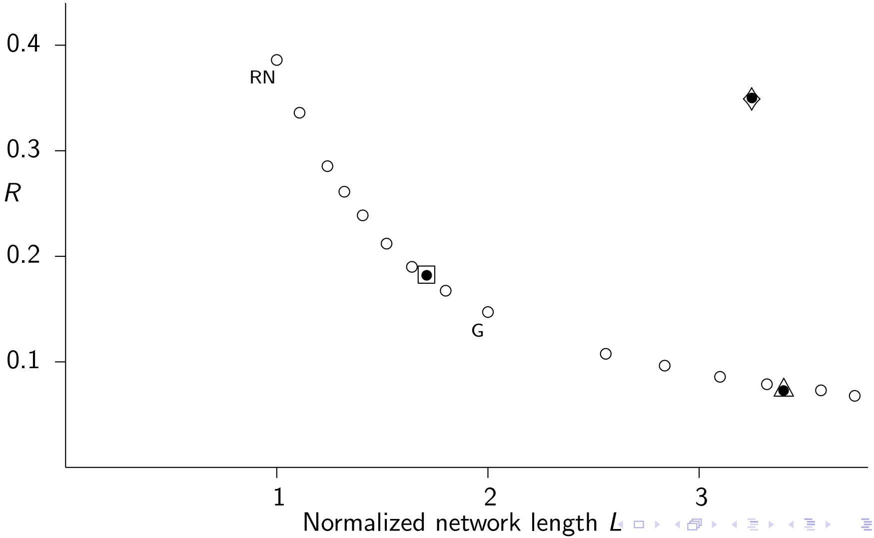
0.4
0.3
0.2
0.1
R
RN
G
1
2
3
Normalized network length L

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

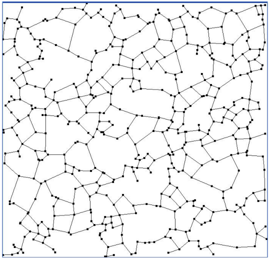

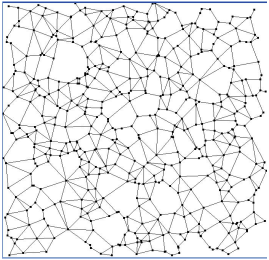

Open Problem: We conjecture that this family is almost optimal in our R versus L sense, but have no theorems.

Comment: Intuitively, this family of models should behave similarly for varying city density, because the network is defined by **relative** distances.

No interesting actual theorem within the preceding material . . . . . .
If you insist on a theorem then we need a cleaner set-up. Instead of measuring shortness of routes by

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

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

we can use the worst-case over city pairs (x, y):

$$\sigma := \max_{x,y} \frac{\text{route-length}(x, y)}{||y - x||}.$$

This σ is called **stretch** and studied in computational geometry, but along with other statistics and emphasizing orders of magnitude for worst-case data. Instead, let's be a mathematician and consider the most basic question we can think of . . .

Consider a configuration of n cities at arbitrary positions $\mathrm{x}_n = (x_1, \ldots, x_n)$ in a square of area n. For a network $\mathcal{N}$ connecting these cities, write $S(\mathcal{N})$ for the *stretch* and write

$$L(\mathcal{N}) = \frac{1}{n} \times (\text{network length of } \mathcal{N})$$

for normalized network length. We then define

$$\psi_n(\mathrm{x}_n, s) := \inf\{L(\mathcal{N}) : \; S(\mathcal{N}) \leq s\}$$

the infimum over all networks $\mathcal{N}$ connecting the cities x_n. So this quantifies the optimal trade-off between length and stretch for the given configuration. We can now consider in parallel the worst-case, that is $\sup_{\mathrm{x}_n} \psi_n(\mathrm{x}_n, s)$, and the average case $\mathbb{E}\psi_n(\mathrm{X}_n, s)$, where X_n consists of n independent uniform random positions in the area-n square. The purpose of this set-up is that it is intuitively obvious (and true) that there exist limit functions

$$\Psi^{\text{worst}}(s) = \lim_{n\to\infty} \sup_{\mathrm{x}_n} \psi_n(\mathrm{x}_n, s)$$

$$\Psi^{\text{ave}}(s) = \lim_{n\to\infty} \mathbb{E}\psi_n(\mathrm{X}_n, s)$$

where $0 < \Psi^{\text{ave}}(s) \leq \Psi^{\text{worst}}(s) \leq \infty$ for $0 < s < \infty$.

Aldous-Lando (2015) derive crude but explicit bounds for all 4 cases, via 4 different arguments.

- Derive upper bounds on $\Psi^{\text{worst}}(s)$ from elementary constructions where one first lays down a regular network of roads without paying attention to city positions, and then adds local links from cities to the network.
- Derive upper bounds on $\Psi^{\text{ave}}(s)$ from constructions (similar to the Θ-graphs in geometric spanner networks) in which from every city and every cone of given angle there exists a road leaving the city within the cone.
- Derive lower bounds on $\Psi^{\text{ave}}(s)$ for small s, based on the stochastic geometry relationship between network length and rate of intersections with a typical line.
- Derive lower bounds on $\Psi^{\text{worst}}(s)$ based on a notion of "local optimality" for specific networks on specific configurations.

There is much scope for improving these explicit bounds!

As a **maybe do-able theory question**, analogous to the earlier Theorem, the "large network" limit is the $s \downarrow 1$ limit, and one expects scaling exponents

$$\Psi(s) \asymp (s-1)^{-\alpha} \text{ as } s \downarrow 1$$

where the value of α does not depend on any detailed assumptions in the model (worst-case or average-case; whether or not Steiner points are allowed) but instead depends only on the fact the we are studying the length-stretch trade-off in two-dimensional space. Our results imply crude bounds on α:
an upper bound of $\frac{3}{4}$ for Ψ^{ave} and $\frac{5}{4}$ for Ψ^{worst}
a lower bound of $\frac{3}{8}$ for Ψ^{ave} and hence for Ψ^{worst} also.

Conclusions from Topic 2: inter-city networks.

1. Graphic is plausible as qualitative trade-off between length and "short routes" in such networks. Indeed it 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 roughly $2\sqrt{NA}$.
2. But these networks are visually wrong. Real world networks have long freeways (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.

Open Problem. Find a model (e.g. with speeds) in which this hierarchy is an emergent property of optimal networks.

Major practical issue: I don't have an algorithm to optimize over **all** possible networks on N points. Ad hoc attempts to tweak the networks above have shown little improvement.

Topic 3: subway networks.

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

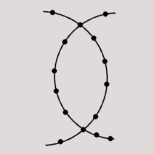

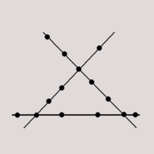

This is another context where variation of population density (source/destination) is important. What do we expect to see in a large network?

We typically see a well-connected core, often delineated by a circular line, from which branching lines spread away.

Can we reproduce these qualitative features as an optimal network within some toy model?

Model:
Gaussian or power-law density of source/destination.
Travel fast on subway network, slow off network.
Seek to minimize mean journey time.

One parameter S =fast/slow ratio.
What are the optimal networks for different total lengths L?
[Previous theorem was the $S = \infty$ case.]

In next figure, dashed circle is 1 s.d. of Gaussian, contains 40% of population. (Similar results for power-law density).

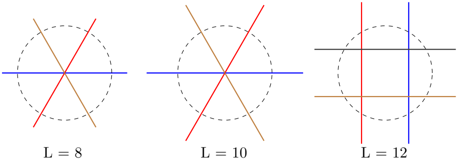

Figure 10: General model, S = 8 and W = 0.05, Gaussian population density; for L around 12 the **star** and **hashtag** shapes are almost equally good.

Figure 11: General model, S = 4 and W = 0: for L = 18 the star is optimal but for L = 21 the **3x3 grid** is optimal for the Gaussian density,

Figure from unfinished project with Marc.

Open problem. Write some efficient code to study large L for different networks.

Published paper Aldous - Barthélemy (2019) discussed tree-networks. Not exciting.

Topic 4: The KPZ relation.

Consider some translation-invariant model for routes in the infinite plane. For the route between two typical points at Euclidean distance r, consider

$$\begin{aligned} T(r) &= \text{route length} \\ D(r) &= \text{max deviation from straight line.} \end{aligned}$$

One expects

$$\mathbb{E}[T(r)] \sim cr, \quad \mathrm{var}[T(r)] \asymp r^{2\chi}, \quad \mathbb{E}[D(r)] \asymp r^{\xi}$$

where exponents χ, ξ are model-dependent but should satisfy the KPZ relation

$$\chi = 2\xi - 1.$$

Extensive study of this relation in statistical physics. Some aspects are rigorously proved in the lattice setting, e.g. Sourav Chatterjee (2013), but **hard Open Problem** to adapt to models based on Poisson points.

$$\mathbb{E}[T(r)] \sim cr, \quad \mathrm{var}[T(r)] \asymp r^{2\chi}, \quad \mathbb{E}[D(r)] \asymp r^{\xi}$$

$$\chi = 2\xi - 1.$$

Kartun-Giles – Barthelemy – Dettmann (2019) do a simulation study of a variety of *proximity graphs* in which it appears that $\chi = 1/5$ or $2/5$ in each model studied. Proving this would be very impressive!

[OK, not actually relevant to real road networks].

Here's one back-of-envelope indication of the relation.

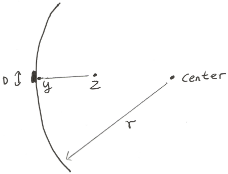

As y moves through arc length D, distance $||y-z||$ varies by order D^2/r

Routes in a random network

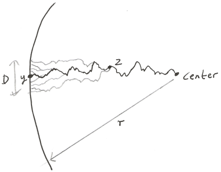

- arc-length D is order $D(r)$.
- route-length L is $\approx$ linear in distance.
- If we ignore randomness:
- $y \to L(y, \text{center}) \approx$ constant
- $y \to L(y,z)$ varies by $O(D^2(r)/r)$ by last slide
- Contradiction: $L(y, \text{center}) = L(y,z) + L(z, \text{center})$
- This difference must be masked by randomness
- s.d.(T_r) at least $O(D^2(r)/r)$
- $r^{\chi} \gtrsim r^{2\zeta - 1}$
- $\chi \geq 2\zeta - 1$

This is my attempt at sophisticated math theory, albeit non-realistic.

(Final) Topic 5: Scale-invariant random spatial networks

I will show a simulation of the following type of process.

- Start with an arbitrary network on the infinite plane (see a window).
- New vertices arrive as "Poisson rain" in space-time.
- Each arriving vertex is then linked to the existing network by new edges defined by some rule that is "scale invariant" in the sense of depending only on **relative** distances. For instance "link to the 2 closest vertices".
- Now "zoom in", that is continually expand the plane, to maintain a constant mean number of vertices within the window.

`https://www.stat.berkeley.edu/~aldous/Research/SInetwork-4.mp4`

Claim: under minimal assumptions on the "rule" and the initial network, this process converges in distribution to a random network on the plane which is invariant under this "zoom in, and add new vertices" procedure.

Comment 1. I have not tried to write a general proof – looks similar to standard methods for random geometric graphs – **student project?**

Comment 2. Can one do any quantitative study of this invariant distribution (in terms of the rule)? For instance, distribution of edge-lengths at a typical vertex?

I want to make a model for inter-city road networks, in which roads have different speeds.

- Consider the time-invariant distribution as the time-0 configuration of the dynamic process run over time $-\infty < t \leq 0$.
- On an edge (road) appearing at time $t < 0$, the speed is $e^{-\beta t} > 1$.
- Define the *route* between two vertices as the shortest-**time** route.

Now imagine (sorry, no graphics) the simulated process, scaled to have 1,000,000 vertices in the unit square window.

Fix k positions in the square – say $k = 7$.
Choose k vertices close to these k positions and draw the $\binom{k}{2}$ routes between each pair.

If we didn't have the hierarchy of different speeds, these routes would be almost straight lines between each pair. But now our routes depend on relative speeds along edges. From the time-invariance of the dynamic construction, we get (heuristics now) a **scale-invariance** property in the ' density of vertices $\to \infty$ limit.

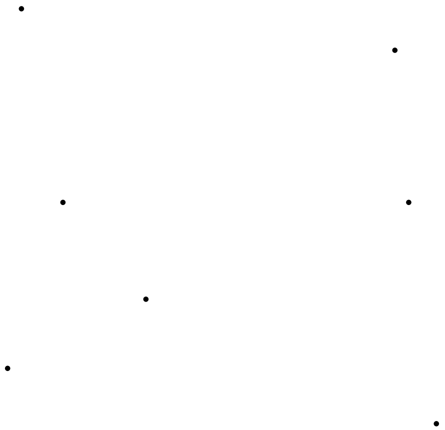

7 points in a window.

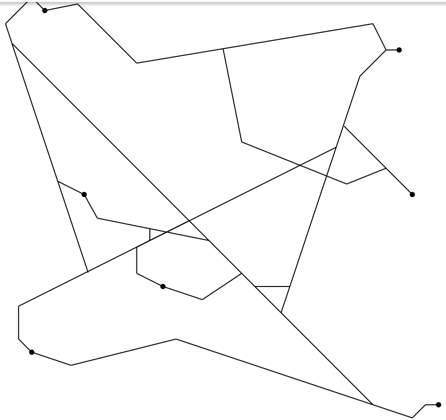

Scale-invariance means: doing this within a randomly positioned window, the statistics of the subnetwork observed don't depend on the scale, i.e. don't depend on whether the side length is 10 km or 100 km.

As undergraduate project we have looked at real-world subnetwork topologies (for $k = 4$ vertices, roughly at corners of a square).

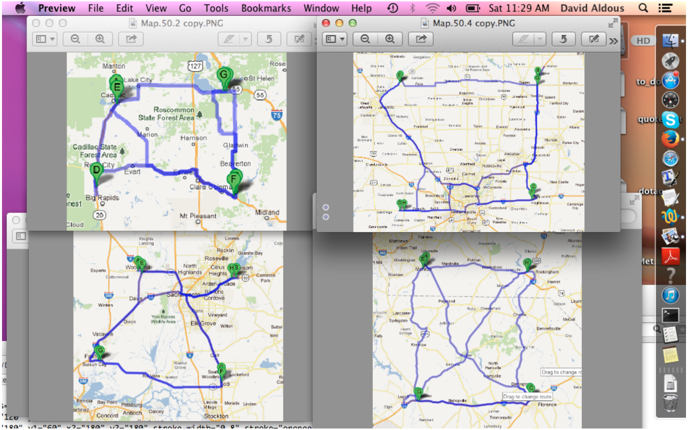

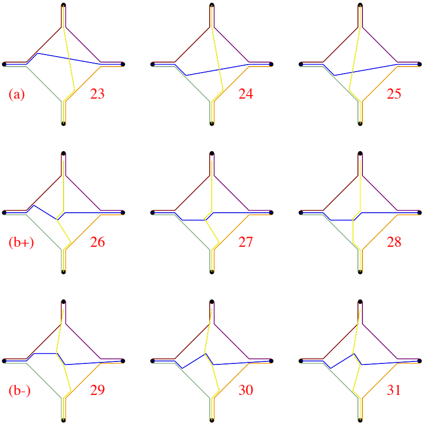

https://www.stat.berkeley.edu/~aldous/Research/all-types.html 3/7

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.

The dynamic construction in the video is very artificial, but the heuristics suggest existence of a large class **SIRSN** of processes satisfying certain axioms. Our old papers contain (only) two rigorous explicit constructions of SIRSN models, based on a rectangular grid or a Poisson line process for the different speed edges.

- (Aldous, David and Karthik Ganesan). True scale-invariant random spatial networks. Proc. Natl. Acad. Sci. USA 110 (2013).
- (Aldous, David). Scale-Invariant Random Spatial Networks. Electronic J. Probability 19 (2014) article 15: 1–41.
- (Kendall, Wilfrid S). From random lines to metric spaces. Ann. Probab. 45 (2017), no. 1, 469–517.

Conceptual starting point:

Idea behind our mathematical set-up: start with **routes** between **addresses** instead of **roads**, and work in the continuum plane.

We abstract *Google maps* as an "oracle" that for any start/destination 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 earlier.

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 $\mathbb{E}[D_1] < \infty$ 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.

Technically 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 $\ell < \infty$.

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 ones?
- Which SIRSNs optimize the trade-off between $\mathbb{E}[D_1]$ and ℓ, that is "short routes" versus "cost"?
- What are their mathematical properties? Similar or different from first-passage percolation paths? Doubly-infinite geodesics?
- Any realistic aspects?

$$\mathbb{E}[T(r)] \sim cr, \quad \mathrm{var}[T(r)] \asymp r^{2\chi}, \quad \mathbb{E}[D(r)] \asymp r^{\xi}$$

$$\chi = 2\xi - 1.$$

SIRSN models satisfy KPZ trivially because by scale-invariance $\chi = \xi = 1$.