J. R. Landers April 2026

Path-Length Distributions, Inverse Geometry, and Designed Neighborhood Spectra

1. Flat Euclidean Space as a Baseline

Let $x, y \in \mathbb{R}^2$ be two fixed points separated by Euclidean distance

$$d = \|x - y\|.$$

In ordinary deterministic geometry, the distance between x and y is the length of the shortest path, namely d. But one can instead ask how the *collection of possible paths* is distributed by length. To make the counting problem finite and concrete, restrict to one-waypoint paths

$$x \to z \to y, \qquad z \in \mathbb{R}^2.$$

The length of such a path is

$$L_{xy}(z) = \|x - z\| + \|z - y\|.$$

The minimum occurs when z lies on the line segment from x to y, giving $L_{xy}(z) = d$. For any $\ell > d$, the level set

$$\{z : \|x - z\| + \|z - y\| = \ell\}$$

is an ellipse with foci x and y. Therefore the cumulative number of one-waypoint paths with length at most ℓ is proportional to the area enclosed by that ellipse.

The ellipse has semi-major axis $a = \ell/2$ and focal distance $c = d/2$, so its semi-minor axis is

$$b = \sqrt{a^2 - c^2} = \frac{1}{2}\sqrt{\ell^2 - d^2}.$$

Thus the enclosed area is

$$A(\ell) = \pi ab = \frac{\pi}{4}\ell\sqrt{\ell^2 - d^2}, \qquad \ell \geq d.$$

Writing $N_{xy}(\leq \ell) \propto A(\ell)$, the corresponding path-count density is

$$\rho_{xy}^{\mathrm{Euc}}(\ell) \propto \frac{d}{d\ell}\left[\ell\sqrt{\ell^2 - d^2}\right] = \frac{2\ell^2 - d^2}{\sqrt{\ell^2 - d^2}}, \qquad \ell > d.$$

In dimensionless form, with $u = \ell/d$,

$$\rho^{\mathrm{Euc}}(u) \propto \frac{2u^2 - 1}{\sqrt{u^2 - 1}}, \qquad u > 1.$$

This is not a normalized probability density over an infinite domain. It is better interpreted as a *path-volume density*: the amount of one-waypoint path volume available at each length. The important point is that this form is not arbitrary. Flat Euclidean space gives elliptical equal-length sets, and elliptical equal-length sets give this specific path-volume law.

2. The Forward Map and the Inverse Problem

The Euclidean example suggests a general construction. Let a geometry be represented by a metric-measure space

$$\mathcal{G} = (M, d, \nu),$$

where M is a space, d is a distance function, and ν measures the amount of available intermediate space. For $x, y, z \in M$, define the one-waypoint length

$$L_{xy}(z) = d(x, z) + d(z, y).$$

The cumulative path-volume function is

$$F_{xy}(\ell) = \nu\big(\{z \in M : d(x, z) + d(z, y) \leq \ell\}\big),$$

and the associated path-length density is

$$\rho_{xy}(\ell) = \frac{d}{d\ell} F_{xy}(\ell).$$

Thus the forward problem is

$$(M, d, \nu) \longmapsto \{\rho_{xy}(\ell) : x, y \in M\}.$$

The inverse problem asks for the reverse reconstruction:

$$\text{given } \widehat{\rho}_{xy}(\ell) \text{ for many pairs } x, y, \quad \text{infer } (M, d, \nu).$$

A variational formulation is

$$(\widehat{M}, \widehat{d}, \widehat{\nu}) = \arg \min_{(M,d,\nu)} \sum_{x,y} D\left(\widehat{\rho}_{xy}(\ell) \,\middle\|\, \frac{d}{d\ell}\nu\{z : d(x,z) + d(z,y) \le \ell\}\right),$$

where D may be squared L^2 error, Wasserstein distance, KL divergence, or another discrepancy between distributions.

The ordinary shortest-path distance is recovered from the left edge of the support:

$$d(x, y) = \inf\{\ell : \rho_{xy}(\ell) > 0\}.$$

But ρ_{xy} contains more information than $d(x, y)$. It tells us how the volume of near-shortest, moderately long, and very long detours accumulates. It is therefore a local path spectrum: a geometric signature based not only on optimal paths, but on the abundance of alternatives around them.

3. Starting from a Lognormal Path-Volume Law

Now suppose the observed or desired path-volume law is lognormal. Since all paths from x to y have length at least $d(x, y)$, it is natural to model the excess length

$$\epsilon_{xy} = \ell - d(x, y)$$

rather than ℓ itself. A shifted lognormal law is

$$\epsilon_{xy} \sim \text{LogNormal}(\mu_{xy}, \sigma_{xy}^2), \qquad \text{equivalently} \qquad \log \epsilon_{xy} \sim \mathcal{N}(\mu_{xy}, \sigma_{xy}^2).$$

The corresponding path-length density is

$$\rho_{xy}^{\text{LN}}(\ell) = \frac{1}{(\ell - d(x,y))\sigma_{xy}\sqrt{2\pi}} \exp\left[-\frac{(\log(\ell - d(x,y)) - \mu_{xy})^2}{2\sigma_{xy}^2}\right], \qquad \ell > d(x, y).$$

Geometrically, this says that the shell

$$\{z : d(x,z) + d(z,y) - d(x,y) \approx \epsilon\}$$

has measure approximately lognormal in ϵ. This is not the path-volume structure of ordinary flat space. Euclidean geometry produces elliptical volume growth; a lognormal law suggests a geometry in which detours accumulate multiplicatively.

A heuristic mechanism is

$$\epsilon \approx \prod_{k=1}^{m} A_k,$$

where the $A_k > 0$ are local distortion factors. Taking logarithms gives

$$\log \epsilon \approx \sum_{k=1}^{m} \log A_k.$$

If these local log-distortions are weakly dependent and satisfy a central-limit-like effect, then $\log \epsilon$ becomes approximately Gaussian and ϵ becomes approximately lognormal. The inverse problem is then:

$$\boxed{\text{Find } (M, d, \nu) \text{ such that } \frac{d}{d\ell}\nu\{z : d(x,z) + d(z,y) \le \ell\} \approx \rho_{xy}^{\text{LN}}(\ell).}$$

This may be called a *lognormal path-volume geometry* or a *multiplicative metric-measure geometry*. The essential point is that the geometry is characterized not merely by shortest paths, but by the full distribution of possible detours around shortest paths.

4. Designed Neighborhood Spectra for Engineering Systems

The same idea gives a practical design principle: learn a distance whose local neighborhoods have a prescribed shape. Suppose we have engineering objects

$$\mathcal{X} = \{x_1, \dots, x_n\},$$

such as sensors in an industrial system, substations in a power grid, components in a manufacturing process, machines in a fleet, or operating regimes of a dynamical system. Learn a distance $d_\theta(x_i, x_j)$ so that each object has a controlled radial neighborhood profile.

For an anchor x_i, define the empirical radial neighbor density

$$\widehat{\rho}_i(r) = \sum_{j \neq i} K_h(r - d_\theta(x_i, x_j)),$$

where K_h is a smoothing kernel. Choose a target spectrum $\rho_i^\star(r)$, for example a lognormal profile with a few very close analogues, several near alternatives, and a long tail of more weakly related states:

$$\rho_i^\star(r) = \frac{1}{(r - r_{0,i})\sigma_i \sqrt{2\pi}} \exp\left[-\frac{(\log(r - r_{0,i}) - \mu_i)^2}{2\sigma_i^2}\right], \qquad r > r_{0,i}.$$

The design problem is

$$\theta^\star = \arg\min_\theta \sum_i D(\widehat{\rho}_i(r) \| \rho_i^\star(r)) + \lambda \mathcal{L}_{\text{task}}(\theta) + \Omega(\theta),$$

where $\mathcal{L}_{\text{task}}$ may encode forecasting accuracy, anomaly-detection performance, clustering quality, or physical consistency, and Ω prevents degenerate solutions.

A discrete shell version is even more direct. For annuli $A_m = [a_m, a_{m+1})$ and desired counts c_m, impose

$$\#\{j : d_\theta(x_i, x_j) \in A_m\} \approx c_m.$$

Equivalently,

$$\theta^\star = \arg\min_\theta \sum_{i,m} \left(\sum_{j \neq i} \mathbf{1}[d_\theta(x_i, x_j) \in A_m] - c_m \right)^2.$$

This turns geometry into a design variable. The learned space is not only optimized for task performance; it is shaped so that every object has a prescribed arrangement of close, intermediate, and distant neighbors. In engineering systems, this could support stable analogues for prediction, redundant near alternatives, and structured long-tail behavior for extrapolation or anomaly detection.

A small numerical illustration makes the design objective concrete. Take fifteen random points in $\mathbb{R}^3$, and let $Y = \{y_1, \dots, y_n\}$ denote optimized coordinates. A simple discrete surrogate for the Section 4 objective is

$$\mathcal{L}(Y) = \sum_{i=1}^n \frac{1}{n-1} \left\| \operatorname{sort}\big(\{\|y_i - y_j\| : j \neq i\}\big) - q^{\text{LN}} \right\|_2^2 + \eta\, R(Y),$$

where q^{LN} is a target vector of lognormal quantiles and $R(Y)$ is a mild regularizer preventing degenerate collapse. This is a direct finite-sample analogue of matching each empirical neighborhood spectrum to a prescribed lognormal profile. In a simple experiment with $n = 15$, the mean spectrum mismatch decreases from approximately 1.914 in the original Euclidean configuration to approximately 0.949 after optimization.

Figure 1 shows the initial Euclidean point cloud. Figure 2 shows the optimized geometry obtained by minimizing the global neighborhood-spectrum loss. The optimized configuration is not intended to be visually “lognormal” in a literal spatial sense; rather, what becomes lognormal-like is the distribution of distances seen from each anchor. Figure 3 shows these anchor-wise radial distributions. The dashed curve in each panel is the target lognormal law, while the solid curve is the fitted empirical density in the learned space. The point of the example is not that the fit is perfect in such a small system, but that the Section 4 objective has a clear computational realization and visibly reshapes the geometry in the intended distributional direction.

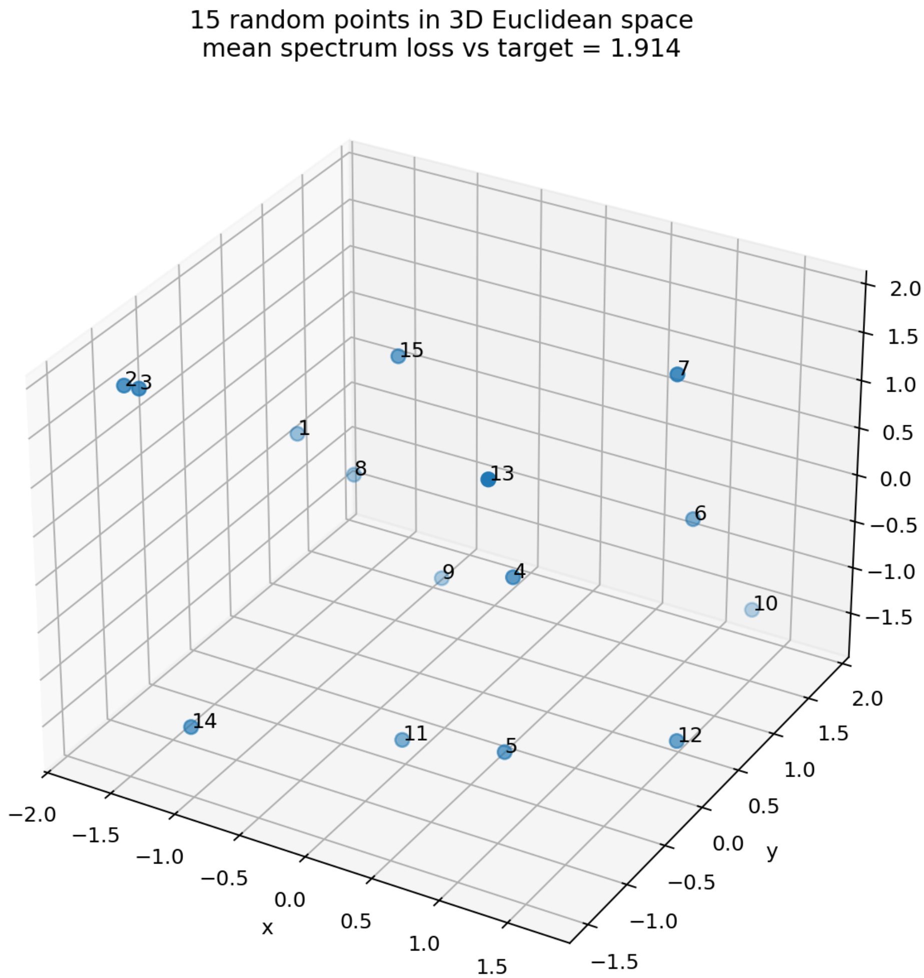

Figure 1: A toy dataset of fifteen random points in ordinary Euclidean $\mathbb{R}^3$. The title reports the mean mismatch between the anchor-wise neighborhood spectra and a common target lognormal spectrum.

5. Lognormal Spectrum Sparsification

The construction in Sections 1–3 concerns path-volume geometry. For fixed endpoints x, y, one varies an intermediate point z and studies the induced distribution of detour lengths

$$L_{xy}(z) = d(x, z) + d(z, y).$$

Section 4 shifts the viewpoint. Instead of studying detours between a fixed pair of endpoints, it designs the radial neighborhood spectrum around each data point. For an anchor x_i, the empirical object is

$$\widehat{\rho}_i(r) = \sum_{j \neq i} K_h(r - d_\theta(x_i, x_j)).$$

Thus the path-volume idea becomes a local metric-design principle: choose d_θ so that each point sees the rest of the system through a prescribed radial distribution. In particular, a shifted lognormal spectrum imposes a geometry with a few close analogues, a structured intermediate region, and a controlled long tail of weaker relations. This is no longer literally a distribution of paths from x to y; it is a distribution of neighbors around each anchor, induced by the learned metric.

Let the target radial law around x_i be

$$\rho_i^\star(r) = \frac{1}{(r - r_{0,i})\sigma_i\sqrt{2\pi}} \exp\left[-\frac{(\log(r - r_{0,i}) - \mu_i)^2}{2\sigma_i^2}\right], \qquad r > r_{0,i}.$$

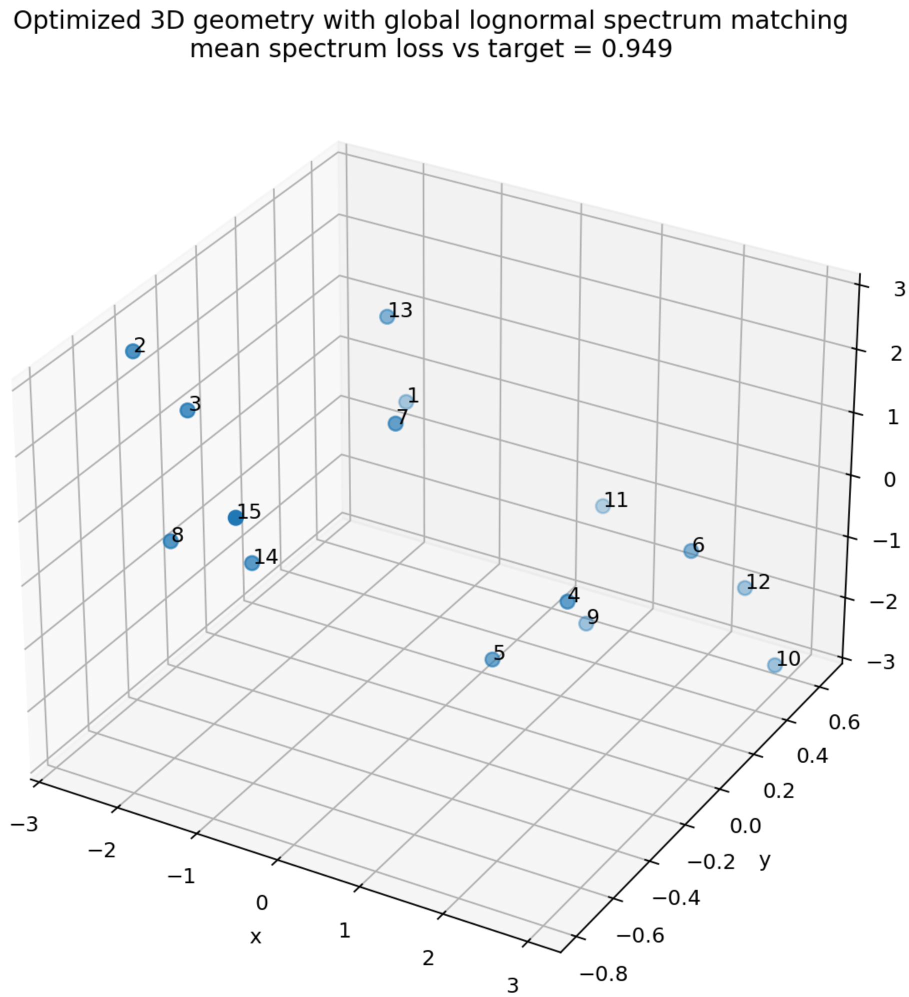

Figure 2: The same fifteen points after optimizing the coordinates to reduce the global lognormal neighborhood-spectrum loss. The optimization does not prescribe a particular visual shape in $\mathbb{R}^3$; it prescribes the radial distance distributions around anchors.

The global design objective is

$$\theta^{\star} = \arg\min_{\theta} \sum_{i} D(\widehat{\rho}_i(r) \| \rho_i^{\star}(r)) + \lambda \mathcal{L}_{\text{task}}(\theta) + \Omega(\theta).$$

The summation over anchors is essential. A lognormal spectrum around one point does not imply a lognormal spectrum around every other point. The condition is imposed anchor by anchor, producing a learned geometry whose local radial spectra are globally constrained.

This gives a natural graph-search interpretation. Let G be the full graph on the data points, with edge weights

$$w_{ij} = d_{\theta}(x_i, x_j).$$

Dijkstra's algorithm is unchanged: it computes shortest paths using these learned edge weights. The algorithmic gain comes instead from using the induced lognormal geometry to construct a sparse graph $H \subseteq G$.

Partition the radial distances around x_i into shells

$$A_m = [a_m, a_{m+1})$$

and define the target lognormal shell mass

$$p_{i,m} = \int_{A_m} \rho_i^{\star}(r)\, dr.$$

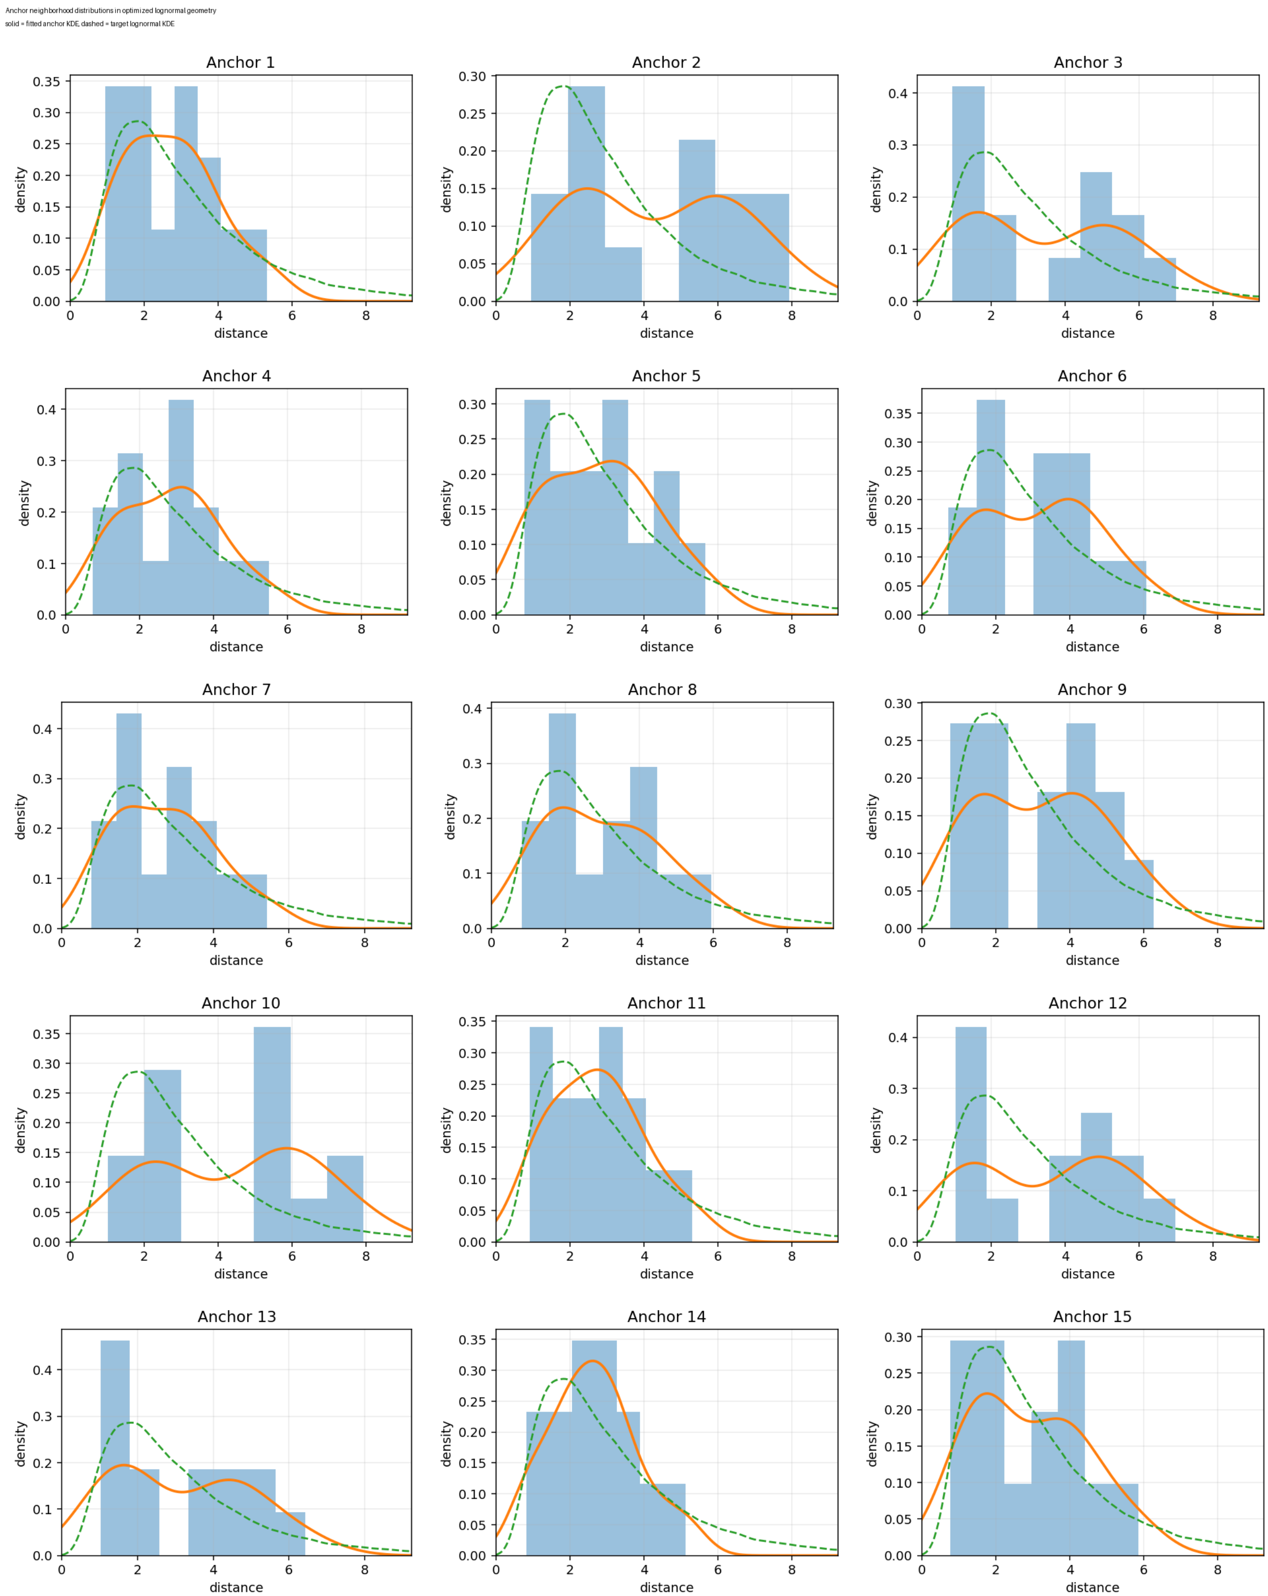

Figure 3: Anchor-wise neighborhood distributions in the optimized geometry. Each panel corresponds to one anchor. The solid curve is the fitted empirical density in the learned geometry, and the dashed curve is the target lognormal density. Collectively, the panels illustrate the global, anchor-by-anchor character of the Section 4 objective.

The learned geometry gives these shell masses a direct meaning: they describe how much of the local neighborhood structure lies at each distance scale. A sparse graph can then be formed by retaining representative edges from the important shells and discarding edges from low-mass regions. The sparsification is not merely an arbitrary deletion of edges. It is controlled by the radial distribution that the geometry was designed to realize.

For example, if all edges beyond radius R_i are removed around anchor x_i, then the discarded lognormal mass is

$$\beta_i(R_i) = \int_{R_i}^{\infty} \rho_i^\star(r)\, dr.$$

For the shifted lognormal spectrum,

$$\beta_i(R_i) = 1 - \Phi\left(\frac{\log(R_i - r_{0,i}) - \mu_i}{\sigma_i}\right).$$

Equivalently, to retain approximately $1-\beta$ of the local radial mass, choose

$$R_i = r_{0,i} + \exp(\mu_i + z_{1-\beta}\sigma_i),$$

where $z_{1-\beta}$ is the corresponding standard-normal quantile. Thus the lognormal geometry provides an explicit relation between a radius cutoff, the amount of discarded neighborhood mass, and the resulting edge budget. This is the distributional guarantee supplied by the lognormal spectrum.

This guarantee is distinct from a shortest-path guarantee. The discarded mass $\beta_i(R_i)$ measures how much local radial structure has been pruned. By itself, it does not guarantee that shortest paths are preserved. For that, one needs a replacement-path condition.

Proposition. Let G be the full learned-geometry graph with weights $w_{ij} = d_\theta(x_i, x_j)$, and let $H \subseteq G$ be a sparse graph obtained by lognormal spectrum sparsification. Suppose every dropped edge $(u, v) \in G \setminus H$ has a replacement path in H satisfying

$$\delta_H(u, v) \leq \tau d_\theta(u, v)$$

for some $\tau \geq 1$. Then for all source-target pairs s, t,

$$\delta_G(s,t) \leq \delta_H(s,t) \leq \tau\delta_G(s,t).$$

Consequently,

$$\frac{\delta_H(s,t) - \delta_G(s,t)}{\delta_G(s,t)} \leq \tau - 1.$$

Proof. Since $H \subseteq G$, every path in H is also a path in G. Removing edges can only increase shortest-path distances, so

$$\delta_G(s,t) \leq \delta_H(s,t).$$

Now take a shortest path in G,

$$s = v_0 \to v_1 \to \cdots \to v_q = t.$$

Each edge on this path is either retained in H, or has a replacement path in H of length at most τ times its original weight. Replacing every full-graph edge by its sparse-graph substitute gives a path in H of total length at most

$$\sum_{r=1}^{q} \tau d_\theta(v_{r-1}, v_r) = \tau\delta_G(s,t).$$

Since $\delta_H(s,t)$ is the shortest such path in H,

$$\delta_H(s,t) \leq \tau\delta_G(s,t).$$

Combining the two inequalities proves the claim.

The full interpretation is therefore

$$\begin{gathered}\text{lognormal neighborhood geometry} \Rightarrow \text{distribution-controlled sparsification} \\ \Rightarrow \text{sparse search} \Rightarrow \text{bounded distortion when replacement paths exist.}\end{gathered}$$

The lognormal law controls what portion of local geometry is discarded. The replacement-path condition controls how much shortest-path error this deletion can introduce. The first guarantee is specific to the induced lognormal geometry; the second is the usual stretch guarantee for sparse graph search. Together they give a principled way to prune a learned-geometry graph: not by arbitrary thresholding, but by removing edges according to the distributional structure the geometry was designed to have.

Summary. Flat Euclidean geometry implies an elliptical path-volume law. A general path-volume law defines an inverse problem for reconstructing a metric-measure geometry. A lognormal law suggests a multiplicative geometry of detours. In applications, the same formalism becomes a way to learn or design distances with controlled neighborhood spectra. The lognormal neighborhood view also induces a spectrum-controlled sparsification principle: discarded edge mass can be measured by the missing lognormal tail, while shortest-path distortion is bounded when dropped edges admit short replacement paths.