Geometric theory of cosmological structure via entropic curvature in Wasserstein space

Tsutomu T. Takeuchi [*]
Division of Particle and Astrophysical Science, Nagoya University, Furo-cho, Chikusa-ku, Nagoya 464-8602, Japan
and The Research Center for Statistical Machine Learning, The Institute of Statistical Mathematics, 10-3 Midori-cho, Tachikawa, Tokyo 190-8562, Japan

(Received 1 April 2026; revised 15 August 2026; accepted 27 October 2026; published 3 November 2026)

We formulate a weighted-Wasserstein framework for curvature statistics of cosmological large-scale structure. A coarse-grained density defines a weighted reference measure whose relative-entropy response separates into kinematic and density-induced Bakry-Émery contributions. On a Euclidean background, the latter is governed locally by $H_R = -\nabla^2 \ln \rho_R$; finite windows define density-weighted directional responses, while finite synthetic dimension subtracts a non-negative score-alignment deficit. The controlled local limit recovers the eigenvalue-threshold classification of the smoothed log-density Hessian, not a general equivalence with density-Hessian or gravitational-tidal classifiers. Independent smoothing and probe scales define a resolution flow. Controlled tests recover quadratic tensors to relative accuracy 10^{-9} or better and confirm $O(\varepsilon)$ and $O(\varepsilon^2)$ convergence. In a periodic three-dimensional Zel'dovich mock, the specified six-direction transport tensor matches its direct finite-window reference with subpercent median error and directional correlation above 0.999 995. Rotated and enriched probe designs reveal appreciable finite-window eigenvalue and classification dependence, distinct from estimator error; a converged angular projection still has a nonquadratic directional residual. Under thinning to $f_{\mathrm{samp}} = 1/16$, the six-direction internal error remains below 0.3%, while density reconstruction dominates the total error. At the fiducial radius, this finite-window observable has approximately 30% lower normalized sampling error than the pointwise log-density Hessian, each relative to its own full-sampling reference. In the Gaussian-linear limit, curvature spectra are differential transforms of two-point statistics; finite-window and finite-dimensional observables retain additional profile and directional information beyond that limit.

DOI: 10.1103/yrgj-rb72

I. INTRODUCTION

Cosmological large-scale structure forms through the amplification of density fluctuations by gravitational instability and their nonlinear evolution, and exhibits characteristic geometric structures such as voids, sheets, filaments, and halos (e.g., Refs. [1–3]). These structures are closely related to the statistical properties of the matter distribution as well as to the environmental

[*]Contact author: tsutomu.takeuchi.ttt@gmail.com

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

dependence of galaxy formation and baryonic physics, and their quantification and classification remain central problems in cosmology. Various theoretical and numerical methods have been proposed for this purpose (e.g., Refs. [4–6]).

Many conventional approaches are based on the local differential structure of the density or gravitational potential fields. For example, cosmic web classifications using the Hessian of the density field or the tidal tensor identify voids, sheets, filaments, and halos through the sign structure of eigenvalues [7–9]. Methods based on velocity fields and shear tensors are also widely used [10,11]. While these approaches characterize geometric structures, they do not explicitly incorporate global mass rearrangement or transport effects. Correlation functions and power spectra likewise encode amplitude correlations rather than the transport paths by which matter is rearranged [1,3,5].

In recent years, the Wasserstein distance from optimal transport theory has emerged as a powerful tool for comparing probability distributions and analyzing their geometry [12,13]. In cosmology, it has been applied to reconstruct initial conditions from galaxy distributions and to Lagrangian-based structure analysis (e.g., Refs. [14–17]). Combined with topological data analysis, it has also been used to compare summaries from persistent homology (e.g., Refs. [18–20]). These studies show that Wasserstein distance effectively captures morphological differences between distributions (e.g., Ref. [21]). On a finite domain, normalizing a positive matter density defines an element of $\mathcal{P}_2(M)$ without altering its spatial morphology. For the quadratic cost, W_2^2 is the minimum mean-squared displacement between normalized mass profiles, while their total mass or mean density can be retained separately as an amplitude observable.

However, existing approaches largely use the Wasserstein distance as a comparison metric without exploiting the intrinsic geometric structure of the probability measure space $\mathcal{P}_2(M)$, in particular, its curvature. In optimal transport theory, this space possesses a rich geometry in which Ricci curvature is defined via the geodesic convexity of entropy [22–24]. This formulation is closely related to Otto's formal Riemannian structure [25] and the gradient-flow formulation of the Jordan-Kinderlehrer-Otto scheme [26], providing a unified description of diffusion and curvature. The Wasserstein Hessian of relative entropy separates into a non-negative kinematic term generated by velocity gradients and a weighted Ricci-curvature term determined by the background metric-measure geometry. Isolating the latter is essential when entropy convexity is used to define a curvature observable.

Discrete analogues such as Ollivier-Ricci curvature define curvature via Wasserstein distances between local measures [27–29]. While useful for network and discrete data analysis, they differ fundamentally from curvature defined on continuous measure spaces via entropy convexity and its second variation. Morphological analysis that directly uses the continuous geometric structure of $\mathcal{P}_2(M)$ remains largely unexplored.

From this viewpoint, cosmological large-scale structure can be regarded not merely as a density field, but as the deformation of mass distributions within Wasserstein geometry (e.g., Ref. [21]). Its morphology can then be characterized through responses to controlled rearrangements as well as through static configurations. This provides a transport-geometric description of morphology distinct from conventional differential or statistical approaches.

Figure 1 summarizes the construction. A normalized coarse-grained density defines the weighted background and local-probe measures, while a controlled deformation generates a diagnostic Wasserstein geodesic whose relative-entropy response separates into kinematic and density-induced components.

Gravitational dynamics supplies the cosmological density field ρ_R, whereas the controlled Wasserstein geodesic probes the fixed density-weighted geometry and is not identified with physical time evolution.

We isolate the density-induced Bakry-Émery contribution from the total normalized entropy response. For the weighted measure defined by the coarse-grained density, the Euclidean curvature tensor is $H_R = -\nabla^2 \ln \rho_R$. Controlled probes define its finite-transport response $K^{(\infty)}_{\rho,\varepsilon}$ and infinitesimal directional limit $K^{(\infty)}_{\mathrm{eff}}$, while the N-Bakry-Émery tensor yields a non-negative score-alignment deficit $\Delta K^{(N)}$.

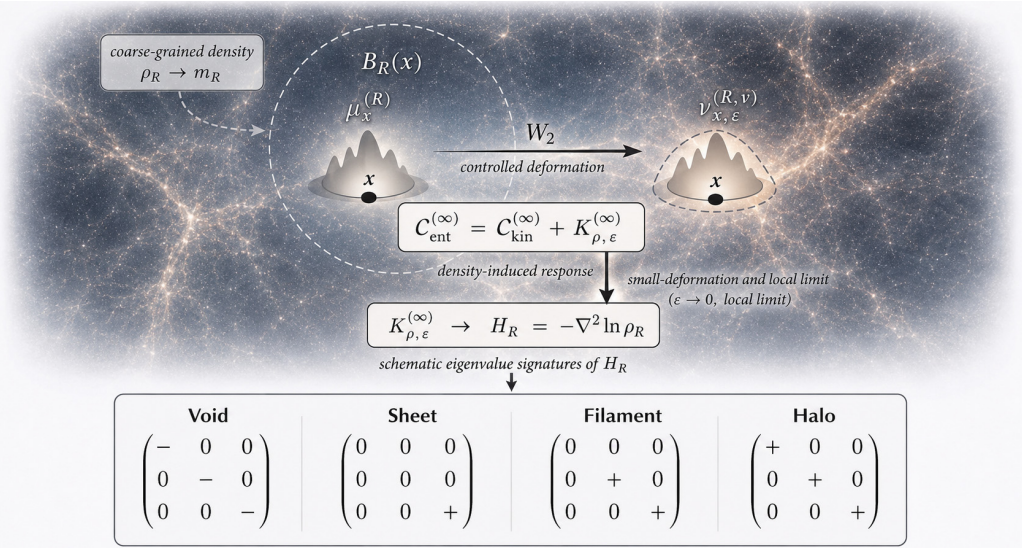


FIG. 1. Overview of the finite-scale curvature construction. A coarse-grained density ρ_R defines the weighted background measure m_R and its normalized local probe $\mu_x^{(R)}$. A controlled deformation generates the end point measure $\nu_{x,\varepsilon}^{(R,v)}$ and a diagnostic Wasserstein geodesic. The response decomposes as $C_{\mathrm{ent}}^{(\infty)} = C_{\mathrm{kin}}^{(\infty)} + K_{\rho,\varepsilon}^{(\infty)}$ and approaches the density-induced tensor $H_R = -\nabla^2 \ln \rho_R$ in the small-deformation and local-window limit. The schematic eigenvalue signatures of H_R yield the log-density-Hessian classification into voids, sheets, filaments, and halos.

The controlled small-deformation and local-window limit recovers the classification defined by the smoothed log-density Hessian, rather than classifiers based on the density-Hessian or gravitational-tidal tensor in general. Their relation is made explicit in Sec. V B. At finite displacement and probe radius, the transport and finite-dimensional responses retain information about the surrounding density profile. Their tensor representation and threshold classification also depend on the chosen probe directions and weights. We validate estimator fidelity, finite-window interpretation, sampling robustness, and this separate directional-design dependence using controlled potentials and a cosmological mock.

To our knowledge, no previous work has directly applied entropy-based Ricci curvature in Wasserstein geometry to the quantitative analysis of cosmological large-scale structure. The present framework places log-density-Hessian classification in a weighted-Wasserstein geometry and extends it to controlled finite-window transport responses.

The organization of this paper is as follows. In Sec. II, we review the foundations of optimal transport and measure-space geometry. Sections III–VI separate the entropy and curvature responses, localize them to finite windows, construct morphology indicators, and develop their operational estimation and validation. In Sec. VII, we discuss the implications of the framework and present our conclusions and outlook. The Appendixes provide the comparison-geometric derivation of finite-dimensional volume distortion and technical results on the local limit, scale dependence, curvature statistics, and supplementary numerical diagnostics.

II. FUNDAMENTALS OF OPTIMAL TRANSPORT AND GEOMETRY OF MEASURE SPACES

A. The optimal transport problem: Monge and Kantorovich

The optimal transport problem provides a framework for describing the rearrangement of probability distributions [12,30,31]. Given probability measures $\mu, \nu \in \mathcal{P}_2(M)$, Monge's formulation

seeks a map $T : M \to M$ satisfying

$$\nu = T_{\#}\mu \tag{1}$$

that minimizes

$$\int c(x, T(x))\, \mu(\mathrm{d}x). \tag{2}$$

Here, $T_{\#}\mu$ denotes the pushforward measure defined by

$$(T_{\#}\mu)(A) = \mu(T^{-1}(A)). \tag{3}$$

Kantorovich introduced a relaxed formulation based on a coupling $\pi \in \Pi(\mu, \nu)$:

$$\inf_{\pi \in \Pi(\mu,\nu)} \int_{M\times M} c(x, y)\, \pi(\mathrm{d}x\, \mathrm{d}y). \tag{4}$$

Here, $\Pi(\mu, \nu)$ is the set of probability measures on $M \times M$ with marginals μ and ν, so that $\pi(A \times M) = \mu(A)$ and $\pi(M \times B) = \nu(B)$ for measurable sets $A, B \subset M$. This formulation allows mass splitting and ensures the existence of an optimal plan under standard lower semicontinuity and moment assumptions. In the quadratic cost case $c(x, y) = d(x, y)^2$, the optimal value defines the squared Wasserstein distance. For $M = \mathbb{R}^d$ with quadratic cost and absolutely continuous μ, the optimal plan is induced by an almost-everywhere unique transport map [12].

B. Probability measures and the Wasserstein distance

Let $\mathcal{P}(M)$ denote the set of probability measures on a separable metric space (M, d). The subset with finite second moment is

$$\mathcal{P}_2(M) = \left\{ \mu \in \mathcal{P}(M) \;\middle|\; \int_M d(x_0, x)^2\, \mu(\mathrm{d}x) < \infty \right\}. \tag{5}$$

The reference point $x_0 \in M$ is arbitrary, and the definition is independent of its choice.

For $\mu, \nu \in \mathcal{P}_2(M)$, the Wasserstein distance is defined by

$$W_2(\mu, \nu)^2 = \inf_{\pi \in \Pi(\mu,\nu)} \int_{M\times M} d(x, y)^2\, \pi(\mathrm{d}x\, \mathrm{d}y). \tag{6}$$

This distance represents the minimal cost required to rearrange μ into ν and defines a nontrivial geometric structure on $\mathcal{P}_2(M)$ [12]. When d has dimensions of length, W_2^2 is the minimum mean-squared displacement. For locally normalized cosmological measures, it characterizes profile rearrangement rather than total mass, which may be retained separately; no interpretation as a physical action or gravitational energy is assumed.

C. Wasserstein geodesics and displacement interpolation

If (M, d) is a complete separable geodesic space, then $(\mathcal{P}_2(M), W_2)$ is also geodesic. For $\mu_0, \mu_1 \in \mathcal{P}_2(M)$, there exists a geodesic $\{\mu_t\}_{t\in[0,1]}$ satisfying

$$W_2(\mu_s, \mu_t) = |t - s|\, W_2(\mu_0, \mu_1). \tag{7}$$

When $M = \mathbb{R}^n$ and an optimal map T exists,

$$\mu_t = ((1 - t)\mathrm{Id} + tT)_{\#}\mu_0 \tag{8}$$

defines the displacement interpolation. On a Riemannian manifold,

$$\mu_t = (F_t)_{\#}\mu_0, \tag{9}$$

$$F_t(x) = \gamma_x(t), \tag{10}$$

with γ_x a constant-speed minimizing geodesic connecting x and $T(x)$.

These geodesics provide a natural notion of interpolation in measure space and play a central role in defining curvature through entropy convexity. In the present application, a prescribed local deformation defines a diagnostic Wasserstein interpolation rather than the physical trajectory generated by gravity.

D. Entropy and its gradient flow

Let (M, d, m) be a metric-measure space with fixed reference measure m, and write $\mu = q\,m$. Its Kullback–Leibler (KL)-type relative entropy is

$$\mathcal{E}_\infty(\mu \mid m) = \int_M q \ln q \,\mathrm{d}m. \tag{11}$$

For $M = \mathbb{R}^d$, $m = \mathrm{d}x$, and $\mu(\mathrm{d}x) = \rho(x)\mathrm{d}x$, this becomes $\mathcal{E}_\infty(\mu \mid \mathrm{d}x) = \int_{\mathbb{R}^d} \rho \ln \rho \,\mathrm{d}x$. We suppress the fixed reference measure when no confusion can arise.

Its Wasserstein gradient flow satisfies

$$\frac{\partial \mu_t}{\partial t} = -\nabla_{W_2} \mathcal{E}_\infty(\mu_t \mid m). \tag{12}$$

In the unweighted Euclidean case $m = \mathrm{d}x$, the continuity equation

$$\frac{\partial \rho_t}{\partial t} + \nabla \cdot (\rho_t v_t) = 0 \tag{13}$$

and velocity field

$$v_t = -\nabla \ln \rho_t \tag{14}$$

give

$$\frac{\partial \rho_t}{\partial t} = \Delta \rho_t. \tag{15}$$

Thus, the gradient flow of entropy coincides with the heat equation [26,32].

More generally, let $\mathrm{d}m_V(x) = e^{-V(x)}\mathrm{d}x$ and write $\mu_t = q_t m_V = r_t \mathrm{d}x$. The gradient velocity is $v_t = -\nabla \ln q_t = -\nabla(\ln r_t + V)$, and the corresponding equations are

$$\frac{\partial q_t}{\partial t} = \Delta q_t - \nabla V \cdot \nabla q_t, \tag{16}$$

$$\frac{\partial r_t}{\partial t} = \Delta r_t + \nabla \cdot (r_t \nabla V). \tag{17}$$

Thus, the same relative-entropy gradient flow is the Witten heat flow for q_t and the Fokker-Planck flow for r_t; it describes variational measure-space geometry rather than gravitational evolution.

E. Curvature-dimension condition and entropy convexity

The curvature-dimension condition $\mathrm{CD}(K, N)$ encodes a Ricci lower bound K and a synthetic dimension bound N. In the smooth weighted setting, it is represented by the weighted Ricci tensor introduced in Sec. III.

In metric-measure spaces, curvature is defined via entropy convexity along Wasserstein geodesics [22–24]. We use the relative entropy in Eq. (11). The $\mathrm{CD}(K, \infty)$ condition is as follows.

Definition II.1 [CD(K, ∞)]. A metric-measure space (M, d, m) satisfies $\mathrm{CD}(K, \infty)$ if, for any $\mu_0, \mu_1 \in \mathcal{P}_2(M)$ with finite relative entropy, there exists a Wasserstein geodesic $\{\mu_t\}_{t\in[0,1]}$ such that

$$\mathcal{E}_\infty(\mu_t \mid m) \leqslant (1-t)\mathcal{E}_\infty(\mu_0 \mid m) + t\,\mathcal{E}_\infty(\mu_1 \mid m) - \frac{K}{2} t(1-t) W_2(\mu_0, \mu_1)^2. \tag{18}$$

This inequality quantifies curvature through entropy variation. For finite N, the Rényi-type entropy is

$$\mathcal{E}_N(\mu \mid m) = -\int_M q^{1-\frac{1}{N}}\,\mathrm{d}m. \tag{19}$$

In the smooth or nonbranching setting, write $\mu_i = q_i m$ for optimally coupled end points and let γ_t be their geodesic. The CD(K, N) interpolation inequality is

$$q_t^{-1/N}(\gamma_t) \geqslant (1-t)\beta_{1-t}^{K,N}(r)^{1/N} q_0^{-1/N}(\gamma_0) + t\,\beta_t^{K,N}(r)^{1/N} q_1^{-1/N}(\gamma_1), \tag{20}$$

where $r = d(\gamma_0, \gamma_1)$. The coefficient $\beta_t^{K,N}(r)$ describes curvature-dependent transverse-volume distortion; its comparison-geometric derivation is given in Appendix A.

Unlike CD(K, ∞), the finite-dimensional condition retains this distance-dependent coefficient and is not obtained by copying the quadratic entropy-convexity deficit. Together, the two conditions link entropy convexity, weighted Ricci curvature, and volume distortion. In the next section, we separate the total entropy response from the weighted-curvature component of its second variation.

III. ENTROPY-CONVEXITY RESPONSE AND EFFECTIVE RICCI CURVATURE

Section II introduced the Wasserstein geometry and entropy functionals used below. We now distinguish the normalized geodesic-convexity response of entropy from the Ricci-curvature component of its second variation, which also contains a kinematic deformation term. This formulation uses, in the reverse direction, the characterization of Ricci curvature in optimal transport theory [22–24], and provides a framework in which curvature is regarded as a deformation property of probability distributions.

The KL-type entropy is associated with CD(K, ∞), whereas the Rényi-type entropy encodes the volume distortion entering CD(K, N); the finite-dimensional condition is therefore not obtained by copying the quadratic infinite-dimensional formula.

A. Wasserstein geometry and KL-type entropy

We consider the probability measure space $(\mathcal{P}_2(M), W_2)$ introduced in the previous section. Let $\mu = q\,m$, with $q = \mathrm{d}\mu/\mathrm{d}m$, and use the relative entropy $\mathcal{E}_\infty(\mu \mid m)$ defined in Eq. (11). We suppress the fixed reference measure when no confusion can arise; the distinction between q and m becomes essential when the cosmological density defines the weighted background geometry.

When viewed along a Wasserstein geodesic $\{\mu_t\}_{t\in[0,1]}$ connecting $\mu_0, \mu_1 \in \mathcal{P}_2(M)$, the change of entropy reflects not only the values of the density but also the geometry of mass rearrangement by optimal transport. What expresses this point most clearly is Otto's formal Riemannian geometry. From Otto's standpoint, $\mathcal{P}_2(M)$ is regarded as a formal infinite-dimensional Riemannian manifold, and geodesics are described by a velocity field v_t satisfying the continuity equation

$$\frac{\partial}{\partial t}\mu_t + \nabla\cdot(\mu_t v_t) = 0. \tag{21}$$

In particular, along the geodesic direction one can write $v_t = \nabla\phi_t$, and the norm of its tangent vector is given by

$$\|\dot{\mu}_t\|_{W_2}^2 = \int_M |v_t|^2\,\mathrm{d}\mu_t. \tag{22}$$

Under this formal metric, the fact that the Hessian of the entropy functional $\mathcal{E}_\infty$ carries curvature information constitutes the core of the notion of curvature in measure space.

At a general geodesic parameter $t \in (0, 1)$, the convexity deficit $(1-t)\mathcal{E}_\infty(\mu_0) + t\,\mathcal{E}_\infty(\mu_1) - \mathcal{E}_\infty(\mu_t)$ is a natural quantity that measures the strength of the geodesic convexity of entropy. Normalization by $W_2^2(\mu_0, \mu_1)$ gives a curvature-scale response whose finite-displacement value can retain higher-order dependence on the displacement amplitude.

Definition III.1 (KL-type normalized entropic convexity response). For $\mu_0, \mu_1 \in \mathcal{P}_2(M)$ and a Wasserstein geodesic $\{\mu_t\}_{t\in[0,1]}$ connecting them,

$$\mathcal{C}^{(\infty)}_{\mathrm{ent}}(\mu_0, \mu_1; t) \equiv \frac{2}{t(1-t)W_2^2(\mu_0, \mu_1)}[(1-t)\mathcal{E}_\infty(\mu_0) + t\,\mathcal{E}_\infty(\mu_1) - \mathcal{E}_\infty(\mu_t)]. \tag{23}$$

This is the total entropy-convexity response along the chosen geodesic; it is not yet the Ricci-curvature contribution because its short-geodesic limit also contains a non-negative kinematic term.

Furthermore, under the curvature-dimension condition CD(K, ∞),

$$\mathcal{E}_\infty(\mu_t \mid m) \leqslant (1-t)\mathcal{E}_\infty(\mu_0 \mid m) + t\,\mathcal{E}_\infty(\mu_1 \mid m) - \frac{K}{2}t(1-t)W_2^2(\mu_0, \mu_1) \tag{24}$$

holds, and therefore

$$\mathcal{C}^{(\infty)}_{\mathrm{ent}}(\mu_0, \mu_1; t) \geqslant K \tag{25}$$

follows. Thus, the response inherits the CD curvature bound, while its kinematic and weighted-curvature components must still be separated.

B. Otto structure, second variation, and Bochner-type representation

To isolate the curvature component, specialize to a smooth d-dimensional weighted Riemannian manifold (M, g, m_V) with $\mathrm{d}m_V = e^{-V}\mathrm{dvol}_g$. The reference measure m_V is held fixed along the diagnostic Wasserstein geodesic, and we write $\mu_t = q_t m_V$. Using the continuity equation, one obtains

$$\frac{\partial \mathcal{E}_\infty(\mu_t \mid m_V)}{\partial t} = \int_M \langle \nabla \ln q_t, v_t \rangle \,\mathrm{d}\mu_t. \tag{26}$$

Differentiating along the Wasserstein geodesic and integrating by parts gives

$$\frac{\partial^2 \mathcal{E}_\infty(\mu_t \mid m_V)}{\partial t^2} = \int_M [\|\nabla v_t\|^2 + \mathrm{Ric}_V(v_t, v_t)]\mathrm{d}\mu_t, \tag{27}$$

where ∇v_t is the covariant derivative of the velocity and $\mathrm{Ric}_V \equiv \mathrm{Ric} + \nabla^2 V$ is the Bakry-Émery Ricci tensor [12,22,33]. The two terms are, respectively, the kinematic deformation and the curvature of the weighted background geometry.

For a nonvanishing tangent velocity v, define their normalized contributions by

$$\mathcal{S}[\mu; v] \equiv \frac{\int_M \|\nabla v\|^2 \,\mathrm{d}\mu}{\int_M |v|^2 \,\mathrm{d}\mu}, \tag{28}$$

$$K^{(\infty)}_{\mathrm{eff}}[\mu; v] \equiv \frac{\int_M \mathrm{Ric}_V(v, v)\,\mathrm{d}\mu}{\int_M |v|^2 \,\mathrm{d}\mu}. \tag{29}$$

Here, $\mathcal{S}$ is the normalized strain contribution and $K^{(\infty)}_{\mathrm{eff}}$ is the averaged Bakry-Émery curvature in the selected transport direction.

Let $\{\mu_s\}_{s\in[0,\varepsilon]}$ be a Wasserstein geodesic with $\mu_0 = \mu$ and initial velocity v_0, and reparametrize it on $[0, 1]$ as $\widehat{\mu}_t = \mu_{t\varepsilon}$. For fixed $t \in (0, 1)$,

$$\lim_{\varepsilon\to 0} \mathcal{C}^{(\infty)}_{\mathrm{ent}}(\mu_0, \mu_\varepsilon; t) = \mathcal{S}[\mu; v_0] + K^{(\infty)}_{\mathrm{eff}}[\mu; v_0]. \tag{30}$$

The total response and the effective Ricci curvature are therefore distinct; the latter is obtained only after removing the kinematic term.

Furthermore, this structure is closely connected with the Bochner inequality and the Γ_2 operator. On a smooth Riemannian manifold, if one takes the generator to be

$$L = \Delta - \nabla V \cdot \nabla, \tag{31}$$

then in the notation of Bakry-Émery,

$$\Gamma(f) = \tfrac{1}{2}[L(f^2) - 2fLf] = |\nabla f|^2, \tag{32}$$

$$\Gamma_2(f) = \tfrac{1}{2}[L\Gamma(f) - 2\Gamma(f, Lf)] \tag{33}$$

are defined. In this case, the Bochner identity gives

$$\Gamma_2(f) = \|\nabla^2 f\|^2 + \mathrm{Ric}_V(\nabla f, \nabla f). \tag{34}$$

Here, $\mathrm{Ric}_V = \mathrm{Ric} + \nabla^2 V$ is the Bakry-Émery-type Ricci curvature. Therefore, the inequality $\Gamma_2 \geqslant K\Gamma$ simultaneously expresses the lower bound of Ricci curvature and the convexity control of the diffusion semigroup [12,33].

For $v = \nabla f$, the two terms in the Bochner identity coincide with the kinematic and weighted-curvature contributions in Eq. (30); the total entropy Hessian must therefore not be identified with its Ricci component.

C. Cosmological density as a weighted reference measure

Let $\rho_R(x) > 0$ be a smooth cosmological density field coarse-grained at scale R. The same symbol is used here for the coarse graining and local-probe scales; they are separated as (ℓ, R) in the operational analysis. On a finite or periodic domain M, define

$$Z_R \equiv \int_M \rho_R(x)\, \mathrm{dvol}_g(x), \tag{35}$$

$$p_R(x) \equiv \frac{\rho_R(x)}{Z_R}, \tag{36}$$

$$\mathrm{d}m_R(x) \equiv p_R(x)\, \mathrm{dvol}_g(x) = e^{-V_R(x)} \mathrm{dvol}_g(x), \tag{37}$$

$$V_R(x) \equiv -\ln p_R(x). \tag{38}$$

At fixed R, these background quantities remain unchanged while the probe measure is transported. Since Z_R is spatially constant, the derivatives of V_R equal those of $-\ln \rho_R$.

For the Euclidean cosmological background, $\mathrm{Ric} = 0$, and the density-weighted Bakry-Émery tensor is

$$\mathrm{Ric}_{V_R} = -\nabla^2 \ln \rho_R \equiv H_R. \tag{39}$$

Thus, the log-density Hessian is the weighted Ricci tensor of the normalized cosmological density, rather than an additional assumption applied to the entropy response.

For a local-probe measure μ and nonzero direction v, define

$$K_{\mathrm{eff}}^{(\infty)}[\mu; v; R] \equiv \frac{\int_M H_R(v, v)\, \mathrm{d}\mu}{\int_M |v|^2\, \mathrm{d}\mu}. \tag{40}$$

Equation (30) then applies with this density-induced curvature. It is the weighted Ricci term already present in the Otto-Bochner decomposition, not an *a posteriori* subtraction introduced to recover H_R.

D. Finite-dimensional extension: Rényi-type entropy and CD(K, N)

The KL-type entropy $\mathcal{E}_\infty$ corresponds to the infinite-dimensional limit. In contrast, in order to incorporate a finite dimension N explicitly, it is natural to introduce the Rényi-type entropy $\mathcal{E}_N$ defined in Eq. (19), with $q = \mathrm{d}\mu/\mathrm{d}m$. Here, the notation $\mathcal{E}_N$ is adopted in order to make explicit that $\mathcal{E}_\infty$ and $\mathcal{E}_N$ are situated within the same theory of entropy convexity. Its role is to encode finite-dimensional volume distortion through CD(K, N), not to define an independent high-density weighting of the cosmological field.

The curvature-dimension condition CD(K, N) asserts that density deformation along a Wasserstein geodesic is controlled by the volume distortion coefficient $\beta_t^{K,N}$. For a geodesic γ_t matched by optimal transport,

$$q_t^{-1/N}(\gamma_t) \geqslant (1-t)\beta_{1-t}^{K,N}(r)^{1/N} q_0^{-1/N}(\gamma_0) + t\,\beta_t^{K,N}(r)^{1/N} q_1^{-1/N}(\gamma_1) \tag{41}$$

holds. Here,

$$r = d(\gamma_0, \gamma_1) \tag{42}$$

is the distance between the pair of points matched by transport, and $t \in [0, 1]$ is the geodesic parameter. Therefore, $\beta_t^{K,N}(r)$ represents the distortion of the volume element at position t along a geodesic of length r. For example, in the case $K > 0$,

$$\beta_t^{K,N}(r) = \left[\frac{\sin(tr\sqrt{K/(N-1)})}{t\,\sin(r\sqrt{K/(N-1)})}\right]^{N-1} \tag{43}$$

holds. This is the model-space transverse-volume distortion; its comparison-geometric derivation is given in Appendix A.

What rewrites this comparison-geometric information in the language of entropy is the geodesic convexity of $\mathcal{E}_N$. Because the finite-dimensional condition retains $\beta_t^{K,N}(r)$ explicitly, its curvature is not obtained from an uncorrected quadratic midpoint deficit parallel to the KL-type formula.

For a smooth d-dimensional weighted Riemannian manifold with $\mathrm{d}m_V = e^{-V}\mathrm{dvol}_g$ and $N > d$, the corresponding tensor is

$$\mathrm{Ric}_{V,N} \equiv \mathrm{Ric} + \nabla^2 V - \frac{1}{N-d}\nabla V \otimes \nabla V. \tag{44}$$

Under the usual regularity assumptions, CD(K, N) is equivalent to $\mathrm{Ric}_{V,N} \geqslant Kg$, and $\mathrm{Ric}_{V,N} \to \mathrm{Ric}_V$ as $N \to \infty$ [22,24,34].

Definition III.2 (Finite-dimensional directional effective Ricci curvature). For a nonvanishing tangent velocity v and $N > d$, define

$$K_{\mathrm{eff}}^{(N)}[\mu; v] \equiv \frac{\int_M \mathrm{Ric}_{V,N}(v, v)\,\mathrm{d}\mu}{\int_M |v|^2\,\mathrm{d}\mu}. \tag{45}$$

The finite- and infinite-dimensional curvatures satisfy

$$K_{\mathrm{eff}}^{(\infty)}[\mu; v] - K_{\mathrm{eff}}^{(N)}[\mu; v] = \frac{1}{N-d}\frac{\int_M \langle \nabla V, v\rangle^2\,\mathrm{d}\mu}{\int_M |v|^2\,\mathrm{d}\mu} \geqslant 0. \tag{46}$$

Finite dimensionality therefore subtracts a non-negative directional score-alignment term rather than enhancing high-density regions generically.

For the cosmological weighted measure of Sec. III C,

$$\mathrm{Ric}_{V_R,N} = H_R - \frac{1}{N-d}\nabla \ln \rho_R \otimes \nabla \ln \rho_R. \tag{47}$$

Accordingly,

$$K_{\mathrm{eff}}^{(N)}[\mu; v; R] = K_{\mathrm{eff}}^{(\infty)}[\mu; v; R] - \frac{1}{N-d}\frac{\int_M \langle \nabla \ln \rho_R, v\rangle^2\,\mathrm{d}\mu}{\int_M |v|^2\,\mathrm{d}\mu}. \tag{48}$$

The correction vanishes pointwise at a stationary density point but can remain nonzero after finite-window averaging because the score field need not vanish throughout the surrounding region.

E. Entropy convexity as a unified framework

The KL-type and Rényi-type entropies encode complementary aspects of curvature-dimension geometry: The former yields the Otto-Bochner decomposition, while the latter retains finite-dimensional volume distortion through $\mathrm{Ric}_{V,N}$ and $\beta_t^{K,N}$. More specifically, $\mathcal{E}_\infty$ is associated with relative-entropy gradient flow, the Otto metric, the Bochner identity, and $\mathrm{CD}(K,\infty)$, whereas $\mathcal{E}_N$ is associated with Jacobi comparison, finite-dimensional volume distortion, and $\mathrm{CD}(K,N)$.

Equation (30) separates the total normalized response into $\mathcal{S} + K_{\mathrm{eff}}^{(\infty)}$, while $K_{\mathrm{eff}}^{(N)}$ is obtained by replacing Ric_V with $\mathrm{Ric}_{V,N}$. Thus, $\mathcal{C}_{\mathrm{ent}}^{(\infty)}$, $K_{\mathrm{eff}}^{(\infty)}$, and $K_{\mathrm{eff}}^{(N)}$ are, respectively, the total response and its infinite- and finite-dimensional curvature components. Their cosmological tensors are given in Eqs. (39) and (47).

Furthermore, this standpoint naturally connects to the issues of coarse graining, scale dependence, and finite resolution that will be considered later. Changing R changes the weighted geometry (m_R, V_R) and hence its Bakry-Émery tensors; the resulting scale dependence concerns directional curvature observables rather than abstract entropy deficits.

IV. LOCAL EFFECTIVE CURVATURE AND OBSERVABLES

We now localize the entropy and curvature responses of Sec. III to finite spatial probes. At fixed background geometry, the total finite-transport response contains both kinematic deformation and density-induced curvature, whereas the latter is isolated by the potential-energy component of relative entropy.

A. Local-probe measures and controlled transport directions

At fixed analysis scale R, the coarse-grained density ρ_R, weighted reference measure m_R, and potential V_R are held fixed. The same symbol is used here for density coarse graining and probe radius; the two scales are separated in the operational analysis.

For a point $x \in M$ and a scale $R > 0$, consider the geodesic ball with respect to the distance function d,

$$B_R(x) = \{y \in M \mid d(x,y) < R\}. \tag{49}$$

Its normalized restriction defines the local-probe measure

$$\mu_x^{(R)}(\mathrm{d}y) = \frac{\mathbf{1}_{B_R(x)}(y)}{m_R(B_R(x))}\,\mathrm{d}m_R(y). \tag{50}$$

Here, m_R specifies the background geometry and $\mu_x^{(R)}$ the region over which it is probed. A normalized smooth non-negative window may replace the top-hat restriction.

An arbitrary end point measure can compare local distributions but does not isolate a prescribed directional curvature. We therefore choose a smooth velocity $v = \nabla\phi$ near $B_R(x)$ and, in the flat spatial setting used below, define

$$T_{\varepsilon,v}(y) = y + \varepsilon v(y). \tag{51}$$

We take $\varepsilon > 0$ sufficiently small that this is an optimal quadratic-cost map; a sufficient local condition is $I + \varepsilon\nabla^2\phi \succ 0$. The end point measure is

$$\nu_{x,\varepsilon}^{(R,v)} = (T_{\varepsilon,v})_{\#}\mu_x^{(R)}. \tag{52}$$

Denote the connecting Wasserstein geodesic by $\{\mu_{x,t}^{(R,v,\varepsilon)}\}_{t\in[0,1]}$. For the displacement interpolation generated by Eq. (51),

$$\mu_{x,t}^{(R,v,\varepsilon)} = (T_{t\varepsilon,v})_{\#}\mu_x^{(R)}, \tag{53}$$

$$T_{t\varepsilon,v}(y) = y + t\varepsilon v(y). \tag{54}$$

The background quantities (ρ_R, m_R, V_R) remain fixed along this diagnostic interpolation. Define the squared transport distance by

$$\mathcal{W}^2_{x,R}(v,\varepsilon) \equiv W_2^2\big(\mu_x^{(R)}, \nu_{x,\varepsilon}^{(R,v)}\big). \tag{55}$$

B. Finite-scale entropic and density-induced responses

At the midpoint $t = 1/2$, the total normalized KL-type response is

$$\mathcal{C}^{(\infty)}_{\mathrm{ent}}(x;R,v,\varepsilon) \equiv \frac{8}{\mathcal{W}^2_{x,R}(v,\varepsilon)}\left[\frac{1}{2}\mathcal{E}_\infty\big(\mu_x^{(R)}\big|m_R\big) + \frac{1}{2}\mathcal{E}_\infty\big(\nu_{x,\varepsilon}^{(R,v)}\big|m_R\big) - \mathcal{E}_\infty\big(\mu_{x,1/2}^{(R,v,\varepsilon)}\big|m_R\big)\right]. \tag{56}$$

This is the total finite-scale entropy response and still contains the deformation-field contribution.

On the flat background, write $\eta(\mathrm{d}y) = r_\eta(y)\,\mathrm{dvol}_g(y)$ and define

$$\mathcal{E}_{\mathrm{vol}}(\eta) \equiv \int_M r_\eta(y)\ln r_\eta(y)\,\mathrm{dvol}_g(y), \tag{57}$$

$$\mathcal{V}_R(\eta) \equiv \int_M V_R(y)\,\eta(\mathrm{d}y). \tag{58}$$

The relative entropy then has the canonical decomposition

$$\mathcal{E}_\infty(\eta \mid m_R) = \mathcal{E}_{\mathrm{vol}}(\eta) + \mathcal{V}_R(\eta). \tag{59}$$

The midpoint response of $\mathcal{E}_{\mathrm{vol}}$ defines the finite-scale kinematic contribution

$$\mathcal{C}^{(\infty)}_{\mathrm{kin}}(x;R,v,\varepsilon) \equiv \frac{8}{\mathcal{W}^2_{x,R}(v,\varepsilon)}\left[\frac{1}{2}\mathcal{E}_{\mathrm{vol}}\big(\mu_x^{(R)}\big) + \frac{1}{2}\mathcal{E}_{\mathrm{vol}}\big(\nu_{x,\varepsilon}^{(R,v)}\big) - \mathcal{E}_{\mathrm{vol}}\big(\mu_{x,1/2}^{(R,v,\varepsilon)}\big)\right], \tag{60}$$

which is non-negative whenever the relevant entropies are finite.

The potential-energy component defines the density-induced response

$$K^{(\infty)}_{\rho,\varepsilon}(x;R,v) \equiv \frac{8}{\mathcal{W}^2_{x,R}(v,\varepsilon)}\left[\frac{1}{2}\mathcal{V}_R\big(\mu_x^{(R)}\big) + \frac{1}{2}\mathcal{V}_R\big(\nu_{x,\varepsilon}^{(R,v)}\big) - \mathcal{V}_R\big(\mu_{x,1/2}^{(R,v,\varepsilon)}\big)\right]. \tag{61}$$

Equations (59)–(61) yield the exact decomposition

$$\mathcal{C}^{(\infty)}_{\mathrm{ent}}(x;R,v,\varepsilon) = \mathcal{C}^{(\infty)}_{\mathrm{kin}}(x;R,v,\varepsilon) + K^{(\infty)}_{\rho,\varepsilon}(x;R,v). \tag{62}$$

Thus, $K^{(\infty)}_{\rho,\varepsilon}$ is selected by the relative-entropy decomposition rather than by an *ad hoc* subtraction.

Define the velocity normalization; under the optimality assumption, the corresponding transport cost is

$$\mathcal{D}_v(x;R) \equiv \int_{B_R(x)} |v(y)|^2\,\mu_x^{(R)}(\mathrm{d}y), \tag{63}$$

$$\mathcal{W}^2_{x,R}(v,\varepsilon) = \varepsilon^2\mathcal{D}_v(x;R). \tag{64}$$

The kinematic response then satisfies

$$\lim_{\varepsilon\to 0}\mathcal{C}^{(\infty)}_{\mathrm{kin}}(x;R,v,\varepsilon) = \mathcal{S}\big[\mu_x^{(R)};v\big]. \tag{65}$$

Assuming sufficient regularity and using $\nabla^2 V_R = -\nabla^2\ln\rho_R \equiv H_R$, expansion of the potential midpoint deficit gives

$$K^{(\infty)}_{\rho,\varepsilon}(x;R,v) = \frac{1}{\mathcal{D}_v(x;R)}\int_{B_R(x)} v(y)^\top H_R(y)v(y)\mu_x^{(R)}(\mathrm{d}y) + O(\varepsilon). \tag{66}$$

Hence, the infinitesimal density-induced curvature is

$$K_{\mathrm{eff}}^{(\infty)}(x;R,v) \equiv \lim_{\varepsilon\to 0} K_{\rho,\varepsilon}^{(\infty)}(x;R,v) = \frac{1}{\mathcal{D}_v(x;R)} \int_{B_R(x)} v(y)^\top H_R(y) v(y) \mu_x^{(R)}(\mathrm{d}y). \tag{67}$$

The finite-transport form Eq. (61) can be evaluated from $V_R = -\ln p_R$ without estimating H_R pointwise; the Hessian average appears only in the controlled small-deformation limit.

C. Finite-dimensional curvature correction and sensitivity

For $N > d$, the finite-dimensional tensor is

$$\mathrm{Ric}_{V_R,N} = H_R - \frac{1}{N-d} \nabla \ln \rho_R \otimes \nabla \ln \rho_R. \tag{68}$$

Define the non-negative score-alignment functional

$$\mathcal{G}_v(x;R) \equiv \int_{B_R(x)} [v(y)\cdot \nabla \ln \rho_R(y)]^2 \mu_x^{(R)}(\mathrm{d}y). \tag{69}$$

The directional finite-dimensional curvature is therefore

$$\begin{aligned} K_{\mathrm{eff}}^{(N)}(x;R,v) &\equiv \frac{\int_{B_R(x)} \mathrm{Ric}_{V_R,N}(v,v)\, \mu_x^{(R)}(\mathrm{d}y)}{\mathcal{D}_v(x;R)} \\ &= K_{\mathrm{eff}}^{(\infty)}(x;R,v) - \frac{\mathcal{G}_v(x;R)}{(N-d)\mathcal{D}_v(x;R)}. \end{aligned} \tag{70}$$

Finite dimensionality thus subtracts a non-negative directional score-alignment term. It can be enhanced by steep gradients, and it may remain nonzero over a finite window even when $\nabla \ln \rho_R(x) = 0$ at its center.

D. Indicators as observables

Define the non-negative finite-dimensional deficit by

$$\Delta K^{(N)}(x;R,v) \equiv K_{\mathrm{eff}}^{(\infty)}(x;R,v) - K_{\mathrm{eff}}^{(N)}(x;R,v) = \frac{\mathcal{G}_v(x;R)}{(N-d)\mathcal{D}_v(x;R)} \geqslant 0. \tag{71}$$

When the infinite-dimensional curvature is bounded away from zero, its relative magnitude is

$$\mathcal{R}^{(N)}(x;R,v) \equiv \frac{\Delta K^{(N)}(x;R,v)}{\left|K_{\mathrm{eff}}^{(\infty)}(x;R,v)\right|}, \quad \left|K_{\mathrm{eff}}^{(\infty)}(x;R,v)\right| > K_{\min}. \tag{72}$$

The threshold $K_{\min} > 0$ avoids an unstable ratio near zero curvature.

The observable hierarchy is now explicit: $\mathcal{C}_{\mathrm{ent}}^{(\infty)}$ is the total finite-transport response, $\mathcal{C}_{\mathrm{kin}}^{(\infty)}$ its kinematic part, $K_{\rho,\varepsilon}^{(\infty)}$ its density-induced finite-transport part, $K_{\mathrm{eff}}^{(\infty)}$ the corresponding infinitesimal curvature, and $\Delta K^{(N)}$ the finite-dimensional score-alignment deficit. The ratio $\mathcal{R}^{(N)}$ measures the last quantity relative to the infinite-dimensional curvature.

For point data, these observables require a positive reconstructed reference density. The potential expectation entering $K_{\rho,\varepsilon}^{(\infty)}$ is directly evaluable for weighted samples, whereas the entropy responses require either a regularized density or a discrete relative entropy defined with respect to a common reference measure. The operational construction is developed in Sec. VI.

The next section specializes v to affine and rank-one probes for morphology classification.

V. CLASSIFICATION OF COSMOLOGICAL LARGE-SCALE STRUCTURE BY LOCAL EFFECTIVE CURVATURE

The primary classification uses the density-induced finite-transport response $K^{(\infty)}_{\rho,\varepsilon}$ and its infinitesimal curvature limit $K^{(\infty)}_{\mathrm{eff}}$. The total entropy response is retained only as a diagnostic because it also contains the kinematic term, while $\Delta K^{(N)}$ and $\mathcal{R}^{(N)}$ provide secondary finite-window score-alignment indicators. The construction is summarized in Fig. 1.

A. Basic principles and simple local models

Voids, sheets, filaments, and halos are distinguished by the principal expansion and contraction directions of the density-weighted geometry. For each structure, we approximate the local density distribution in a form in which geometric characteristics appear explicitly.

In voids, the density is low at the center and increases toward the surroundings. As a local model, $\rho_R(x) \simeq \rho_0(1+\epsilon|x|^2)$ with $\epsilon > 0$ is adopted. This is a distribution in which the density is minimal at the origin and increases isotropically.

Sheets have one concentrated normal direction and two approximately extended in-plane directions. Writing the normal coordinate as $x_\perp$, a simple model is

$$\rho_R(x) \simeq \rho_0 \exp\left(-\frac{x_\perp^2}{\sigma_\perp^2}\right). \tag{73}$$

It is approximately uniform within the sheet.

Filaments have anisotropic structures elongated in one direction. Decomposing the coordinates as

$$x = (x_\parallel, x_\perp), \tag{74}$$

one can approximate $\rho_R(x) \simeq \rho_0 \exp(-|x_\perp|^2/\sigma_\perp^2)$. This is a structure that is nearly uniform along the filament direction and has contraction only in the perpendicular directions.

Since halos have strongly concentrated structures, $\rho_R(x) \simeq \rho_0 \exp(-|x|^2/\sigma^2)$ is used as an approximation. This distribution contracts isotropically, and the density gradient is steep toward the center.

These are local curvature models, not globally normalized cosmological profiles.

B. Local quadratic approximation and density-induced curvature

At a general point x, expand the coarse-grained log-density as

$$\ln \rho_R(y) = \ln \rho_R(x) + g_R(x)^\top (y-x) - \tfrac{1}{2}(y-x)^\top H_R(x)(y-x) + O(|y-x|^3), \tag{75}$$

where

$$g_R(x) \equiv \nabla \ln \rho_R(x), \tag{76}$$

$$H_R(x) \equiv -\nabla^2 \ln \rho_R(x). \tag{77}$$

The score term permits evaluation away from extrema and vanishes at a stationary point.

The field whose Hessian is classified must be specified. For the same smooth positive density at fixed smoothing scale,

$$-\partial_i \partial_j \ln \rho_R = -\frac{\partial_i \partial_j \rho_R}{\rho_R} + \frac{\partial_i \rho_R \, \partial_j \rho_R}{\rho_R^2}. \tag{78}$$

The gradient term generally prevents equivalence with the sign-reversed density Hessian. At a stationary point it vanishes, so the two tensors are related by the positive factor ρ_R^{-1} and have the same zero-threshold inertia. Their eigenvalue amplitudes, and classifications using the same

nonzero numerical threshold, need not agree. The gravitational-tidal tensor is instead the Hessian of the gravitational potential obtained from the field equation, not of the log-density. Thus, the four morphology names describe a common eigenvalue-counting rule, not an identity between these different classifiers throughout the volume.

For the local measure $\mu_x^{(R)}$, consider the affine deformation

$$T_{\varepsilon,A}(y) = x + (I + \varepsilon A)(y - x). \tag{79}$$

Assume

$$A = A^\top, \tag{80}$$

$$I + \varepsilon A \succ 0. \tag{81}$$

Then, $T_{\varepsilon,A}$ is the gradient of a convex quadratic potential and defines a controlled optimal-transport probe, with $\nu_{x,\varepsilon}^{(R,A)} = (T_{\varepsilon,A})_\# \mu_x^{(R)}$.

Writing $z = y - x$, define the second moment about the prescribed probe center by

$$\Sigma_x^{(R)} \equiv \int_{B_R(x)} zz^\top \, \mu_x^{(R)}(\mathrm{d}y). \tag{82}$$

It equals the covariance only when x is the probe barycenter.

Define the positive normalization

$$\mathcal{D}_A(x;R) \equiv \mathrm{Tr}\big(A\Sigma_x^{(R)}A^\top\big). \tag{83}$$

For a nondegenerate probe, $\mathcal{D}_A(x;R) > 0$, and optimality of the affine map gives

$$W_2^2\big(\mu_x^{(R)}, \nu_{x,\varepsilon}^{(R,A)}\big) = \varepsilon^2 \mathcal{D}_A(x;R). \tag{84}$$

The density-induced effective curvature is

$$K_{\mathrm{eff}}^{(\infty)}(x;R;A) = \frac{1}{\mathcal{D}_A(x;R)} \int_{B_R(x)} z^\top A^\top H_R(y) A z \mu_x^{(R)}(\mathrm{d}y). \tag{85}$$

It is therefore a transport-weighted finite-window average of the curvature tensor.

If $H_R(y) = H_R(x) + O(R)$ within the probe, then

$$K_{\mathrm{eff}}^{(\infty)}(x;R;A) = \frac{\mathrm{Tr}\big[H_R(x)A\Sigma_x^{(R)}A^\top\big]}{\mathcal{D}_A(x;R)} + O(R). \tag{86}$$

The finite-transport response correspondingly satisfies

$$K_{\rho,\varepsilon}^{(\infty)}(x;R;A) = \frac{\mathrm{Tr}\big[H_R(x)A\Sigma_x^{(R)}A^\top\big]}{\mathcal{D}_A(x;R)} + O(R+\varepsilon). \tag{87}$$

Thus, the quotient arises from the density-induced response, not from the total entropy response containing the kinematic term.

In particular, if one chooses

$$A = \mathbf{e}\mathbf{e}^\top \tag{88}$$

for a unit vector $\mathbf{e}$, the directional curvature is defined by

$$K_{\mathrm{eff}}^{(\infty)}(x;R;\mathbf{e}) \equiv K_{\mathrm{eff}}^{(\infty)}(x;R;A = \mathbf{e}\mathbf{e}^\top). \tag{89}$$

Since $A\Sigma_x^{(R)}A^\top = (\mathbf{e}^\top\Sigma_x^{(R)}\mathbf{e})\mathbf{e}\mathbf{e}^\top$, the finite-window response is

$$K_{\mathrm{eff}}^{(\infty)}(x;R;\mathbf{e}) = \frac{1}{\mathbf{e}^\top\Sigma_x^{(R)}\mathbf{e}} \int_{B_R(x)} \langle z, \mathbf{e}\rangle^2 [\mathbf{e}^\top H_R(y)\mathbf{e}] \mu_x^{(R)}(\mathrm{d}y). \tag{90}$$

It averages directional curvature over the local mass profile.

Whenever $\mathbf{e}^\top \Sigma_x^{(R)} \mathbf{e} > 0$,

$$K_{\mathrm{eff}}^{(\infty)}(x;R;\mathbf{e}) = \mathbf{e}^\top H_R(x)\mathbf{e} + O(R). \tag{91}$$

No isotropy assumption on $\Sigma_x^{(R)}$ is required. The finite-transport counterpart is $K_{\rho,\varepsilon}^{(\infty)}(x;R;\mathbf{e}) = \mathbf{e}^\top H_R(x)\mathbf{e} + O(R+\varepsilon)$; it provides a finite-displacement realization without requiring a pointwise Hessian estimate.

C. Directional finite-dimensional curvature correction

For the same affine probe, define the score-alignment functional

$$\mathcal{G}_A(x;R) \equiv \int_{B_R(x)} [(Az)\cdot \nabla \ln \rho_R(y)]^2 \mu_x^{(R)}(\mathrm{d}y). \tag{92}$$

For $N > d$,

$$K_{\mathrm{eff}}^{(N)}(x;R;A) = K_{\mathrm{eff}}^{(\infty)}(x;R;A) - \frac{\mathcal{G}_A(x;R)}{(N-d)\mathcal{D}_A(x;R)}. \tag{93}$$

Equivalently, the non-negative deficit is

$$\Delta K^{(N)}(x;R;A) \equiv K_{\mathrm{eff}}^{(\infty)}(x;R;A) - K_{\mathrm{eff}}^{(N)}(x;R;A) = \frac{\mathcal{G}_A(x;R)}{(N-d)\mathcal{D}_A(x;R)} \geqslant 0. \tag{94}$$

For $A = \mathbf{e}\mathbf{e}^\top$, define

$$\mathcal{D}_{\mathbf{e}}(x;R) \equiv \int_{B_R(x)} \langle y-x, \mathbf{e}\rangle^2 \mu_x^{(R)}(\mathrm{d}y), \tag{95}$$

$$\mathcal{G}_{\mathbf{e}}(x;R) \equiv \int_{B_R(x)} \langle y-x, \mathbf{e}\rangle^2 [\mathbf{e}\cdot \nabla \ln \rho_R(y)]^2 \mu_x^{(R)}(\mathrm{d}y). \tag{96}$$

The directional deficit is

$$\Delta K^{(N)}(x;R;\mathbf{e}) = \frac{\mathcal{G}_{\mathbf{e}}(x;R)}{(N-d)\mathcal{D}_{\mathbf{e}}(x;R)}. \tag{97}$$

It measures finite-window alignment between the probe direction and the log-density score, not density amplitude alone; it can remain nonzero even when the score vanishes at the window center.

D. Characteristics of each structure

In the local quadratic approximation, the curvature matrix is written as

$$H_R(x) = \mathrm{diag}(\lambda_1, \lambda_2, \lambda_3), \quad \lambda_1 \leqslant \lambda_2 \leqslant \lambda_3. \tag{98}$$

Introduce a common threshold $\lambda_{\mathrm{th}} \geqslant 0$ and define

$$N_+(x;R) \equiv \sum_{a=1}^{3} \mathbf{1}[\lambda_a(x;R) > \lambda_{\mathrm{th}}]. \tag{99}$$

The four classes are

$$\begin{aligned} N_+ &= 0: \quad \text{void},\\ N_+ &= 1: \quad \text{sheet},\\ N_+ &= 2: \quad \text{filament},\\ N_+ &= 3: \quad \text{halo or cluster}. \end{aligned} \tag{100}$$

The sign classification has $\lambda_{\rm th} = 0$; for a positive threshold, schematic zeros denote directions that do not exceed the threshold rather than exact zero eigenvalues.

Here, the counted eigenvalues belong specifically to H_R; the labels do not assert agreement with a density-Hessian or tidal-web classification.

Void. Since the density is minimal at the center, the local model gives $H_R(0) = -2\epsilon I$. For $\lambda_{\rm th} = 0$, all principal curvatures are negative and $N_+ = 0$. A nearly spherical interior has weak anisotropy and score deficit, whereas a steep boundary can increase $\Delta K^{(N)}$.

Sheet. For Eq. (73), $H_R = \mathrm{diag}(0, 0, 2/\sigma_\perp^2)$. Hence, $N_+ = 1$, with strong anisotropy and a deficit concentrated along the sheet normal.

Filament. For anisotropic structures,

$$H_R = \mathrm{diag}\left(0, \frac{2}{\sigma_\perp^2}, \frac{2}{\sigma_\perp^2}\right), \tag{101}$$

and the curvature differs significantly depending on direction. Along the filament direction it is nearly zero, whereas in the perpendicular directions, it takes finite positive values. Thus, $N_+ = 2$; both curvature anisotropy and the score deficit are primarily transverse.

Halo. For a centrally concentrated structure,

$$H_R = \frac{2}{\sigma^2} I, \tag{102}$$

and all eigenvalues are positive and large. Therefore, the curvature is strong in all directions and isotropic. Thus, $N_+ = 3$. For a spherical finite window, $\Delta K^{(N)}$ is approximately isotropic; it vanishes at the stationary center only in the strict pointwise limit.

The eigenvalue count determines the primary morphology, while $\Delta K^{(N)}$ and $\mathcal{R}^{(N)}$ quantify finite-window profile steepness.

E. Construction of classification indicators

Let $\{\mathbf{e}_a\}_{a=1}^m$ sample the unit sphere and satisfy

$$\frac{1}{m}\sum_{a=1}^{m} \mathbf{e}_a \mathbf{e}_a^\top = \frac{1}{d} I. \tag{103}$$

For this set of directions,

$$\overline{K}^{(\infty)}(x;R) = \frac{1}{m}\sum_{a=1}^{m} K_{\rm eff}^{(\infty)}(x;R;\mathbf{e}_a), \tag{104}$$

$$\overline{K}^{(N)}(x;R) = \frac{1}{m}\sum_{a=1}^{m} K_{\rm eff}^{(N)}(x;R;\mathbf{e}_a) \tag{105}$$

are defined as the mean curvatures. Furthermore, anisotropy is defined by

$$A^{(\infty)}(x;R) = \max_a K_{\rm eff}^{(\infty)}(x;R;\mathbf{e}_a) - \min_a K_{\rm eff}^{(\infty)}(x;R;\mathbf{e}_a), \tag{106}$$

$$A^{(N)}(x;R) = \max_a K_{\rm eff}^{(N)}(x;R;\mathbf{e}_a) - \min_a K_{\rm eff}^{(N)}(x;R;\mathbf{e}_a). \tag{107}$$

The direction-averaged deficit is

$$\overline{\Delta K}^{(N)}(x;R) \equiv \frac{1}{m}\sum_{a=1}^{m} \Delta K^{(N)}(x;R;\mathbf{e}_a) = \overline{K}^{(\infty)}(x;R) - \overline{K}^{(N)}(x;R) \geqslant 0. \tag{108}$$

Its anisotropy is

$$A_\Delta^{(N)}(x;R) \equiv \max_a \Delta K^{(N)}(x;R;\mathbf{e}_a) - \min_a \Delta K^{(N)}(x;R;\mathbf{e}_a). \tag{109}$$

When $|\overline{K}^{(\infty)}|$ is bounded away from zero, define

$$\overline{\mathcal{R}}^{(N)}(x;R) \equiv \frac{\overline{\Delta K}^{(N)}(x;R)}{|\overline{K}^{(\infty)}(x;R)|}, \quad |\overline{K}^{(\infty)}(x;R)| > K_{\min}. \tag{110}$$

This ratio of averaged observables need not equal the average of the directional ratios.

The local limits are

$$\overline{K}^{(\infty)}(x;R) = \frac{1}{d}\mathrm{Tr}H_R(x) + O(R), \tag{111}$$

$$A^{(\infty)}(x;R) = \lambda_{\max}(x;R) - \lambda_{\min}(x;R) + O(R), \tag{112}$$

where the second relation assumes extrema over the full unit sphere; a finite direction set approximates the eigenvalue spread.

Thus, N_+ sets the primary class, $\overline{K}^{(\infty)}$ records the mean expansive or contractive signature, $A^{(\infty)}$ measures directional anisotropy, and $\overline{\Delta K}^{(N)}$ with $A_\Delta^{(N)}$ characterizes the magnitude and directionality of the finite-window score correction.

When directional curvature is evaluated by the direction-dependent deformation

$$T_{\varepsilon,\mathbf{e}}(y) = y + \varepsilon\langle y - x, \mathbf{e}\rangle\mathbf{e}, \tag{113}$$

sheets and filaments produce strongly directional responses, while approximately spherical voids and halos are nearly isotropic. The finite-transport response converges to the Hessian signature locally but retains finite-window potential information.

A full-rank direction set defines the finite-scale tensor representation in Sec. VI, but algebraic completeness does not imply isotropic angular quadrature. In particular, the minimal six-direction design below does not satisfy Eq. (103); its raw directional mean is not a uniform angular mean. At finite radius, the reconstructed tensor's eigenvalues and classes inherit the design dependence. The total entropy response may be retained as a deformation diagnostic, but its sign is not used alone for morphology because it includes the non-negative kinematic term.

VI. OPERATIONAL ESTIMATION AND CONTROLLED VALIDATION

This section translates the density-induced curvature into an operational estimator for gridded fields, particle realizations, and reconstructed tracers. We distinguish the density-smoothing scale ℓ from the probe radius R; the earlier one-scale notation is recovered by setting $\ell = R$.

A. From density contrast to local-probe measures

Let $\widehat{\delta}_\ell(y)$ be a density-contrast field reconstructed or smoothed at scale ℓ, with positive density

$$\widehat{\rho}_\ell(y) = \bar{\rho}[1 + \widehat{\delta}_\ell(y)] > 0. \tag{114}$$

Here, $\bar{\rho}$ is the mean matter density for a simulation; for a tracer catalog, selection correction and positive reconstruction are understood. Any regularization or density floor must be documented and varied in stability tests.

On a finite or periodic domain M, define

$$\widehat{Z}_\ell \equiv \int_M \widehat{\rho}_\ell(y)\,\mathrm{d}y, \tag{115}$$

$$\widehat{p}_\ell(y) \equiv \frac{\widehat{\rho}_\ell(y)}{\widehat{Z}_\ell}, \tag{116}$$

$$\widehat{V}_\ell(y) \equiv -\ln \widehat{p}_\ell(y). \tag{117}$$

Because additive constants cancel from the midpoint deficit, $-\ln\widehat{\rho}_\ell$ may be used equivalently when only the potential response is required.

For quadrature nodes $\{y_i\}_{i=1}^{n_x}$, cell weights ΔV_i, and a non-negative window W_R, define

$$\omega_i(x;R,\ell)\equiv\frac{W_R(y_i-x)\widehat{p}_\ell(y_i)\Delta V_i}{\sum_{j=1}^{n_x}W_R(y_j-x)\widehat{p}_\ell(y_j)\Delta V_j}. \tag{118}$$

The discrete local probe is

$$\widehat{\mu}_x^{(R;\ell)}\equiv\sum_{i=1}^{n_x}\omega_i(x;R,\ell)\delta_{y_i}. \tag{119}$$

A top-hat reproduces the theoretical ball restriction, while a smooth compact or Gaussian window reduces discontinuities as nodes enter or leave the probe.

Since local normalization removes the window mass from the transport geometry, retain the complementary amplitude observables

$$\widehat{M}_x^{(R;\ell)}\equiv\sum_{i=1}^{n_x}W_R(y_i-x)\widehat{\rho}_\ell(y_i)\Delta V_i, \tag{120}$$

$$\widehat{\overline{\delta}}_x^{(R;\ell)}\equiv\frac{\sum_i W_R(y_i-x)\widehat{\delta}_\ell(y_i)\Delta V_i}{\sum_i W_R(y_i-x)\Delta V_i}. \tag{121}$$

B. Finite-transport curvature estimator

Let $A=A^\top$ satisfy $I+\varepsilon A\succ 0$, set $z_i=y_i-x$, and define

$$y_i(t)\equiv x+(I+t\varepsilon A)z_i,\quad 0\leqslant t\leqslant 1. \tag{122}$$

The end point measure transports the same weights to $y_i(1)$.

The discrete transport normalization is

$$\widehat{\mathcal{D}}_A(x;R,\ell)\equiv\sum_{i=1}^{n_x}\omega_i(x;R,\ell)|Az_i|^2. \tag{123}$$

For the prescribed optimal affine pairing, the quadrature transport cost is

$$\widehat{W}_{2,A}^2(x;R,\ell,\varepsilon)=\varepsilon^2\widehat{\mathcal{D}}_A(x;R,\ell). \tag{124}$$

For an empirical measure without a certified pairing, this cost must be checked against, or replaced by, a numerical optimal-transport solution.

Define the potential midpoint deficit by

$$\widehat{\mathcal{Q}}_A(x;R,\ell,\varepsilon)\equiv\sum_{i=1}^{n_x}\omega_i(x;R,\ell)\left[\frac{\widehat{V}_\ell(y_i)+\widehat{V}_\ell(y_i(1))}{2}-\widehat{V}_\ell(y_i(1/2))\right]. \tag{125}$$

Interpolating $\widehat{V}_\ell$ at the deformed positions gives

$$\widehat{K}_{\rho,\varepsilon}^{(\infty)}(x;R,\ell;A)\equiv\frac{8\,\widehat{\mathcal{Q}}_A(x;R,\ell,\varepsilon)}{\widehat{W}_{2,A}^2(x;R,\ell,\varepsilon)}=\frac{8\,\widehat{\mathcal{Q}}_A(x;R,\ell,\varepsilon)}{\varepsilon^2\widehat{\mathcal{D}}_A(x;R,\ell)}. \tag{126}$$

For completeness, the affine Jacobian gives the kinematic response

$$\widehat{\mathcal{C}}_{\mathrm{kin}}^{(\infty)}(x;R,\ell;A)=\frac{8}{\varepsilon^2\widehat{\mathcal{D}}_A(x;R,\ell)}\left[\ln\det\left(I+\frac{\varepsilon}{2}A\right)-\frac{1}{2}\ln\det(I+\varepsilon A)\right], \tag{127}$$

and hence

$$\widehat{\mathcal{C}}^{(\infty)}_{\rm ent} = \widehat{\mathcal{C}}^{(\infty)}_{\rm kin} + \widehat{K}^{(\infty)}_{\rho,\varepsilon}. \tag{128}$$

A symmetric estimator reduces the leading displacement bias. With $y_i^{\pm} = x + (I \pm \varepsilon A/2)z_i$, define

$$\widehat{K}^{(\infty)}_{\rho,\varepsilon,\rm sym}(x;R,\ell;A) \equiv \frac{8}{\varepsilon^2 \widehat{\mathcal{D}}_A(x;R,\ell)} \sum_i \omega_i \left[\frac{\widehat{V}_\ell(y_i^-) + \widehat{V}_\ell(y_i^+)}{2} - \widehat{V}_\ell(y_i) \right]. \tag{129}$$

For a smooth potential, the one-sided and symmetric truncation errors are $O(\varepsilon)$ and $O(\varepsilon^2)$, respectively.

A direct short-geodesic reference, evaluated from the same smooth interpolant, is

$$\widehat{K}^{(\infty)}_{\rm eff}(x;R,\ell;A) \equiv \lim_{\varepsilon\to 0} \widehat{K}^{(\infty)}_{\rho,\varepsilon}(x;R,\ell;A) = \frac{\sum_{i=1}^{n_x} \omega_i(x;R,\ell)(Az_i)^\top [\nabla\nabla \widehat{V}_\ell](y_i)(Az_i)}{\widehat{\mathcal{D}}_A(x;R,\ell)}. \tag{130}$$

Unlike the finite-ε estimator, this reference uses no midpoint difference. Equation (130) assumes differentiation and transport evaluation of the same smooth potential. In the gridded mock, the direct reference uses spectral derivatives of the reconstructed potential, whereas off-grid transport values use cubic interpolation. Their finite-resolution difference is included in the measured internal error; exact equality at fixed grid spacing is not assumed. The implementation is specified in Appendix E 1.

Define the score-alignment numerator

$$\widehat{\mathcal{G}}_A(x;R,\ell) \equiv \sum_{i=1}^{n_x} \omega_i(x;R,\ell)[Az_i \cdot \widehat{\nabla \ln \rho_\ell}(y_i)]^2. \tag{131}$$

For $N > d$, the finite-dimensional estimator is

$$\widehat{\Delta K}^{(N)}(x;R,\ell;A) \equiv \frac{\widehat{\mathcal{G}}_A(x;R,\ell)}{(N-d)\widehat{\mathcal{D}}_A(x;R,\ell)} \geqslant 0, \tag{132}$$

$$\widehat{K}^{(N)}_{\rm eff}(x;R,\ell;A) \equiv \widehat{K}^{(\infty)}_{\rm eff}(x;R,\ell;A) - \widehat{\Delta K}^{(N)}(x;R,\ell;A). \tag{133}$$

The Hessian and score must be derived consistently from the same reconstructed density and smoothing scale, with the differentiation and interpolation operations stated explicitly. In the periodic tests, both are obtained by spectral differentiation. The Bakry-Émery parameter N is unrelated to the number of particles in the probe.

C. Reconstruction of the finite-scale response tensor

Define the pointwise log-density Hessian by

$$H_{\rm point}(x;\ell) \equiv \nabla\nabla \widehat{V}_\ell(x) = -\nabla\nabla \ln \widehat{\rho}_\ell(x). \tag{134}$$

The same differentiation and interpolation operator must be used at the probe center and quadrature nodes.

For unit directions $\mathcal{E} = \{\mathbf{e}_a\}_{a=1}^m$, set $A_a = \mathbf{e}_a\mathbf{e}_a^\top$. The finite-transport and direct finite-window responses are

$$\widehat{k}_{{\rm transport},a}(x;R,\ell,\varepsilon) \equiv \widehat{K}^{(\infty)}_{\rho,\varepsilon}(x;R,\ell;A_a),$$

$$k_{{\rm window},a}(x;R,\ell) \equiv \widehat{K}^{(\infty)}_{\rm eff}(x;R,\ell;A_a) = \frac{1}{\widehat{\mathcal{D}}_{A_a}(x;R,\ell)} \sum_{i=1}^{n_x} \omega_i(x;R,\ell)\left(\mathbf{e}_a^\top z_i\right)^2 \mathbf{e}_a^\top H_{\rm point}(y_i;\ell)\mathbf{e}_a. \tag{135}$$

The second quantity is the direct reference for the same weighted probe and must have positive denominator.

Using common directional weights w_a, reconstruct

$$\widehat{H}^{(\mathcal{E})}_{\text{transport}}(x; R, \ell, \varepsilon) \equiv \underset{H=H^\top}{\arg\min} \sum_{a=1}^{m} w_a \left[\widehat{k}_{\text{transport},a} - \mathbf{e}_a^\top H \mathbf{e}_a\right]^2,$$

$$H^{(\mathcal{E})}_{\text{window}}(x; R, \ell) \equiv \underset{H=H^\top}{\arg\min} \sum_{a=1}^{m} w_a \left[k_{\text{window},a} - \mathbf{e}_a^\top H \mathbf{e}_a\right]^2. \tag{136}$$

The directional design must have full rank on symmetric tensors. Because the direction enters both the spatial weighting and the normalization in Eq. (135), the complete finite-window response need not be a single quadratic form. Consequently, $H^{(\mathcal{E})}_{\text{window}}$, its eigenvalues, and its threshold classifications depend in general on the specified directions and weights. This is distinct from a joint change of spatial coordinates, under which the same tensor transforms covariantly and retains its eigenvalues. Comparing transport and direct tensors for the same design establishes estimator fidelity, not invariance under replacing that design. The fixed-field tests in Appendix E 6 quantify the resulting eigenvalue and classification sensitivity and separately check angular convergence and estimator error.

In three dimensions, a minimal algebraically determined design is

$$\mathbf{e}_1 = (1, 0, 0)^\top, \quad \mathbf{e}_2 = (0, 1, 0)^\top, \quad \mathbf{e}_3 = (0, 0, 1)^\top, \tag{137}$$

$$\mathbf{e}_{12} = \frac{\mathbf{e}_1 + \mathbf{e}_2}{\sqrt{2}}, \quad \mathbf{e}_{13} = \frac{\mathbf{e}_1 + \mathbf{e}_3}{\sqrt{2}}, \quad \mathbf{e}_{23} = \frac{\mathbf{e}_2 + \mathbf{e}_3}{\sqrt{2}}. \tag{138}$$

For either response set $q = \{q_i, q_{ij}\}$, define

$$[\mathcal{R}_6(q)]_{ii} \equiv q_i, \tag{139}$$

$$[\mathcal{R}_6(q)]_{ij} \equiv q_{ij} - \frac{q_i + q_j}{2}, \quad i < j. \tag{140}$$

Thus, $\widehat{H}^{(6)}_{\text{transport}} = \mathcal{R}_6(\widehat{k}_{\text{transport}})$ and $H^{(6)}_{\text{window}} = \mathcal{R}_6(k_{\text{window}})$. The reconstruction is exact for the six supplied values and represents all directions only when the response is a quadratic form.

At fixed ℓ, the two controlled limits are

$$\lim_{\varepsilon \to 0} \widehat{H}^{(\mathcal{E})}_{\text{transport}}(x; R, \ell, \varepsilon) = H^{(\mathcal{E})}_{\text{window}}(x; R, \ell),$$

$$\lim_{R \to 0} H^{(\mathcal{E})}_{\text{window}}(x; R, \ell) = H_{\text{point}}(x; \ell). \tag{141}$$

They require smooth interpolation, nondegenerate normalization, and adequate resolution; a constant quadratic Hessian is recovered exactly in the noiseless construction.

The local limit is specifically the log-density Hessian in Eq. (134), with the qualifications in Eq. (78). At finite spatial resolution, the derivative and interpolation approximations must also be controlled, as described above.

For any of these tensors, write the ordered eigenvalues as $\lambda_1[X] \leqslant \lambda_2[X] \leqslant \lambda_3[X]$ and define

$$C_X(x) \equiv \sum_{a=1}^{3} \mathbf{1}[\lambda_a[X(x)] > \lambda_{\text{th}}]. \tag{142}$$

The values 0, 1, 2, 3 denote void, sheet, filament, and halo or cluster, using one common threshold for all comparisons.

The reconstructed transport mean curvature and anisotropy are

$$\widehat{\overline{K}}_{\text{tr}} \equiv \frac{1}{d} \text{Tr} \widehat{H}^{(\mathcal{E})}_{\text{transport}}, \tag{143}$$

$$\widehat{A}_{\text{tr}} \equiv \lambda_{\max}\left[\widehat{H}^{(\mathcal{E})}_{\text{transport}}\right] - \lambda_{\min}\left[\widehat{H}^{(\mathcal{E})}_{\text{transport}}\right]. \tag{144}$$

Replacing the tensor gives the corresponding finite-window and pointwise invariants; the two limits in Eq. (141) must remain distinct.

D. Simulation benchmark and quantitative diagnostics

A benchmark must hold fixed the reconstructed density, ℓ, the probe window, direction set, interpolation and differentiation operators, classification threshold, and evaluation sample. It then compares the three objects defined above: the transport reconstruction, the direct finite-window response tensor, and the pointwise Hessian.

Layer I: Estimator fidelity. Within one reconstructed field, compare $\widehat{k}_{\mathrm{transport},a}$ with $k_{\mathrm{window},a}$ and $\widehat{H}^{(\mathcal{E})}_{\mathrm{transport}}$ with $H^{(\mathcal{E})}_{\mathrm{window}}$, using the same reconstruction operator. The residual measures finite-ε, interpolation, quadrature, conditioning, and tensor-fit errors. For an overdetermined design, the raw directional responses should also be compared because a tensor fit can conceal nonquadratic residuals.

Layer II: Observable departure. The finite-window departure from the pointwise Hessian is

$$\mathcal{D}^{(\mathcal{E})}_{\text{window-point}}(x;R,\ell) \equiv H^{(\mathcal{E})}_{\mathrm{window}}(x;R,\ell) - H_{\mathrm{point}}(x;\ell). \tag{145}$$

This is not estimator error. Directionally, it compares $k_{\mathrm{window},a}$ with $k_{\mathrm{point},a} = \mathbf{e}_a^{\top} H_{\mathrm{point}} \mathbf{e}_a$. A stable nonzero departure after Layer-I convergence is finite-window information and may be examined with the local amplitude and score-alignment observables.

Replacing $\mathcal{E}$ at fixed field and scales defines another observable comparison, not a Layer-I error. Its eigenvalue and class changes must be interpreted using the margin of the relevant finite-window tensor, not the pointwise-Hessian margin.

Layer III: Sampling robustness. For sampled field s and a fixed full-sampling reconstruction, the internal comparison is $\widehat{H}^{(s)}_{\mathrm{transport}}$ versus $H^{(s)}_{\mathrm{window}}$; field degradation is $H^{(s)}_{\mathrm{window}}$ versus $H^{(\mathrm{full})}_{\mathrm{window}}$; the end-to-end comparison is $\widehat{H}^{(s)}_{\mathrm{transport}}$ versus $H^{(\mathrm{full})}_{\mathrm{window}}$; and pointwise-Hessian sampling compares $H^{(s)}_{\mathrm{point}}$ with $H^{(\mathrm{full})}_{\mathrm{point}}$. Thus, each observable uses its own full-sampling reference, which is numerical rather than continuum truth.

Let $N_{\mathrm{eval}} = |\mathcal{I}|$. For directional families X and Y, define

$$r_{\mathrm{dir}}(X,Y;\mathcal{I},\mathcal{E}) \equiv \operatorname{Corr}_{\substack{i\in\mathcal{I}\\ a=1,\dots,m}} [k_{X,a}(x_i), k_{Y,a}(x_i)]. \tag{146}$$

Layer I uses transport and window responses; Layer II uses window and pointwise responses.

For tensor observables, with Y as reference, define

$$E(X \to Y;\mathcal{I}) \equiv \left[\frac{\sum_{i\in\mathcal{I}} \|X(x_i) - Y(x_i)\|_{\mathrm{F}}^2}{\sum_{i\in\mathcal{I}} \|Y(x_i)\|_{\mathrm{F}}^2} \right]^{1/2}. \tag{147}$$

The arrow points to the reference. We use the shorthand $E_{\mathrm{transport}\to\mathrm{window}}$ and $E_{\mathrm{window}\to\mathrm{point}}$ for Layers I and II; the sampled-field specializations are defined below.

The median regularized per-point discrepancies reported for Stage 3 are distinct from this aggregate norm; their normalization is specified in Appendix E 1. The additional design comparisons use Eq. (147).

The class-agreement fraction is

$$f_{\mathrm{class}}(X,Y;\mathcal{I}) \equiv \frac{1}{N_{\mathrm{eval}}} \sum_{i\in\mathcal{I}} \mathbf{1}[C_X(x_i) = C_Y(x_i)]. \tag{148}$$

The compared pair should always be named explicitly. The jointly normalized transition matrix is

$$\mathsf{M}^{Y\to X}_{ab}(\mathcal{I}) \equiv \frac{1}{N_{\mathrm{eval}}} \sum_{i\in\mathcal{I}} \mathbf{1}[C_Y(x_i) = a]\mathbf{1}[C_X(x_i) = b], \quad a,b \in \{0,1,2,3\}. \tag{149}$$

Rows are the reference classes. The transport-window matrix assesses Layer I, whereas the pointwise-to-window matrix records Layer-II morphology transitions.

A convincing validation distinguishes three regimes. Controlled local convergence takes $\varepsilon \to 0$ and $R \to 0$ at fixed ℓ. A resolved finite-window response has $E_{\text{transport}\to\text{window}} \ll 1$ while $E_{\text{window}\to\text{point}}$ is stable and may be nonzero. Sampling degradation increases E_{field}, E_{total}, and E_{Hessian}, even if the internal estimator error remains small.

E. Finite samples, boundaries, and uncertainty

For a regularized field and certified convex affine probe, the estimator requires only local quadrature and interpolation. Arbitrary empirical couplings instead require a discrete or regularized optimal-transport solve.

The density must remain positive, the potential and derivatives must be adequately interpolated, and each probe direction must satisfy

$$\widehat{\mathcal{D}}_A(x; R, \ell) < D_{\min} \tag{150}$$

only at rejected locations. The directional design condition number should also be monitored, with the same accepted directions used for the transport and direct-reference tensors.

Periodic simulations permit wrapping. For finite surveys, one may restrict probes to the mask interior, renormalize truncated windows, or use matched random catalogs; the same boundary operator must be applied to transport, finite-window, and pointwise branches.

Bootstrap, Poisson, posterior, or jackknife realizations may be propagated through all three tensors and their classes. For $X \in \{\text{transport}, \text{window}, \text{point}\}$, define

$$\widehat{P}_X(C = c \mid x) \equiv \frac{1}{B} \sum_{b=1}^{B} \mathbf{1}\big[C_X^{(b)}(x) = c\big]. \tag{151}$$

The observable label separates uncertainty within one classifier from systematic transitions between finite-window and pointwise classifications. Operationally, one density reconstruction supplies the finite-transport estimator, the direct finite-window reference, and the pointwise-Hessian branch; repeating them for sampled and full-sampling fields preserves the three-layer comparison.

F. Controlled validation on analytic and nonquadratic mocks

Controlled three-dimensional potentials test exact quadratic recovery, finite-displacement convergence, grid refinement, transition stability, and rotational covariance. Stage 1 uses quadratic fields; Stage 2 uses a quartic perturbation to separate estimator convergence from finite-window observable departure; and Stage 2b tests grid and rotation stability. In Stages 2 and 2b, "exact finite-window reference" means the analytic six-direction reference, not a direction-independent tensor for an arbitrary field. The complete results are summarized in Table I.

The exact polynomial potentials, coefficient families, rotations, grid and displacement scans, and classification thresholds are given in Appendix E 1, including Eq. (E1) and Table II.

The tests recover quadratic curvature exactly to numerical precision, verify the expected one-sided and symmetric convergence orders, and identify a stable displacement-resolution regime. They also confirm that finite-window class transitions are observable departures rather than estimator failures. The cosmological test below replaces the analytic target with a direct numerical reference.

G. Cosmological mock and finite-window curvature

We next used a periodic three-dimensional Zel'dovich mock with $N_g = 128$, $N_p = 128^3$, $L_{\text{box}} = 200\,h^{-1}\,\text{Mpc}$, and $\Delta x = 1.5625\,h^{-1}\,\text{Mpc}$. The smoothing scale was $\ell = \sigma_G = 4\Delta x = 6.25\,h^{-1}\,\text{Mpc}$, the probe radii were $R/\Delta x = (8, 12, 16)$, and the displacement parameters were $\eta = (1, 1.5, 2)$.

TABLE I. Controlled validation of the finite-window transport estimator. Stage 2 convergence is measured against the exact analytic finite-window reference for the six-direction design, whereas the finite-window-observable row compares that reference with the central pointwise Hessian.

Stage	Validation target	Setting	Quantitative result	Interpretation
1	Exact quadratic response	Analytic potential	256/256 classes; maximum directional error 6.99×10^{-11}; maximum tensor relative error 8.64×10^{-11}	Finite-transport normalization and six-direction tensor reconstruction recover the exact quadratic curvature.
1	Exact quadratic response	Cubic-grid potential	256/256 classes; maximum directional error 9.41×10^{-10}; maximum tensor relative error 1.03×10^{-9}	The quadratic response remains accurate after cubic-grid interpolation, including rotated tensors and near-threshold eigenvalue configurations.
2	One-sided estimator convergence	Analytic quartic potential	$p_K = 1.030$ and $p_H = 1.030$; transport-window class agreement: 98.25%	The directional and tensor biases relative to the exact analytic finite-window reference follow the expected first-order asymptotic scaling.
2	Symmetric estimator convergence	Analytic quartic potential	$p_K = 2.000$ and $p_H = 2.000$; transport-window class agreement: 99.75%	Symmetrization removes the leading odd contribution and yields the expected second-order asymptotic scaling relative to the exact analytic finite-window reference.
2	Finite-window observable departure	$\alpha = 0.20$	$R/L_0 = 0.8$: window-point class agreement 0.875, sheet → filament; $R/L_0 = 1$: window-point class agreement 0.500, sheet → filament and filament → halo	Finite-window curvature contains scale-dependent geometric information beyond the central pointwise Hessian. These class transitions are properties of the exact analytic finite-window reference, not estimator failures.
2b	Grid and displacement calibration	$N_g = 65, 97, 129$	$N_g = 65$: $\varepsilon = 0.20$, $\eta = 1.71$–2.13, maximum tensor relative error 5.48×10^{-4}; $N_g = 97$: $\varepsilon = 0.10$, $\eta = 1.28$–1.60, maximum tensor relative error 1.94×10^{-4}; $N_g = 129$: $\varepsilon = 0.10$, $\eta = 1.71$–2.13, maximum tensor relative error 1.36×10^{-4}	The practical displacement-resolution regime is $\eta = O(1)$, approximately $1 \lesssim \eta \lesssim 2$. The interpolation floor decreases under grid refinement.
2b	Transition stability	$N_g \geqslant 97$	Stable class signatures: 6/6; robust transitioned rows: 12/12; all fine-grid transitions exceed ten threshold-margin units	The sheet-filament transition near $R/L_0 = 0.8$ and the sheet-filament and filament-halo transitions at $R/L_0 = 1$ are stable on the two finer grids.
2b	Rotation covariance	Axis-aligned vs covariantly rotated control	Maximum tensor residual 3.63×10^{-3} over all grids and 1.20×10^{-3} for $N_g \geqslant 97$	Rotating both the curvature tensor and the linear score yields the expected covariant response to the grid-accuracy level.

The initial spectrum, seed, density reconstruction, probe window, and derivative/interpolation settings are specified in Appendix E 1. The periodic mock uses $\lambda_{th} = 0$ and $N = 6$; N is the synthetic Bakry-Émery dimension, not a particle count.

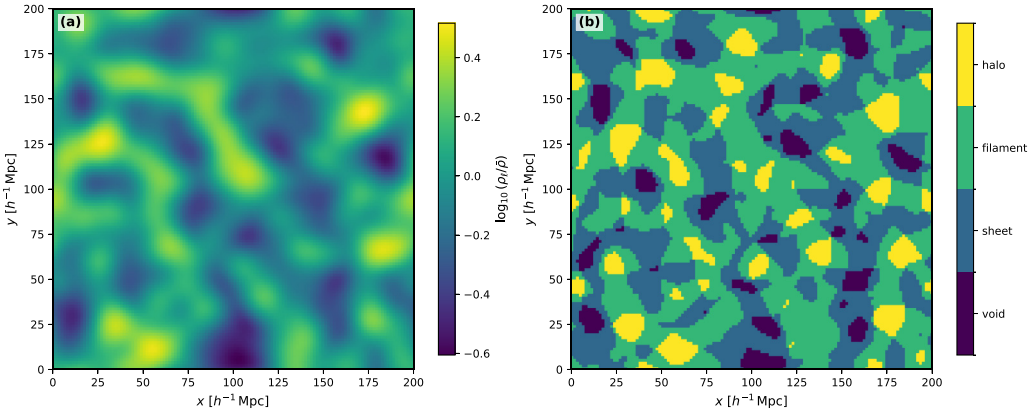


FIG. 2. Periodic Zel'dovich mock used for validation: a representative smoothed-density slice and its pointwise log-density-Hessian void, sheet, filament, and halo classification.

The evaluation set contained 96 pointwise-stable points, balanced among the four pointwise log-density-Hessian classes, and 32 additional pointwise-boundary points for threshold diagnostics. These names refer to selection by the pointwise-Hessian margin and do not imply stability of a finite-window or alternative-design classification. Figure 2 shows the smoothed field and pointwise log-density-Hessian classification.

Using the six-direction design, we evaluated $H_{\mathrm{point}}(\boldsymbol{x};\ell)$, $H^{(6)}_{\mathrm{window}}(\boldsymbol{x};R,\ell)$, and $\widehat{H}^{(6)}_{\mathrm{transport}}(\boldsymbol{x};R,\ell,\eta)$. Here, "direct" means evaluation from the reconstructed curvature field rather than a finite-ε midpoint difference; it does not denote continuum truth.

Transport-window agreement was uniformly high: The directional correlation exceeded 0.999 995, the stable-sample median tensor error was 0.19%–0.37%, and class agreement was approximately 97.9%–100%. The smallest error occurred at $\eta = 2$ (Fig. 3).

The direct finite-window tensor nevertheless differed from the pointwise Hessian. For $R/\Delta x = 8$, 12, and 16, class agreement was 0.6458, 0.2917, and 0.2083, while the median normalized tensor departure was 0.5795, 0.9322, and 1.1287. This radius-dependent difference belongs to Layer II.

Figure 4 shows the structured pointwise-to-window transitions and the complementary finite-dimensional score-alignment deficit.

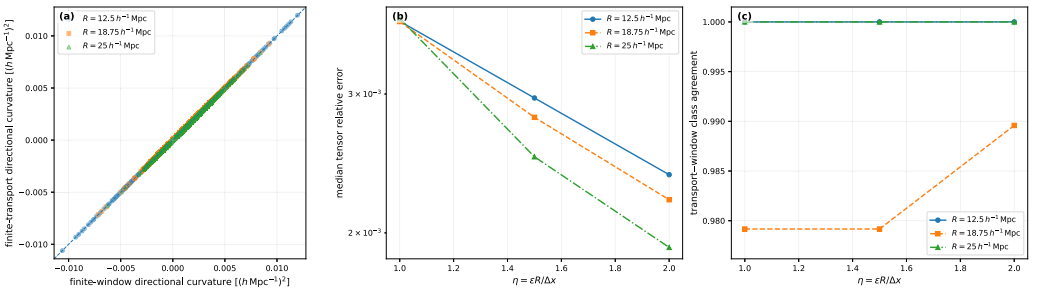


FIG. 3. Estimator fidelity on the cosmological mock. (a) Directional transport vs direct finite-window response. (b) Tensor error vs η. (c) Transport-window class agreement. All comparisons use the same six-direction design; they do not test invariance under changing that design.

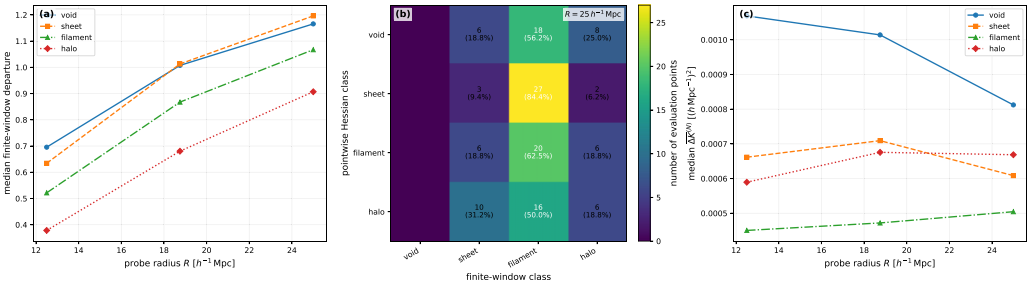


FIG. 4. Finite-window information beyond the pointwise Hessian. (a) Window-point tensor departure. (b) Pointwise-to-window class transitions. (c) Finite-dimensional score-alignment deficit. The tensor and class comparisons refer to the specified six-direction construction.

The central result is

$$\widehat{H}^{(6)}_{\text{transport}}(\boldsymbol{x}; R, \ell, \eta) \simeq H^{(6)}_{\text{window}}(\boldsymbol{x}; R, \ell),$$
$$H^{(6)}_{\text{window}}(\boldsymbol{x}; R, \ell) \neq H_{\text{point}}(\boldsymbol{x}; \ell). \tag{152}$$

Thus, the transport construction estimates the six-direction finite-window response, not the pointwise Hessian at finite R; the latter is recovered only in the resolved local limit of Eq. (141).

This high estimator fidelity does not make the finite-window classification independent of its directional design. The fixed-field audit in Appendix E 6 finds an aggregate tensor difference of 0.2141 and class agreement of 88/96 between the 96-axis and original six-direction direct fits at $R/\Delta x = 8$, with larger differences at the two larger radii. In contrast, transport errors relative to each tested design's own direct reference remain 0.2707%–0.3238% at $R/\Delta x = 8$ and $\eta = 2$. These results separate observable-design dependence from numerical estimator error.

H. Finite-sampling robustness

We thinned the full-particle catalog with nested Bernoulli samples $f_{\text{samp}} = (1, 1/2, 1/4, 1/8, 1/16)$. Each nonunit fraction used eight replicates; ℓ, the evaluation points, and the full catalog were fixed. Rebuilding the full catalog reproduced all Stage 3 tensors and classes to machine precision before thinning.

All tensors below use $\mathcal{E}_6$, with common spatial and scale arguments suppressed. Superscripts (s) and (full) denote sampled and full-sampling reconstructions. Define

$$\begin{aligned} D_{\text{est}} &= \widehat{H}^{(s)}_{\text{transport}} - H^{(s)}_{\text{window}}, \\ D_{\text{field}} &= H^{(s)}_{\text{window}} - H^{(\text{full})}_{\text{window}}, \\ D_{\text{Hess}} &= H^{(s)}_{\text{point}} - H^{(\text{full})}_{\text{point}}, \\ D_{\text{total}} &= \widehat{H}^{(s)}_{\text{transport}} - H^{(\text{full})}_{\text{window}} \\ &= D_{\text{est}} + D_{\text{field}}. \end{aligned} \tag{153}$$

For a difference D_i and reference $H_{\text{ref},i}$, use

$$E(D; H_{\text{ref}}) \equiv E(H_{\text{ref}} + D \to H_{\text{ref}}; \mathcal{I}) = \left[\frac{\sum_{i\in\mathcal{I}} \|D_i\|_{\text{F}}^2}{\sum_{i\in\mathcal{I}} \|H_{\text{ref},i}\|_{\text{F}}^2} \right]^{1/2}. \tag{154}$$

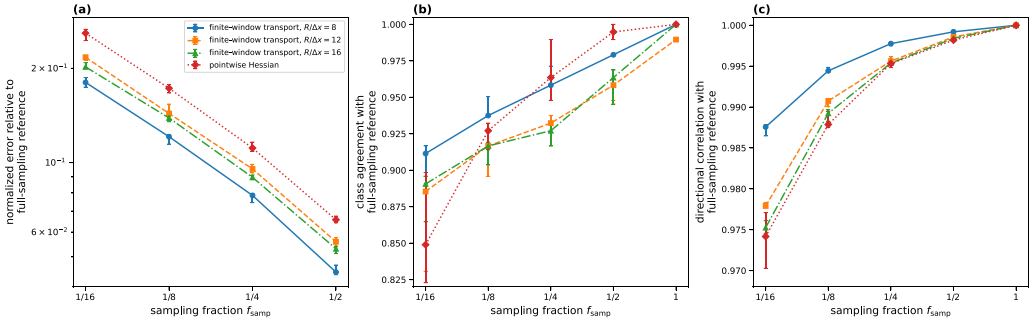


FIG. 5. Finite-sampling robustness for the stable sample and $\eta = 2$. Medians and interquartile ranges summarize eight replicates. (a) End-to-end transport and pointwise-Hessian errors, each relative to its own full-sampling reference. (b) Sampled-full class agreement. (c) Sampled-full directional correlation.

Accordingly,

$$
\begin{aligned}
E_{\text{internal}} &\equiv E\left(D_{\text{est}}; H^{(s)}_{\text{window}}\right), \\
E_{\text{field}} &\equiv E\left(D_{\text{field}}; H^{(\text{full})}_{\text{window}}\right), \\
E_{\text{Hessian}} &\equiv E\left(D_{\text{Hess}}; H^{(\text{full})}_{\text{point}}\right), \\
E_{\text{total}} &\equiv E\left(D_{\text{total}}; H^{(\text{full})}_{\text{window}}\right).
\end{aligned}
$$

The tensor identity is exact, although these normalized norms need not add linearly.

At $f_{\text{samp}} = 1/16$, the internal errors for $R/\Delta x = (8, 12, 16)$ were 2.78×10^{-3}, 2.66×10^{-3}, and 2.30×10^{-3}. Directional correlation remained above 0.999 995 and transport-window class agreement was unity, showing that the estimator floor did not amplify sampling noise.

Figure 5 summarizes the sampling dependence of the end-to-end errors, sampled-full classification agreement, and sampled-full directional correlations. The end-to-end error was instead dominated by field reconstruction: At the same fraction, E_{field} and E_{total} were both approximately (0.180, 0.217, 0.202) across the three radii (Fig. 6).

For every nonunit fraction and radius, $E_{\text{total}} < E_{\text{Hessian}}$. At $f_{\text{samp}} = 1/16$, their ratios were 0.700, 0.834, and 0.792 for $R/\Delta x = 8$, 12, and 16; the fiducial radius therefore gave an approximately 30% reduction. At that radius, the correlations were 0.988 and 0.974, and class agreements were 0.911 and 0.849, for the finite-window and pointwise observables, respectively.

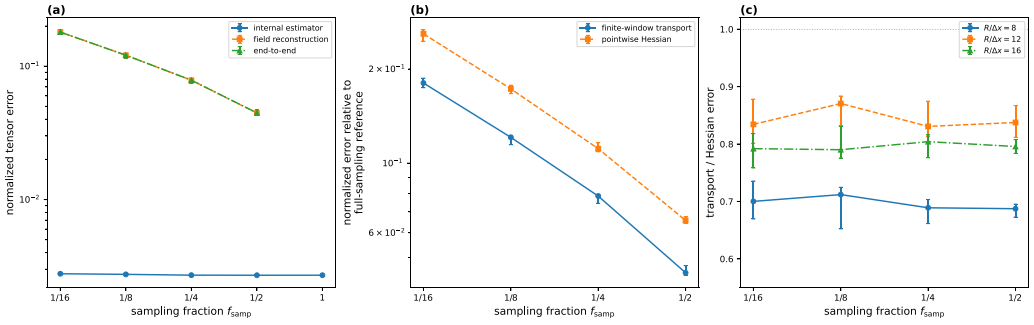


FIG. 6. Error decomposition at $R/\Delta x = 8$ and $\eta = 2$. (a) Internal, field-reconstruction, and total errors. (b) E_{total} and E_{Hessian}. (c) Their ratio for all probe radii.

Combining nonunit fractions selected $R/\Delta x = 8$ as the fiducial radius. Increasing η from 1.5 to 2 lowered the internal error by 15%–22% but changed the total error by less than 0.3%. The empirical scaling $E \propto f_{\rm samp}^{-\alpha}$ with $\alpha \simeq 0.65$ and median coefficients of determination near 0.99 characterizes this mock and reconstruction pipeline, not a universal Poisson law.

The Stage 4 result is

$$E_{\rm total} \simeq E_{\rm field} \gg E_{\rm internal}, \quad E_{\rm total} < E_{\rm Hessian}. \tag{155}$$

The second relation compares two distinct observables with their own full-sampling references. Supplementary displacement, boundary, morphology, finite-dimensional, and spectral-margin tests are presented in Appendix E.

The sampling-error reduction applies to the original six-direction observable under this reconstruction and thinning protocol. The multidirection audit does not repeat the thinning experiment and does not establish the same advantage for every directional design.

I. Practical parameter choices and scope

For this mock, the adopted operating point is

$$\begin{aligned} \eta_{\rm fid} &= 2, \\ R_{\rm fid} &= 8\Delta x = 12.5\, h^{-1}\,{\rm Mpc}, \\ \ell_{\rm fid} &\equiv \sigma_{\rm G} = 4\Delta x = 6.25\, h^{-1}\,{\rm Mpc}. \end{aligned} \tag{156}$$

The value $\eta_{\rm fid} = 2$ optimizes numerical accuracy and has no dynamical interpretation. The selected radius minimizes the combined normalized errors over the tested grid but is conditional on ℓ, resolution, realization, and radius range.

The full catalog is a numerical reference, not continuum truth. The stable sample supplies the primary statistics, while the boundary sample tests threshold sensitivity and was selected using the pointwise-Hessian margin. No density-floor activation occurred, and the Bakry-Émery dimension N was held independent of particle number.

The test assumes a periodic box, uniform thinning, fixed Gaussian smoothing, and common evaluation points. Survey masks, variable completeness, radial selection, redshift-space distortions, anisotropic uncertainties, and finite-volume boundaries require a survey-specific forward operator and renewed full-sampling calibration.

Within this scope and for the specified six-direction design, the transport estimator accurately recovers its direct finite-window response; end-to-end uncertainty is dominated by density reconstruction; and that finite-window observable is more robust to thinning than the pointwise log-density Hessian under their respective full-sampling comparisons. Directional-design sensitivity is quantified separately in Appendix E 6.

VII. DISCUSSION AND CONCLUSIONS

A. Local limit and Hessian structure

The framework separates the total entropy-convexity response from the density-induced Ricci-curvature contribution. The exact finite-transport and short-geodesic decompositions are given in Eqs. (30) and (62): The kinematic term depends on the chosen deformation, while the curvature term probes the fixed weighted background geometry.

For a flat Euclidean or periodic background, the cosmological density defines the Bakry-Émery tensor $H_R = -\nabla^2 \ln \rho_R$ through Eq. (39). A rank-one probe averages $\mathbf{e}^\top H_R \mathbf{e}$ over the complete local mass profile [Eq. (90)] and approaches the pointwise Rayleigh quotient as $R \to 0$ [Eq. (91)]. The eigenvalue-threshold classifier of the smoothed log-density Hessian is therefore recovered in the controlled small-deformation and local-window limit. Equation (78) distinguishes it from density-Hessian classifiers, and no general equivalence with gravitational-tidal classifiers is implied. At

finite R, the response is a distinct observable whose tensor representation and classification also depend on the specified directional design.

The corresponding mean curvature and anisotropy approach the trace and eigenvalue spread of H_R through Eqs. (111) and (112). Finite synthetic dimension further subtracts the non-negative deficit defined in Eq. (71). This term measures directional alignment with the log-density score over the probe and supplements, rather than replaces, the Hessian eigenvalue classification.

B. Scale-dependent curvature and multiscale structure

The construction contains two independent scales: the density-smoothing scale ℓ, which fixes the weighted background geometry, and the probe radius R, which selects the region over which that geometry is sampled. The formal one-parameter notation corresponds to the diagonal choice $\ell = R$. In the cosmological validation, $\ell = \sigma_G = 4\Delta x$ was held fixed while $R/\Delta x = (8, 12, 16)$ was varied, thereby isolating the probe-window dependence.

Varying either scale defines a diagnostic resolution flow for a fixed cosmological realization and must be distinguished from the physical dynamics that generated the density field. Unlike Ricci flow, this does not describe temporal evolution of geometry, but rather a change of resolution for a fixed cosmological realization. Changing ℓ modifies the Bakry-Émery tensor itself, whereas changing R alters the density-weighted finite-window response of that tensor.

When component weights, component responses, or tensor eigendirections vary sufficiently with scale, the reconstructed eigenvalues can evolve nonmonotonically and cross the classification threshold. As a result, curvature and anisotropy may exhibit nonmonotonic behavior and even sign changes due to the mixing of different structures. This implies that structure classification is intrinsically scale dependent. A monotone mixture of two scale-independent scalar responses is not, by itself, sufficient to generate nonmonotonicity. The tensor scale flow and its threshold-crossing conditions are derived in Appendix C. From this viewpoint, cosmological structures are characterized not only by fixed geometric types, but also by their trajectories under coarse graining and finite-window probing.

C. Curvature statistics and observable content

By treating the local effective curvature as a random field, the present framework naturally leads to a statistical description in terms of curvature. At linear order in a density contrast smoothed at scale ℓ, the normalized pointwise trace is $\overline{K}^{(1)}_{\text{point}} = -(1/d)\nabla^2\delta_\ell$ and has zero ensemble mean on a statistically homogeneous domain. Its density cross-spectrum is $P_{\delta_\ell \overline{K}}(k) = k^2 P_\ell(k)/d$, whereas its autospectrum is $P_{\overline{K}}(k) = k^4 P_\ell(k)/d^2$, where P_ℓ is the smoothed-density power spectrum. The full tensor and finite-window transfer functions are given in Appendix D.

Thus, in the Gaussian-linear limit, curvature statistics are differential reweightings of conventional two-point information rather than independent statistics. Beyond that limit, the logarithmic transformation, density-weighted window normalization, and finite-transport response generate nonlinear mode coupling. The finite-dimensional deficit begins quadratically in the score field and lowers the effective curvature, so its statistics probe directional profile steepness rather than a generic high-density enhancement.

Operationally, the finite-transport response is evaluated from the reconstructed potential $\widehat{V}_\ell$ along the controlled interpolation and does not require a pointwise Hessian estimate. The direct finite-window and pointwise branches are used as distinct references in the three-layer validation of Sec. VI.

D. Numerical validation and operational robustness

The controlled tests summarized in Table I recover quadratic tensors to a relative accuracy of 10^{-9} or better, confirm the expected $O(\varepsilon)$ and $O(\varepsilon^2)$ convergence laws for the one-sided and symmetric estimators, and identify the stable resolution range $1 \lesssim \eta \lesssim 2$. They also show that resolved finite-window class transitions are observable departures rather than estimator failures.

In the periodic Zel'dovich mock, the transport reconstruction matched the direct finite-window reference with directional correlation above 0.999 995 and median tensor error of 0.19%–0.37%. The finite-window tensor nevertheless differed substantially from the pointwise Hessian as the probe radius increased, establishing the two relations summarized in Eq. (152).

The additional direction tests show that replacing the design can produce substantially larger tensor and classification changes than these internal errors. A uniformly weighted angular projection converges under angular refinement, but retains a nonzero residual relative to the complete directional response. Thus, the finite-window tensor is a specified quadratic summary, not a direction-independent representation of that full response (Appendix E 6).

Under thinning to $f_{\mathrm{samp}} = 1/16$, the internal estimator error remained below 0.3%, while the end-to-end error was dominated by density reconstruction. At the fiducial radius, $E_{\mathrm{total}} = 0.180$ compared with $E_{\mathrm{Hessian}} = 0.259$, an approximately 30% reduction when each observable is evaluated relative to its own full-sampling reference. The score-alignment deficit also retained strong sampled-full correlation. These findings are summarized by Eq. (155). The comparison is restricted to the original six-direction observable and the pointwise log-density Hessian, not every directional design or cosmic-web classifier.

The numerical results therefore support the intended hierarchy: Layer I validates the finite-transport estimator against the direct finite-window reference, Layer II measures finite-window departure from the pointwise Hessian, and Layer III isolates degradation caused by sampling and density reconstruction. Classification agreement must be interpreted jointly with tensor error and eigenvalue margins.

E. Scope, implementation, and limitations

The framework is a diagnostic geometry of a simulated or reconstructed density field. It neither replaces the gravitational dynamics that generated the field nor identifies the controlled Wasserstein geodesic with a physical matter trajectory. The smoothing scale, probe radius, deformation, probe directions and directional weights, interpolation scheme, and eigenvalue threshold are part of the observable definition and must be reported and calibrated.

A positive density reconstruction is required because $V = -\ln p$ must remain finite. Local normalization isolates spatial profile geometry and removes the window mass, which should be retained separately when amplitude information is required. The potential response can be evaluated directly for weighted discrete probes, whereas the total and kinematic entropy responses require a regularized density or a discrete relative entropy defined with respect to a common reference measure.

The present calibration uses a periodic box, uniform Bernoulli thinning, fixed Gaussian smoothing, and the full catalog as a numerical reference rather than continuum truth. At finite ε, the response depends on the chosen transport path and numerical resolution. Survey masks, varying completeness, radial selection, redshift-space distortions, tracer bias, and anisotropic uncertainties must therefore be incorporated through a survey-specific forward operator and a new full-sampling calibration.

F. Outlook

The controlled program has established the basic consistency and numerical viability of the estimator. The immediate next step is to apply the same hierarchy to realistic survey mocks in which

masks, selection functions, redshift-space distortions, tracer bias, and measurement uncertainties are propagated through the density reconstruction and reference-measure construction.

A systematic exploration of the two-scale plane (ℓ, R) is also important. Independent variation of ℓ and R would separate changes in the weighted background tensor from changes in its finite-window sampling and would identify structures that persist or undergo resolved threshold crossings under coarse graining.

Because the curvature is defined as a response, the framework may also connect naturally with symmetry-response and separate-universe-type analyses. Long-wavelength perturbations can modify the reference potential, local-probe measure, and directional response separately, suggesting a route to environmental response and large-scale bias.

The connection with peak statistics requires care. For a smooth positive density, stationary points of ρ_R and $\ln \rho_R$ coincide, and at such a point $H_R = -\rho_R^{-1}\nabla^2\rho_R$; hence, H_R and the sign-reversed density Hessian $-\nabla^2\rho_R$ have the same inertia, with the zero-threshold qualification discussed in Sec. V B. The unweighted Kac-Rice operator is correspondingly invariant:

$$|\det H_R|\delta_{\mathrm{D}}^{(d)}(\nabla \ln \rho_R)\Theta(H_R) = |\det(-\nabla^2\rho_R)|\delta_{\mathrm{D}}^{(d)}(\nabla\rho_R)\Theta(-\nabla^2\rho_R).$$

The logarithmic representation therefore does not by itself define a new unweighted peak point process at fixed smoothing scale. New peak observables must instead arise from finite-window transport, scale evolution, explicit weighting, or changes in the underlying probability law.

In conclusion, the work establishes a consistent separation of total entropy response, kinematic deformation, density-induced Bakry-Émery curvature, and finite-dimensional score alignment. The log-density-Hessian classifier is recovered as a controlled local limit. At finite radius, the response retains surrounding-profile information, while its tensor summary requires an explicit directional design and can be complemented by a converged angular projection and its nonquadratic residual. For the original six-direction construction, the finite-window observable is more robust than the pointwise log-density Hessian under the tested thinning conditions, each relative to its own full-sampling reference.

These results provide a mathematically defined and numerically validated bridge between weighted-Wasserstein geometry and cosmological morphology. Extending that bridge to observational data now requires survey-specific forward modeling and calibration rather than a change in the underlying estimator hierarchy.

ACKNOWLEDGMENTS

We thank the referees for their constructive comments that improved the paper much. We are deeply grateful to Shiro Ikeda for helpful insights for this work. This work was supported by JSPS Grant-in-Aid in Scientific Research (24H00247), JST CREST Grant No. JPMJCR24Q1, and by the Joint Research program of the Institute of Statistical Mathematics (General Research 2) entitled "Machine-Learning Cosmogony: From Structure Formation to Galaxy Evolution."

DATA AVAILABILITY

There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

APPENDIX A: DERIVATION OF THE VOLUME DISTORTION COEFFICIENT $\beta_t^{K,N}$ BY JACOBI FIELDS

This Appendix summarizes the comparison-geometric origin of the volume distortion coefficient $\beta_t^{K,N}$ used in Sec. III D. The geometric dimension $d = \dim M$ determines the actual $d-1$ transverse Jacobi fields, whereas the synthetic dimension N enters the weighted curvature-dimension comparison and need not be an integer [34].

1. Jacobi comparison and transverse volume deformation

Let $\gamma : [0, r] \to M$ be a unit-speed minimizing geodesic on a d-dimensional Riemannian manifold. On the orthogonal bundle $\dot\gamma^{\perp}$, define $\mathsf{R}_\gamma(s)w = R(w, \dot\gamma(s))\dot\gamma(s)$. The transverse Jacobi tensor Y satisfies

$$\frac{\mathrm{D}^2\mathsf{Y}}{\mathrm{d}s^2} + \mathsf{R}_\gamma \mathsf{Y} = 0, \quad \mathsf{Y}(0) = 0, \quad \frac{\mathrm{D}\mathsf{Y}}{\mathrm{d}s}(0) = I_\perp. \tag{A1}$$

Before the first conjugate point, $j_\gamma(s) = \det \mathsf{Y}(s) > 0$ is the transverse volume density. Writing

$$\mathsf{S} \equiv \frac{\mathrm{D}\mathsf{Y}}{\mathrm{d}s}\mathsf{Y}^{-1}, \quad \theta \equiv \mathrm{Tr}\mathsf{S} = \frac{\mathrm{d}}{\mathrm{d}s}\ln j_\gamma, \quad \Sigma \equiv \mathsf{S} - \frac{\theta}{d-1} I_\perp, \tag{A2}$$

the Jacobi equation gives the Riccati and trace equations

$$\begin{aligned} \frac{\mathrm{D}\mathsf{S}}{\mathrm{d}s} + \mathsf{S}^2 + \mathsf{R}_\gamma &= 0, \\ \frac{\mathrm{d}\theta}{\mathrm{d}s} + \frac{\theta^2}{d-1} + \|\Sigma\|^2_{\mathrm{HS}} + \mathrm{Ric}(\dot\gamma, \dot\gamma) &= 0. \end{aligned} \tag{A3}$$

Thus, Ricci curvature controls volume deformation together with the non-negative shear contribution. For $u_\gamma = j_\gamma^{1/(d-1)}$,

$$\begin{aligned} \frac{\mathrm{d}^2 u_\gamma}{\mathrm{d}s^2} &= -\frac{\mathrm{Ric}(\dot\gamma, \dot\gamma) + \|\Sigma\|^2_{\mathrm{HS}}}{d-1} u_\gamma, \\ \mathrm{Ric} \geqslant K g \Longrightarrow \frac{\mathrm{d}^2 u_\gamma}{\mathrm{d}s^2} &+ \frac{K}{d-1} u_\gamma \leqslant 0. \end{aligned} \tag{A4}$$

2. Model spaces and the distortion coefficient

For $\kappa \in \mathbb{R}$, let

$$s_\kappa(r) \equiv \begin{cases} \frac{\sin(\sqrt{\kappa}\, r)}{\sqrt{\kappa}}, & \kappa > 0, \\ r, & \kappa = 0, \\ \frac{\sinh(\sqrt{-\kappa}\, r)}{\sqrt{-\kappa}}, & \kappa < 0, \end{cases} \tag{A5}$$

so that $s_\kappa'' + \kappa s_\kappa = 0$, $s_\kappa(0) = 0$, and $s_\kappa'(0) = 1$. The scalar comparison from Eq. (A4) yields, before the first model conjugate point,

$$\frac{j_\gamma(tr)}{t^{d-1} j_\gamma(r)} \geqslant \left[\frac{s_{K/(d-1)}(tr)}{t\, s_{K/(d-1)}(r)}\right]^{d-1}, \quad 0 < t \leqslant 1. \tag{A6}$$

The corresponding synthetic comparison coefficient for $N > 1$ is

$$\beta_t^{K,N}(r) \equiv \left[\frac{s_{K/(N-1)}(tr)}{t\, s_{K/(N-1)}(r)}\right]^{N-1}. \tag{A7}$$

The Euclidean radial factor has been divided out, so $\beta_t^{0,N}(r) = 1$. For $K > 0$, the formula is used for $r < \pi\sqrt{(N-1)/K}$ and extended by the standard $+\infty$ convention beyond the first model conjugate point; at $r = 0$, it is defined by continuity. When N is noninteger, $\beta_t^{K,N}$ is a comparison coefficient and not the determinant of $N - 1$ literal Jacobi fields.

3. Weighted transport Jacobians and CD(K, N)

Consider a smooth weighted manifold $\mathrm{d}m = e^{-V}\mathrm{dvol}_g$ of geometric dimension d, with $N > d$ and

$$\mathrm{Ric}_{V,N} \equiv \mathrm{Ric} + \nabla^2 V - \frac{1}{N-d}\nabla V \otimes \nabla V. \tag{A8}$$

Let $F_t(x) = \exp_x(t\xi_x)$ be the regular interpolation maps of a smooth optimal transport and set $r_x = |\xi_x| = d(x, F_1(x))$. The Jacobian relative to m is

$$\mathcal{J}_t^m(x) \equiv \exp\left[V(x) - V(F_t(x))\right] \det \mathrm{d}F_t(x). \tag{A9}$$

If $\mu_t = q_t m = (F_t)_{\#}\mu_0$, weighted mass conservation gives

$$q_t(F_t(x))\mathcal{J}_t^m(x) = q_0(x). \tag{A10}$$

Under $\mathrm{Ric}_{V,N} \geqslant Kg$, the sharp weighted Jacobi comparison is expressed through

$$\tau_{K,N}^{(t)}(r) \equiv t^{1/N}\left[\frac{s_{K/(N-1)}(tr)}{s_{K/(N-1)}(r)}\right]^{(N-1)/N} = t\left[\beta_t^{K,N}(r)\right]^{1/N}, \tag{A11}$$

and yields

$$\left[\mathcal{J}_t^m(x)\right]^{1/N} \geqslant \tau_{K,N}^{(1-t)}(r_x)\left[\mathcal{J}_0^m(x)\right]^{1/N} + \tau_{K,N}^{(t)}(r_x)\left[\mathcal{J}_1^m(x)\right]^{1/N}. \tag{A12}$$

Combining this inequality with Eq. (A10) gives

$$q_t(F_t(x))^{-1/N} \geqslant (1-t)\left[\beta_{1-t}^{K,N}(r_x)\right]^{1/N} q_0(x)^{-1/N} + t\left[\beta_t^{K,N}(r_x)\right]^{1/N} q_1(F_1(x))^{-1/N}. \tag{A13}$$

This is the smooth-map form of the CD(K, N) density interpolation used in Sec. III D; integration with respect to the optimal end point coupling gives the corresponding Rényi-entropy inequality. Thus, $\beta_t^{K,N}$ compares weighted transport-volume deformation with the constant-curvature model. It should not be identified with the local score-alignment deficit $\Delta K^{(N)}$, and N is neither the number of Jacobi fields nor the number of particles.

APPENDIX B: LOCAL LIMIT OF THE FINITE-WINDOW CURVATURE RESPONSE

This Appendix establishes the local relation among the finite-transport estimator, the direct finite-window response, and the pointwise log-density Hessian. The density-smoothing scale ℓ and probe radius R remain distinct. On a curved background, H_{point} below is replaced by the corresponding Bakry-Émery tensor.

1. Local expansion and probe moments

Let $\rho_\ell > 0$ be smoothed at scale ℓ, and define

$$\begin{aligned} V_\ell &\equiv -\ln p_\ell = -\ln \rho_\ell + \text{constant}, \\ s_x^{(\ell)} &\equiv \nabla \ln \rho_\ell(x), \\ H_{\mathrm{point}}(x;\ell) &\equiv \nabla\nabla V_\ell(x) = -\nabla\nabla \ln \rho_\ell(x). \end{aligned} \tag{B1}$$

For $z = y - x$,

$$V_\ell(x+z) = V_\ell(x) - \left[s_x^{(\ell)}\right]^\top z + \tfrac{1}{2} z^\top H_{\mathrm{point}}(x;\ell) z + O(|z|^3). \tag{B2}$$

The linear score term is retained at a general point; it cancels from the quadratic midpoint response and therefore does not obstruct the Hessian limit.

Let $W_R(z) = R^{-d}W(z/R)$, where W is non-negative, centered, and isotropic, with

$$\begin{aligned}\int_{\mathbb{R}^d} W(u)\,\mathrm{d}u &= 1,\\ \int_{\mathbb{R}^d} uW(u)\,\mathrm{d}u &= 0,\\ \int_{\mathbb{R}^d} uu^\top W(u)\,\mathrm{d}u &= c_2 I, \quad c_2 > 0.\end{aligned} \tag{B3}$$

The density-weighted probe is

$$\begin{aligned}\mathrm{d}\mu_x^{(R;\ell)}(x+z) &\equiv \frac{W_R(z)\rho_\ell(x+z)}{Z_x^{(R;\ell)}}\,\mathrm{d}z,\\ Z_x^{(R;\ell)} &\equiv \int_{\mathbb{R}^d} W_R(z)\rho_\ell(x+z)\,\mathrm{d}z.\end{aligned} \tag{B4}$$

Writing $m_x^{(R;\ell)} = \int z\,\mathrm{d}\mu_x^{(R;\ell)}$ and $\Sigma_x^{(R;\ell)} = \int zz^\top\,\mathrm{d}\mu_x^{(R;\ell)}$, the centered even window gives

$$\begin{aligned}m_x^{(R;\ell)} &= c_2R^2 s_x^{(\ell)} + O(R^4),\\ \Sigma_x^{(R;\ell)} &= c_2R^2 I + R^4 M_x^{(\ell)} + O(R^6).\end{aligned} \tag{B5}$$

Thus, the score shifts the probe barycenter at order R^2, while the leading second moment is isotropic; no Gaussian approximation is required.

2. Direct response and the two-step local limit

For a fixed affine probe A, define the normalization and curvature numerator by

$$\begin{aligned}\mathcal{D}_A(x;R,\ell) &\equiv \int |Az|^2\,\mathrm{d}\mu_x^{(R;\ell)},\\ \mathcal{N}_A(x;R,\ell) &\equiv \int (Az)^\top H_{\mathrm{point}}(x+z;\ell)(Az)\mathrm{d}\mu_x^{(R;\ell)}.\end{aligned} \tag{B6}$$

Whenever $\mathcal{D}_A(x;R,\ell) > 0$, the direct response is

$$K_{\mathrm{eff}}^{(\infty)}(x;R,\ell;A) \equiv \frac{\mathcal{N}_A(x;R,\ell)}{\mathcal{D}_A(x;R,\ell)}. \tag{B7}$$

For a constant Hessian, this is the exact quadratic reference. For a smooth field and the centered even window,

$$K_{\mathrm{eff}}^{(\infty)}(x;R,\ell;A) = \frac{\mathrm{Tr}[H_{\mathrm{point}}(x;\ell)AA^\top]}{\mathrm{Tr}(AA^\top)} + O(R^2).$$

For the rank-one probe $A = \mathbf{e}\mathbf{e}^\top$, this becomes

$$K_{\mathrm{eff}}^{(\infty)}(x;R,\ell;\mathbf{e}) = \mathbf{e}^\top H_{\mathrm{point}}(x;\ell)\mathbf{e} + O(R^2). \tag{B8}$$

An asymmetric or off-centered window can instead produce an $O(R)$ correction.

For a sufficiently smooth potential and nondegenerate transport normalization, the one-sided and symmetric finite-transport responses satisfy

$$K_{\rho,\varepsilon}^{(\infty)} = K_{\mathrm{eff}}^{(\infty)} + O(\varepsilon), \quad K_{\rho,\varepsilon,\mathrm{sym}}^{(\infty)} = K_{\mathrm{eff}}^{(\infty)} + O(\varepsilon^2). \tag{B9}$$

Applying a fixed full-rank linear tensor reconstruction to the same directional design therefore gives

$$\widehat{H}^{(\mathcal{E})}_{\text{transport}}(x; R, \ell, \varepsilon) = H^{(\mathcal{E})}_{\text{window}}(x; R, \ell) + O(\varepsilon),$$
$$H^{(\mathcal{E})}_{\text{window}}(x; R, \ell) = H_{\text{point}}(x; \ell) + O(R^2). \tag{B10}$$

The first remainder becomes $O(\varepsilon^2)$ for the symmetric estimator. The finite-transport error and the finite-window departure are therefore distinct even when both limits are ultimately taken.

3. Directional invariants and finite-dimensional correction

Averaging rank-one responses uniformly over $\mathbb{S}^{d-1}$ gives

$$\overline{K}^{(\infty)}(x; R, \ell) = \frac{1}{d}\text{Tr}H_{\text{point}}(x; \ell) + O(R^2). \tag{B11}$$

The corresponding directional variance is

$$\mathbb{V}_{\mathbf{e}}\big[K^{(\infty)}_{\text{eff}}\big] = \frac{2}{d(d+2)}\left[\text{Tr}\big(H^2_{\text{point}}\big) - \frac{(\text{Tr}H_{\text{point}})^2}{d}\right] + O(R^2). \tag{B12}$$

The variance measures the squared norm of the traceless Hessian and vanishes precisely for an isotropic pointwise tensor.

For $N > d$, define the score-alignment numerator by

$$\mathcal{G}_A(x; R, \ell) \equiv \int [Az \cdot \nabla \ln \rho_\ell(x+z)]^2 \mathrm{d}\mu^{(R;\ell)}_x. \tag{B13}$$

The finite-dimensional deficit is then

$$\Delta K^{(N)}(x; R, \ell; A) \equiv \frac{\mathcal{G}_A(x; R, \ell)}{(N-d)\mathcal{D}_A(x; R, \ell)} \geqslant 0. \tag{B14}$$

For $A = \mathbf{e}\mathbf{e}^\top$,

$$\Delta K^{(N)}(x; R, \ell; \mathbf{e}) = \frac{\big[\mathbf{e} \cdot s^{(\ell)}_x\big]^2}{N-d} + O(R^2),$$
$$K^{(N)}_{\text{eff}}(x; R, \ell; \mathbf{e}) = \mathbf{e}^\top\left[H_{\text{point}}(x; \ell) - \frac{s^{(\ell)}_x[s^{(\ell)}_x]^\top}{N-d}\right]\mathbf{e} + O(R^2). \tag{B15}$$

The bracketed tensor is the Euclidean N-Bakry-Émery tensor. The correction is therefore a non-negative score-alignment deficit that lowers the infinite-dimensional curvature.

4. Classification stability

Let $\lambda_a(x; \ell)$ be the eigenvalues of $H_{\text{point}}(x; \ell)$ and use the same threshold λ_{th} for the pointwise and finite-window tensors. Define the pointwise margin by

$$m_{\text{point}}(x; \ell) \equiv \min_{1 \leqslant a \leqslant d} |\lambda_a(x; \ell) - \lambda_{\text{th}}|. \tag{B16}$$

Weyl's inequality gives the sufficient stability condition

$$\big\|H^{(\mathcal{E})}_{\text{window}}(x; R, \ell) - H_{\text{point}}(x; \ell)\big\|_2 < m_{\text{point}}(x; \ell). \tag{B17}$$

Under this condition, $C_{\text{window}}(x; R, \ell) = C_{\text{point}}(x; \ell)$. Hence, the two classifications agree in the resolved local limit whenever the pointwise spectrum is separated from the threshold. Near a classification boundary, a small finite-window departure can change the discrete class even when the tensor difference is small.

APPENDIX C: ASYMPTOTIC SCALE DEPENDENCE AND STRUCTURAL TRANSITIONS

This Appendix describes the dependence of the direct finite-window curvature response on the probe radius R at fixed density-smoothing scale ℓ. Varying ℓ changes the weighted background geometry, whereas varying R changes the region over which that fixed geometry is sampled. The scale flow considered here is therefore diagnostic rather than dynamical.

1. Small-radius response and tensor scale flow

Let $\mu_x^{(R;\ell)}$ be the probe measure defined in Eq. (B4). For a nonzero affine probe matrix A, suppress the common arguments $(x; R, \ell)$ and write the numerator and denominator of Eq. (B7) as

$$\mathcal{N}_A \equiv \int (Az)^\top H_{\text{point}}(x+z;\ell)(Az)\, \mathrm{d}\mu_x^{(R;\ell)}(x+z),$$
$$\mathcal{D}_A \equiv \int |Az|^2\, \mathrm{d}\mu_x^{(R;\ell)}(x+z). \tag{C1}$$

Then, $K_{\text{eff}}^{(\infty)}(x; R, \ell; A) = \mathcal{N}_A/\mathcal{D}_A$. For a centered even window and a sufficiently smooth density,

$$\mathcal{D}_A = R^2\mathcal{D}_{A,2} + R^4\mathcal{D}_{A,4} + O(R^6),$$
$$\mathcal{N}_A = R^2\mathcal{N}_{A,2} + R^4\mathcal{N}_{A,4} + O(R^6). \tag{C2}$$

The leading coefficients are $\mathcal{D}_{A,2} = c_2\mathrm{Tr}(AA^\top)$ and $\mathcal{N}_{A,2} = c_2\mathrm{Tr}[H_{\text{point}}(x;\ell)AA^\top]$. The fourth-order terms contain both the density-weighted deformation of the probe moments and spatial derivatives of the pointwise curvature.

Expanding the ratio gives

$$K_{\text{eff}}^{(\infty)}(x; R, \ell; A) = K_A^{(0)}(x;\ell) + R^2\mathcal{C}_A(x;\ell) + O(R^4),$$
$$K_A^{(0)}(x;\ell) \equiv \frac{\mathrm{Tr}[H_{\text{point}}(x;\ell)AA^\top]}{\mathrm{Tr}(AA^\top)}, \tag{C3}$$
$$\mathcal{C}_A(x;\ell) \equiv \frac{\mathcal{N}_{A,4} - K_A^{(0)}\mathcal{D}_{A,4}}{\mathcal{D}_{A,2}}.$$

The occurrence of $AA^\top$ is essential. For the rank-one probe $A_{\mathbf{e}} = \mathbf{e}\mathbf{e}^\top$,

$$K_{\text{eff}}^{(\infty)}(x; R, \ell; \mathbf{e}) = \mathbf{e}^\top H_{\text{point}}(x;\ell)\mathbf{e} + R^2\mathcal{C}_{\mathbf{e}}(x;\ell) + O(R^4). \tag{C4}$$

A constant Hessian gives $\mathcal{C}_{\mathbf{e}} = 0$; a nonzero correction therefore reflects spatial variation of the curvature field and its density weighting.

Let $\mathcal{E} = \{\mathbf{e}_a\}_{a=1}^m$ be a fixed full-rank directional design and let $\mathcal{R}_\mathcal{E}$ be its linear tensor-reconstruction operator. Applying it to Eq. (C4) yields

$$H_{\text{window}}^{(\mathcal{E})}(x; R, \ell) = H_{\text{point}}(x;\ell) + R^2 B_x^{(\mathcal{E},\ell)} + O(R^4), \quad B_x^{(\mathcal{E},\ell)} \equiv \mathcal{R}_\mathcal{E}\left(\left\{\mathcal{C}_{\mathbf{e}_a}(x;\ell)\right\}_{a=1}^m\right). \tag{C5}$$

At finite R, the correction may depend on the directional design because the complete response need not be a single quadratic form. The design dependence disappears at leading order in the local limit.

2. Scale derivative and component mixing

Whenever the window family is differentiable in R, Eq. (C1) gives

$$\frac{\partial}{\partial R}K_{\text{eff}}^{(\infty)}(x; R, \ell; A) = \frac{(\partial_R\mathcal{N}_A)\mathcal{D}_A - \mathcal{N}_A(\partial_R\mathcal{D}_A)}{\mathcal{D}_A^2} = \frac{\partial_R\mathcal{N}_A - K_{\text{eff}}^{(\infty)}\partial_R\mathcal{D}_A}{\mathcal{D}_A}. \tag{C6}$$

A stationary scale is obtained when the numerator of this expression vanishes. Near $R = 0$, the derivative is $2R\mathcal{C}_A + O(R^3)$, but this local sign does not determine the behavior at larger radius. For hard windows or discrete samples, apparent extrema must also be checked against radius refinement.

A component-resolved representation follows from any non-negative partition of unity $\{\chi_q\}_{q=1}^{Q}$ within the probe. Let $\mathcal{N}_{A,q}$ and $\mathcal{D}_{A,q}$ denote the numerator and denominator restricted by χ_q, and set $K_{A,q} = \mathcal{N}_{A,q}/\mathcal{D}_{A,q}$. Then,

$$K_{\mathrm{eff}}^{(\infty)}(x; R, \ell; A) = \sum_{q=1}^{Q} \alpha_{A,q}(R) K_{A,q}(R), \quad \alpha_{A,q}(R) \equiv \frac{\mathcal{D}_{A,q}(R)}{\sum_{p=1}^{Q} \mathcal{D}_{A,p}(R)}, \quad \sum_{q=1}^{Q} \alpha_{A,q}(R) = 1. \tag{C7}$$

The weights are therefore determined by directional transport energy, not by component mass fractions alone. For two components,

$$\frac{\mathrm{d}}{\mathrm{d}R} K_{\mathrm{eff}}^{(\infty)}(R) = \alpha'(R)[K_1(R) - K_2(R)] + \alpha(R) K_1'(R) + [1 - \alpha(R)] K_2'(R). \tag{C8}$$

A monotone weight mixing two scale-independent responses remains monotone. Genuine nonmonotonicity requires scale dependence of a component response, a nonmonotone effective weight, more than two components, or direction-dependent mixing that alters the reconstructed tensor and its eigendirections.

3. Eigenvalue flow and structural transitions

For fixed x, ℓ, and $\mathcal{E}$, write $\mathsf{H}(R) = H_{\mathrm{window}}^{(\mathcal{E})}(x; R, \ell)$. If $\lambda_a(R)$ is a simple eigenvalue with normalized eigenvector $u_a(R)$, then

$$\frac{\mathrm{d}\lambda_a}{\mathrm{d}R} = u_a(R)^{\top} \mathsf{H}'(R) u_a(R). \tag{C9}$$

Thus, an eigenvalue need not follow any single fixed directional response because its eigenvector may rotate. If the pointwise spectrum is simple, Eq. (C5) gives

$$\begin{aligned} \lambda_a(R) &= \lambda_a^{\mathrm{point}} + R^2 b_a^{(\mathcal{E},\ell)} + O(R^4), \\ b_a^{(\mathcal{E},\ell)} &\equiv \left[u_a^{\mathrm{point}}\right]^{\top} B_x^{(\mathcal{E},\ell)} u_a^{\mathrm{point}}. \end{aligned} \tag{C10}$$

Near a degeneracy, the leading splitting is obtained by restricting $B_x^{(\mathcal{E},\ell)}$ to the degenerate eigenspace rather than by using the nondegenerate formula.

The finite-window morphology is

$$C_{\mathrm{window}}(x; R, \ell) \equiv \sum_{a=1}^{d} \mathbf{1}[\lambda_a(R) > \lambda_{\mathrm{th}}]. \tag{C11}$$

Under continuous tensor flow, the class can change only when

$$\lambda_a(R_*) = \lambda_{\mathrm{th}} \tag{C12}$$

for at least one eigenvalue. In the small-radius regime, a simple crossing has the leading estimate

$$R_*^2 \simeq \frac{\lambda_{\mathrm{th}} - \lambda_a^{\mathrm{point}}}{b_a^{(\mathcal{E},\ell)}}, \tag{C13}$$

provided the right-hand side is positive and the expansion remains valid. A zero crossing of a fixed directional response is neither necessary nor sufficient unless the direction follows the relevant eigenvector and $\lambda_{\mathrm{th}} = 0$.

Define the finite-window margin by

$$m_{\mathrm{window}}(x; R, \ell) \equiv \min_{1 \leqslant a \leqslant d} |\lambda_a(R) - \lambda_{\mathrm{th}}|. \tag{C14}$$

Writing $\Delta H_{12}^{(\mathcal{E})} = H_{\mathrm{window}}^{(\mathcal{E})}(x; R_2, \ell) - H_{\mathrm{window}}^{(\mathcal{E})}(x; R_1, \ell)$, Weyl's inequality gives the sufficient stability condition

$$\left\| \Delta H_{12}^{(\mathcal{E})} \right\|_2 < m_{\mathrm{window}}(x; R_1, \ell). \tag{C15}$$

Under this condition, the class is unchanged between R_1 and R_2. A resolved transition should be inferred from tensor eigenvalues and threshold margins only after the transport estimator has converged to the direct finite-window response. Finite-ε, interpolation, and sparse-sampling effects belong to the other validation layers.

APPENDIX D: LINEARIZED CURVATURE STATISTICS AND FINITE-DIMENSIONAL CORRECTIONS

This Appendix relates the pointwise and finite-window curvature fields to conventional density statistics. The smoothing scale ℓ and probe radius R remain distinct. Linear curvature is a differential and windowed transform of the density contrast; nonlinear log-density terms, probe normalization, and the finite-dimensional score-alignment deficit introduce higher-order mode coupling.

1. Linear pointwise-curvature statistics

Let δ be a statistically homogeneous and isotropic density contrast with zero mean and Fourier convention

$$\delta(x) = \int \frac{\mathrm{d}^d k}{(2\pi)^d} \widetilde{\delta}(k) e^{ik\cdot x},$$
$$\mathbb{E}\widetilde{\delta}(k)\widetilde{\delta}(k') = (2\pi)^d \delta_{\mathrm{D}}(k + k') P_\delta(k). \tag{D1}$$

For a smoothing kernel $\mathcal{W}_\ell$,

$$\delta_\ell(x) \equiv \int \frac{\mathrm{d}^d k}{(2\pi)^d} \widetilde{\mathcal{W}}_\ell(k) \widetilde{\delta}(k) e^{ik\cdot x}, \tag{D2}$$

$$\sigma_j^2(\ell) \equiv \int \frac{\mathrm{d}^d k}{(2\pi)^d} k^{2j} |\, \widetilde{\mathcal{W}}_\ell(k)|^2 P_\delta(k). \tag{D3}$$

Gaussianity is unnecessary for the linear two-point relations below; it is needed only when higher-order moments are reduced using Wick's theorem.

For $\rho_\ell = \overline{\rho}(1 + \delta_\ell)$, define

$$H_{\mathrm{point},ij}(x; \ell) \equiv -\partial_i \partial_j \ln [1 + \delta_\ell(x)],$$
$$\overline{K}_{\mathrm{point}}(x; \ell) \equiv \frac{1}{d} \mathrm{Tr} H_{\mathrm{point}}(x; \ell) = -\frac{1}{d} \nabla^2 \ln [1 + \delta_\ell(x)]. \tag{D4}$$

Statistical homogeneity implies

$$\mathbb{E} H_{\mathrm{point},ij}(x; \ell) = 0, \quad \mathbb{E}\overline{K}_{\mathrm{point}}(x; \ell) = 0. \tag{D5}$$

The same spatial average vanishes on a periodic domain. Boundaries, masks, or inhomogeneous selection can invalidate this translation-invariant statement.

To first order in δ_ℓ,

$$H^{(1)}_{\mathrm{point},ij}(x; \ell) = -\partial_i \partial_j \delta_\ell(x), \quad \overline{K}^{(1)}_{\mathrm{point}}(x; \ell) = -\frac{1}{d} \nabla^2 \delta_\ell(x), \tag{D6}$$

$$\widetilde{\overline{K}}^{(1)}_{\mathrm{point}}(k; \ell) = \frac{k^2}{d} \widetilde{\mathcal{W}}_\ell(k) \widetilde{\delta}(k). \tag{D7}$$

Hence, the auto- and density-curvature cross-spectra are

$$P^{(1)}_{\overline{K}_{\rm point}}(k;\ell) = \frac{k^4}{d^2}|\,\widetilde{\mathcal{W}}_\ell(k)|^2 P_\delta(k), \tag{D8}$$

$$P^{(1)}_{\delta_\ell,\overline{K}_{\rm point}}(k;\ell) = \frac{k^2}{d}|\,\widetilde{\mathcal{W}}_\ell(k)|^2 P_\delta(k). \tag{D9}$$

Thus, $\mathbb{V}(\overline{K}^{(1)}_{\rm point}) = \sigma_2^2/d^2$ and $\mathbb{C}\mathbb{o}\mathbb{v}(\delta_\ell, \overline{K}^{(1)}_{\rm point}) = \sigma_1^2/d$. The positive k^2-weighted integral is a cross-covariance, not a one-point mean. In real space,

$$\xi_{\delta_\ell\overline{K}}(r) = -\frac{1}{d}\nabla_r^2\xi_\ell(r), \quad \xi_{\overline{K}}(r) = \frac{1}{d^2}\big(\nabla_r^2\big)^2\xi_\ell(r). \tag{D10}$$

2. Finite-window transfer and tensor reconstruction

For a centered even probe window and a unit direction $\mathbf{e}$, define

$$M_{2,\mathbf{e}}(R) \equiv \int_{\mathbb{R}^d} W_R(z)(\mathbf{e}\cdot z)^2\,\mathrm{d}z, \quad Q_{R,\mathbf{e}}(z) \equiv \frac{W_R(z)(\mathbf{e}\cdot z)^2}{M_{2,\mathbf{e}}(R)}, \quad \int_{\mathbb{R}^d} Q_{R,\mathbf{e}}(z)\,\mathrm{d}z = 1. \tag{D11}$$

At linear order, density weighting of the normalized probe contributes only beyond first order because the background curvature vanishes. Suppressing common Fourier-space arguments where clear, the directional response and its transform are

$$k^{(1)}_{\text{window},\mathbf{e}}(x;R,\ell) = \int_{\mathbb{R}^d} Q_{R,\mathbf{e}}(z)\mathbf{e}^\top H^{(1)}_{\rm point}(x+z;\ell)\mathbf{e}\,\mathrm{d}z,$$
$$\widetilde{k}^{(1)}_{\text{window},\mathbf{e}} = (\mathbf{e}\cdot k)^2\,\widetilde{\mathcal{W}}_\ell(k)\widetilde{Q}_{R,\mathbf{e}}(k)\widetilde{\delta}(k). \tag{D12}$$

The smoothing and probe scales thus enter through distinct filters. The corresponding spectra are

$$P^{(1)}_{\text{window},\mathbf{e}} = (\mathbf{e}\cdot k)^4|\,\widetilde{\mathcal{W}}_\ell(k)|^2|\widetilde{Q}_{R,\mathbf{e}}(k)|^2 P_\delta(k), \tag{D13}$$

$$P^{(1)}_{\text{window-point},\mathbf{e}}(k;R,\ell) = (\mathbf{e}\cdot k)^4|\,\widetilde{\mathcal{W}}_\ell(k)|^2\widetilde{Q}_{R,\mathbf{e}}(k)P_\delta(k), \tag{D14}$$

$$P^{(1)}_{\Delta\text{wp},\mathbf{e}} = (\mathbf{e}\cdot k)^4|\,\widetilde{\mathcal{W}}_\ell(k)|^2|\widetilde{Q}_{R,\mathbf{e}}(k) - 1|^2 P_\delta(k). \tag{D15}$$

The second expression is the power spectrum of the linear window-point departure. It vanishes as $R \to 0$ and represents finite-window observable information rather than transport-estimator error.

For a full-rank directional design, let R^a_{ij} denote the linear tensor-reconstruction coefficients. Define

$$\mathcal{T}^{(\mathcal{E})}_{ij}(k;R,\ell) \equiv\ \widetilde{\mathcal{W}}_\ell(k)\sum_{a=1}^{m}\mathsf{R}^a_{ij}(\mathbf{e}_a\cdot k)^2\widetilde{Q}_{R,\mathbf{e}_a}(k). \tag{D16}$$

The reconstructed field and its spectral covariance are

$$\widetilde{H}^{(\mathcal{E},1)}_{\text{window},ij}(k;R,\ell) = \mathcal{T}^{(\mathcal{E})}_{ij}(k;R,\ell)\widetilde{\delta}(k), \quad P^{(\mathcal{E},1)}_{ijmn}(k;R,\ell) = \mathcal{T}^{(\mathcal{E})}_{ij}\big[\mathcal{T}^{(\mathcal{E})}_{mn}\big]^* P_\delta(k). \tag{D17}$$

Because the reconstruction reproduces quadratic forms on the adopted design, $\mathcal{T}^{(\mathcal{E})}_{ij} \to k_i k_j\,\widetilde{\mathcal{W}}_\ell$ as $R \to 0$. At finite R, the transfer can retain design dependence because the complete directional response need not be an exact quadratic form.

3. Nonlinear and finite-dimensional contributions

Beyond linear order, $-\nabla\nabla\ln(1+\delta_\ell)$ contains quadratic and higher powers of δ_ℓ. Density weighting and normalization of the finite-window probe add further mode coupling. Their spectra therefore depend on convolutions of the power spectrum and, for non-Gaussian fields, on higher

polyspectra. Statistical homogeneity still enforces a zero mean for the exact pointwise log-density Hessian.

For $N > d$, define the score-alignment numerator and deficit by

$$\mathcal{G}_A(x;R,\ell) \equiv \int [Az\cdot\nabla \ln \rho_\ell(x+z)]^2 \mathrm{d}\mu_x^{(R;\ell)}(x+z),$$
$$\Delta K^{(N)}(x;R,\ell;A) \equiv \frac{\mathcal{G}_A(x;R,\ell)}{(N-d)\mathcal{D}_A(x;R,\ell)} \geqslant 0. \tag{D18}$$

The finite-dimensional curvature is $K_{\mathrm{eff}}^{(N)} = K_{\mathrm{eff}}^{(\infty)} - \Delta K^{(N)}$. The correction is a non-negative score-alignment deficit and is not generated by smoothing a Rényi-weighted density.

For $A = \mathbf{e}\mathbf{e}^\top$, the leading weak-field term and its homogeneous isotropic mean are

$$\Delta K_{\mathbf{e}}^{(N,2)}(x;R,\ell) = \frac{1}{N-d}\int_{\mathbb{R}^d} Q_{R,\mathbf{e}}(z)[\mathbf{e}\cdot\nabla\delta_\ell(x+z)]^2\mathrm{d}z. \tag{D19}$$

Its homogeneous isotropic mean is $\mathbb{E}\Delta K_{\mathbf{e}}^{(N,2)} = \sigma_1^2(\ell)/[d(N-d)]$. The deficit begins at quadratic order and therefore does not alter the linear curvature spectrum. For a Gaussian density field, Wick's theorem gives an $O(P_\delta^2)$ convolution with two gradient weights; for a non-Gaussian field, its cross-spectrum with the linear curvature also receives bispectrum contributions.

In the pointwise limit, statistical homogeneity gives

$$\mathbb{E}k_{\mathrm{point},\mathbf{e}}^{(N)} = -\frac{1}{N-d}\mathbb{E}[\mathbf{e}\cdot\nabla\ln\rho_\ell]^2 \leqslant 0, \tag{D20}$$
$$\mathbb{E}\overline{K}_{\mathrm{point}}^{(N)} = -\frac{1}{d(N-d)}\mathbb{E}|\nabla\ln\rho_\ell|^2 \leqslant 0. \tag{D21}$$

The negative finite-dimensional mean is distinct from the vanishing mean of the infinite-dimensional pointwise curvature.

4. Interpretation and scope

At linear order, pointwise curvature statistics contain the same two-point information as the smoothed density field, reweighted by spatial derivatives: The field and density cross-spectrum carry k^2, while the curvature auto-spectrum carries k^4. The finite-window observable adds the direction-dependent transfer $\widetilde{Q}_{R,\mathbf{e}}$, and the finite-dimensional deficit adds a nonlinear score statistic beginning at quadratic order.

Derivative weighting increases sensitivity to short wavelengths, but the ultraviolet response is controlled by the smoothing and probe kernels. On a finite survey, masks, selection, and boundary treatment convolve these spectra with the observational operator. The homogeneous relations above therefore apply to observed data only after the same forward modeling and full-sampling calibration used for the operational estimator.

APPENDIX E: SUPPLEMENTARY NUMERICAL DIAGNOSTICS

This Appendix specifies the numerical implementation and supplements the validation in Sec. VI. After the setup, four diagnostics retain the original six-direction thinning experiment, with the full-particle catalog as numerical reference and eight replicates at each nonunit sampling fraction. The final subsection separately changes the probe directions in a fixed reconstructed field. The comparisons preserve the distinction between internal estimator error, finite-window observable dependence, and sampled-full reconstruction degradation.

TABLE II. Parameter scans for the controlled tests. N_g is the number of grid nodes along each side, not the synthetic dimension.

Stage	N_g	R/L_0	ε
1	65	0.4, 0.6, 0.8, 1.0	0.02, 0.05, 0.10, 0.20
2	65	0.25, 0.40, 0.60, 0.80, 1.00	0.025, 0.05, 0.10, 0.20, 0.30
2b	65, 97, 129	0.80, 1.00	0.025, 0.05, 0.10, 0.20, 0.30

1. Numerical setup and reproducibility

a. Controlled potentials. The following specifies the archived implementation settings for the controlled tests and the periodic mock. Writing $\boldsymbol{z} = \boldsymbol{y} - \boldsymbol{x}$, the controlled fields use

$$V(\boldsymbol{z}) = V_0 + \boldsymbol{b}^\top \boldsymbol{z} + \frac{1}{2}\boldsymbol{z}^\top H_0 \boldsymbol{z} + \frac{\alpha}{4L_0^4}|\boldsymbol{z}|^4. \tag{E1}$$

The positive density is proportional to e^{-V}. Here, $L_0 = 1$ in computational units, $\boldsymbol{b} = (0.15, -0.10, 0.05)/L_0$, and the grid covers $[-3L_0, 3L_0]^3$, including both end points. Stage 1 sets $\alpha = 0$; Stage 2 uses $\alpha = 0.05$ and 0.20, and Stage 2b uses $\alpha = 0.20$. The central Hessian is $H_0 = Q\,\mathrm{diag}(\boldsymbol{\lambda})Q^\top$. In the order void, sheet, filament, and halo, the dimensionless canonical spectra $L_0^2\boldsymbol{\lambda}$ are $(-1, -1, -1)$, $(0, 0, 1)$, $(0, 1, 1)$, and $(1, 1, 1)$; the finite-margin spectra are $(-1.2, -0.8, -0.4)$, $(-0.4, -0.2, 0.8)$, $(-0.3, 0.6, 1.0)$, and $(0.4, 0.8, 1.2)$. Stage 1 uses both families, while Stage 2 uses the finite-margin family. Stage 2b restricts the main grid-refinement study to the finite-margin sheet and filament cases. The orientations are $Q = I$ and $Q = R_z(15^\circ)R_y(40^\circ)R_x(25^\circ)$. In the main rotated cases, $\boldsymbol{b}$ is fixed; the separate Stage-2b covariance control rotates both H_0 and $\boldsymbol{b}$, using the sheet case at $\varepsilon = 0.10$.

The local controlled-test measure is a spherical top-hat restriction with grid weights proportional to ρ and uniform cell volumes. Analytic evaluation uses the polynomial potential and its exact Hessian. Grid evaluation uses cubic B-spline interpolation of V; all requested nodes remain inside the computational box. The controlled classification threshold is $\lambda_{\mathrm{th}} = 10^{-4}L_0^{-2}$. The finite-dimensional score diagnostic uses $N = 6$ where evaluated.

For Stage 2b, the sheet cases use both listed radii and the filament cases use $R/L_0 = 1$. The covariance control uses both sheet radii.

b. Periodic mock and reconstruction. The initial Gaussian field is generated from real standard-normal grid noise using NumPy's default random generator with seed 20 260 807. Its Fourier amplitudes are multiplied by the square root of the test spectrum

$$P_{\mathrm{test}}(k) \propto \frac{k^{0.96}}{1 + (k/k_0)^4}\exp[-(k/k_c)^4], \tag{E2}$$

with $k_0 = 0.15\,h\,\mathrm{Mpc}^{-1}$, $k_c = 0.70\pi/\Delta x$, and the zero Fourier mode removed. This is a controlled spectrum with a cold dark matter (CDM)-like shape, not a precision cosmological transfer function. Its amplitude is set so that smoothing with the Fourier multiplier $\exp[-k^2(4\Delta x)^2/2]$ gives a linear density rms of 0.55. The displacement is $\widetilde{\boldsymbol{\Psi}} = i\boldsymbol{k}\widetilde{\delta}/k^2$, and the growth multiplier is unity. One particle at each of the 128^3 periodic lattice nodes is displaced and wrapped into a box of side $200\,h^{-1}\,\mathrm{Mpc}$.

Periodic cloud-in-cell assignment is followed by Gaussian smoothing of width four cells, truncated at four standard deviations, and normalization to unit mean density. A floor of 10^{-6} times the mean is specified, followed by renormalization; it is inactive in the reported tests. The potential is $V = -\ln\rho$ with its grid mean removed. The pointwise Hessian and score are computed by spectral differentiation of that grid potential, whereas off-grid transport values use a periodic cubic B-spline interpolant. The direct numerical reference and the transport estimator therefore use the same reconstructed field but different numerical derivative/interpolation operations; their measured internal residual includes the associated discretization effects. Density, potential, and derivative

grids are stored in single precision in the archived implementation; Fourier operations, interpolation coefficients, local sums, and tensor fits use double precision.

The cosmological probe uses the compact radial window $W_R(z) \propto (1-q)^4(1+4q)$ for $0 \leqslant q = |z|/R < 1$, and zero otherwise. Integer grid offsets supply the quadrature, with weights proportional to $W_R\rho$. The original direction set is the six-direction design specified in the main text. The scans are $R/\Delta x = (8, 12, 16)$ and $\eta = (1, 1.5, 2)$, using the symmetric transport estimator with $\varepsilon = \eta/(R/\Delta x)$. The classification threshold is zero and the synthetic dimension is $N = 6$. Neither is determined by the particle number.

The reference curvature scale is $H_{\rm rms} = [N_{\rm g}^{-3} \sum_{\rm grid} \|H_{\rm point}\|_{\rm F}^2]^{1/2}$. Stable points have a minimum eigenvalue margin of at least 0.05 times this scale; 24 points per class are drawn without replacement. Eight boundary points per class are drawn from the 32 smallest-margin candidates not already selected. Point selection uses seed $20\,260\,807 + 17$. These definitions refer to the pointwise classifier, not to every finite-window or angular-design classifier. The regularized per-point tensor discrepancies summarized by their median in Stage 3 have denominator $\|H_{\rm ref}(x)\|_{\rm F} + 0.05H_{\rm rms}$. They differ from the aggregate normalized Frobenius error in Eq. (147), used for the thinning and additional angular-design comparisons.

Thinning. Stage 4 holds the full-catalog realization, threshold, smoothing, and all evaluation points fixed. In replicate $b = 1, \ldots, 8$, seed $20\,260\,808 + b$ generates one uniform variate per particle; retaining variates below $f_{\rm samp}$ gives nested catalogs with $f_{\rm samp} = (1/2, 1/4, 1/8, 1/16)$. The full catalog supplies a deterministic reference. Both $\eta = 1.5$ and 2 are evaluated at the three radii, and the subsequent operating-point choice is $\eta = 2$. Each sampled catalog is normalized and reconstructed through the same pipeline, with the full-catalog threshold and curvature scale retained. Error statistics and classifications are compared with the corresponding full-sampling observable, not with continuum truth.

Tensor-error and class-agreement statistics are first computed within each thinning replicate and then summarized by medians and interquartile ranges. The transport-to-Hessian error ratio is formed within each replicate before aggregation. At full sampling, the pointwise-Hessian reconstruction error vanishes up to roundoff, so the error ratio is not used there. The angular-design audit below changes neither this thinning protocol nor its original six-direction target.

2. Controlled-displacement and boundary sensitivity

The controlled-displacement parameter $\eta = \varepsilon R/\Delta x$ balances finite-displacement truncation against interpolation resolution. Its direct diagnostic is the internal estimator error $E_{\rm internal}$ defined in Sec. VI H, whereas the end-to-end error is dominated by sampled-field reconstruction.

The stable sample supplies the primary results. The boundary sample was selected by proximity to the pointwise-Hessian threshold and need not lie near the finite-window classification boundary; it therefore tests threshold and sample-selection sensitivity rather than a universal ordering of the two observables.

Figure 7(a) isolates displacement sensitivity within the estimator, whereas panels (b) and (c) also contain density-reconstruction and classification-margin effects.

3. Morphology dependence

The morphology-conditioned results in Fig. 8 use the same stable subsets as the Stage 3 analysis. The errors $E_{\rm total}$ and $E_{\rm Hessian}$ are each evaluated relative to the corresponding full-sampling reference defined in Sec. VI H. Their ratio therefore compares normalized sampling robustness and is not a direct tensor difference between the finite-window response and the pointwise Hessian.

Because the normalization uses the Frobenius energy of each reference tensor, a smaller relative error can partly reflect a larger reference amplitude and does not by itself imply a smaller absolute fluctuation.

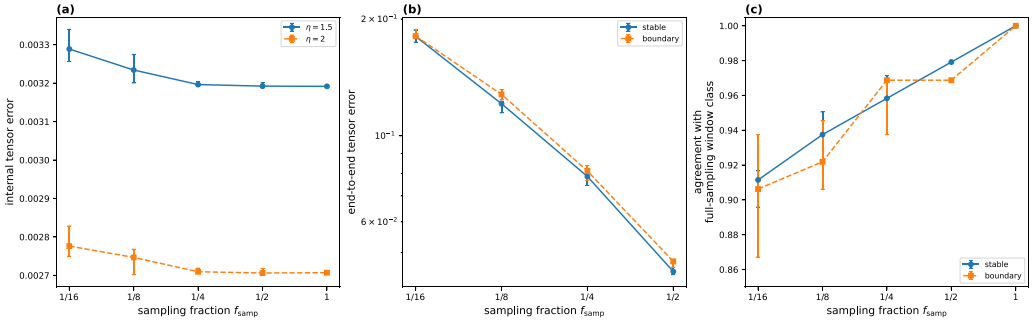


FIG. 7. Sensitivity to the controlled-displacement amplitude and evaluation-sample definition for $R/\Delta x = 8$. (a) Internal estimator error for $\eta = 1.5$ and $\eta = 2$. (b) End-to-end tensor error for the stable and boundary samples at $\eta = 2$. (c) Agreement of the finite-window transport classification with its full-sampling finite-window reference. The boundary sample is selected using the pointwise-Hessian margin and therefore does not define the finite-window classification boundary.

4. Finite-dimensional score-alignment robustness

The sampled score-alignment deficit $\widehat{\Delta K}^{(N)}$ is defined in Eq. (133). Because both its log-density gradient and transport normalization are reconstructed from the sampled field, its sampling sensitivity is distinct from the finite-transport midpoint error and the pointwise-Hessian error.

Figure 9 compares the sampled and full-sampling deficits using the normalized root-mean-square error (NRMSE), correlation, and gain through the origin, which respectively diagnose total relative deviation, preservation of spatial ordering, and multiplicative amplitude bias. The Bakry-Émery parameter N is fixed throughout the thinning sequence.

5. Classification margins and spectral stability

For $X \in \{\mathrm{window}, \mathrm{point}\}$, let $m_X^{(\mathrm{full})}(x)$ denote the minimum absolute distance of the eigenvalues of $H_X^{(\mathrm{full})}(x)$ from the common threshold λ_{th}. At a fixed sampling fraction, the plotted agreement probability is the fraction of the B thinning replicates that preserve the corresponding full-sampling class.

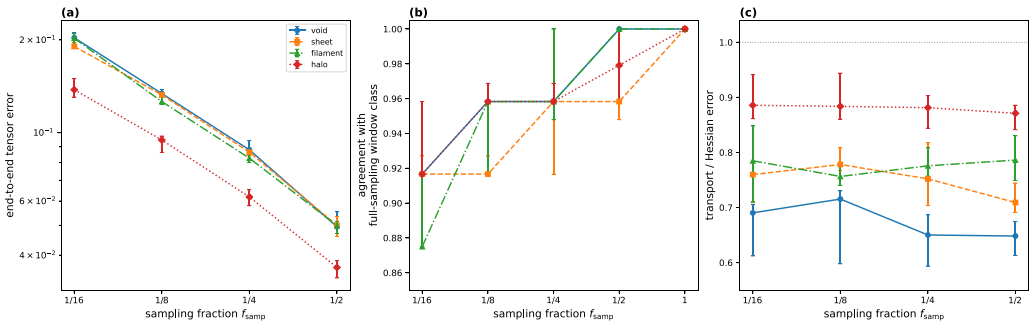


FIG. 8. Morphology dependence of the finite-sampling results for stable evaluation points, $R/\Delta x = 8$, and $\eta = 2$. (a) End-to-end finite-window transport error E_{total}. (b) Agreement with the full-sampling finite-window class. (c) Replicate-wise ratio $E_{\mathrm{total}}/E_{\mathrm{Hessian}}$, aggregated across the thinning realizations, with each observable compared with its own full-sampling reference.

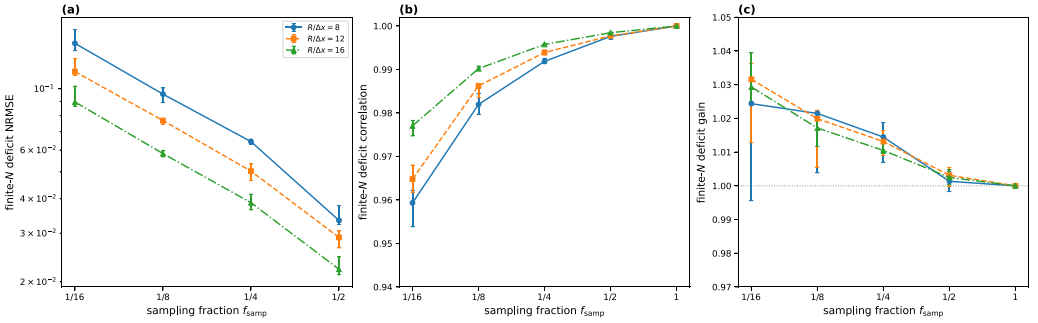


FIG. 9. Robustness of the finite-dimensional score-alignment deficit under particle thinning for the stable evaluation sample. (a) Normalized root-mean-square error (NRMSE). (b) Correlation with the full-sampling reference. (c) Gain through the origin; the horizontal line marks unit gain. The full-sampling point is omitted from panel (a), where the NRMSE vanishes identically. Here, N is the Bakry-Émery dimension parameter, not the number of retained particles.

Using the end-to-end tensor difference $D^{(b)}_{\mathrm{total}}$ defined in Eq. (153), define the worst-case spectral certificate ratio by

$$q_{\max}(x) \equiv \frac{\max_{1\leqslant b\leqslant B} \left\| D^{(b)}_{\mathrm{total}}(x) \right\|_2}{m^{(\mathrm{full})}_{\mathrm{window}}(x)}. \tag{E3}$$

Weyl's eigenvalue perturbation inequality implies that $q_{\max}(x) < 1$ is sufficient to preserve the full-sampling finite-window class in every replicate. The condition is not necessary, so points with $q_{\max} \geqslant 1$ may also remain correctly classified.

The margin analysis in Fig. 10 separates continuous tensor error from the discrete response of the threshold classifier: Near-threshold points can change class under small perturbations, whereas large-margin points tolerate larger tensor errors.

Together, the four diagnostics isolate displacement calibration, morphology-dependent sampling, reconstruction of the finite-dimensional deficit, and threshold stability. They reinforce the three-layer validation without treating the finite-window transport tensor as an estimator of the pointwise Hessian.

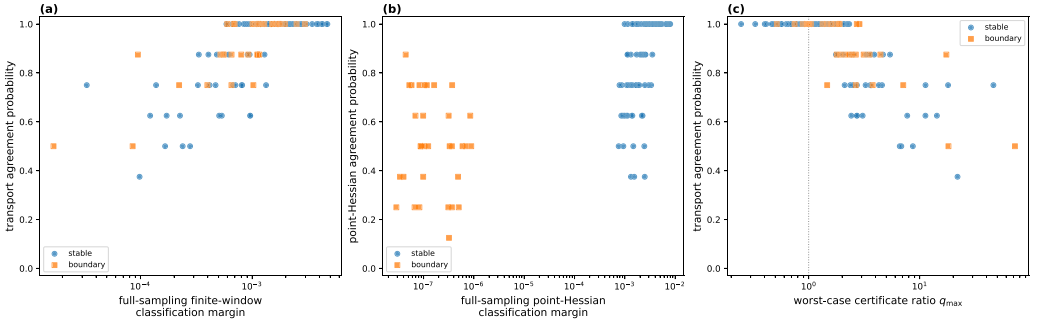


FIG. 10. Classification margins and spectral-norm stability at $f_{\mathrm{samp}} = 1/16$, $R/\Delta x = 8$, and $\eta = 2$. (a) Transport agreement probability vs the full-sampling finite-window margin. (b) Pointwise-Hessian agreement probability vs the corresponding full-sampling margin. (c) Transport agreement probability vs the worst-case certificate $q_{\max}$ of Eq. (E3). The vertical line marks $q_{\max} = 1$; points to its left are guaranteed to preserve their class in all replicates, whereas the bound makes no converse claim.

TABLE III. Fixed-field directional-design sensitivity on the 96 pointwise-stable locations. The tensor difference is $E(H^{(\mathcal{A}_{96})}_{\text{window}} \to H^{(6)}_{\text{window}}; \mathcal{I})$, normalized by the aggregate Frobenius energy of the original six-direction direct tensor. Class agreement also uses that original direct tensor as reference. The last column gives the range over 12 rotated six-direction designs. These are observable-design differences, not transport-estimator errors.

$R/\Delta x$	Tensor difference	96-axis class agreement	Changed locations	Rotated six-direction agreement range
8	0.2141	$88/96 = 0.9167$	8/96	0.7917–0.8958
12	0.4371	$74/96 = 0.7708$	22/96	0.6458–0.7917
16	0.5161	$66/96 = 0.6875$	30/96	0.6042–0.7604

6. Directional-design dependence and angular resolution

A coordinate change and a change of probe directions in a fixed density field are different operations. A covariant coordinate transformation preserves the eigenvalues of the same tensor; replacing the directional design in Eq. (136) need not preserve its finite-window eigenvalues or classes. We therefore tested design sensitivity separately from transport-estimator accuracy.

For this audit, the baseline was regenerated from the archived Stage 3 seed and settings. All 102 comparisons with the archived printed log, including the four grid-wide class counts, field diagnostics, and radius-displacement summaries, agreed at the stored print precision. This verifies the recorded baseline statistics, not bitwise identity with the original particle archive. The regenerated field, smoothing, window, threshold, and 128 evaluation locations were then fixed. The 96 pointwise-stable locations form the primary sample; the 32 pointwise-boundary locations are analyzed separately. At $R/\Delta x = (8, 12, 16)$, we compared the original six-direction design $\mathcal{E}_6$, 12 proper rotations of $\mathcal{E}_6$, and angular quadratures with 24, 96, 384, and 1536 distinct unoriented probe axes. All tensors were fitted in the same Cartesian frame. Rotations were generated using NumPy's `default_rng` with seed 20 261 021: Each four-component standard-normal vector was normalized and interpreted, in the order (w, x, y, z), as a unit quaternion. The first three rotations were also applied to the 24- and 96-axis designs.

a. Angular projection and diagnostics. For positive nodes μ_a and weights w_a of an even-order n_μ-point Gauss-Legendre rule on $[-1, 1]$, set $\phi_b = 2\pi b/n_\phi$, with $b = 0, \ldots, n_\phi - 1$. The probe axes are $\boldsymbol{e}_{ab} = (\sqrt{1-\mu_a^2}\cos\phi_b, \sqrt{1-\mu_a^2}\sin\phi_b, \mu_a)$ with weights w_a/n_ϕ. The pairs $(n_\mu, n_\phi) = (6, 8), (12, 16), (24, 32), (48, 64)$ give the four quadrature sizes, denoted $\mathcal{A}_{24}$ through $\mathcal{A}_{1536}$. Restriction to positive μ_a uses the exact evenness $k_{\text{window}}(-\boldsymbol{e}) = k_{\text{window}}(\boldsymbol{e})$; the retained weights sum to unity. The fits approximate

$$H_{\text{ang}} \equiv \underset{H=H^\top}{\arg\min} \int_{\mathbb{S}^2} [k_{\text{window}}(\boldsymbol{e}) - \boldsymbol{e}^\top H \boldsymbol{e}]^2 \mathrm{d}\sigma(\boldsymbol{e}), \tag{E4}$$

where σ is normalized surface measure and the spatial and scale arguments are suppressed. The 1536-axis fit is an auxiliary angular integration reference, not continuum-density truth or a replacement for the original six-direction observable.

Design differences use E and f_{class} from Eqs. (147) and (148), with the original direct six-direction tensor as reference. Angular convergence instead uses the direct $\mathcal{A}_{1536}$ fit as reference. To test whether a converged tensor represents the full directional response, define

$$\mathcal{R}_{\text{ang}}[H] \equiv \left[\frac{\sum_{i\in\mathcal{I}} \int_{\mathbb{S}^2} [k_i(\boldsymbol{e}) - \boldsymbol{e}^\top H_i \boldsymbol{e}]^2 \mathrm{d}\sigma}{\sum_{i\in\mathcal{I}} \int_{\mathbb{S}^2} k_i(\boldsymbol{e})^2 \mathrm{d}\sigma} \right]^{1/2}, \tag{E5}$$

where $k_i(\boldsymbol{e}) = k_{\text{window}}(x_i; R, \ell; \boldsymbol{e})$. All residuals use the same 1536-axis evaluation quadrature.

b. Design dependence and angular convergence. Table III shows appreciable design dependence that grows over the tested radius range. Averaged over the 12 rotations, the original six-direction

TABLE IV. Transport fidelity on the 96 pointwise-stable locations at $R/\Delta x = 8$ and $\eta = 2$. Each design uses its own direct reference; the error is the aggregate normalized Frobenius error.

Design	Internal error	Class agreement
$\mathcal{E}_6$	0.2707%	96/96
Rotation 1	0.3238%	96/96
Rotation 2	0.3102%	96/96
Rotation 3	0.3097%	96/96
$\mathcal{A}_{24}$	0.2969%	94/96
$\mathcal{A}_{96}$	0.2956%	95/96

class is retained in 84.64%, 71.18%, and 67.80% of the point-rotation comparisons, respectively. The pointwise-stable selection does not guarantee a large finite-window classification margin. For a reference design $\mathcal{E}$, write $m_{\mathcal{E}}(x) = \min_a |\lambda_a[H^{(\mathcal{E})}_{\text{window}}(x)] - \lambda_{\text{th}}|$. The same Weyl argument used in Appendix B guarantees class preservation under a design change whenever the spectral norm of the tensor difference is less than $m_{\mathcal{E}}(x)$. No tested pair satisfying this sufficient condition changed class; the converse is not required.

The 96-axis direct fits differ from the 1536-axis projection by aggregate relative tensor differences of 1.43×10^{-9}, 8.44×10^{-8}, and 6.11×10^{-7} at the three radii. The 384-to-1536 differences are below 2.4×10^{-14}. The unrotated and three rotated 96-axis designs yield the high-resolution projection's classes at all 96 primary locations and all three radii. These are angular-integration convergence results within the fixed reconstructed field, not accuracies relative to an underlying continuum density. Nevertheless, $\mathcal{R}_{\text{ang}}[H_{\text{ang}}]$ remains approximately 0.1094, 0.2378, and 0.2868, compared with 0.2066, 0.4536, and 0.5562 for the original six-direction fit. A stable quadratic projection therefore need not reproduce the complete directional response.

c. Estimator fidelity for each design. At $R/\Delta x = 8$ and $\eta = 2$, we separately evaluated the symmetric transport estimator for $\mathcal{E}_6$, its first three rotations, $\mathcal{A}_{24}$, and $\mathcal{A}_{96}$. Each transport tensor was compared with its own design's direct finite-window reference using the aggregate Frobenius error of Eq. (147), not the median regularized per-point error used in the Stage 3 summary.

The internal errors remain 0.2707%–0.3238% (Table IV). The three mismatched design-location pairs fail the corresponding spectral-norm class-preservation certificate; no certified pair changes class. All six tested designs preserve their own direct-reference classes at all 32 pointwise-boundary locations. Small estimator error and appreciable design dependence thus coexist without contradiction.

The original thinning results remain statements about the fixed six-direction observable and its full-sampling reference. Their sampling advantage has not been established for every angular design, since the multidirection thinning experiments were not repeated. Applications seeking a less orientation-sensitive quadratic summary should specify and check an angular quadrature while retaining the nonquadratic residual as a separate diagnostic.

[1] P. J. E. Peebles, *The Large-Scale Structure of the Universe* (Princeton University Press, Princeton, 1980).
[2] A. P. Fairall, *Large-Scale Structures in the Universe*, Wiley-Praxis Series in Astronomy and Astrophysics (Wiley, Chichester, 1998).
[3] T. T. Takeuchi, *Physics of the Formation and Evolution of Galaxies: Multiwavelength Point of View*, Springer Series in Astrophysics and Cosmology (Springer Nature, Singapore, 2025).
[4] B. J. T. Jones, V. J. Martínez, E. Saar, and V. Trimble, Scaling laws in the distribution of galaxies, Rev. Mod. Phys. **76**, 1211 (2005).

[5] V. J. Martínez and E. Saar, *Statistics of the Galaxy Distribution* (Chapman & Hall/CRC, Boca Raton, 2002).
[6] T. T. Takeuchi, *Applications of Big Data and Machine Learning in Galaxy Formation and Evolution*, Series in Astronomy and Astrophysics (CRC Press, Boca Raton, 2025).
[7] O. Hahn, C. Marcella Carollo, C. Porciani, and A. Dekel, The evolution of dark matter halo properties in clusters, filaments, sheets and voids, Mon. Not. R. Astron. Soc. **381**, 41 (2007).
[8] J. E. Forero-Romero, Y. Hoffman, S. Gottlöber, A. Klypin, and G. Yepes, A dynamical classification of the cosmic web, Mon. Not. R. Astron. Soc. **396**, 1815 (2009).
[9] R. van de Weygaert and E. Platen, Cosmic voids: Structure, dynamics and galaxies, Int. J. Mod. Phys.: Conf. Ser. **01**, 41 (2011).
[10] Y. Hoffman, O. Metuki, G. Yepes, S. Gottlöber, J. E. Forero-Romero, N. I. Libeskind, and A. Knebe, A kinematic classification of the cosmic web, Mon. Not. R. Astron. Soc. **425**, 2049 (2012).
[11] F.-S. Kitaura and F. Sinigaglia, Spectral hierarchy of the cosmic web, arXiv:2603.15834 [astro-ph.CO].
[12] C. Villani, *Optimal Transport: Old and New*, *Grundlehren der mathematischen Wissenschaften* (Springer, Berlin, 2009).
[13] F. Santambrogio, *Optimal Transport for Applied Mathematicians: Calculus of Variations, PDEs, and Modeling*, Progress in Nonlinear Differential Equations and Their Applications (Springer International Publishing, Cham, Switzerland, 2015).
[14] U. Frisch, S. Matarrese, R. Mohayaee, and A. Sobolevski, A reconstruction of the initial conditions of the universe by optimal mass transportation, Nature (London) **417**, 260 (2002).
[15] Y. Brenier, U. Frisch, M. Hénon, G. Loeper, S. Matarrese, R. Mohayaee, and A. Sobolevskiĭ, Reconstruction of the early universe as a convex optimization problem, Mon. Not. R. Astron. Soc. **346**, 501 (2003).
[16] F. Nikakhtar, R. K. Sheth, N. Padmanabhan, B. Lévy, and R. Mohayaee, Displacement field analysis via optimal transport: Multitracer approach to cosmological reconstruction, Phys. Rev. D **109**, 123512 (2024).
[17] S. Sartori, S. Contarini, E. Sarpa, G. Degni, F. Marulli, S. Escoffier, and L. Moscardini, The Back-in-time Void Finder: dynamical identification of cosmic voids through optimal transport reconstruction, Astron. Astrophys. **714**, A293 (2026).
[18] H. Edelsbrunner and J. L. Harer, *Computational Topology: An Introduction* (American Mathematical Society, Providence, RI, 2010).
[19] M. Kerber, D. Morozov, and A. Nigmetov, Geometry helps to compare persistence diagrams, ACM J. Exp. Algorithmics **22**, 1 (2017).
[20] M. Tsizh, V. Tymchyshyn, and F. Vazza, Wasserstein distance as a new tool for discriminating cosmologies through the topology of large-scale structure, Mon. Not. R. Astron. Soc. **522**, 2697 (2023).
[21] T. T. Takeuchi, Wasserstein distance in cosmological structure formation: An optimal transport perspective, Phys. Rev. D **114**, 023555 (2026).
[22] J. Lott and C. Villani, Ricci curvature for metric-measure spaces via optimal transport, Ann. Math. **169**, 903 (2009).
[23] K.-T. Sturm, On the geometry of metric measure spaces, Acta Math. **196**, 65 (2006).
[24] K.-T. Sturm, On the geometry of metric measure spaces. II, Acta Math. **196**, 133 (2006).
[25] F. Otto, The geometry of dissipative evolution equations: The porous medium equation, Commun. Partial Differ. Equations **26**, 101 (2001).
[26] R. Jordan, D. Kinderlehrer, and F. Otto, The variational formulation of the Fokker–Planck equation, SIAM J. Math. Anal. **29**, 1 (1998).
[27] Y. Ollivier, Ricci curvature of Markov chains on metric spaces, J. Funct. Anal. **256**, 810 (2009).
[28] Y. Ollivier, A survey of Ricci curvature for metric spaces and Markov chains, in *Probabilistic Approach to Geometry*, edited by M. Kotani, M. Hino, and T. Kumagai, Advanced Studies in Pure Mathematics Vol. 57 (Mathematical Society of Japan, Tokyo, 2010), pp. 343–381.
[29] N. Garcia Trillos and M. Weber, Continuum Limits of Ollivier's Ricci Curvature on data clouds: pointwise consistency and global lower bounds, Discrete Comput. Geom. **76**, 1572 (2026).
[30] G. Monge, Mémoire sur la théorie des déblais et des remblais, *Histoire de l'Académie Royale des Sciences de Paris*, avec les Mémoires de Mathématique et de Physique pour la même année (1781), pp. 666–704.

[31] L. V. Kantorovich, On the translocation of masses, Dokl. Akad. Nauk SSSR **37**, 227 (1942).
[32] L. Ambrosio, N. Gigli, and G. Savaré, *Gradient Flows: In Metric Spaces and in the Space of Probability Measures*, Lectures in Mathematics ETH Zürich, 2nd ed. (Birkhäuser, Basel, Switzerland, 2008).
[33] D. Bakry, I. Gentil, and M. Ledoux, *Analysis and Geometry of Markov Diffusion Operators*, *Grundlehren der mathematischen Wissenschaften* (Springer, Cham, 2014), Vol. 348.
[34] S.-I. Ohta, Ricci curvature, entropy, and optimal transport, in *Optimal Transport: Theory and Applications*, edited by Y. Ollivier, H. Pajot, and C. Villani, London Mathematical Society Lecture Note Series Vol. 413 (Cambridge University Press, Cambridge, 2014), pp. 145–200.