Operational Curvature and Geodesic Structure in the VERSF Framework

Refinement Transport, Distinguishability-Density Deformation, and the Dynamical Bridge to Emergent Macroscopic Geometry

Keith Taylor VERSF Theoretical Physics Programme versf-eos.com

---

A General Reader's Summary

Where does curvature come from?

In general relativity, the geometry of the macroscopic continuum bends in the presence of mass and energy. Trajectories that would be straight in flat geometry curve toward concentrations of matter. Light rays bend around the Sun; clocks run slower in deeper gravitational wells. The Einstein equations relate this bending of geometry to the local density of energy and momentum. But the equations take the geometry itself as a primitive given — a smooth 4-manifold on which everything else lives. Why should there be a geometric continuum at all? And why should it bend in response to what it contains?

The previous paper of this programme — *Operational Geometry in the VERSF Framework* — established that the admissible subspace $\mathcal{A} = \mathrm{Im}(\Omega_{\max})$ of VERSF closure architecture is intrinsically a finite-dimensional real Hilbert space, derived from admissibility alone. That paper supplied the **kinematics** of admissibility geometry: metric, dimension, measure, volume, compactness. What it did *not* supply was the **dynamics** — how transport happens on this geometry, what role curvature plays, and how the bending of the macroscopic continuum could arise from the substrate.

The present paper develops that dynamics. The central claim is structural rather than computational:

The flat reference Hilbert geometry of admissibility is the unperturbed substrate. Local variation in distinguishability density deforms this reference conformally — by a specific factor forced by operational distinguishability counting — generating a curved Riemannian geometry whose curvature is bounded by finite-packing constraints. Refinement transport on this curved geometry obeys natural geodesic and continuity equations. Under coarse-graining at scales much larger than the operational distinguishability quantum, this curved geometry approaches a $C^{1,\alpha}$ Riemannian limit — the structural candidate for the emergent 4-manifold of macroscopic physics.

In short: the *flat* Hilbert geometry of admissibility is the substrate; *distinguishability-density gradients* are what bend it; *finite packing* is what bounds the bending; and *coarse-graining* is the bridge to the smooth macroscopic continuum. Curvature is not assumed — it is generated by inhomogeneity in the local distinguishability density on the admissibility substrate, and the form of the operational metric is forced by what "operational distance" means.

This does not derive general relativity. The conjecture that the macroscopic continuum is the coarse-grained limit of operational distinguishability geometry, and that the Einstein equations are the coarse-grained limit of refinement-transport dynamics, is left open (§14). What this paper establishes is that the *structural ingredients* for such a derivation — operational metric (derived), geodesics, connection, curvature tensor, bounded curvature, continuity equation, conformal deformation, coarse-graining bridge — are all available on the admissibility substrate, derived from prior VERSF results, without importing macroscopic geometry as input.

The results are conditional on a single load-bearing structural hypothesis — the **kernel-coverage hypothesis** — which requires that the admissible sector set $\Sigma(M)$ covers the operational image M_{op} densely enough that local distinguishability density never vanishes. This is the precise mathematical condition under which the operational geometry is well-defined; closure manifolds violating it lie outside the scope of the framework.

The technical body of the paper follows. Readers without a mathematical-physics background may wish to read §1 (Introduction), §5 (Distinguishability Density and Conformal Deformation), and §12 (Relation to the Emergent 4-Manifold), which together convey the essential content in continuous prose, before consulting the theorems and proofs.

Abstract

The previous paper *Operational Geometry in the VERSF Framework* established that the admissible subspace $\mathcal{A} = \mathrm{Im}(\Omega_{max})$ of the closure architecture is intrinsically a finite-dimensional real Hilbert space $(\mathcal{A}, \langle \cdot,\cdot \rangle_{op})$, with metric D_{op}, finite operational dimension d_{op}, operational Hausdorff measure $\mathcal{H}^{d_{op}}_{op}$, finite operational volume $\mathrm{Vol}_{op}(M)$, non-expansive refinement projection Ω_{max}, total boundedness of the operational image $M_{op} = \Omega_{max}(M)$, and the packing ceiling $P(M) \leq \mathrm{Vol}_{op}(M) / \Delta_{op}^{d_{op}}$ as a one-step Euclidean sphere-packing corollary.

That paper supplied the **kinematics** of admissibility geometry. The present paper develops the **dynamics**. The structural framing in four phrases: the flat reference Hilbert geometry of admissibility is the substrate; distinguishability-density gradients bend it conformally; finite packing bounds the bending; coarse-graining is the bridge to the macroscopic continuum.

The structural backbone is a two-step derivation. First (Theorem 5.0): operational distance — the count of local distinguishability quanta along a path — *forces* the operational metric to be the conformally rescaled form

$$g_{op}(x) = \rho_{op}(x)^{2/d_{op}} \cdot \langle \cdot,\cdot \rangle_{op},$$

where ρ_{op} is the **operational distinguishability density**, defined as the smoothed packing density on M_{op} at coarse-graining scale $\geq \Delta_{op}$. Under three structural principles — $\mathbb{Z}_7$-equivariance (substrate symmetry), channel-uniformity (scalar source ⇒ channel-symmetric response), and locality (Remark 5.0.3) — the factor $\rho_{op}^{2/d_{op}}$ is forced, not postulated: it is the unique conformal factor consistent with local distinguishability counting within the conformal class, and the conformal restriction itself is supplied by the first two principles. Second (Theorem 5.2): the conformally rescaled g_{op} is a smooth Riemannian metric whose Riemann curvature vanishes when ρ_{op} is constant, and is generated by gradients of ρ_{op} when it varies.

Every other dynamical structure of the paper follows from these two structural inputs:

1. Refinement-stable admissible transport — flows $\Phi_\tau : \mathcal{A} \to \mathcal{A}$ indexed by refinement parameter τ — is defined as the 1-Lipschitz Ω_{max}-stable semigroup on $(\mathcal{A}, \langle\cdot,\cdot\rangle_{op})$, with the entropy-gradient flow as natural physical realisation under the S_{op}-concavity hypothesis (Definition 4.1, Theorem 4.2).
2. Operational geodesics under g_{op} exist locally and minimise g_{op}-length between sufficiently close interior points; full geodesic completeness holds on closed completions (Theorem 6.2, Theorem 6.3).
3. The Levi-Civita connection ∇_{op}^{LC} of g_{op} exists by the fundamental theorem of Riemannian geometry; its Christoffel symbols are given in closed form by derivatives of $\ln \rho_{op}$ (Theorem 7.1).
4. The conformal-flat curvature formula gives R_{op} in closed form as a function of $\nabla^2(\ln \rho_{op})$ and $\nabla(\ln \rho_{op}) \otimes \nabla(\ln \rho_{op})$ (Theorem 8.1). Curvature is generated by distinguishability-density gradients; constant ρ_{op} gives $R_{op} = 0$, while non-constant ρ_{op} generically produces non-trivial curvature (Theorem 8.2). The converse direction $R_{op} = 0 \Rightarrow \rho_{op}$ constant requires excluding non-trivial Möbius-type Liouville solutions via additional admissibility conditions.
5. Finite packing $P(M) \leq Vol_{op}(M) / \Delta_{op}^{d_{op}}$ bounds the maximum gradient $\|\nabla \ln \rho_{op}\| \leq C_1/\Delta_{op}$ pointwise on M_{op}, hence bounds the operational curvature pointwise: $\|R_{op}(x)\| \leq C / \Delta_{op}^2$ (Theorem 9.1), where the constant C depends on dimension, kernel choice, and the density contrast ratio $\rho_{op}^{max}/\rho_{op}^{min}$ across M_{op}.
6. Refinement transport obeys the operational continuity equation $\partial_\tau \rho_{op} + \nabla_{op} \cdot J_{op} = 0$, with $J_{op} = \rho_{op} v_{op}$ the natural pushforward flux (Theorem 10.2).
7. Under coarse-graining at scales $L \gg \Delta_{op}$, the conformally-deformed g_{op}-geometry satisfies the three Cheeger–Gromov hypotheses (bounded curvature, bounded diameter, volume non-collapsing — the last derived from kernel-coverage in Lemma 11.0), giving subsequential GH-convergence to a $C^{1,\alpha}$ Riemannian limit with the substrate-level curvature bound preserved (Theorem 11.2).

Standing hypotheses. All structural results of this paper are conditional on the **kernel-coverage hypothesis** (Definition 3.3): every $x \in M_{op}$ lies within D_{op}-distance $\varepsilon - \varepsilon'$ of some admissible sector $s \in \Sigma(M)$, where ε is the coarse-graining scale and $\varepsilon' \in (0, \varepsilon)$ is the inner radius of the kernel-positive sub-ball. This is the precise structural condition (strictly stronger than topological ε-covering) needed to bound the operational distinguishability density ρ_{op} below by a positive constant uniformly on M_{op}, supporting smoothness of g_{op}, finite curvature bounds, and

Cheeger–Gromov non-collapsing. Closure manifolds violating kernel-coverage admit regions of vanishing ρ_{op} and degenerate g_{op}, falling outside the scope of the present framework.

The result upgrades the VERSF programme from **operational geometry** to **operational geometric dynamics**: admissibility-induced geometry is not merely a kinematic backdrop but a substrate carrying its own dynamical structure (transport, geodesics, connection, curvature) — a structure whose coarse-grained limit is the natural candidate for the emergent 4-manifold geometry of macroscopic physics.

Contents

---

1. Introduction

The previous paper of this programme established that admissibility induces a finite-dimensional real Hilbert geometry on the closure manifold. The structural backbone of that paper was the inner-product theorem (Theorem 4.0 of *Operational Geometry*): the $\mathbb{Z}_7$ Fourier architecture of VERSF closure dynamics endows the admissible subspace $\mathcal{A} = \mathrm{Im}(\Omega_{max})$ with a natural real inner product $\langle\cdot,\cdot\rangle_{op}$, under which $(\mathcal{A}, \langle\cdot,\cdot\rangle_{op})$ is a finite-dimensional real Hilbert space, the spectral-channel decomposition $\mathcal{A} = \bigoplus_\alpha V_\alpha$ is an orthogonal direct sum, and Ω_{max} is the orthogonal projection from the ambient closure-state-space onto $\mathcal{A}$.

Every kinematic geometric property of the closure manifold — operational metric, finite operational dimension, operational Hausdorff measure, finite operational volume, total boundedness, finite covering numbers, finite distinguishability packing — followed as a standard consequence of this Hilbert structure.

That paper established the **kinematics**. It did not address the **dynamics**: how transport happens on the operational geometry, what role curvature plays, and whether the geometric structure of macroscopic physics could arise as a coarse-grained limit of admissibility dynamics.

The present paper develops those questions. The structural picture has three layers:

Layer 1: Flat reference Hilbert geometry. The unperturbed admissible subspace $(\mathcal{A}, \langle\cdot,\cdot\rangle_{op})$ is a flat finite-dimensional real Hilbert space with vanishing curvature. This is the substrate.

Layer 2: Refinement-deformed Riemannian geometry. The previous paper measured operational distance using the flat reference inner product directly. The present paper observes that this is appropriate only when admissible sectors are uniformly distributed across $\mathcal{A}$. When they are not — when the **operational distinguishability density** $\rho_{op}(x)$ varies across M_{op} — operational distance is no longer the flat $\langle\cdot,\cdot\rangle_{op}$ distance, because two points separated by the same flat distance count *more* distinguishability quanta in regions of denser packing. The **Operational Distinguishability Metric Theorem** (Theorem 5.0) makes this precise: the operational metric forced by local distinguishability counting is

$$g_{op}(x) = \rho_{op}(x)^{2/d_{op}} \cdot \langle\cdot,\cdot\rangle_{op}.$$

The factor $\rho_{op}^{2/d_{op}}$ is not a free choice. Under the substrate-level principles of $\mathbb{Z}_7$-equivariance, channel-uniformity, and locality, it is the unique conformal factor under which g_{op}-distance counts local distinguishability quanta correctly. The Riemann curvature tensor of g_{op} is non-trivial precisely when ρ_{op} varies, and is given in closed form by the standard conformal-flat curvature formula. Finite packing — Theorem 10.1 of the previous paper, $P(M) \leq \mathrm{Vol}_{op}(M) / \Delta_{op}^{d_{op}}$ — bounds the gradient of $\ln \rho_{op}$ pointwise, hence bounds the curvature pointwise on M_{op}.

Layer 3: Coarse-grained emergent macroscopic geometry. Averaging over scales $L \gg \Delta_{op}$ smooths the discrete packing into a continuous Riemannian limit. The coarse-grained operational geometry (M_{op}, g_{op}) converges to a $C^{1,\alpha}$ Riemannian limit (Cheeger–Gromov) with bounded curvature, geodesic transport, and continuity conservation — the natural structural candidate for the emergent 4-manifold of macroscopic physics. The full derivation of the macroscopic continuum from this limit is left open (§14); what this paper establishes is the availability of the structural ingredients.

The paper proceeds in seven stages:

§3–§4 set up definitions and the refinement-flow semigroup Φ_τ on the operational Hilbert geometry, with refinement parameter τ replacing time (which is emergent in VERSF). §5 derives the operational metric g_{op} from distinguishability counting (Theorem 5.0), introduces ρ_{op} (Definitions 3.2–3.3), and establishes the conformal-deformation theorem (Theorem 5.2). §6–§8 develop geodesics, Levi-Civita connection, and the curvature tensor under g_{op}, each as a one-step consequence of finite-dim Riemannian standard results. §9 derives the pointwise curvature bound from finite packing. §10 establishes the refinement continuity equation. §11 outlines the coarse-graining limit toward macroscopic geometry. §12 sketches the relation to the emergent 4-

manifold of macroscopic physics, with §13–§14 listing falsifiable predictions and open problems.

The previous paper said: *admissibility geometrises*. The present paper says: *and operational distance, measured by distinguishability count, bends the geometry — generating dynamics*.

2. Prior Structural Results

The argument depends on the seven structural results of the previous paper *Operational Geometry in the VERSF Framework*, summarised here for self-containment. We refer to that paper as **OG** below for brevity.

2.1 Operational Hilbert Geometry (OG Theorem 4.0)

The $\mathbb{Z}_7$ Fourier architecture of VERSF closure dynamics endows the admissible subspace $\mathcal{A} = \mathrm{Im}(\Omega_\text{max})$ with a natural real inner product $\langle\cdot,\cdot\rangle_\text{op}$ under which:

- $(\mathcal{A}, \langle\cdot,\cdot\rangle_\text{op})$ is a finite-dimensional real Hilbert space of dimension d_op;
- the spectral-channel decomposition $\mathcal{A} = \bigoplus_\alpha V_\alpha$ is an orthogonal direct sum;
- Ω_max is the orthogonal projection from the ambient closure-state-space onto $\mathcal{A}$.

2.2 Operational Metric (OG Theorem 4.1, Corollary 4.2)

The norm-induced metric

$$D_\text{op}(x, y) := \|x - y\|_\text{op} = \sqrt{\langle x - y, x - y\rangle_\text{op}}$$

is a metric on $\mathcal{A}$. $(\mathcal{A}, D_\text{op})$ is a d_op-dimensional Euclidean Hilbert space satisfying the standard ball-volume identity $\mathrm{vol}_\text{op}(B(x, r)) = \omega_{d_\text{op}} \cdot r^{d_\text{op}}$.

2.3 Non-Expansive Refinement Projection (OG Theorem 5.1)

Ω_max is non-expansive on the ambient closure-state-space, with equality on $\mathcal{A}$ where it acts as the identity.

2.4 Finite Operational Volume (OG Theorem 8.1)

For closure manifolds M with bounded closure support, the operational image $M_\text{op} := \Omega_\text{max}(M) \subseteq \mathcal{A}$ is bounded in $(\mathcal{A}, D_\text{op})$, and

$$\mathrm{Vol}_\text{op}(M) = (2^{d_\text{op}} / \omega_{d_\text{op}}) \cdot \mathcal{H}^{d_\text{op}}_\text{op}(M_\text{op}) < \infty.$$

2.5 Total Boundedness and Precompactness (OG Theorem 9.1)

M_op is totally bounded under D_op; its closure M̄_op in 𝒜 is compact. Finite covering numbers N(ε) ≲ Vol_op(M) / ε^{d_op}.

2.6 Finite Distinguishability Packing (OG Theorem 10.1)

The admissible packing number satisfies

P(M) ≤ Vol_op(M) / Δ_op^{d_op},

with Δ_op > 0 the universal distinguishability quantum (OG §2.3).

2.7 Continuous–Discrete Distinction (OG §3)

We carry forward the OG distinction:

- M_op = Ω_max(M) ⊆ 𝒜 is the **continuous** bounded operational image, carrying ℋ^{d_op}_op and Vol_op;
- Σ(M) ⊆ M_op is the **finite discrete** subset of admissible sectors, with |Σ(M)| ≤ P(M).

The Hausdorff measure and the conformal deformation of this paper live on continuous M_op; the discrete Σ(M) is what gets counted by packing bounds.

---

3. Definitions

We adopt all definitions of **OG §3** (closure manifold M, operational quotient Q(M), admissible sector set Σ(M), admissible subspace 𝒜, operational image M_op, distinguishability quantum Δ_op, operational topology τ_op, operational Hausdorff measure ℋ^{d_op}_op, operational volume Vol_op(M)).

The following definitions are new to the present paper.

Definition 3.1 (Refinement parameter). The **refinement parameter** τ ∈ [0, ∞) is the substrate-internal scalar parameter indexing commitment ledger advance under admissible refinement dynamics. τ is a substrate primitive, not the macroscopic time coordinate (which is emergent in VERSF). Throughout this paper, all dynamical statements refer to τ-evolution.

Remark 3.1.1. *On the term "refinement".* Within VERSF, "refinement" denotes substrate-internal commitment ledger advance under admissible dynamics — *not* increase of operational resolution in the colloquial sense. Indeed, under the entropy-gradient refinement flow (Theorem 4.2), the operational distinguishability density ρ_op evolves toward the uniform equilibrium 1/Vol_op(M), at which the operational geometry becomes *flat* (Theorem 8.2(i), Remark 10.4.1). Refinement-flow τ-evolution therefore drives the operational geometry toward flatness, not toward finer-grained structure. The terminology is VERSF-internal and refers to ledger-advance dynamics in the substrate architecture.

Definition 3.2 (Operational density of states). Fix a coarse-graining radius $\varepsilon \geq \Delta_{op}$. The **operational density of states** at coarse-graining scale ε is the function $\eta_{op}^{\varepsilon} : M_{op} \to \mathbb{R}_{\geq 0}$ defined by the smoothed packing count

$$\eta_{op}^{\varepsilon}(x) := |\Sigma(M) \cap B_{op}(x, \varepsilon)| / \mathrm{vol}_{op}(B_{op}(x, \varepsilon)) = |\Sigma(M) \cap B_{op}(x, \varepsilon)| / (\omega_{d_{op}} \cdot \varepsilon^{d_{op}}),$$

where $B_{op}(x, \varepsilon)$ is the open D_{op}-ball of radius ε around $x \in M_{op}$.

Remark 3.2.1. η_{op}^{ε} counts the number of admissible sectors per unit reference operational volume in a ε-neighbourhood of x. At the natural choice $\varepsilon = \Delta_{op}$, finite packing gives the uniform upper bound

$$\eta_{op}^{\varepsilon}(x) \leq P(B_{op}(x, \varepsilon)) / (\omega_{d_{op}} \cdot \Delta_{op}^{d_{op}}) \leq C_{pack}(d_{op}) / \Delta_{op}^{d_{op}}$$

with $C_{pack}(d_{op})$ an order-unity geometric constant (the d_{op}-dimensional sphere-packing density).

Definition 3.3 (Kernel-coverage hypothesis). A closure manifold M satisfies the **kernel-coverage hypothesis** at coarse-graining scale ε and inner radius $\varepsilon' \in (0, \varepsilon)$ — relative to a smooth mollifier kernel K_{ε} bounded below by $K_* > 0$ on the sub-ball $B(0, \varepsilon') \subset \mathrm{supp}\, K_{\varepsilon}$ — if every $x \in M_{op}$ lies within D_{op}-distance $\varepsilon - \varepsilon'$ of some $s \in \Sigma(M)$:

$$M_{op} \subseteq \bigcup_{s \in \Sigma(M)} B_{op}(s, \varepsilon - \varepsilon').$$

Equivalently: every $x \in M_{op}$ has at least one sector $s \in \Sigma(M)$ such that $K_{\varepsilon}(x - s) \geq K_*$ (since $|x - s| \leq \varepsilon - \varepsilon'$ implies $x - s$ lies in the kernel-positive sub-ball $B(0, \varepsilon')$).

This is the structural condition needed to bound ρ_{op} below by a positive constant uniformly on M_{op}. It is strictly stronger than the bare topological covering condition "every x is within ε of some s," because compactly supported mollifiers vanish at the boundary of their support; a positive lower bound on ρ_{op} requires sectors to lie inside the kernel-positive interior of K_{ε}, not merely inside its full support. We assume this hypothesis throughout the paper at working scale $\varepsilon \geq \Delta_{op}$, with inner radius ε' a kernel-dependent parameter, and note explicitly where it is invoked.

Remark 3.3.1. *Why the topological ε-coverage is not enough.* If every $x \in M_{op}$ were only within ε of some s (the weaker topological condition), then $x - s$ could lie at the boundary of the kernel support where K_{ε} vanishes. The mollified density $\rho_{op}(x)$ would then receive contributions only from sectors *outside* the kernel's positive interior, possibly summing to zero. The kernel-coverage hypothesis with $\varepsilon' > 0$ closes this gap by requiring sectors to lie inside the kernel-positive sub-ball, where $K_{\varepsilon} \geq K_* > 0$.

Definition 3.4 (Operational distinguishability density). Under the kernel-coverage hypothesis (Def 3.3), the **operational distinguishability density** is the mollified, normalised density

$$\rho_{op}(x) := (1 / \mathcal{Z}) \cdot \int_{M_{op}} K_\varepsilon(x, y) \cdot \eta_{op}^{\varepsilon}(y) \cdot d\mathcal{H}^{d_{op}}_{op}(y),$$

where K_ε is a smooth radial mollifier of scale $\varepsilon \geq \Delta_{op}$ (e.g. a Gaussian truncated at radius 3ε, or any standard d_{op}-dimensional smoothing kernel of unit $\int K_\varepsilon \, d\mathcal{H} = 1$), and $\mathcal{Z}$ is the normalisation constant given by integration of η_{op}^{ε} against $\mathcal{H}^{d_{op}}_{op}$. By Fubini:

$$\int_{M_{op}} \eta_{op}^{\varepsilon}(y) \cdot d\mathcal{H}^{d_{op}}_{op}(y) = \int_{M_{op}} \left[(1 / \mathrm{vol}_{op}(B(y, \varepsilon))) \cdot \sum_{s \in \Sigma(M)} \mathbb{1}_{B(s, \varepsilon)}(y)\right] \cdot d\mathcal{H}^{d_{op}}_{op}(y) = \sum_{s \in \Sigma(M)} \int_{B(s, \varepsilon) \cap M_{op}} d\mathcal{H}^{d_{op}}_{op}(y) / \mathrm{vol}_{op}(B(s, \varepsilon)) = |\Sigma(M)|$$

(the last equality assuming $B(s, \varepsilon) \subseteq M_{op}$ for all $s \in \Sigma(M)$, so that each sector ball contributes a full $\omega_{d_{op}} \varepsilon^{d_{op}}$ volume to the integral; this **interior-sector condition** is technically distinct from the kernel-coverage hypothesis of Def 3.3 — which requires every $x \in M_{op}$ to be near some sector, but not conversely that every sector be interior to M_{op}. For sectors s near ∂M_{op}, $B(s, \varepsilon)$ extends outside M_{op} and η_{op}^{ε} picks up boundary truncation, giving $\mathcal{Z} = |\Sigma(M)| - \delta_\partial$ for a boundary correction $\delta_\partial \geq 0$. The boundary correction is negligible when M_{op} is large compared to ε in the relevant boundary geometry, and we work in this regime henceforth; otherwise $\mathcal{Z}$ is replaced by its precise integral value in the normalisation, leaving the structural results unchanged). Hence $\mathcal{Z} = |\Sigma(M)|$ (up to negligible boundary correction), and

$$\int_{M_{op}} \rho_{op}(x) \cdot d\mathcal{H}^{d_{op}}_{op}(x) = 1,$$

so ρ_{op} is a smooth probability-density-like function on M_{op}.

Remark 3.4.1. The mollification scale $\varepsilon \geq \Delta_{op}$ ensures ρ_{op} is smooth at scales $\geq \Delta_{op}$ — i.e. at all *resolved* operational scales. The discrete packing structure of $\Sigma(M)$ is invisible to ρ_{op} at scales below ε; above ε, ρ_{op} encodes the local density of admissible sectors as a smooth function. This is the natural and physical density: where admissibility packs sectors tightly, ρ_{op} is large; where admissibility is sparse, ρ_{op} is small.

Remark 3.4.2. *Convention on the smoothing scale.* Throughout the remainder of the paper, the choice $\varepsilon = \Delta_{op}$ is the *default* smoothing scale unless explicitly stated otherwise. With this convention, all derivative bounds, curvature bounds, and the resulting pointwise estimates of §9 take their sharpest (worst-case) form. Specifically, mollifier-derivative estimates of the form $\|\nabla K_\varepsilon\| \leq C/\varepsilon$ and $\|\nabla^2 K_\varepsilon\| \leq C/\varepsilon^2$ become $\|\nabla K_{\Delta_{op}}\| \leq C/\Delta_{op}$ and $\|\nabla^2 K_{\Delta_{op}}\| \leq C/\Delta_{op}^2$ at the default scale, and the curvature bound of Theorem 9.1 takes the form $\|R_{op}\| \leq C/\Delta_{op}^2$. At larger $\varepsilon > \Delta_{op}$, the corresponding bounds soften to $\|R_{op}\| \leq C/\varepsilon^2$ accordingly; the C/Δ_{op}^2 form represents the substrate-cutoff worst case at the finest resolved scale. The coarse-graining family of §11 carries each $L > \Delta_{op}$ explicitly, with L-dependent bounds; the uniform-in-L family bound is achieved as $L \downarrow \Delta_{op}^{+}$.

Definition 3.5 (Operational gradient and Laplacian). The **operational gradient** ∇_{op} of a smooth function $f : \mathcal{A} \to \mathbb{R}$ is the gradient with respect to the reference Hilbert structure $(\mathcal{A}, \langle \cdot,\cdot \rangle_{op})$:

$$\langle \nabla_{op} f, v \rangle_{op} := df(v) \text{ for all } v \in T_x \mathcal{A} \cong \mathcal{A}.$$

In coordinates dual to an orthonormal basis of $\mathcal{A}$, $\nabla_{op} = \sum_i e_i \partial_i$. The corresponding **operational Laplacian** with respect to the flat reference is denoted

$$\bar{\Delta}_{op} f := \sum_i \partial_i^2 f.$$

The notation $\bar{\Delta}_{op}$ (with overbar) distinguishes the Laplacian from the scalar distinguishability quantum Δ_{op} (no overbar) of OG §2.3.

Definition 3.6 (Conformal factor). The **operational conformal factor** is

$$\omega_{op}(x) := \rho_{op}(x)^{1/d_{op}}.$$

Definition 3.7 (Conformally deformed operational metric). The **conformally deformed operational metric** is the Riemannian metric on M_{op} defined by

$$g_{op}(x)(v, w) := \omega_{op}(x)^2 \cdot \langle v, w \rangle_{op} = \rho_{op}(x)^{2/d_{op}} \cdot \langle v, w \rangle_{op},$$

for $v, w \in T_x \mathcal{A}$.

Remark 3.7.1. g_{op} is a smooth, position-dependent rescaling of the reference Hilbert inner product. The volume form of g_{op} is

$$dV_{g_{op}}(x) = \omega_{op}(x)^{d_{op}} \cdot d\mathcal{H}^{d_{op}}_{op}(x) = \rho_{op}(x) \cdot d\mathcal{H}^{d_{op}}_{op}(x),$$

so integrating against $dV_{g_{op}}$ is integrating against the natural measure with density ρ_{op}. The motivation for the specific exponent $2/d_{op}$ is supplied by Theorem 5.0 below: it is forced by operational distinguishability counting, not chosen for volume-form convenience.

Definition 3.8 (Operational length functional). For a piecewise-C^1 path $\gamma : [0, 1] \to M_{op}$, the **$g_{op}$-length** of γ is

$$L_{op}(\gamma) := \int_0^1 \sqrt{g_{op}(\gamma(\tau))(\dot{\gamma}(\tau), \dot{\gamma}(\tau))}\, d\tau = \int_0^1 \omega_{op}(\gamma(\tau)) \cdot \|\dot{\gamma}(\tau)\|_{op}\, d\tau.$$

When ρ_{op} is constant, ω_{op} is constant and L_{op} reduces to a constant multiple of the flat Hilbert length.

4. Refinement Flow on the Operational Hilbert Geometry

We introduce the substrate dynamics. The refinement-parameter τ generates a one-parameter family of admissible self-maps of the operational geometry — the **refinement flow**. Following the operational-admissibility convention of the previous paper, we take *non-expansiveness* as the defining property of refinement-stable admissible transport.

Definition 4.1 (Refinement flow). A **refinement flow** on $(\mathcal{A}, \langle\cdot,\cdot\rangle_{\mathrm{op}})$ is a one-parameter semigroup

$$\{\Phi_\tau : \mathcal{A} \to \mathcal{A}, \tau \in [0, \infty)\}$$

satisfying:

- **Identity at $\tau = 0$:** $\Phi_0 = \mathrm{id}$.
- **Semigroup property:** $\Phi_{\tau_1 + \tau_2} = \Phi_{\tau_2} \circ \Phi_{\tau_1}$.
- **Continuity in τ:** $\tau \mapsto \Phi_\tau(x)$ is continuous for every $x \in \mathcal{A}$.
- **Non-expansiveness:** $\|\Phi_\tau(x) - \Phi_\tau(y)\|_{\mathrm{op}} \leq \|x - y\|_{\mathrm{op}}$ for all $x, y \in \mathcal{A}$, $\tau \geq 0$.
- **$\Omega_{\max}$-stability:** $\Phi_\tau(\mathcal{A}) \subseteq \mathcal{A}$ for all τ; equivalently, $\Omega_{\max} \circ \Phi_\tau = \Phi_\tau \circ \Omega_{\max}$.
- **M_{op}-stability:** $\Phi_\tau(M_{\mathrm{op}}) \subseteq M_{\mathrm{op}}$ for all $\tau \geq 0$.

The non-expansiveness condition is the substrate-level expression of operational admissibility: refinement transport cannot generate new distinguishability finer than already present in the source configuration. The M_{op}-stability condition is the substrate-level expression of operational-image conservation: admissible configurations cannot flow out of the operational image $M_{\mathrm{op}} = \Omega_{\max}(M)$ into the ambient closure-state-space, since doing so would correspond to refinement producing configurations not realizable as projections of the original closure manifold.

Remark 4.1.1. *Why M_{op}-stability is needed in addition to $\Omega_{\max}$-stability.* The $\Omega_{\max}$-stability condition $\Phi_\tau(\mathcal{A}) \subseteq \mathcal{A}$ preserves the ambient admissible *subspace* $\mathcal{A}$, but $M_{\mathrm{op}} = \Omega_{\max}(M)$ is a bounded *subset* of $\mathcal{A}$. A flow can preserve $\mathcal{A}$ while moving points off M_{op} into the complement $\mathcal{A} \setminus M_{\mathrm{op}}$ — for example, by extrapolating beyond the operational image. The continuity equation (Theorem 10.3) requires conservation of the normalisation $\int_{M_{\mathrm{op}}} \rho_{\mathrm{op}} \, d\mathcal{H} = 1$ under the flow, which holds only if $\Phi_\tau(M_{\mathrm{op}}) \subseteq M_{\mathrm{op}}$. M_{op}-stability is therefore the structural condition needed for refinement transport to preserve the operational-image-normalised density ρ_{op} as a well-defined probability density. Both natural realisations of refinement flow — the identity flow $\Phi_\tau = \mathrm{id}_{\mathcal{A}}$ and the entropy-gradient flow (Theorem 4.2) — preserve M_{op} automatically.

Theorem 4.2 (Existence of refinement flow). *On every operational Hilbert geometry $(\mathcal{A}, \langle\cdot,\cdot\rangle_{\mathrm{op}})$ with bounded admissible image M_{op}, there exists a non-trivial refinement flow Φ_τ. The trivial choice $\Phi_\tau = \mathrm{id}_{\mathcal{A}}$ (the identity flow on $\mathcal{A}$) is always admissible, since by OG Theorem 4.0(iii) $\Omega_{\max}$ acts as the identity on $\mathcal{A}$ and so $\Phi_\tau = \mathrm{id}_{\mathcal{A}}$ coincides with the constant-projector flow on the admissible subspace. The natural physical choice is the entropy-gradient flow*

$$\frac{d}{d\tau} \Phi_\tau(x) = \nabla_{\mathrm{op}} S_{\mathrm{op}}(\Phi_\tau(x)),$$

provided the operational entropy functional $S_{op} : \mathcal{A} \to \mathbb{R}$ *of OG §2.3 is concave on* $\mathcal{A}$ *(a hypothesis whose substrate-level derivation is left open; see §14, item 1).*

Proof. For the trivial choice $\Phi_\tau = \mathrm{id}_{\mathcal{A}}$, all conditions of Definition 4.1 are immediate: $\Phi_0 = \mathrm{id}$, the semigroup property holds trivially, continuity in τ holds trivially, $\|\mathrm{id}(x) - \mathrm{id}(y)\|_{op} = \|x - y\|_{op}$ gives non-expansiveness with equality, Ω_{max}-stability follows from $\Omega_{max}(\mathcal{A}) = \mathcal{A}$, and M_{op}-stability is trivial since $\mathrm{id}(M_{op}) = M_{op}$.

For the entropy-gradient flow, the standard Hilbert-space gradient-flow theory of Brezis (*Opérateurs maximaux monotones*, §3.2) gives: if $F : \mathcal{A} \to \mathbb{R}$ is a concave smooth function on a finite-dimensional real Hilbert space, then the gradient ascent flow $d\Phi_\tau/d\tau = \nabla_{op} F(\Phi_\tau)$ generates a non-expansive semigroup on $\mathcal{A}$ — equivalently, the flow of $-F$ (a convex function) is the resolvent of a maximal monotone operator, which is 1-Lipschitz. Applying this with $F = S_{op}$ (concave by hypothesis) gives a non-expansive semigroup. Ω_{max}-stability follows from S_{op} being Ω_{max}-invariant (operational entropy depends only on the admissible projection, OG §2.3): $\nabla_{op} S_{op}$ points along $\mathcal{A}$, so the flow preserves $\mathcal{A}$.

M_{op}-stability follows from boundedness of M_{op} (OG Lemma 8.0) combined with the entropy-gradient flow being attractive toward the entropy-maximum set, which lies in the interior of M_{op} (uniform-density configurations realise maximum operational entropy on the bounded operational image, OG §2.3). The flow therefore drives configurations *toward* the interior of M_{op} rather than out of it; the trajectory of every $x_0 \in M_{op}$ remains in M_{op} for all $\tau \geq 0$. Formally, this is the standard contraction property of gradient flows of concave functions on convex sublevel sets containing the maximum.

The concavity of S_{op} on $\mathcal{A}$ is the standard concavity of Shannon-type entropy functionals on probability simplices (or, more generally, on convex sets of operational configurations), and is taken as a hypothesis here. A full derivation from substrate primitives — establishing concavity of S_{op} as a corollary of admissibility — is one of the open problems of the broader programme (§14, item 1).

Remark 4.2.1. Two sources of non-expansiveness coexist for the entropy-gradient refinement flow: (i) it is the gradient flow of $-(-S_{op}) = S_{op}$, a concave function on a Hilbert space (Brezis); (ii) S_{op}-maximising configurations lie in the admissible attractor, toward which the flow converges. Both produce 1-Lipschitz semigroup behaviour.

Theorem 4.3 (Asymptotic projector recovery). *Under the concavity hypothesis of Theorem 4.2, the* $\tau \to \infty$ *limit of the entropy-gradient refinement flow projects onto the set of operational entropy maxima:*

$$\lim_{\tau \to \infty} \Phi_\tau(x) \in \operatorname{argmax}_{y \in \mathcal{A}} S_{op}(y),$$

for all $x \in \mathcal{A}$ *with bounded operational trajectory.*

Proof. By concavity of S_{op} and boundedness of M_{op} (OG Lemma 8.0), the entropy-gradient flow is gradient ascent on a concave functional with bounded sublevel sets, which converges to $\operatorname{argmax}_{y \in \mathcal{A}} S_{op}$ (standard convergence theorem; Brezis op. cit. §3.2).

---

5. Operational Metric, Distinguishability Density, and Conformal Deformation

This section contains the central structural derivation of the paper. We first derive the operational metric g_{op} from local distinguishability counting (Theorem 5.0), then establish its smoothness (Theorem 5.1), then derive the curvature-generation theorem (Theorem 5.2).

5.1 Derivation of g_{op} from Operational Distinguishability Counting

The operational metric of OG measured admissible distance using the flat reference $\langle\cdot,\cdot\rangle_{op}$ directly. That choice is appropriate when admissible sectors are uniformly distributed across $\mathcal{A}$ — i.e. when ρ_{op} is constant. When ρ_{op} varies, it is no longer appropriate: two operational points separated by the same flat distance can be separated by *different* numbers of distinguishability quanta, depending on the local density of admissible sectors.

Theorem 5.0 (Operational Distinguishability Metric Theorem). *Suppose admissible sectors of M are distributed across M_{op} with smooth density $\rho_{op}(x)$ at coarse-graining scale $\varepsilon \geq \Delta_{op}$ (Definition 3.4). Within the class of conformal rescalings*

$$g_{op}(x) = f(\rho_{op}(x)) \cdot \langle\cdot,\cdot\rangle_{op}$$

of the flat reference inner product by a smooth positive function f, the unique form under which g_{op}-length equals the local distinguishability-quantum count along paths is

$$g_{op}(x)(v, v) = C \cdot \rho_{op}(x)^{2/d_{op}} \cdot \langle v, v\rangle_{op},$$

for some overall positive constant C fixed by the convention that constant $\rho_{op} \equiv 1/Vol_{op}(M)$ recovers the flat reference metric. The broader uniqueness — within the class of all smooth Riemannian deformations of $\langle\cdot,\cdot\rangle_{op}$ — is separately supplied by the $\mathbb{Z}_7$-equivariance argument of Remark 5.2.2, which forces the conformal restriction itself.

Proof. The operational metric must encode the operational interpretation of distance: distance between two points is the count of distinguishability quanta separating them. We derive the conformal factor $f(\rho_{op})$.

Inter-sector spacing. ρ_{op} is a probability density ($\int_{M_{op}} \rho_{op}\, d\mathcal{H}^{d_{op}}_{op} = 1$), not a sector count per reference volume. The sector count density at x is therefore

$$n_{op}(x) := |\Sigma(M)| \cdot \rho_{op}(x),$$

so that $\int n_{op}\, d\mathscr{H}^{d_{op}}_{op} = |\Sigma(M)|$ recovers the total sector count. The typical inter-sector flat-reference spacing in a region of count density n_{op} is

$$\delta_{local}(x) = c_d \cdot n_{op}(x)^{-1/d_{op}} = c_d \cdot |\Sigma(M)|^{-1/d_{op}} \cdot \rho_{op}(x)^{-1/d_{op}},$$

where c_d is a dimension-dependent geometric constant (the d_{op}-dimensional unit-cell scale).

Quantum count. For a tangent vector $v \in T_x M_{op}$, the number of local distinguishability quanta along v is

$$N_{quanta}(v) = \|v\|_{op} / \delta_{local}(x) = (|\Sigma(M)|^{1/d_{op}} / c_d) \cdot \rho_{op}(x)^{1/d_{op}} \cdot \|v\|_{op}.$$

The factor $|\Sigma(M)|^{1/d_{op}} / c_d$ is a *global* constant — independent of x — and therefore induces only an overall multiplicative rescaling of g_{op}. The form $g_{op} = C \cdot \rho_{op}^{2/d_{op}} \cdot \langle\cdot,\cdot\rangle_{op}$ is preserved; only the constant C is affected.

Overall constant. Any overall constant rescaling of g_{op} is unobservable at the level of curvature, geodesics-as-curves, or any structural claim of the present paper: curvature is invariant under constant conformal rescaling; geodesic equations are invariant; relative distances and angles are preserved. We fix C by requiring that constant $\rho_{op} \equiv 1/\mathrm{Vol}_{op}(M)$ — uniform packing — recovers the flat reference metric $\langle\cdot,\cdot\rangle_{op}$ of the previous paper. This determines $C = \mathrm{Vol}_{op}(M)^{2/d_{op}}$, but the structural results below depend only on the position-dependence $\rho_{op}^{2/d_{op}}$, not on this absolute normalisation.

Squaring the quantum count. With C absorbed,

$$g_{op}(v, v) = \rho_{op}^{2/d_{op}} \cdot \langle v, v\rangle_{op}.$$

Uniqueness within the conformal class. Within the class of conformal rescalings $g_{op} = f(\rho_{op}) \cdot \langle\cdot,\cdot\rangle_{op}$, the requirement $N_{quanta}(v) = \sqrt{g_{op}(v, v)}$ forces $f(\rho_{op}) \propto \rho_{op}^{2/d_{op}}$. Any other functional dependence would either miscount local quanta (if the exponent differs from $2/d_{op}$) or introduce dimensional inconsistency in the relation $N_{quanta} = \|v\|_{op} / \delta_{local}$.

Beyond the conformal class. Non-conformal Riemannian deformations of $\langle\cdot,\cdot\rangle_{op}$ (e.g. position-dependent channel-mixing of the spectral decomposition $\mathcal{A} = \bigoplus_\alpha V_\alpha$, or channel-weighted direct sums with non-uniform weights) are excluded separately, by the $\mathbb{Z}_7$-equivariance + channel-uniformity argument of Remark 5.2.2. $\mathbb{Z}_7$-equivariance restricts the allowed deformations to channel-weighted direct sums; the channel-uniformity principle (the scalar source ρ_{op} generates only channel-symmetric response) restricts further to the conformal sub-class. The conformal restriction is itself supplied by these two substrate-level principles, not an independent assumption.

Remark 5.0.1. Theorem 5.0 promotes g_{op} from a postulated definition to a derived consequence. The uniqueness argument is three-layered: $\mathbb{Z}_7$-equivariance restricts the admissible geometric deformations to the *channel-weighted direct sum* class (Remark 5.2.2); the channel-uniformity principle restricts further to the *conformal* sub-class (Remark 5.2.2); the counting argument of Theorem 5.0 fixes the conformal factor within the conformal class to ρ_op^{2/d_op}. All three layers are independently necessary; together they determine g_op = ρ_op^{2/d_op} · ⟨·,·⟩_op up to the global rescaling fixed by flat-reference recovery.

Remark 5.0.2. *Relation to information-geometric metrics.* Theorem 5.0's derivation is structurally similar to the construction of the Fisher information metric on probability simplices — both are conformal rescalings of an ambient flat metric by a power of a local density that encodes distinguishability scale. Differences: (i) the Fisher information metric is constructed from the expected Hessian of log-likelihood (not the Hessian itself), whereas Theorem 5.0 rescales by the density itself raised to 2/d_op; (ii) the underlying counting is over admissible sectors (substrate-discrete) rather than over distinguishable probability distributions (continuum-statistical). Whether the substrate-level operational metric reduces to a Fisher metric on the emergent macroscopic continuum is an open question of the broader programme (§14, item 14).

Remark 5.0.3. *Locality principle.* Theorem 5.0's counting argument implicitly assumes a **locality principle**: the operational metric g_op(x) at the point x is determined by operational data *at* x — specifically by ρ_op(x) and the local distinguishability spacing δ_local(x) — rather than by ρ_op weighted against a non-local reference density or by global functionals of ρ_op. The substrate-level justification is straightforward: operational distance between two infinitesimally close configurations is a *local* quantity (it counts local quanta along an infinitesimal segment), and a local metric is the natural geometric object encoding such local distances. Non-local metric responses (e.g. depending on ∫_{M_op} ρ_op dℋ in some sub-region) would correspond to refinement transport being influenced by distant admissible-sector density, which would violate the substrate-level locality of the closure architecture (OG §2.5: Ω_max is finite-rank and channel-local). Within this locality principle, the position-dependence ρ_op(x)^{2/d_op} is the unique conformal factor consistent with operational counting.

5.2 Properties of ρ_op and g_op

Theorem 5.1 (Smoothness and boundedness of ρ_op). *Under the kernel-coverage hypothesis (Def 3.3) with inner radius ε' and kernel lower bound K_*, the operational distinguishability density ρ_op of Definition 3.4 is C^∞ on M_op and uniformly bounded:*

0 < ρ_op^{min} ≤ ρ_op(x) ≤ ρ_op^{max} < ∞ for all x ∈ M_op.

Proof.

Smoothness. ρ_op is the convolution of the bounded function η_op^ε against a C^∞ mollifier K_ε. Convolution of an L^∞ function with a smooth compactly-supported kernel produces a smooth function (standard mollification theorem). Hence ρ_op ∈ C^∞(M_op).

Upper bound. By Remark 3.2.1, η_op^ε ≤ C_{pack}(d_op) / Δ_op^{d_op} uniformly on M_op. Convolution against a unit-mass kernel cannot increase the L^∞ norm, so

ρ_op^{max} = ‖ρ_op‖*{L^∞}* ≤ ‖*η_op^ε*‖{L^∞} / 𝒵 ≤ C_{pack}(d_op) / (|Σ(M)| · Δ_op^{d_op}).

Lower bound. Working directly from the convolution definition, 𝒵 · ρ_op(x) = ∫_{M_op} K_ε(x − y) · η_op^ε(y) · d𝓗^{d_op}*op(y). We bound this integral from below by restricting to the sub-ball B(x, ε') on which the kernel is bounded below by* K*.

Claim: for every y ∈ B(x, ε'), η_op^ε(y) ≥ 1/(ω_{d_op} · ε^{d_op}). By kernel-coverage, there exists s_x ∈ Σ(M) with |x − s_x| ≤ ε − ε'. For any y ∈ B(x, ε'), the triangle inequality gives

|y − s_x| ≤ |y − x| + |x − s_x| ≤ ε' + (ε − ε') = ε,

so s_x ∈ B(y, ε). Therefore |Σ(M) ∩ B(y, ε)| ≥ 1, and η_op^ε(y) = |Σ(M) ∩ B(y, ε)| / (ω_{d_op} · ε^{d_op}) ≥ 1/(ω_{d_op} · ε^{d_op}).

The integral bound. Restricting the convolution to B(x, ε') and using K_ε ≥ K_* on this sub-ball:

𝒵 · ρ_op(x) ≥ K_* · (1 / (ω_{d_op} · ε^{d_op})) · vol_op(B(x, ε')) = K_* · (1 / (ω_{d_op} · ε^{d_op})) · ω_{d_op} · (ε')^{d_op} = K_* · (ε' / ε)^{d_op}.

With 𝒵 = |Σ(M)|, this gives the clean closed-form lower bound

ρ_op(x) ≥ K_* · (ε' / ε)^{d_op} / |Σ(M)| > 0.

The bound is strictly positive, uniform in x ∈ M_op, and depends explicitly on the kernel lower-bound K_*, the inner-radius-to-support ratio ε'/ε ∈ (0, 1), the sector count |Σ(M)|, and the dimension d_op (entering through the (ε'/ε)^{d_op} factor).

Remark 5.1.1. The kernel-coverage hypothesis is the precise structural condition required for the lower bound. The weaker topological ε-covering (every x within ε of some s) is *not* sufficient, because a compactly supported mollifier vanishes at the boundary of its support — sectors at distance close to ε contribute zero to ρ_op(x). The kernel-coverage hypothesis with inner radius ε' < ε strictly bounded away from ε ensures sectors lie in the kernel-positive interior. Under kernel-coverage, ρ_op is bounded away from zero uniformly on M_op, and g_op is non-degenerate everywhere on M_op. Closure manifolds violating kernel-coverage admit regions of vanishing ρ_op and degenerate g_op.

Theorem 5.2 (Conformal-deformation theorem). *Under the hypotheses of Theorem 5.1 (kernel-coverage), the conformally deformed operational metric*

g_op(x) = ρ_op(x)^{2/d_op} · ⟨·,·⟩_op

is a smooth Riemannian metric on M_op. The Riemann curvature tensor R_op of g_op vanishes identically when ρ_op is constant, and is generated by gradients of ρ_op when it varies (Theorem 8.1 supplies the closed-form expression).

Proof. g_op is smooth because ρ_op is smooth (Theorem 5.1) and the conformal factor ω_op = ρ_op^{1/d_op} is smooth and bounded away from zero (Theorem 5.1 lower bound). g_op is positive-definite for the same reason. Hence g_op is a smooth Riemannian metric on the d_op-dimensional smooth manifold M_op.

When ρ_op is constant, g_op is a constant multiple of the flat reference inner product. The Riemann tensor of a constant multiple of a flat metric vanishes identically. Hence R_op = 0 when ρ_op is constant.

When ρ_op varies, the closed-form curvature formula of Theorem 8.1 shows R_op ≠ 0 generically. The structural characterisation of the converse — when R_op = 0 forces ρ_op constant — is established in Theorem 8.2 in the form: vanishing R_op forces ψ to solve a Liouville-type equation, whose only solutions excluding non-trivial Möbius transformations are constant.

Remark 5.2.1. This is the structural backbone theorem of the paper. The flat reference Hilbert geometry of admissibility has no curvature. Distinguishability-density gradients — i.e. inhomogeneity in the packing of admissible sectors across M_op — conformally deform this flat reference, generating a curved Riemannian metric whose Riemann tensor is determined in closed form by derivatives of ln ρ_op. Curvature is not assumed; it is generated by inhomogeneity in the admissibility substrate, and the specific form of the deformation is forced by Theorem 5.0.

Remark 5.2.2. *Why conformal? The $\mathbb{Z}_7$-equivariance and channel-uniformity arguments.* The choice of a conformal deformation g_op = ω_op^2 · ⟨·,·⟩_op (rescaling by a single scalar function) — rather than a general position-dependent Riemannian deformation — is forced by two complementary substrate-level principles.

$\mathbb{Z}_7$-equivariance: the weighted direct sum form. The spectral decomposition 𝒜 = ⊕_α V_α of OG Theorem 4.0 is an orthogonal direct sum of channels carrying distinct $\mathbb{Z}_7$-characters. By OG §2.5, Ω_max is $\mathbb{Z}_7$-equivariant, and so is the operational closure dynamics. Any operational geometric deformation must inherit this $\mathbb{Z}_7$-equivariance. By Schur's lemma applied to the pairwise inequivalent $\mathbb{Z}_7$-irreps V_α, the most general $\mathbb{Z}_7$-equivariant position-dependent inner product on 𝒜 is the *channel-weighted direct sum*

g(x)(v, w) = ∑_α λ_α(x) · ⟨v_α, w_α⟩_op,

with channel-specific positive smooth weights λ_α : M_op → ℝ_{>0} and v_α, w_α the V_α-projections of v, w. $\mathbb{Z}_7$-equivariance alone permits channel-asymmetric deformations (distinct λ_α for distinct α) and does not by itself force conformality.

Channel-uniformity from scalar source. The conformal restriction λ_α(x) = ω_op(x)^2 for all α is selected by an additional **channel-uniformity principle**: the operational distinguishability density ρ_op is a *scalar* function on M_op, carrying no channel structure (it counts admissible sectors irrespective of spectral channel), so the metric response it generates must be channel-uniform. Formally: if g_op is determined locally by ρ_op alone (locality principle, Remark 5.0.3 below), and ρ_op carries no channel structure, then no channel-asymmetric coupling λ_α(x) = F(c_α · ρ_op(x)) with channel-dependent fixed couplings c_α can arise — such couplings would require a substrate-level mechanism distinguishing channels, which the scalar source does not supply.

Under both $\mathbb{Z}_7$-equivariance and channel-uniformity, the most general operational geometric deformation by a scalar density is conformal:

$$g_{op}(x) = \omega_{op}(x)^2 \cdot \langle\cdot,\cdot\rangle_{op}.$$

The counting argument of Theorem 5.0 then fixes $\omega_{op}(x)^2 = \rho_{op}(x)^{2/d_{op}}$ within this conformal class. The two principles are complementary: $\mathbb{Z}_7$-equivariance restricts to the channel-weighted direct sum class; channel-uniformity restricts to the conformal sub-class; counting fixes the conformal factor.

Whether the channel-uniformity principle is itself a corollary of substrate primitives (rather than an independent input) is one of the open problems of the broader programme (§14, item 16).

---

6. Operational Geodesics

The conformally deformed Riemannian metric g_{op} on M_{op} induces a natural notion of geodesic — a path that locally minimises the g_{op}-length L_{op} of Definition 3.8.

Theorem 6.1 (Operational geodesic equation). *On (M_{op}, g_{op}), the g_{op}-geodesic equation is the smooth second-order ODE*

$$\ddot{\gamma}^a + \Gamma_{op}{}^a{}_{bc}(\gamma) \cdot \dot{\gamma}^b \cdot \dot{\gamma}^c = 0,$$

with Christoffel symbols of the conformally-rescaled metric ($g_{op} = e^{2\psi} \langle\cdot,\cdot\rangle_{op}$ with $\psi = (1/d_{op}) \ln \rho_{op}$) given in closed form by

$$\Gamma_{op}{}^a{}_{bc} = \delta^a{}_b\, \psi_c + \delta^a{}_c\, \psi_b - \delta_{bc}\, \psi^a,$$

where $\psi_a := \partial_a \psi$ and indices are raised with the flat reference inner product.

Proof. g_{op} is a smooth Riemannian metric on M_{op} (Theorem 5.2). The Christoffel symbols of a conformally rescaled flat metric $g = e^{2\psi} \cdot g_{flat}$ are standard (Besse, *Einstein Manifolds*, §1.J, Eq. 1.159):

$$\tilde{\Gamma}^a{}_{bc} = \Gamma^a{}_{bc} + \delta^a{}_b\, \psi_c + \delta^a{}_c\, \psi_b - g_{bc}\, \psi^a,$$

which reduces to the stated form when $\Gamma^a{}_{bc} = 0$ (flat reference) and $g_{bc} = \delta_{bc}$. Smoothness of $\Gamma_{op}{}^a{}_{bc}$ follows from smoothness of $\psi = (1/d_{op}) \ln \rho_{op}$ (Theorem 5.1, plus strict positivity of ρ_{op} which makes $\ln \rho_{op}$ smooth). The geodesic equation is then a smooth second-order ODE.

Theorem 6.2 (Local geodesic existence). *For every x_0 ∈ M_op and every tangent vector v_0 ∈ T_{x_0} M_op = $\mathcal{A}$, there exists ε(x_0, v_0) > 0 and a unique smooth geodesic γ : (−ε, ε) → M_op with γ(0) = x_0 and $\dot{\gamma}(0) = v_0$.*

Proof. Immediate from Theorem 6.1 (smooth ODE) and Picard–Lindelöf.

Theorem 6.3 (Local minimising geodesics on the interior). *For every interior point x_0 ∈ int(M_op), there exists a neighbourhood U ⊆ int(M_op) of x_0 such that any two points x, y ∈ U are joined by a unique minimising g_op-geodesic γ : [0, 1] → U.*

Proof. By Theorem 5.1, the conformal factor ω_op is uniformly bounded above and below on M_op. Hence the g_op-metric and the flat D_op-metric are bilipschitz-equivalent on M_op:

ω_op^{min} · D_op(x, y) ≤ d_{g_op}(x, y) ≤ ω_op^{max} · D_op(x, y),

where d_{g_op} is the g_op-distance. On the interior, the standard local geodesic-convexity theorem for smooth Riemannian metrics (see e.g. do Carmo, *Riemannian Geometry*, Theorem 3.7) gives the existence of geodesically convex neighbourhoods U of every interior point, on which any two points are joined by a unique minimising geodesic.

Remark 6.3.1. Full geodesic completeness of (M_op, g_op) — in the sense of Hopf–Rinow — requires either (i) closedness of M_op in $\mathcal{A}$ (which does not hold in general, since open precompact subsets of Hilbert spaces are not closed), or (ii) a specific boundary condition on ∂M_op. On the **closed completion** ($\overline{M}$_op, g_op) — where g_op extends continuously to the closure under the kernel-coverage hypothesis — compactness plus Hopf–Rinow gives full geodesic completeness. On the open interior int(M_op), only local geodesic existence and local minimising-geodesic existence (Theorem 6.3) hold without further hypotheses. Trajectories of interest in physical applications typically remain in the interior; if a trajectory approaches ∂M_op, additional boundary conditions become relevant.

Remark 6.3.2. In the flat-reference limit ρ_op ≡ const, g_op-geodesics reduce to straight lines in ($\mathcal{A}$, ⟨·,·⟩_op). In the curved case, geodesics bend toward regions of higher ρ_op: where admissibility packs sectors more densely, the local g_op-distance per unit ⟨·,·⟩_op-distance is larger, so g_op-shortest paths take detours through low-density regions. This is the substrate-level analogue of light bending around mass concentrations in general relativity.

---

7. Operational Connection

The conformally deformed metric g_op induces a Levi-Civita connection on M_op.

Theorem 7.1 (Existence and uniqueness of operational Levi-Civita connection). *On (M_op, g_op), there exists a unique torsion-free metric-compatible connection ∇_op^{LC} — the Levi-Civita connection — characterised by:*

- *Torsion-freeness:* $\nabla_{op}^{LC}{}_X Y - \nabla_{op}^{LC}{}_Y X = [X, Y]$;
- *Metric compatibility:* $X \cdot g_{op}(Y, Z) = g_{op}(\nabla_{op}^{LC}{}_X Y, Z) + g_{op}(Y, \nabla_{op}^{LC}{}_X Z)$.

Its Christoffel symbols are

$$\Gamma_{op}{}^a{}_{bc} = \delta^a_b \psi_c + \delta^a_c \psi_b - \delta_{bc} \psi^a,$$

with $\psi = (1/d_{op}) \ln \rho_{op}$ (cf. Theorem 6.1).

Proof. By the fundamental theorem of Riemannian geometry, every smooth Riemannian manifold carries a unique torsion-free metric-compatible connection. g_{op} is a smooth Riemannian metric on M_{op} (Theorem 5.2), so ∇_{op}^{LC} exists and is unique. The closed-form Christoffel symbols are standard for a conformally rescaled flat metric (Besse op. cit. §1.J, Eq. 1.159).

Corollary 7.2 (Operational parallel transport). *Operational parallel transport along a g_{op}-geodesic $\gamma : [0, 1] \to \mathrm{int}(M_{op})$ is defined, for each tangent vector $v \in T_{\gamma(0)} M_{op}$, as the unique solution $V(\tau) \in T_{\gamma(\tau)} M_{op}$ of*

$$\nabla_{op}^{LC}{}_{\dot\gamma} V = 0, \quad V(0) = v.$$

Parallel transport is an isometry $T_{\gamma(0)} M_{op} \to T_{\gamma(1)} M_{op}$ of the corresponding g_{op}-inner products at the endpoints, by metric compatibility.

Remark 7.2.1. In the flat-reference limit $\rho_{op} \equiv \text{const}$, ψ is constant and $\Gamma_{op}{}^a{}_{bc} = 0$. Parallel transport reduces to translation in $(\mathcal{A}, \langle\cdot,\cdot\rangle_{op})$ — the trivial flat connection. In the curved case, parallel transport around closed loops generically returns vectors rotated by an amount determined by the curvature tensor R_{op} of §8 — the holonomy.

---

8. Operational Curvature

We now develop the operational curvature tensor in closed form.

Theorem 8.1 (Operational curvature tensor in closed form). *The Riemann curvature tensor of the operational Levi-Civita connection ∇_{op}^{LC}, expressed in the (1,3) component form $R_{op}{}^a{}_{bcd}$, is given by the standard conformal-flat formula for $g_{op} = e^{2\psi} \cdot \langle\cdot,\cdot\rangle_{op}$ (Besse op. cit. §1.J):*

$$R_{op}{}^a{}_{bcd} = \delta^a_c (\psi_{bd} - \psi_b \psi_d) - \delta^a_d (\psi_{bc} - \psi_b \psi_c) + \delta_{bd} (\psi^a{}_c - \psi^a \psi_c) - \delta_{bc} (\psi^a{}_d - \psi^a \psi_d) - \|\nabla\psi\|^2 \cdot (\delta_{bd} \delta^a_c - \delta_{bc} \delta^a_d),$$

where:

- $\psi = (1/d\_op) \ln \rho\_op$,
- $\psi\_a = \partial\_a \psi = \nabla\_a \psi$ *(flat-reference gradient),*
- $\psi\_\{ab\} = \partial\_a \partial\_b \psi$ *(flat-reference Hessian),*
- *indices raised with the flat reference δ-tensor,*
- $\| \nabla\psi\|^2 = \delta\^\{ab\} \psi\_a \psi\_b$.

Proof. This is the conformal-flat Riemann tensor formula for $g\_op = e\^\{2\psi\} \cdot \langle\cdot,\cdot\rangle\_op$ with ψ smooth (Theorem 5.1). The derivation is standard (Besse, *Einstein Manifolds*, §1.J). The formula expresses R_op entirely in terms of ψ and its first two derivatives.

Remark 8.1.1. *Reconciliation of conventions.* The curvature formula uses the sign convention in which the last term carries a *minus* $\|\nabla\psi\|^2$ factor without an ω_op^{-1} prefactor, consistent with the (1,3)-component form of the Riemann tensor for the rescaled metric. The factor ω_op^2 appearing in indexes raised with $g\_op\^\{ab\} = \omega\_op^{-2} \delta\^\{ab\}$ introduces ω_op-prefactors in the (0,4)-form $R\_op\_\{abcd\} = \omega\_op^2 \cdot$ (the expression above with δ_{ab} replacing g_{ab}); we work with the (1,3) form throughout to avoid these prefactors in the basic curvature expression.

Theorem 8.2 (Curvature vanishing). *On a connected open subset $U \subseteq M\_op$:*

(i) *If ρ_op is constant on U, then R_op = 0 on U.*

(ii) *Conversely, in dimension $d\_op \geq 3$, R_op = 0 on U forces $\psi = (1/d\_op) \ln \rho\_op$ to satisfy the Liouville-type equation*

$$\psi\_\{ab\} = \psi\_a \psi\_b - (1/2) \|\nabla\psi\|^2 \cdot \delta\_\{ab\},$$

whose solutions on simply connected U include — beyond the constant ψ — non-trivial conformal factors corresponding to restrictions of Möbius transformations of Euclidean R^{d_op}. Such non-constant Möbius solutions cannot in general be excluded from R_op = 0 by curvature data alone.

Proof. (i) If ρ_op is constant on U, then ψ is constant, all ψ-derivatives vanish, and R_op = 0 by Theorem 8.1.

(ii) For a conformally flat metric in dimension $d \geq 3$, the Weyl tensor vanishes identically. The vanishing of the full Riemann tensor R_op = 0 then forces the Schouten tensor

$$P\_\{ab\} = - \psi\_\{ab\} + \psi\_a \psi\_b - (1/2) \|\nabla\psi\|^2 \cdot \delta\_\{ab\}$$

to vanish (Besse op. cit. Eq. 1.159), giving the stated Liouville-type equation. The space of smooth solutions to this equation on simply connected R^{d_op} ($d\_op \geq 3$) is precisely the conformal factors of Möbius transformations — i.e. compositions of translations, rotations, dilations, and inversions — which include the constant solution but also non-constant inversion-type factors. Restrictions of such Möbius factors to simply connected $U \subseteq M\_op$ are smooth (away from any inversion centres outside U) and remain solutions of the Liouville equation. Therefore R_op = 0 on U does not, by curvature data alone, force ρ_op constant on U.

Remark 8.2.1. *Excluding Möbius solutions via admissibility.* The Möbius non-constant solutions to $R_{op} = 0$ correspond to inversion-type conformal maps. Excluding them — and recovering the converse direction "$R_{op} = 0 \Rightarrow \rho_{op}$ constant" — requires additional admissibility-level conditions on ρ_{op}. Candidate conditions include:

- *Boundedness with no inversion-centre singularity.* Möbius conformal factors corresponding to inversions through centres outside M_{op} are bounded on M_{op}, but their gradients $\|\nabla\psi\| \to \infty$ as one approaches the inversion centre. Restricting to ρ_{op} with kernel-coverage-bounded gradient (Lemma 9.0) excludes inversion centres at finite distance from M_{op} but does not exclude inversion centres at arbitrary distance from M_{op}.
- *Normalisation invariance.* The condition $\int_{M_{op}} \rho_{op} \, d\mathscr{H}^{d_{op}}_{op} = 1$ must hold; Möbius transformations generically do not preserve this normalisation on a fixed bounded domain.
- *$\mathbb{Z}_7$-equivariance.* Möbius transformations of $\mathbb{R}^{d_{op}}$ need not respect the spectral channel decomposition $\mathcal{A} = \bigoplus_\alpha V_\alpha$; restricting to $\mathbb{Z}_7$-equivariant ρ_{op} restricts the allowed solutions.

Under a combination of these admissibility conditions, the only solution of $R_{op} = 0$ on M_{op} is ρ_{op} constant. A full structural characterisation of which admissibility conditions are necessary and sufficient is one of the open problems of the broader programme (§14, item 15).

Remark 8.2.2. *Two-dimensional exception.* In $d_{op} = 2$, every metric is conformally flat locally (uniformisation), and the Gaussian curvature of $g = e^{2\psi} \cdot g_{flat}$ is $K = -e^{-2\psi} \cdot \Delta_{op} \psi$. $R_{op} = 0$ in 2D is equivalent to $\Delta_{op} \psi = 0$, i.e. ψ harmonic — admitting many non-constant solutions. For the closure architecture of VERSF, $d_{op} \geq N_{spec} \geq 3$ (since the $\mathbb{Z}_7$ closure architecture contributes at least three orthogonal spectral channels carrying admissibility), so $d_{op} \geq 3$ in physically relevant cases.

Corollary 8.3 (Sectional curvature). *In dimension $d_{op} \geq 2$, the sectional curvature $K_{op}(\sigma)$ of a 2-plane $\sigma = \mathrm{span}(v, w) \subseteq T_x M_{op}$ (with v, w g_{op}-orthonormal) is given by the standard expression*

$$K_{op}(\sigma) = g_{op}(R_{op}(v, w)\, w, v),$$

explicitly in terms of ψ-derivatives via Theorem 8.1. K_{op} is bounded by Theorem 9.1.

Remark 8.3.1. The closed-form curvature formula gives operational curvature as an explicit function of the operational distinguishability density ρ_{op}. The curvature is not an independent dynamical degree of freedom — it is determined entirely by the local distinguishability density on M_{op}.

---

9. Curvature Bounds from Finite Packing

The finite-packing constraint of OG Theorem 10.1 — P(M) ≤ Vol_op(M) / Δ_op^{d_op} — bounds the operational curvature pointwise on M_op.

Lemma 9.0 (Bounded operational gradient of ln ρ_op). *Under the kernel-coverage hypothesis (Def 3.3), on M_op:*

‖∇_op ln ρ_op(x)‖_op ≤ C$_1$(d_op, K_ε) · (ρ_op^{max} / ρ_op^{min}) / Δ_op, uniformly in x ∈ M_op,

for a constant C$_1$(d_op, K_ε) of order unity depending on dimension and the choice of mollifier kernel.

Proof. By Definition 3.4, ∇_op ρ_op = (∇_op K_ε) ∗ η_op^ε. The mollifier K_ε is a C^∞ kernel of length-scale ε ≥ Δ_op, so

‖∇*op K_ε*‖{L^∞} ≤ C_K / ε ≤ C_K / Δ_op,

(standard mollifier-derivative estimate with C_K a kernel-dependent constant). Convolving against L^∞-bounded η_op^ε:

‖∇*op ρ_op*‖{L^∞} ≤ C(d_op, K_ε) · ρ_op^{max} / Δ_op.

Dividing by ρ_op^{min} > 0:

‖∇*op ln ρ_op*‖{L^∞} = ‖∇*op ρ_op / ρ_op*‖{L^∞} ≤ C$_1$(d_op, K_ε) · (ρ_op^{max} / ρ_op^{min}) / Δ_op.

Lemma 9.0' (Bounded operational Hessian of ln ρ_op). *Similarly,*

‖Hess_op ln ρ_op(x)‖_{op} ≤ C$_2$(d_op, K_ε) · (ρ_op^{max} / ρ_op^{min})^2 / Δ_op^2,

uniformly in x ∈ M_op, for a constant C$_2$(d_op, K_ε) of order unity.

Proof. Direct calculation gives

Hess_op (ln ρ_op) = ρ_op^{-1} · Hess_op ρ_op − ρ_op^{-2} · (∇_op ρ_op)(∇_op ρ_op)^{T}.

We bound the two terms separately.

Term 1: ρ_op^{-1} · Hess_op ρ_op. By the mollifier-derivative estimate ‖∇^{2}_op K_ε‖_{L^∞} ≤ C / ε^2 ≤ C / Δ_op^2 and convolution against the L^∞-bounded η_op^ε:

‖Hess_op ρ_op‖_{L^∞} ≤ C(d_op, K_ε) · ρ_op^{max} / Δ_op^2.

Dividing by ρ_op^{min}:

$\|\rho_{op}^{-1} \cdot \mathrm{Hess}_{op}\, \rho_{op}\|_{L^\infty} \le C(d_{op}, K_\varepsilon) \cdot (\rho_{op}^{\max} / \rho_{op}^{\min}) / \Delta_{op}^2$.

Term 2: $\rho_{op}^{-2} \cdot (\nabla_{op} \rho_{op})(\nabla_{op} \rho_{op})^T$. By Lemma 9.0's estimate $\|\nabla_{op} \rho_{op}\| \le C \cdot \rho_{op}^{\max} / \Delta_{op}$:

$\|\rho_{op}^{-2} \cdot (\nabla_{op} \rho_{op})(\nabla_{op} \rho_{op})^T\|_{L^\infty} \le \rho_{op}^{-2}{}_{\min} \cdot (C \cdot \rho_{op}^{\max} / \Delta_{op})^2 = C^2 \cdot (\rho_{op}^{\max} / \rho_{op}^{\min})^2 / \Delta_{op}^2$.

Comparison. Term 2 scales as $(\rho_{op}^{\max}/\rho_{op}^{\min})^2/\Delta_{op}^2$ and Term 1 as $(\rho_{op}^{\max}/\rho_{op}^{\min})/\Delta_{op}^2$. Since the density-contrast ratio $\rho_{op}^{\max}/\rho_{op}^{\min} \ge 1$, Term 2 dominates the bound when the contrast is non-trivial. The combined bound is therefore

$\|\mathrm{Hess}_{op} \ln \rho_{op}\|_{L^\infty} \le C_2(d_{op}, K_\varepsilon) \cdot (\rho_{op}^{\max} / \rho_{op}^{\min})^2 / \Delta_{op}^2$,

with C_2 absorbing both contributions.

Theorem 9.1 (Bounded operational curvature). *Under the kernel-coverage hypothesis, the operational Riemann curvature tensor satisfies*

$\|R_{op}(x)\|_{op} \le C(d_{op}, K_\varepsilon, \rho_{op}^{\max}/\rho_{op}^{\min}) / \Delta_{op}^2$, uniformly in $x \in M_{op}$,

where the constant C is explicit in dimension, kernel choice, and the density contrast ratio across M_op. In particular, sectional curvatures are uniformly bounded:

$|K_{op}(\sigma)| \le C'(d_{op}, K_\varepsilon, \rho_{op}^{\max}/\rho_{op}^{\min}) / \Delta_{op}^2$

for all 2-planes $\sigma \subseteq T_x M_{op}$, $x \in M_{op}$.

Proof. By Theorem 8.1, R_{op} is a sum of terms involving:

(i) **First-derivative quadratics** $\psi_a \psi_b$: by Lemma 9.0, $\|\nabla \psi\|^2 = (1/d_{op}^2) \cdot \|\nabla \ln \rho_{op}\|^2 \le C_1^2 \cdot (\rho_{op}^{\max}/\rho_{op}^{\min})^2 / (d_{op}^2 \Delta_{op}^2)$. Bounded by C / Δ_{op}^2 with C depending on $(d_{op}, \rho_{op}^{\max}/\rho_{op}^{\min}, K_\varepsilon)$.

(ii) **Second derivatives** ψ_{ab}: by Lemma 9.0', $\|\nabla^2 \psi\| = (1/d_{op}) \cdot \|\mathrm{Hess} \ln \rho_{op}\| \le C_2 \cdot (\rho_{op}^{\max}/\rho_{op}^{\min})^2 / (d_{op} \Delta_{op}^2)$. Bounded by C / Δ_{op}^2 with the same constant dependence.

(iii) **Tensor combinations** $\delta_{bd} \delta^a{}_c$, $\delta_{bc} \delta^a{}_d$: dimension-only constants.

Combining, $\|R_{op}\|_{op}$ is bounded by a constant multiple of $1/\Delta_{op}^2$ with explicit dependence on d_{op}, K_ε, and the density contrast ratio $\rho_{op}^{\max}/\rho_{op}^{\min}$ across M_{op}.

Remark 9.1.1. *Explicit constant dependence.* The constant in Theorem 9.1 depends not only on dimension but also on the density contrast ratio $\rho_{op}^{\max}/\rho_{op}^{\min}$ across M_{op} — a manifold-dependent quantity. For closure manifolds with uniform admissibility

($\rho_{op}^{max}/\rho_{op}^{min} \sim 1$), the constant is purely dimension-dependent. For closure manifolds with large density contrast, the constant grows polynomially in the contrast ratio. The bound is genuinely $1/\Delta_{op}^2$ in scaling but not manifold-independent in prefactor.

Remark 9.1.2. *Connection to a Planck-scale curvature cutoff — with caveat.* In macroscopic geometric theory, classical curvature is expected to break down at the Planck scale ℓ_P, with a natural upper bound $K \lesssim 1/\ell_P^2$. Theorem 9.1 gives the *substrate-level* analogue: Δ_{op} plays the role of a substrate cutoff scale, and operational curvature is bounded by $1/\Delta_{op}^2$ substrate-side. However, the identification with a manifold-independent Planck cutoff requires the prefactor to be manifold-independent — which holds only under uniform-admissibility conditions ($\rho_{op}^{max}/\rho_{op}^{min} \sim 1$). Closure manifolds with significant density contrast would have larger prefactors, indicating that the Planck-cutoff identification is sharpest for uniform admissibility regimes and softer otherwise. Whether Δ_{op} corresponds quantitatively to ℓ_P under coarse-graining is one of the open problems of the broader programme (§14).

---

10. Refinement Transport Equations

Refinement flow on M_{op} induces dynamical equations on the distinguishability density ρ_{op}.

Definition 10.1 (Operational velocity field). Let Φ_τ be a refinement flow on $(\mathcal{A}, \langle\cdot,\cdot\rangle_{op})$, $\tau \mapsto \Phi_\tau$ smoothly differentiable in τ. The **operational velocity field** at $\tau = 0$ is

$$v_{op}(x) := \left(\frac{d}{d\tau} \Phi_\tau(x)\right)\Big|_{\tau=0}.$$

For the entropy-gradient flow of Theorem 4.2, $v_{op} = \nabla_{op} S_{op}$.

Definition 10.2 (Operational transport current). The **operational transport current** associated with velocity field v_{op} is

$$J_{op}(x) := \rho_{op}(x) \cdot v_{op}(x).$$

When $v_{op} = \nabla_{op} S_{op}$ *(entropy-gradient flow) and* S_{op} *admits the density-functional form* $S_{op}[\rho_{op}] = \int_{M_{op}} s(\rho_{op}(x)) \cdot d\mathscr{H}^{d_{op}}_{op}(x)$ *for some smooth scalar function* $s : \mathbb{R}_{>0} \to \mathbb{R}$ (so that the variational derivative $\delta S_{op}/\delta\rho_{op}$ exists pointwise as $s'(\rho_{op})$), the chain rule gives the Fickian form

$$J_{op} = -\kappa_{op}(x) \cdot \nabla_{op}\rho_{op},$$

with $\kappa_{op}(x) := \rho_{op}(x) \cdot s''(\rho_{op}(x))$ (assumed positive by convexity-of-density-functional). In coordinate-free language, S_{op} is a *local functional of ρ_{op} alone*, with no explicit position-dependence and no non-local kernel.

Remark 10.2.1. The density-functional assumption $S_{op} = \int s(\rho_{op})\, d\mathscr{H}$ is stronger than mere concavity of S_{op} on $\mathcal{A}$ (Theorem 4.2 hypothesis): it requires S_{op} to be expressible as an integral of a *local* scalar function of ρ_op, ruling out non-local entropy functionals such as Tsallis-like or relative-entropy forms. The standard Shannon-type functional S_op = −∫ ρ_op ln ρ_op dℋ satisfies this with s(ρ) = −ρ ln ρ; under this choice, κ_op = ρ_op · (1/ρ_op²) · ρ_op = 1, giving the standard heat equation in Corollary 10.4 with κ_op ≡ 1. Whether this density-functional form is forced by substrate primitives, or must be taken as an additional assumption beyond concavity, is closely related to Open Problem 13.

Theorem 10.3 (Operational continuity equation). *Let Φ_τ be a smooth refinement flow — i.e. a refinement flow (Definition 4.1) with τ-differentiable trajectories and smooth velocity field v_op(x). Then the operational distinguishability density evolves according to the operational continuity equation*

$$\partial\rho_{op} / \partial\tau + \nabla_{op} \cdot (\rho_{op} \cdot v_{op}) = 0.$$

Proof. This is the standard Eulerian mass-conservation argument for a smooth transport semigroup on a Riemannian manifold.

Let ρ_op(x, τ) be the τ-dependent density obtained by pushing ρ_op(x, 0) forward along Φ_τ. For any compactly-supported test region Ω ⊆ M_op:

d/dτ ∫_Ω ρ_op(x, τ) · dℋ^{d_op}*op = − flux of (ρ_op · v_op) through ∂Ω* = − ∫{∂Ω} ρ_op · v_op · n · dℋ^{d_op−1},

by the operational analogue of the Reynolds transport theorem applied to the τ-pushforward. The integrand on the left is ∫_Ω ∂_τ ρ_op · dℋ^{d_op}_op (smooth ρ_op-dependence on τ for smooth v_op). The right is by the divergence theorem ∫_Ω ∇_op · (ρ_op · v_op) · dℋ^{d_op}_op. Equality on every test region Ω forces the local form:

$$\partial\rho_{op} / \partial\tau + \nabla_{op} \cdot (\rho_{op} \cdot v_{op}) = 0.$$

This holds independently of whether Φ_τ is measure-preserving — the non-expansiveness of Definition 4.1 is *not* used here; what is used is the Eulerian conservation of the pushforward density, which holds for any smooth semigroup of self-maps.

Remark 10.3.1. *Smoothness hypothesis.* The theorem requires τ-differentiability of Φ_τ and smoothness of v_op. The Brezis maximal-monotone-semigroup framework (Theorem 4.2) gives only Lipschitz (not C^1) τ-dependence for general non-expansive semigroups, so the smoothness hypothesis is not automatic. For the entropy-gradient flow of Theorem 4.2, smoothness of v_op = ∇_op S_op holds on the interior of the complement of the S_op-maximum set (where ∇_op S_op is smooth), and the continuity equation holds there. At the S_op-maximum set itself, the flow becomes stationary (∇_op S_op = 0) and the continuity equation reduces trivially to ∂_τ ρ_op = 0. Whether smoothness extends to the boundary of the maximum set is an open regularity question.

Remark 10.3.2. The previous wording suggested non-expansiveness implies measure preservation, which is false in general (non-expansive flows can contract volumes). The correct statement is that the continuity equation expresses Eulerian conservation of the pushforward density along the flow, which holds for any smooth semigroup. Non-expansiveness enters elsewhere — in Definition 4.1 as the admissibility condition on Φ_τ, not in the derivation of the continuity equation.

Corollary 10.4 (Diffusive limit). *For the entropy-gradient refinement flow of Theorem 4.2 with Fickian transport (Definition 10.2), the continuity equation takes the operational diffusion form*

$$\partial\rho_{op} / \partial\tau = \nabla_{op} \cdot (\kappa_{op} \cdot \nabla_{op} \rho_{op}).$$

In the $\kappa_{op} \equiv$ constant limit, this reduces to

$$\partial\rho_{op} / \partial\tau = \kappa_{op} \cdot \bar{\Delta}_{op} \rho_{op},$$

where $\bar{\Delta}_{op}$ is the flat-reference operational Laplacian of Definition 3.5 (distinguished from the scalar distinguishability quantum Δ_{op} of OG §2.3 by the overbar).

Remark 10.4.1. In the asymptotic $\tau \to \infty$ limit, ρ_{op} approaches the constant operational equilibrium density $1/\mathrm{Vol}_{op}(M)$, at which point the conformal factor ω_{op} is constant and operational curvature $R_{op} = 0$ (Theorem 8.2(i), the constant-density-implies-flat direction). Refinement transport drives the operational geometry toward its flat reference limit. The non-equilibrium operational geometry — where $R_{op} \neq 0$ — corresponds to refinement-incomplete distinguishability density: regions where admissibility packing has not yet relaxed to uniform density.

---

11. Coarse-Graining and Macroscopic Emergence

We now sketch the bridge from operational geometric dynamics to the macroscopic geometric continuum.

Definition 11.1 (Coarse-graining map). Fix a coarse-graining length $L > \Delta_{op}$. The **coarse-graining map** at scale L is the map sending the operational geometry $(M_{op}, g_{op}, \rho_{op})$ to its L-smoothed version $(M_{op}, g_{op}^{(L)}, \rho_{op}^{(L)})$, where:

(i) the underlying smooth manifold M_{op} is unchanged — L-averaging acts on the metric data, not on the manifold; (ii) $\rho_{op}^{(L)}(x) := \int K_L(x, y) \cdot \rho_{op}(y) \cdot d\mathcal{H}^{d_{op}}_{op}(y)$, with K_L a smooth radial mollifier of scale L; (iii) $g_{op}^{(L)} := (\rho_{op}^{(L)})^{2/d_{op}} \cdot \langle\cdot,\cdot\rangle_{op}$, the L-smoothed conformally deformed metric.

Lemma 11.0 (Volume non-collapsing of g_{op}-balls). *Under the kernel-coverage hypothesis, there exists $v_0 > 0$ such that for every $x \in M_{op}$ and every $r \in [\Delta_{op}, \mathrm{diameter}(M_{op})]$,*

$$\mathrm{Vol}_{g_{op}}(B_{g_{op}}(x, r)) \geq v_0 \cdot r^{d_{op}}.$$

Equivalently, g_op-balls at resolved scales r ≥ Δ_op carry volume bounded below by a uniform constant times r^{d_op}; no degeneration of M_op to a lower-dimensional limit occurs at resolved scales.

Proof. The g_op-volume of any measurable region E ⊆ M_op is

Vol_{g_op}(E) = ∫_E ρ_op · d𝓗^{d_op}_op ≥ ρ_op^{min} · 𝓗^{d_op}_op(E).

By Theorem 5.1, ρ_op^{min} > 0 under the kernel-coverage hypothesis. Applied to a g_op-ball B_{g_op}(x, r) — which contains a flat reference D_op-ball of radius at least r / ω_op^{max} by the bilipschitz comparison ω_op^{min} · D_op ≤ d_{g_op} ≤ ω_op^{max} · D_op:

Vol_{g_op}(B_{g_op}(x, r)) ≥ ρ_op^{min} · vol_op(B_{D_op}(x, r / ω_op^{max})) = ρ_op^{min} · ω_{d_op} · (r / ω_op^{max})^{d_op} = v_0 · r^{d_op}

with v_0 := ρ_op^{min} · ω_{d_op} · (ω_op^{max})^{-d_op} > 0. The bound is uniform in x and in r ∈ [Δ_op, diameter(M_op)].

Remark 11.0.1. *The role of non-collapsing in Cheeger–Gromov.* Non-collapsing is the third Cheeger–Gromov hypothesis (alongside sectional curvature bound and diameter bound) needed for C^{1,α} subsequential convergence of a family of smooth Riemannian metrics. The bare curvature-plus-diameter hypotheses give Gromov–Hausdorff precompactness, but the limit could collapse to a lower-dimensional space without non-collapsing (e.g. a family of thinner and thinner cylinders collapses to a circle). Lemma 11.0 supplies the missing hypothesis as a direct consequence of kernel-coverage + bounded density contrast: ρ_op^{min} > 0 prevents the g_op-volume of resolved-scale balls from vanishing, hence prevents dimensional collapse in the GH limit. The bound holds at *resolved* scales r ≥ Δ_op; below Δ_op, the substrate-discrete structure of Σ(M) takes over and the continuous-Riemannian framework no longer applies.

Remark 11.0.2. *Uniformity across the coarse-graining family.* For the Cheeger–Gromov assembly of Theorem 11.2 to give a uniform-in-L bound, the constant v_0 in Lemma 11.0 must itself be uniform in L > Δ_op as the smoothing scale varies. Tracking through the proof, v_0 = ρ_op^{(L),min} · ω_{d_op} · (ω_op^{(L),max})^{-d_op}, which depends on the smoothing scale L through ρ_op^{(L),min}. By Theorem 5.1's lower-bound estimate ρ_op^{min} ≥ K_* · (ε'/ε)^{d_op} / |Σ(M)|, uniformity in L requires the **kernel family ratio condition**: the inner-radius-to-support ratio ε'_L / L of the coarse-graining kernel family {K_L : L > Δ_op} is bounded below by a positive constant uniformly in L. This is automatic for standard kernel families (e.g. Gaussian truncated at fixed multiples of L) and is taken as a standing kernel-choice convention henceforth.

Theorem 11.2 (C^{1,α} Riemannian limit). *Under the kernel-coverage hypothesis with bounded density contrast, the family of coarse-grained operational geometries {(M_op, g_op^{(L)}) : L > Δ_op} satisfies the three Cheeger–Gromov hypotheses uniformly in L:*

(i) *bounded sectional curvature: |K_op^{(L)}| ≤ C/Δ_op² (Theorem 9.1);*

(ii) *bounded diameter: diam(M_op, g_op^{(L)}) ≤ ω_op^{max} · diam(M_op, D_op) < ∞ (OG Lemma 8.0 with the bilipschitz comparison);*

(iii) *volume non-collapsing at resolved scales: Vol_{g_op^{(L)}}(B(x, r)) ≥ v_0 · r^{d_op} for r ∈ [Δ_op, diameter] (Lemma 11.0).*

The family is therefore precompact in the Gromov–Hausdorff topology of C^{1,α} Riemannian metrics. Any subsequential limit (M_∞, g_∞) along sequences L_n → Δ_op$^+$ with L_n > Δ_op is a C^{1,α} Riemannian metric on the admissibility substrate, with the bounded-curvature estimate of Theorem 9.1 and the non-collapsing estimate of Lemma 11.0 preserved in the limit.

Proof sketch. Conditions (i)–(iii) are the standard Cheeger–Gromov hypotheses for C^{1,α} compactness of Riemannian manifolds (Cheeger, *Finiteness theorems for Riemannian manifolds*, 1970; Petersen, *Riemannian Geometry*, Chapter 11). With all three uniformly in L, the family is precompact and any subsequential limit is C^{1,α} Riemannian with the bounds preserved.

The interesting sequence is L_n → Δ_op$^+$ from above (refining the coarse-graining toward the substrate scale). Higher regularity (C^2 or smooth) does not follow from Cheeger–Gromov alone; it would require additional control on covariant derivatives of curvature, which we have not established (see §14, item 13).

The L → ∞ limit (full coarse-graining) is uninteresting structurally: it smooths ρ_op to its mean value 1/Vol_op(M), under which g_op^{(L)} → constant rescaling of ⟨·,·⟩_op, recovering the flat reference. The refinement-toward-substrate-scale limit L_n → Δ_op$^+$ is the operationally significant case, since it retains the curvature structure generated by ρ_op-gradients at scales just above the substrate cutoff.

Remark 11.2.1. The earlier draft claimed a smooth (C^∞) limit; this is an overclaim. Cheeger–Gromov compactness produces C^{1,α} regularity under sectional curvature and diameter bounds; achieving higher regularity requires additional control on covariant derivatives of curvature, which would follow from bounds on higher derivatives of ρ_op. Whether such higher-regularity bounds follow from finite packing (perhaps with stronger structural hypotheses on the closure architecture) is an open question of the broader programme (§14, item 13).

Remark 11.2.2. *Conformal-flat structure preserved in the limit.* Each coarse-grained metric g_op^{(L)} = (ρ_op^{(L)})^{2/d_op} · ⟨·,·⟩_op is conformal-flat with respect to the *fixed* ambient inner product ⟨·,·⟩*op of the operational Hilbert geometry. Gromov–Hausdorff convergence along L_n → Δ_op$^+$ takes place within this fixed ambient Hilbert space, so any subsequential limit metric g∞* inherits the conformal-flat form

g_∞(x) = (ρ_op^∞(x))^{2/d_op} · ⟨·,·⟩_op

for some C^{1,α} limiting distinguishability density ρ_op^∞(x) (the GH limit of ρ_op^{(L_n)}). The limit is therefore not merely C^{1,α} Riemannian but *C^{1,α} conformal-flat* — a more restrictive structure that interacts with the Möbius-type Liouville solutions of Theorem 8.2(ii). Whether the macroscopic 4-manifold of physics inherits this conformal-flat structure or whether dimensional reduction (§12(a)) generates non-conformal-flat curvature in the reduced geometry is one of the structural questions of the coarse-graining bridge (§14, item 10).

Remark 11.2.3. Theorem 11.2 is the structural endpoint of the present paper. The coarse-grained operational geometry is a $C^{1,\alpha}$ Riemannian manifold with bounded curvature, substrate-derived dynamics, and geodesic transport on its interior. This is the natural candidate for the emergent 4-manifold geometry of macroscopic physics — but the identification requires further structure (dimensional reduction from d_{op} to 4, emergence of Lorentzian signature from refinement-flow asymmetry, derivation of Einstein-like equations from operational continuity) which is left open (§14).

---

12. Relation to the Emergent 4-Manifold

The operational geometric dynamics developed above suggests a structural hypothesis about macroscopic geometric structure:

The 4-manifold geometry of macroscopic physics is the coarse-grained limit of operational distinguishability geometric dynamics on the admissibility substrate.

Within VERSF, the emergent 4-manifold is constructed from substrate primitives (the fold, irreversible commitment events, the BCB and TPB layers, and the $K = 7$ closure architecture). Operational geometric dynamics sits structurally between these primitives and the emergent 4-manifold:

substrate primitives → $\mathbb{Z}_7$ closure architecture → operational Hilbert geometry (previous paper) → operational geometric dynamics (this paper) → $C^{1,\alpha}$ coarse-grained Riemannian limit → emergent 4-manifold.

The coarse-grained limit of §11 supplies a $C^{1,\alpha}$ Riemannian manifold with bounded curvature, geodesic transport (on the interior), continuity conservation, and substrate-derived dynamical equations. The full derivation of the emergent 4-manifold from this limit requires three further steps, none attempted in the present paper:

(a) **Dimensional reduction $d_{op} \to 4$.** Why does the macroscopic continuum carry the specific dimension 4 rather than the full operational dimension $d_{op} = \sum_\alpha \dim V_\alpha$? Candidate mechanisms: stable-channel selection (only 4 of the d_{op} operational dimensions survive coarse-graining as effective macroscopic dimensions), or hierarchical spectral filtering (4 emerges from $K = 7$ closure architecture via specific representation-theoretic counting).

Dimensional bookkeeping for the conformal exponent. The substrate-level conformal factor is $\rho_{op}^{2/d_{op}}$, with exponent $2/d_{op}$ fixed by counting in d_{op}-dimensional operational geometry (Theorem 5.0). The exponent inherited by the macroscopic 4-dimensional metric depends on where in the bridge the dimensional reduction occurs:

- *Reduction before coarse-graining (substrate-level dimensional selection).* If 4 macroscopic dimensions are selected at substrate level and the conformal-factor derivation is re-run within the selected 4-dimensional subspace, the macroscopic exponent is $2/4 = 1/2$.
- *Reduction after coarse-graining (geometric dimensional reduction in the limit).* If the full d_{op}-dimensional g_{op}-geometry is coarse-grained first and dimensional reduction proceeds within the $C^{1,\alpha}$ limit (M_∞, g_∞), the macroscopic metric inherits the d_{op}-dimensional conformal exponent $2/d_{op}$ directly, with the 4-dimensional structure arising as a sub-manifold or quotient.

Which reduction-ordering is the physically correct one — and the corresponding macroscopic conformal exponent — is part of the open question of dimensional reduction (item 9 of §14).

(b) **Refinement-flow asymmetry and a preferred direction in $T_x M_{op}$.** The operational geometry developed here is Riemannian (positive-definite). Macroscopic geometry is Lorentzian (signature 1, 3). The candidate bridge has two steps, with a sharp boundary between what is established and what is open:

Step 1 — From parameter τ to a preferred tangent ray (established here). The refinement parameter τ is a scalar parameter along trajectories of Φ_τ, not directly a direction in $T_x M_{op}$. However, at every $x \in M_{op}$, the velocity field $v_{op}(x) = (d/d\tau\, \Phi_\tau)|_{\tau=0}$ of Definition 10.1 defines a distinguished tangent vector. The refinement-flow asymmetry — Φ_τ is defined for $\tau \geq 0$ only, not for $\tau < 0$ (commitment is irreversible) — picks out $v_{op}(x)$ with a preferred sign at each x. Hence the refinement flow selects, at each point of M_{op}, a preferred ray in $T_x M_{op}$ rather than a full bidirectional line. This breaks the $O(d_{op})$ symmetry of the Riemannian inner product to the stabiliser subgroup $O(1) \times O(d_{op} - 1)$ of the ray.

Step 2 — From preferred ray to Lorentzian signature (Open Problem 8, entangled with Open Problem 9). The conversion of the preferred ray into a Riemannian-to-Lorentzian signature flip on the ray direction — yielding signature $(1, d_{op} - 1)$ on $T_x M_{op}$ — is left as Open Problem 8. This question is entangled with the dimensional-reduction question of §12(a) / Open Problem 9: a one-direction sign flip on the d_{op}-dimensional substrate gives signature $(1, d_{op} - 1)$, but the macroscopic Lorentzian signature (1, 3) requires the additional reduction $d_{op} \to 4$ with the time direction preserved through reduction. Resolving Open Problem 8 therefore presupposes (or co-resolves) Open Problem 9.

(c) **Einstein-like equations from operational continuity.** Theorem 10.3 establishes the operational continuity equation $\partial_\tau \rho_{op} + \nabla_{op} \cdot J_{op} = 0$. *Under coarse-graining, this should become a constraint relating the macroscopic curvature* R_∞ to the macroscopic energy-momentum density (which encodes operational distinguishability density). The structure of this relation — whether it is the standard Einstein equation $R_{\mu\nu} - (1/2) R\, g_{\mu\nu} = 8\pi T_{\mu\nu}$, or some modified form — is determined by the precise coarse-graining of operational dynamics, and is the central open problem of the bridge.

Remark 12.1. This section deliberately avoids reference to "spacetime" — within VERSF, time is emergent, and the macroscopic geometric continuum is the emergent 4-manifold whose construction from substrate primitives is the subject of the geometric emergence programme.

Remark 12.2. The hypothesis that macroscopic geometric structure is *the coarse-grained limit of operational distinguishability dynamics* is what makes the present paper's structural claims falsifiable: if the coarse-grained operational dynamics does not produce a $C^{1,\alpha}$ Riemannian limit, the hypothesis fails; if the limit does not admit a Lorentzian reduction with bounded sectional curvature, the hypothesis fails; if the operational continuity equation cannot be reduced to a curvature–stress-energy relation under coarse-graining, the hypothesis fails. These are concrete structural tests (§13).

---

13. Predictions and Falsifiability

The operational geometric dynamics framework predicts:

- **Operational metric forced by distinguishability counting.** $g_{op}(x) = \rho_{op}(x)^{2/d_{op}} \cdot \langle\cdot,\cdot\rangle_{op}$ is the unique smooth Riemannian metric consistent with the joint conditions of operational-distance-as-quantum-count (Theorem 5.0), $\mathbb{Z}_7$-equivariance, channel-uniformity, and locality (Remarks 5.2.2 and 5.0.3).
- **Conformal-flat structure.** The operational metric is a conformal rescaling of the flat reference by a scalar density-derived factor.
- **Curvature from density gradients.** Operational curvature is non-trivial precisely when ρ_{op} varies; constant ρ_{op} gives $R_{op} = 0$ (Theorem 8.2(i)). The converse — that $R_{op} = 0$ forces ρ_{op} constant — holds in $d_{op} \geq 3$ only under additional admissibility conditions excluding non-trivial Möbius-type solutions of the Liouville equation (Theorem 8.2(ii), Remark 8.2.1).
- **Bounded substrate-level curvature.** Pointwise on M_{op}, $\|R_{op}\| \leq C(d_{op}, \rho_{op}^{max}/\rho_{op}^{min}) / \Delta_{op}^2$ (Theorem 9.1).
- **Non-expansive refinement transport.** Refinement flows Φ_τ are 1-Lipschitz semigroups on $(\mathcal{A}, \langle\cdot,\cdot\rangle_{op})$ (Definition 4.1).
- **Operational continuity equation.** $\partial_\tau \rho_{op} + \nabla_{op} \cdot J_{op} = 0$ (Theorem 10.3).
- **Refinement-driven flattening.** $\rho_{op} \to$ constant under entropy-gradient refinement flow (Corollary 10.4, Remark 10.4.1) — operational geometry relaxes to flat reference under asymptotic refinement.
- **$C^{1,\alpha}$ Riemannian coarse-grained limit.** Coarse-graining at $L \gg \Delta_{op}$ produces a Cheeger–Gromov-regular Riemannian limit with bounded curvature and bounded-below g_{op}-ball volume at resolved scales (Theorem 11.2, Lemma 11.0).

Potential falsifiers include:

- demonstration that local operational distance is not consistent with the distinguishability-quantum count of Theorem 5.0 (direct violation of the operational metric derivation);

- detection of operational curvature exceeding the bound of Theorem 9.1 (violation of finite-packing curvature bound);
- demonstration of refinement transport that is *not* 1-Lipschitz (violation of non-expansive admissible transport);
- failure of the operational continuity equation under refinement flow (violation of Eulerian density conservation);
- demonstration that operational curvature persists uniformly across regions of constant ρ_op (violation of Theorem 8.2(i): constant ρ_op should yield R_op = 0);
- failure of coarse-graining to produce a C^{1,α} Riemannian limit (violation of Cheeger–Gromov convergence — could arise from violation of any of the three hypotheses: curvature bound, diameter bound, or non-collapsing);
- demonstration that g_op-volume of some resolved-scale ball collapses to zero (violation of Lemma 11.0 non-collapsing, forcing the GH limit to lose dimension);
- demonstration that operational geometric dynamics generates non-conformal deformation of the flat reference (violation of distinguishability-counting derivation Theorem 5.0 or $\mathbb{Z}_7$-equivariance argument of Remark 5.2.2);
- demonstration that ρ_op admits regions of vanishing density (violation of kernel-coverage hypothesis), which would degenerate g_op.

Each falsifier corresponds to a violation of one of the prior structural results of OG §2, the operational-metric derivation of Theorem 5.0, or the conformal-deformation theorem (Theorem 5.2) of the present paper.

14. Open Problems

The structural framework of operational geometric dynamics is now in place. The following problems remain open within the broader programme:

1. **Substrate-level derivation of S_op concavity.** Whether the operational entropy concavity hypothesis of Theorem 4.2 follows from admissibility-level substrate primitives, or must be taken as an independent input. Theorem 4.2 and Corollary 10.4 both rest on this hypothesis; its substrate-level justification is therefore central rather than peripheral.
2. **Explicit construction of v_op for physical refinement flows.** Closed-form derivation of the entropy-gradient velocity field v_op = ∇_op S_op for the physical closure manifold M, in terms of substrate primitives.
3. **Quantitative form of κ_op.** The refinement-mobility coefficient κ_op(x) of Definition 10.2 — its explicit form and physical interpretation in terms of substrate kinetics, including whether S_op admits the local density-functional form S_op = ∫ s(ρ_op) d$\mathscr{H}$ assumed in Definition 10.2.
4. **Beyond conformal: channel-mixing deformations.** Whether refinement dynamics admits non-conformal deformations of g_op (e.g. position-dependent channel-mixing terms or channel-asymmetric weights λ_α(x) in the Schur-decomposed direct sum) under conditions where $\mathbb{Z}_7$-equivariance is broken or refined.

5. **Operational Ricci structure.** Closed-form derivation of the operational Ricci tensor Ric_op^{ab} from R_op^a_{bcd}, and its relation to the operational distinguishability density.
6. **Operational Einstein-like equations.** Whether the continuity equation (Theorem 10.3), combined with the conformal-flat curvature formula (Theorem 8.1) and the coarse-graining limit (Theorem 11.2), yields an Einstein-like field equation relating operational curvature to operational stress-energy.
7. **Operational geodesic deviation.** The d²ξ^a/dτ² = − R_op^a_{bcd} v^b v^d ξ^c geodesic deviation equation on (M_op, g_op), and its coarse-grained limit.
8. **Entropy-flow gravity analogues.** Whether the entropy-gradient refinement flow of Theorem 4.2, under coarse-graining, generates a gravity-like attractive structure between operational sectors of high distinguishability density.
9. **Lorentzian signature from refinement-flow asymmetry (entangled with item 10).** Establishing that the preferred-ray structure on T_x M_op induced by refinement-flow asymmetry (§12(b) Step 1) *forces* a Riemannian-to-Lorentzian signature flip on the ray direction — rather than merely permitting one. The candidate mechanism: time-like directions (those aligned with v_op) carry a *cost* under operational distance reflecting irreversible commitment, distinct from spatial distinguishability count. Formally, the candidate signature flip is g_Lor(v_op, v_op) = −g_op(v_op, v_op). This question is entangled with item 10 (dimensional reduction): a sign flip on the d_op-dimensional substrate gives signature (1, d_op − 1), but the macroscopic Lorentzian signature (1, 3) requires the additional reduction d_op → 4 with the time direction preserved through reduction.
10. **Dimensional reduction d_op → 4.** Mechanism by which d_op operational dimensions reduce to 4 macroscopic continuum dimensions under coarse-graining: stable-channel selection, hierarchical spectral filtering, or other representation-theoretic mechanisms internal to the K = 7 closure architecture. Includes the dimensional-bookkeeping question of §12(a): does the macroscopic conformal exponent inherit the substrate value 2/d_op or the reduced value 2/4 = 1/2, depending on whether reduction occurs before or after coarse-graining.
11. **Quantum operational geometric dynamics.** Whether the Hilbert structure (𝒜, ⟨·,·⟩_op) carries a natural quantum-mechanical operator structure on top of the present classical-Riemannian dynamics, consistent with the VERSF quantum reconstruction programme — and whether the resulting quantum-operational dynamics reduces to standard quantum field theory in the coarse-grained limit.
12. **Identification of Δ_op with the Planck length.** Whether the substrate-level distinguishability quantum Δ_op, as a curvature cutoff (Remark 9.1.2), coincides quantitatively under coarse-graining with the Planck scale ℓ_P, accounting for the density-contrast prefactor caveat.
13. **Higher regularity of the coarse-grained limit.** Whether finite packing, combined with additional structural hypotheses on the closure architecture, controls higher covariant derivatives of curvature sufficiently to upgrade the C^{1,α} Cheeger–Gromov limit of Theorem 11.2 to C^k or C^∞.
14. **Relation to information geometry.** Whether the operational metric g_op = ρ_op^{2/d_op} · ⟨·,·⟩_op (Theorem 5.0) reduces, in some macroscopic limit, to a Fisher

information metric — and whether the substrate-level distinguishability counting that derives g_op admits a continuum-statistical interpretation.

15. **Admissibility conditions excluding Möbius solutions.** Identification of necessary and sufficient admissibility conditions on ρ_op (e.g. normalisation invariance, $\mathbb{Z}_7$-equivariance, exclusion of inversion-centre singularities) under which the converse direction of Theorem 8.2 — R_op = 0 ⇒ ρ_op constant — holds, excluding the non-trivial Möbius-type solutions of the Liouville equation in d_op ≥ 3 (Remark 8.2.1). Also: whether the conformal-flat structure of the coarse-grained limit (Remark 11.2.2) preserves or breaks the relevance of Möbius non-trivial solutions.
16. **Substrate-level derivation of channel-uniformity.** Whether the channel-uniformity principle of Remark 5.2.2 — the conformal restriction within the $\mathbb{Z}_7$-equivariant Schur-decomposed class of metric deformations — follows from substrate primitives, or must be taken as an independent assumption alongside $\mathbb{Z}_7$-equivariance and locality. The principle states that a scalar source (ρ_op without channel structure) generates only channel-symmetric metric response, but this is not automatic: fixed channel-dependent couplings c_α could in principle enter, and excluding them requires a substrate-level "no fixed channel-asymmetric coupling" minimality principle.

Each open problem is now a problem on the dynamical operational Hilbert geometry of M, supplied with definite mathematical structure (g_op, ∇_op^{LC}, R_op, Φ_τ, ρ_op, J_op) and definite open targets.

---

15. Conclusion

The previous paper *Operational Geometry in the VERSF Framework* established that the admissible subspace 𝒜 = Im(Ω_max) carries an intrinsic finite-dimensional real Hilbert geometry, derived from admissibility alone via the $\mathbb{Z}_7$ Fourier architecture and the self-adjoint idempotent structure of Ω_max. That paper supplied the kinematics of admissibility geometry.

The present paper supplies the dynamics. The structural backbone is the operational distinguishability metric theorem (Theorem 5.0) together with the substrate-level structural principles of Remark 5.2.2: under the joint conditions of $\mathbb{Z}_7$-equivariance (substrate symmetry), channel-uniformity (scalar source generates only channel-symmetric metric response), and locality (Remark 5.0.3), the unique smooth Riemannian metric on M_op under which operational length counts local distinguishability quanta is

g_op(x) = ρ_op(x)^{2/d_op} · ⟨·,·⟩_op,

with ρ_op the smoothed packing density on M_op at scale ≥ Δ_op. The metric g_op is therefore not postulated but forced by the operational interpretation of distance. The Riemann curvature tensor R_op of g_op is non-trivial precisely when ρ_op varies; constant ρ_op gives R_op = 0, with the converse — vanishing curvature forcing constant density — holding only under additional admissibility conditions excluding non-trivial Möbius solutions of the Liouville equation in d_op ≥ 3.

Every other dynamical structure of the paper follows as a standard finite-dimensional Riemannian fact:

- refinement flow Φ_τ is the non-expansive Ω_max-stable semigroup on (𝒜, ⟨·,·⟩_op), with the entropy-gradient flow as natural physical realisation under the S_op-concavity hypothesis (Theorem 4.2);
- operational geodesics exist locally and minimise g_op-length between sufficiently close interior points (Theorem 6.2, Theorem 6.3); full geodesic completeness requires closed completions or boundary conditions;
- the operational Levi-Civita connection ∇_op^{LC} has closed-form Christoffel symbols (Theorem 7.1);
- the operational Riemann tensor R_op is given in closed form by the conformal-flat curvature formula (Theorem 8.1), with the 2D case exceptional (Remark 8.2.1);
- finite packing P(M) ≤ Vol_op(M)/Δ_op^{d_op} bounds the operational curvature pointwise: ‖R_op‖ ≤ C(d_op, ρ_op^{max}/ρ_op^{min})/Δ_op² (Theorem 9.1), with explicit dependence on the density contrast across M_op;
- refinement transport obeys the operational continuity equation ∂_τ ρ_op + ∇_op · (ρ_op v_op) = 0 (Theorem 10.3), derived as Eulerian density conservation;
- coarse-graining at L ≫ Δ_op produces a C^{1,α} Riemannian limit via Cheeger–Gromov compactness, with bounded curvature preserved (Theorem 11.2).

The result upgrades the VERSF programme from **operational geometry** to **operational geometric dynamics**: admissibility-induced geometry is not merely a kinematic backdrop but a substrate carrying its own dynamical structure. The flat reference Hilbert geometry of admissibility is the unperturbed substrate; the operational interpretation of distance (distinguishability counting) forces conformal deformation by ρ_op^{2/d_op}; distinguishability-density gradients generate curvature; finite packing bounds the curvature; refinement transport drives flattening; and coarse-graining produces a C^{1,α} Riemannian limit — the natural structural candidate for the emergent 4-manifold geometry of macroscopic physics.

The full derivation of the macroscopic continuum from this coarse-grained limit — including dimensional reduction d_op → 4, Lorentzian signature from refinement-flow asymmetry via the v_op preferred-direction mechanism of §12(b), and Einstein-like field equations from operational continuity — remains open (§14). What the present paper establishes is that the *structural ingredients* for such a derivation are now available, derived from admissibility alone, with no macroscopic geometry imported as input.

Admissibility does not merely geometrise. It generates geometric dynamics — and the operational metric, the curvature, and the transport are all forced by what "operational distance" means. The flat reference Hilbert geometry of the admissible subspace bends conformally under distinguishability-density gradients by an exponent determined by quantum counting; the resulting curvature is bounded by the substrate-level distinguishability quantum Δ_op; and coarse-graining produces a C^{1,α} Riemannian limit that is the structural candidate for the macroscopic geometric continuum. Operational distinguishability geometry, dynamics included, is the substrate-side scaffold from which the emergent 4-manifold of macroscopic physics is constructed.