International Journal of Topology

MDPI

Article

Informational Holonomy Curvature and Its Discrete-to-Continuous Convergence

David Gutierrez Ule

Independent Researcher, 34230 Torquemada, Spain; david.gutierrezule@gmail.com

Abstract

We introduce a notion of curvature based on *informational holonomy*. Let (M, g) be a smooth Riemannian manifold and let $\pi : \mathcal{P} \to M$ be a bundle of state spaces equipped fibrewise with a smooth divergence D_x inducing an information metric $g^{\mathcal{P}_x}$. Assuming a connection on $\mathcal{P}$ compatible with this fibrewise information geometry, we measure the deviation of holonomy around small geodesic triangles by transporting a reference state μ_x and comparing it to its image via the induced informational distance $d_x = \sqrt{2D_x}$. Normalizing the resulting distance defect by the geometric area yields a *continuous informational holonomy (sectional) curvature* $K^{\mathrm{cont}}_{\mathrm{hol}}(x, \Pi)$. We prove that this limit exists for all (x, Π) and equals the norm of a vector $W_x(\Pi; \mu_x) \in T_{\mu_x}\mathcal{P}_x$ depending linearly on the curvature of the connection along Π. In geometric models induced from the Levi–Civita connection via an isometric representation, $K^{\mathrm{cont}}_{\mathrm{hol}}$ becomes a scalar invariant of $R^g|_\Pi$ and, on spaces of constant sectional curvature, reduces to a constant multiple of $|\sec_g|$. On the discrete side, we consider quasi-uniform sampling graphs whose edges carry channels approximating parallel transport. Discrete triangle holonomies define a curvature estimator, and under explicit sampling, area-approximation, and channel-consistency assumptions, we establish a discrete-to-continuum convergence theorem with a quantitative error bound controlled by the sampling scale.

Keywords: informational holonomy; holonomy curvature; sectional curvature; information geometry; Jensen–Shannon divergence; Ehresmann connection; state bundle; sampling graphs; discrete-to-continuous convergence

1. Introduction

Curvature is one of the central notions of Riemannian geometry and plays a distinguished role in both mathematics and physics. In its classical incarnation, sectional curvature measures the second-order deviation of geodesics in a two-dimensional direction, while scalar curvature is obtained by averaging over all such directions. From a more global perspective, curvature can also be understood in terms of *holonomy*: transporting a vector around a small closed loop and comparing the result with the original vector. In a flat space, the transported vector returns to itself, whereas in a curved space a non-trivial discrepancy appears, and this discrepancy encodes the curvature of the underlying connection.

In parallel, information geometry has developed a rich differential-geometric framework on spaces of probability distributions and quantum states, where Riemannian metrics arise from statistical divergences such as the Kullback–Leibler or Jensen–Shannon divergences. The associated Fisher metrics and their variants have been used to define information-geometric analogues of geodesics, connections and curvature, with applica-

Check for updates

Academic Editor: Michel Planat

Received: 15 December 2025
Revised: 4 January 2026
Accepted: 12 June 2026
Published: 18 June 2026

Copyright: © 2026 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

tions ranging from statistics and machine learning to quantum information theory; see, e.g., ref. [1].

These two viewpoints—Riemannian curvature and information geometry—suggest a natural question:

Can curvature be reconstructed, or even defined, purely in terms of informational data and their transformations, without a priori access to a smooth Riemannian structure?

In the present paper, we do not attempt to solve this reconstruction problem in full generality. Instead, we assume a smooth Riemannian manifold (M, g) and a state bundle with a compatible connection as part of the input, and show how their curvature can be encoded and approximated in purely informational terms.

This question has received substantial attention in the context of discrete and non-smooth geometries. In the setting of graphs and Markov chains, several notions of "discrete curvature" have been introduced, most notably the coarse Ricci curvature à la Ollivier [2], Forman's combinatorial curvature [3], and Regge-type curvature concentrated on simplices in piecewise flat manifolds [4]. These constructions either derive curvature from a discrete metric structure or from combinatorial and measure-theoretic data. In all cases, an important theme is the *consistency problem*: under suitable sampling or refinement hypotheses, do these discrete curvatures converge to the classical Riemannian curvatures of a smooth limit space?

In a companion preprint [5], the author studied a purely metric special case based on the Jensen–Shannon distance on sampling graphs and a Regge-type estimator targeting scalar curvature. Under explicit sampling and local metric-compatibility hypotheses, that work establishes weak consistency of the resulting discrete curvature measure with the scalar curvature of (M, g). The present paper departs from the metric-only setting by introducing curvature from the holonomy of informational channels, yielding a sectional-type invariant and a discrete-to-continuum convergence theory for this holonomy-based construction.

The present paper aims at a refinement of a different nature. Instead of constructing curvature from distances alone, we introduce a new notion of curvature that is based on *holonomy of informational channels*. Concretely, we consider a discrete sampling of a manifold in which each vertex carries a space of informational states (classical or quantum), and each oriented edge is equipped with a channel transporting states between neighbouring vertices. By composing these channels along small loops, we obtain a discrete notion of holonomy acting on the state space at a base point. We then measure how far this holonomy deviates from the identity, using an informational divergence such as the Jensen–Shannon divergence and the associated informational distance $d = \sqrt{2D}$, and normalize by the geometric area of the loop. The resulting quantity is what we call the *informational holonomy curvature*.

Intuitively, one can think of this construction as follows. At each point, we choose a reference state μ_x of the system. We then let this state travel along a small triangular loop $x \to y \to z \to x$ by applying, in sequence, the channels assigned to each edge. If the underlying "informational geometry" were flat, we would expect the state to come back unchanged. Any systematic deviation between the initial state and the state obtained after completing the loop signals the presence of curvature. By comparing these two states with a suitable divergence (or equivalently, with the associated informational distance) and normalizing by the area of the loop, we obtain a discrete curvature associated with that loop. By averaging over families of loops that approximate a given two-dimensional direction, we obtain a quantity that plays the role of a *sectional* curvature in an informational setting.

From the Riemannian viewpoint, this is reminiscent of the classical description of sectional curvature via holonomy of the Levi–Civita connection in the tangent bundle. The novelty here is that the objects being transported are not tangent vectors but infor-

mational states, and the parallel transport is implemented by channels rather than by a linear connection on a vector bundle. Nevertheless, we shall show that, under appropriate assumptions, the resulting informational holonomy curvature can be expressed in terms of the curvature of a connection on a state bundle over (M, g), and in geometric models induced by the Levi–Civita connection it reduces, in spaces of constant curvature, to a constant multiple of $|\sec_g|$.

1.1. Main Contributions

We now summarize the main contributions of this work.

- We introduce a general framework in which to study *informational holonomy* on a discrete sampling of a Riemannian manifold. The basic data consist of:
  1. A sampling graph $G_\varepsilon = (V_\varepsilon, E_\varepsilon)$ embedded in a manifold (M, g);
  2. A state space $\mathcal{P}_x$ attached to each vertex $x \in V_\varepsilon$ (for instance, a space of probability distributions or density matrices);
  3. A family of channels $\Phi_{xy} : \mathcal{P}_x \to \mathcal{P}_y$ associated with each oriented edge $x \to y \in E_\varepsilon$.

  The composition of these channels along loops in G_ε gives rise to discrete holonomy operators on the fibres $\mathcal{P}_x$.
- Given a divergence D on each state space $\mathcal{P}_x$ (in particular, the Jensen–Shannon divergence), we define the *informational holonomy defect* of a loop γ based at x as

$$\delta_\gamma(x) := D\big(\mu_x, H_\gamma(\mu_x)\big),$$

  where H_γ is the holonomy operator obtained by composing channels along γ, and μ_x is a reference state at the base point x. We also consider the associated distance defect

$$\Delta_\gamma(x) := d\big(\mu_x, H_\gamma(\mu_x)\big), \qquad d(s,t) = \sqrt{2D(s,t)}.$$

  For sufficiently small loops bounding an area A_γ, we define the *informational holonomy curvature* of the loop by

$$K_{\mathrm{hol}}(\gamma) := \frac{\Delta_\gamma(x)}{A_\gamma}.$$

  By averaging over discrete loops that approximate a two-plane $\Pi \subset T_xM$, we obtain an *informational sectional curvature* $K^{(\varepsilon)}_{\mathrm{hol}}(x, \Pi)$ at scale ε.
- On the continuum side, we consider a smooth Riemannian manifold (M, g) equipped with a smooth bundle of state spaces $\mathcal{P} \to M$ and a connection-like structure which assigns, to nearby points $x, y \in M$, channels $\Phi_{x\to y} : \mathcal{P}_x \to \mathcal{P}_y$ compatible with g. Under natural regularity and compatibility conditions, we define a *continuous informational holonomy curvature* by transporting a reference state around infinitesimal loops tangent to a given two-plane $\Pi \subset T_xM$ and measuring, via a divergence, the infinitesimal deviation from the identity. Our first main result shows that this quantity can be written as

$$K^{\mathrm{cont}}_{\mathrm{hol}}(x, \Pi) = \big\| W_x(\Pi; \mu_x) \big\|_{g^{\mathcal{P}_x}},$$

  where $W_x(\Pi; \mu_x)$ depends linearly on the curvature of the connection on $\mathcal{P}$ along Π. In particular, when the connection is induced from the Levi–Civita connection via a linear isometric representation, $K^{\mathrm{cont}}_{\mathrm{hol}}(x, \Pi)$ is a scalar invariant of the restriction of R^g_x to Π, and in spaces of constant sectional curvature it is proportional to $|\sec_g(x, \Pi)|$.
- Second, building on previous work on scalar curvature estimators derived from informational metrics, we show that the discrete informational sectional curvature

$K^{(\varepsilon)}_{\mathrm{hol}}(x,\Pi)$ converges, as $\varepsilon \to 0$, to the continuous quantity $K^{\mathrm{cont}}_{\mathrm{hol}}(x,\Pi)$ above. More precisely, we obtain an estimate of the form

$$\left|K^{(\varepsilon)}_{\mathrm{hol}}(x,\Pi) - K^{\mathrm{cont}}_{\mathrm{hol}}(x,\Pi)\right| \leq C\,\kappa_\varepsilon,$$

where

$$\kappa_\varepsilon := r_\varepsilon + \eta_\varepsilon + q_\varepsilon + \rho_\varepsilon,$$

is the error scale appearing in Theorem 2 (see (10)). Here, r_ε is the sampling radius, η_ε and q_ε quantify directional anisotropy and area-approximation errors of the triangle families, and $\rho_\varepsilon = r_\varepsilon^{\alpha-1}$ encodes the channel-consistency error from Assumption 9(2). In particular, in spaces of constant curvature, this yields convergence to a constant multiple of $|\sec_g(x,\Pi)|$.

- Finally, we discuss several model constructions and examples. In particular, we consider classical Fisher-type models in which the state bundle is induced by a statistical model on M, and we construct channels by transporting distributions along short geodesic segments. On spaces of constant curvature, we show how the resulting informational holonomy curvature reproduces, up to a constant factor, the expected constant sectional curvature.

Together, these results provide what we view as a natural notion of curvature in a setting where the primary objects are informational states and channels rather than tangent vectors and linear connections. In contrast with purely metric-based discrete curvatures, the informational holonomy curvature fundamentally exploits the *dynamical* aspect of information transport.

1.2. Relation to Previous Work

This work lies at the interface between several active research directions.

First, in Riemannian geometry and general relativity, Regge calculus and its refinements provide discrete models of curvature in piecewise flat manifolds, where curvature is concentrated on codimension-2 simplices and can be described in terms of deficit angles and holonomy [4]. Second, in the study of graphs and Markov chains, discrete Ricci and scalar curvatures have been introduced using optimal transport, entropy convexity, and combinatorial Laplacians [2,3]. Third, information geometry provides natural Riemannian metrics and connections on spaces of probability measures and quantum states, together with associated notions of curvature [1].

Our construction is directly inspired by the Regge and holonomy viewpoints, but it is formulated in an intrinsically informational setting. Rather than starting from a discrete metric or Laplacian, we start from a network of informational channels and a divergence on each fibre. At the technical level, our convergence results build on discrete-to-continuum analysis of curvature estimators on sampling graphs, including previous work on scalar curvature estimators derived from informational metrics [5].

A simple model to keep in mind throughout the paper is provided by the classical Jensen–Shannon divergence on probability simplices (Section 7.1), which satisfies Assumption 1 and induces fibre metrics proportional to the Fisher information metric. This example shows that the abstract hypotheses on the fibre divergences are realised in a familiar information-geometric setting.

1.3. Organization of the Paper

This paper is organized as follows. In Section 2, we recall the necessary background on Riemannian geometry, information geometry, and holonomy, and we set up the continuous model of state bundles and channels. In Section 3, we introduce the discrete sampling

framework, define discrete holonomy operators on state spaces, and state the main assumptions on sampling, metric approximation, and channel regularity. Section 4 contains the precise definitions of the informational holonomy defect and informational holonomy curvature, both at the level of individual loops and in the averaged sectional form.

In Section 5, we formulate and prove our main continuous theorem relating informational holonomy curvature to the curvature of a connection on the state bundle. Section 6 is devoted to the discrete-to-continuum convergence theorem. Finally, Section 7 presents model constructions and examples, and Section 8 discusses possible extensions and applications.

2. Geometric and Informational Background

In this section, we recall the basic geometric notions and fix the continuous framework in which the informational holonomy curvature will be defined. We start with standard Riemannian preliminaries, then introduce the state bundle and divergences, and finally describe a continuous notion of informational transport and holonomy.

2.1. Riemannian Preliminaries

Let (M, g) be a smooth, connected, oriented Riemannian manifold of dimension $n \geq 2$. We denote by d_g the associated geodesic distance, by ∇ the Levi–Civita connection of g, and by R^g the Riemann curvature tensor, taken with the convention

$$R^g(X,Y)Z := \nabla_X \nabla_Y Z - \nabla_Y \nabla_X Z - \nabla_{[X,Y]} Z.$$

For $x \in M$ and a 2-dimensional subspace $\Pi \subset T_x M$, the *sectional curvature* of g at (x, Π) is defined by

$$\sec_g(x,\Pi) := \frac{g\big(R^g(u,v)v,u\big)}{\|u\|_g^2\|v\|_g^2 - g(u,v)^2},$$

where $u, v \in T_x M$ is any basis of Π with $u \wedge v \neq 0$. This definition is independent of the choice of basis.

We shall frequently use normal coordinate charts. For $x \in M$, we denote by

$$\exp_x : U_x \subset T_x M \longrightarrow M$$

the exponential map at x, defined on a maximal star-shaped neighbourhood U_x of 0 where $\exp_x$ is a diffeomorphism onto its image. There exists $r_0 > 0$ such that for every $x \in M$ the open ball $B_g(x, r_0)$ is contained in the normal neighbourhood $\exp_x(B(0, r_0))$ and is geodesically convex: any two points $y, z \in B_g(x, r_0)$ are joined by a unique minimizing geodesic contained in $B_g(x, r_0)$.

Given three points $x, y, z \in M$ sufficiently close to each other and contained in such a convex normal neighbourhood, there is a unique geodesic triangle $\Delta_g(x, y, z)$ formed by the minimizing geodesic segments $[xy], [yz], [zx]$. We denote by

$$A_g(x,y,z)$$

the Riemannian area of $\Delta_g(x, y, z)$ and by $\alpha_x(x, y, z)$, $\alpha_y(x, y, z)$ and $\alpha_z(x, y, z)$ the interior angles at x, y, z, respectively. The *angle defect* of the triangle is

$$\mathrm{def}_g(x,y,z) := \alpha_x(x,y,z) + \alpha_y(x,y,z) + \alpha_z(x,y,z) - \pi.$$

The following classical fact relates angle defect and sectional curvature; see, for example, Chapter 6 of [6].

Lemma 1 (Angle defect and sectional curvature)**.** *Let (M, g) be a smooth Riemannian manifold. For each compact set $K \subset M$ there exist constants $C_K > 0$ and $r_K > 0$ such that the following holds. For any $x \in K$ and any geodesic triangle $\Delta_g(x, y, z)$ contained in $B_g(x, r_K)$, let $\Pi \subset T_xM$ be the plane spanned by the initial velocities of the geodesics from x to y and from x to z. Then,*

$$\mathrm{def}_g(x, y, z) = \mathrm{sec}_g(x, \Pi)\, A_g(x, y, z) + O(\ell^3), \tag{1}$$

where $\ell := \max\{d_g(x, y), d_g(y, z), d_g(z, x)\}$ and the implicit constant in the $O(\cdot)$ term is bounded by C_K.

In particular, for sequences of triangles shrinking to x with diameters of order $\ell \to 0$ and areas of order $A_g \sim \ell^2$, the ratio $\mathrm{def}_g(x, y, z)/A_g(x, y, z)$ converges to $\mathrm{sec}_g(x, \Pi)$.

Holonomy gives an alternative description of sectional curvature. Let $\mathrm{Hol}_x(\nabla) \subset \mathrm{O}(T_xM, g_x)$ denote the holonomy group of the Levi–Civita connection at x. For any piecewise smooth closed loop $\gamma : [0, 1] \to M$ based at x, there is an associated parallel transport operator

$$P_\gamma^g : T_xM \to T_xM,$$

obtained by solving the parallel transport equation along γ. If γ bounds a small geodesic triangle tangent to a plane $\Pi \subset T_xM$, then one has the expansion

$$P_\gamma^g = \mathrm{Id}_{T_xM} + A_\gamma\, \mathcal{R}_x^g(\Pi) + O(A_\gamma^{3/2}), \tag{2}$$

where A_γ is the area of the triangle, and $\mathcal{R}_x^g(\Pi)$ is a linear map depending linearly on the curvature tensor R_x^g restricted to Π; see, e.g., Chapter II of [7]. We do not need the precise form of $\mathcal{R}_x^g(\Pi)$; it will be enough to assume the existence of expansions of the form (2) in the informational setting below.

2.2. State Spaces and Informational Divergences

We now recall basic notions from information geometry. For the purposes of this work it is convenient to formulate the discussion at the level of a general smooth state space, although throughout we keep classical probability distributions as a canonical example.

Definition 1 (State space)**.** *A* state space *is a smooth manifold S whose points represent informational states of a system. Typical examples include:*

(i) An open subset of the simplex of strictly positive probability vectors on a finite set $\Omega = \{1, \ldots, m\}$:

$$S \subset \Big\{ p \in \mathbb{R}^m : p_i > 0,\ \sum_{i=1}^m p_i = 1 \Big\}.$$

(ii) An open submanifold of the space of faithful density matrices on a finite-dimensional Hilbert space (quantum states).

We endow S with two pieces of structure:

- A smooth Riemannian metric g^S on S;
- An *informational divergence*

$$D : S \times S \to [0, \infty),$$

which is differentiable of sufficiently high order and vanishes on the diagonal.

The divergence D is not assumed to be a distance (it may fail to be symmetric and may not satisfy the triangle inequality), but we require that it induces the Riemannian metric g^S in the usual way.

Assumption 1 (Divergence and information metric)**.** *We assume that D is of class C^3 in a neighbourhood of the diagonal $\{(s,s) : s \in S\}$ and satisfies:*

1. *$D(s,s) = 0$ for all $s \in S$;*
2. *$\partial_2 D(s,s) = 0$ for all $s \in S$ (vanishing first derivative in the second argument along the diagonal);*
3. *the Hessian of D in the second argument at the diagonal recovers the Riemannian metric g^S, i.e.,*

$$\mathrm{Hess}_2 D(s,s)[w,w] = g_s^S(w,w), \qquad \forall s \in S,\ \forall w \in T_s S.$$

The last condition means that for t close to s one has the second-order expansion

$$D(s,t) = \frac{1}{2} g_s^S(v,v) + O\big(\|v\|_{g^S}^3\big), \tag{3}$$

where $v \in T_s S$ is any tangent vector such that $\exp_s^S(v) = t$ in a normal coordinate chart on S.

Remark 1 (Jensen–Shannon divergence)**.** *A central example in this work is the* Jensen–Shannon divergence *on the simplex of probability vectors on a finite set; see [8]. Let $\Omega = \{1, \dots, m\}$ and let*

$$\Delta^\circ(\Omega) := \Big\{ p \in \mathbb{R}^m : \ p_i > 0,\ \sum_{i=1}^m p_i = 1 \Big\}$$

denote the open probability simplex. For $p, q \in \Delta^\circ(\Omega)$, the Jensen–Shannon divergence is defined by

$$D_{\mathrm{JS}}(p,q) := H\Big(\frac{p+q}{2}\Big) - \frac{1}{2}H(p) - \frac{1}{2}H(q),$$

where $H(p) = -\sum_{i=1}^m p_i \log p_i$ is the Shannon entropy (with a fixed choice of logarithm). It is well known that D_{JS} is symmetric and non-negative and that

$$\sqrt{D_{\mathrm{JS}}(p,q)}$$

defines a genuine metric on $\Delta^\circ(\Omega)$ [9]. In particular, with the normalization used throughout this paper,

$$d_{\mathrm{JS}}(p,q) := \sqrt{2\, D_{\mathrm{JS}}(p,q)} = \sqrt{2}\, \sqrt{D_{\mathrm{JS}}(p,q)}$$

is also a genuine metric on $\Delta^\circ(\Omega)$.

Moreover, the second-order expansion of D_{JS} around the diagonal induces a multiple of the Fisher information metric on $\Delta^\circ(\Omega)$. In other words, there exists a constant $c_{\mathrm{JS}} > 0$ such that, for p fixed and q close to p,

$$D_{\mathrm{JS}}(p,q) = \frac{c_{\mathrm{JS}}}{2} g_p^{\mathrm{Fisher}}(v,v) + O\big(\|v\|_{g^{\mathrm{Fisher}}}^3\big),$$

where $v \in T_p\Delta^\circ(\Omega)$ is as in (3). Thus, Assumption 1 holds with g^S proportional to the Fisher metric.

Assumption 1 ensures that the divergence D and the metric g^S are compatible in the sense of information geometry: D is a "potential" whose Hessian yields the local quadratic structure. The precise constant relating D to g^S will play no essential role; it will simply be absorbed into the constant c appearing in the main theorems.

2.3. State Bundles over a Riemannian Manifold

We now couple the Riemannian manifold (M, g) with the state space S.

Definition 2 (State bundle)**.** *A* state bundle *over (M, g) is a smooth fibre bundle*

$$\pi : \mathcal{P} \to M$$

with typical fibre S, together with:

- *For each $x \in M$, a smooth identification of the fibre $\mathcal{P}_x := \pi^{-1}(x)$ with S;*
- *A smooth family of fibrewise Riemannian metrics $(g^{\mathcal{P}_x})_{x \in M}$, where each $g^{\mathcal{P}_x}$ is a Riemannian metric on $\mathcal{P}_x$ obtained from a copy of g^S under the identification $\mathcal{P}_x \simeq S$;*
- *A smooth family of divergences $(D_x)_{x \in M}$, where each*

$$D_x : \mathcal{P}_x \times \mathcal{P}_x \to [0, \infty)$$

satisfies Assumption 1 with respect to $g^{\mathcal{P}_x}$.

For notational simplicity, we often suppress the explicit identifications and denote the metric on $\mathcal{P}_x$ by g^S_x and the divergence by D_x, keeping in mind that they vary smoothly with x.

Remark 2. *In the simplest classical setting, one may take $\mathcal{P} = M \times S$ to be the trivial bundle with fibre $S = \Delta^\circ(\Omega)$, endowed with the product smooth structure. Then, each fibre $\mathcal{P}_x$ is naturally identified with S, and one can set $g^{\mathcal{P}_x} = g^S$ and $D_x = D$ independently of x. The non-trivial bundle case is more appropriate, for instance, in quantum settings or in statistical models where the parameterization of states varies with x.*

The state bundle $\pi : \mathcal{P} \to M$ provides the configuration space for informational states over M. To speak about transport and holonomy of states, we need a notion of connection on $\mathcal{P}$.

2.4. Connections and Continuous Informational Transport

The key geometric input in the continuous theory is a connection producing parallel transport on the state bundle. In order to place the holonomy expansion used later on a fully standard footing, we henceforth restrict the continuous framework to the associated-bundle setting of principal connections.

Assumption 2 (Associated-bundle model and regularity)**.** *There exist:*

(i) *A Lie group G acting smoothly on the left on the typical fibre S by Riemannian isometries of (S, g^S). We denote the action by $(g, s) \mapsto g \cdot s$.*

(ii) *A C^3 principal G-bundle $\pi_{\mathcal{Q}} : \mathcal{Q} \to M$ endowed with a C^2 principal connection $\omega \in \Omega^1(\mathcal{Q}; \mathfrak{g})$, with curvature $F_\omega \in \Omega^2(\mathcal{Q}; \mathfrak{g})$.*

(iii) *The state bundle is the associated bundle*

$$\pi : \mathcal{P} := \mathcal{Q} \times_G S \longrightarrow M,$$

and the connection on $\mathcal{P}$ is the Ehresmann connection induced by ω.

Moreover, all geometric structures are understood on compact subsets of M, so that (after choosing local trivializations) the coefficients of ω, F_ω, and their first derivatives are uniformly bounded on compact sets.

Under Assumption 2, the principal connection ω induces an Ehresmann connection on $\pi : \mathcal{P} \to M$ by declaring a vector in $T_{[q,s]}\mathcal{P}$ to be horizontal if it is represented by a pair $(\dot{q}, \dot{s})$ with $\dot{q}$ horizontal in $\mathcal{Q}$ (i.e., $\omega(\dot{q}) = 0$) and $\dot{s} = 0$ in S. This yields a smooth horizontal distribution $T\mathcal{P} = \mathcal{H} \oplus \mathcal{V}$ and, therefore, parallel transport along curves in M.

Definition 3 (Parallel transport in the associated state bundle)**.** *Let $\gamma : [0,1] \to M$ be a piecewise C^1 curve with $\gamma(0) = x$ and $\gamma(1) = y$. Let $q_0 \in \mathcal{Q}_x$ and denote by $q(t)$ the ω-horizontal lift of γ with $q(0) = q_0$. There exists a unique element $g_\gamma(q_0) \in G$ such that $q(1) = q_0 \cdot g_\gamma(q_0)$. Then, the induced parallel transport on $\mathcal{P}$ is the map Here, $[q, s]$ denotes the equivalence class of $(q, s) \in \mathcal{Q} \times S$ under the relation $(q \cdot g,\, s) \sim (q,\, g \cdot s)$, $g \in G$.*

$$\mathrm{PT}_\gamma : \mathcal{P}_x \to \mathcal{P}_y, \qquad \mathrm{PT}_\gamma([q_0, s]) := [q(1),\, s] = [q_0,\, g_\gamma(q_0) \cdot s].$$

Notation 1. *For $x \in M$, we denote by d_x^{R} the Riemannian distance on the fibre $(\mathcal{P}_x, g^{\mathcal{P}_x})$.*

Parallel transport satisfies the functorial properties

$$\mathrm{PT}_{\gamma_2 \star \gamma_1} = \mathrm{PT}_{\gamma_2} \circ \mathrm{PT}_{\gamma_1}, \qquad \mathrm{PT}_{\gamma^{-1}} = \mathrm{PT}_\gamma^{-1}, \qquad \mathrm{PT}_{\mathrm{constant}} = \mathrm{Id}, \tag{4}$$

whenever the concatenations are defined. In particular, any piecewise smooth closed loop γ based at x yields a holonomy map $\mathrm{Hol}_\gamma := \mathrm{PT}_\gamma : \mathcal{P}_x \to \mathcal{P}_x$.

Lemma 2 (Parallel transport is a fibrewise Riemannian isometry)**.** *Assume the associated-bundle setting of Assumption 2, and that the G-action on (S, g^S) is by Riemannian isometries. Then, for every piecewise C^1 curve $\gamma : [0,1] \to M$ with $\gamma(0) = x$, $\gamma(1) = y$, the induced parallel transport map $\mathrm{PT}_\gamma : \mathcal{P}_x \to \mathcal{P}_y$ is a Riemannian isometry between the fibres:*

$$d_y^{\mathrm{R}}\big(\mathrm{PT}_\gamma(s), \mathrm{PT}_\gamma(t)\big) = d_x^{\mathrm{R}}(s,t) \qquad \forall\, s, t \in \mathcal{P}_x.$$

In particular, PT_γ is locally distance-preserving and globally distance-preserving on each fibre.

Proof. Fix $q_0 \in \mathcal{Q}_x$ and let $q(t)$ be the ω-horizontal lift of γ with $q(0) = q_0$. By Definition 3, there is $g_\gamma \in G$ such that $\mathrm{PT}_\gamma([q_0, s]) = [q_0, g_\gamma \cdot s]$. Since g_γ acts by a Riemannian isometry of (S, g^S), it preserves the Riemannian distance on S, and hence PT_γ preserves the Riemannian distance on the fibres $\mathcal{P}_x \simeq S$ and $\mathcal{P}_y \simeq S$. □

Assumption 3 (Compatibility of divergences with transport)**.** *For each compact set $K \subset M$ there exist constants $r_K > 0$ and $C_K > 0$ such that the following holds. For every piecewise C^1 curve $\gamma : [0,1] \to K$ with $\gamma(0) = x$ and $\gamma(1) = y$, and every $s, t \in \mathcal{P}_x$ satisfying $d_x^{\mathrm{R}}(s,t) \le r_K$, one has*

$$\big|D_y(\mathrm{PT}_\gamma(s), \mathrm{PT}_\gamma(t)) - D_x(s,t)\big| \le C_K\, d_x^{\mathrm{R}}(s,t)^3, \tag{5}$$

where d_x^{R} denotes the Riemannian distance in the fibre $(\mathcal{P}_x, g^{\mathcal{P}_x})$.

Remark 3. *Assumption 3 is* local *in the fibre variables: it is only required for pairs (s,t) with $d_x^{\mathrm{R}}(s,t) \le r_K$, and it quantifies how the chosen family of divergences varies under parallel transport on compact subsets of M. If the fibre divergence is induced by a single divergence D on S which is G-invariant (i.e., $D(g \cdot s, g \cdot t) = D(s,t)$ for all $g \in G$), then (5) holds with equality (hence $C_K = 0$). More generally, when D is C^3 near the diagonal and PT_γ is a fibrewise Riemannian isometry, the cubic order on the right-hand side is consistent with the fact that the first and second derivatives of D along the diagonal are fixed by the induced information metric; the inequality can thus be viewed as a mild regularity/modelling axiom rather than a global invariance requirement.*

2.5. Informational Holonomy in the Continuous Setting

In the continuous setting, holonomy of the state bundle is defined exactly as in the Riemannian case: for a piecewise smooth closed loop $\gamma : [0,1] \to M$ based at x, we have a holonomy map

$$\mathrm{Hol}_\gamma : \mathcal{P}_x \to \mathcal{P}_x.$$

We interpret Hol_γ as an *informational channel* (In information theory, a (classical) channel is a map that sends input probability distributions to output distributions (equivalently, a Markov kernel), and in quantum information it is a completely positive trace-preserving map. Here we use the term in this broad sense: parallel transport provides a canonical state-to-state transformation along a path). acting on the state space at x: starting from a state $s \in \mathcal{P}_x$, one transports s along γ using the connection and compares the resulting state $\mathrm{Hol}_\gamma(s)$ with the original one.

To quantify the deviation from trivial holonomy, we introduce reference states and informational defects.

Definition 4 (Reference states)**.** *A reference state field is a smooth section*

$$\mu : M \to \mathcal{P}$$

of the state bundle, i.e., $\pi \circ \mu = \mathrm{id}_M$. For each $x \in M$, the point $\mu_x := \mu(x) \in \mathcal{P}_x$ will serve as the base state with respect to which informational changes are measured.

Definition 5 (Continuous informational holonomy defect)**.** *Let μ be a reference state field and let γ be a piecewise smooth closed loop in M based at x. The informational holonomy defect of γ at x is*

$$\delta_\gamma(x) := D_x\big(\mu_x,\ \mathrm{Hol}_\gamma(\mu_x)\big),$$

where D_x is the divergence on the fibre $\mathcal{P}_x$.

We also consider the associated (local) informational distance defect:

$$\Delta_\gamma(x) := d_x\big(\mu_x,\ \mathrm{Hol}_\gamma(\mu_x)\big), \qquad d_x(s,t) = \sqrt{2D_x(s,t)}.$$

In this continuous framework, we will be interested in loops which are the boundaries of small geodesic triangles. Let $x \in M$ and let $\Pi \subset T_xM$ be a two-dimensional subspace. For $u, v \in \Pi$ sufficiently small, we consider the geodesic triangle with vertices

$$x, \quad y = \exp_x(u), \quad z = \exp_x(v),$$

and denote by $\gamma_{x,u,v}$ the corresponding closed loop obtained by traversing the geodesic segments $x \to y \to z \to x$ in order. Its Riemannian area is $A_g(x,y,z)$.

Definition 6 (Continuous informational holonomy curvature)**.** *Let $(M, g, \mathcal{P}, \mu)$ be as above. For $x \in M$ and a two-dimensional subspace $\Pi \subset T_xM$, consider geodesic triangles based at x and tangent to Π with area $A \to 0$. We say that the continuous informational holonomy curvature at (x, Π) exists if the limit*

$$K^{\mathrm{cont}}_{\mathrm{hol}}(x,\Pi) := \lim_{A \to 0} \frac{\Delta_{\gamma_{x,u,v}}(x)}{A_g(x,y,z)} \tag{6}$$

exists and is independent of the particular way in which the triangle shrinks to x within Π.

Remark 4 (Local versus finite-scale defects)**.** *The limit in* (6) *is an* infinitesimal *(small-loop) quantity, hence only the behaviour of the divergence near the diagonal—equivalently, the induced information metric—enters the leading-order curvature term. For applications at a fixed (non-infinitesimal) scale, one may also consider the* finite-scale *holonomy defects*

$$\delta_\gamma(x) := D_x\big(\mu_x, \mathrm{Hol}_\gamma(\mu_x)\big), \qquad \Delta_\gamma^{\mathrm{sym}}(x) := \sqrt{D_x(\mu_x, \mathrm{Hol}_\gamma(\mu_x)) + D_x(\mathrm{Hol}_\gamma(\mu_x), \mu_x)},$$

and normalize by the enclosed area. When D is symmetric and induces a genuine metric at finite scale (e.g., Jensen–Shannon), Δ_γ can be used non-locally without passing to a second-order approximation.

In later sections, we will show, under appropriate assumptions, that this limit exists and is proportional to $\mathrm{sec}_g(x, \Pi)$, with a constant factor depending on the chosen connection, divergence, and reference state field.

Remark 5. *The definition above is closely analogous to the classical definition of sectional curvature via holonomy of the Levi–Civita connection, with the crucial difference that we compare* states *in a non-linear state space using an informational divergence/distance, rather than tangent vectors using a linear norm. Assumptions 1 and 3 ensure that, to second order, the informational defect behaves quadratically in the infinitesimal holonomy, while the associated distance defect scales linearly in the holonomy displacement, hence is the natural object to normalize by area.*

The continuous framework developed in this section provides the conceptual target for the discrete constructions that follow. In the next section, we introduce sampling graphs, discrete channels, and discrete holonomy operators, which will serve as the basis for our definition of discrete informational holonomy curvature.

3. Discrete Sampling, Channels and Holonomy

In this section, we introduce the discrete framework that will serve as the basis for our definition of informational holonomy curvature. We consider sampling graphs embedded in a Riemannian manifold, discrete state spaces attached to vertices, channels attached to edges, and discrete holonomy operators obtained by composing channels along loops. The assumptions formulated here are discrete counterparts of the continuous structures described in Section 2.

3.1. Sampling Graphs on a Riemannian Manifold

Let (M, g) be a smooth Riemannian manifold of dimension $n \geq 2$. For each small parameter $\varepsilon > 0$, we are given a finite subset

$$V_\varepsilon = \{x_i\}_{i \in I_\varepsilon} \subset M$$

of sampling points and a simple undirected graph

$$G_\varepsilon = (V_\varepsilon, E_\varepsilon),$$

where $E_\varepsilon \subset \{\{x_i, x_j\} : i \neq j\}$ is the set of edges. We denote by $x_i \sim x_j$ the fact that $\{x_i, x_j\} \in E_\varepsilon$.

The vertex sets V_ε are assumed to be asymptotically dense and quasi-uniform in M, in the following sense.

Assumption 4 (Quasi-uniform sampling)**.** *There exist positive constants c_1, c_2 and a sequence $r_\varepsilon \to 0$ as $\varepsilon \to 0$ such that, for all sufficiently small ε:*

1. *(Separation) For all distinct $x_i, x_j \in V_\varepsilon$ one has*

$$d_g(x_i, x_j) \geq c_1 r_\varepsilon.$$

2. *(Covering) For every $x \in M$ there exists $x_i \in V_\varepsilon$ such that*

$$d_g(x, x_i) \leq c_2 r_\varepsilon.$$

Thus, the sampling points form a Delone set at scale r_ε (see, e.g., [10]). We shall refer to r_ε as the *sampling radius*.

The edge set E_ε is assumed to connect points that are at distance of order r_ε.

Assumption 5 (Local connectivity)**.** *There exists a constant $c_3 > 0$ such that, for all sufficiently small ε, one has:*

1. *If $d_g(x_i, x_j) \leq c_3 r_\varepsilon$, then $\{x_i, x_j\} \in E_\varepsilon$.*
2. *The degrees of the graph G_ε are uniformly bounded: there exists $D < \infty$ such that every vertex $x_i \in V_\varepsilon$ satisfies*

$$\deg(x_i) \leq D.$$

Under Assumptions 4 and 5, each connected component of G_ε has uniformly bounded local complexity and provides a reasonable discrete approximation of (M, g) at scale r_ε. In particular, for r_ε sufficiently small, any point $x \in M$ and any direction $\xi \in T_x M$ admit neighbours of x whose geodesic directions approximate ξ up to an error of order r_ε.

In order to average quantities defined at vertices, we will occasionally associate a volume weight $w_i^{(\varepsilon)}$ to each vertex $x_i \in V_\varepsilon$.

Assumption 6 (Volume approximation)**.** *For each ε there exist positive weights $w_i^{(\varepsilon)} > 0$, $x_i \in V_\varepsilon$, such that:*

1. *There exists $C > 0$ independent of ε with*

$$C^{-1} r_\varepsilon^n \leq w_i^{(\varepsilon)} \leq C r_\varepsilon^n \qquad \text{for all } x_i \in V_\varepsilon.$$

2. *For every $f \in C_c(M)$, one has the quadrature convergence*

$$\left| \sum_{x_i \in V_\varepsilon} f(x_i)\, w_i^{(\varepsilon)} - \int_M f(x)\, \mathrm{dvol}_g(x) \right| \longrightarrow 0 \qquad \text{as } \varepsilon \to 0.$$

Assumption 6 is satisfied, for instance, if $\{w_i^{(\varepsilon)}\}$ are defined as the Riemannian volumes of a Voronoi tessellation associated with V_ε.

3.2. Discrete State Spaces and Divergences

We now discretize the state bundle $\pi : \mathcal{P} \to M$ introduced in Section 2.3. For each sampling point $x_i \in V_\varepsilon$, we consider a fibre $\mathcal{P}_i^{(\varepsilon)}$ representing the possible states at x_i. In the simplest setting, one may take

$$\mathcal{P}_i^{(\varepsilon)} = \mathcal{P}_{x_i}$$

to be the fibre of the continuous state bundle at x_i, but we keep the notation $\mathcal{P}_i^{(\varepsilon)}$ to emphasize the discrete nature of the sampling.

Assumption 7 (Discrete fibres and divergences)**.** *For each $\varepsilon > 0$ and each $x_i \in V_\varepsilon$ we are given:*

1. *A smooth manifold* $\mathcal{P}_i^{(\varepsilon)}$ *and a smooth identification*

$$\iota_i^{(\varepsilon)} : \mathcal{P}_i^{(\varepsilon)} \to \mathcal{P}_{x_i}$$

with the continuous fibre at x_i*;*
2. *A Riemannian metric* $g^{\mathcal{P}_i^{(\varepsilon)}}$ *on* $\mathcal{P}_i^{(\varepsilon)}$ *obtained as the pullback of the metric* $g^{\mathcal{P}_{x_i}}$ *under* $\iota_i^{(\varepsilon)}$*;*
3. *An informational divergence*

$$D_i^{(\varepsilon)} : \mathcal{P}_i^{(\varepsilon)} \times \mathcal{P}_i^{(\varepsilon)} \to [0, \infty)$$

such that, after transporting it to $\mathcal{P}_{x_i}$ *via* $\iota_i^{(\varepsilon)}$*, Assumption 1 holds uniformly in* x_i *and* ε*.*

For notational simplicity, we will usually drop the superscript (ε) and write $\mathcal{P}_i$, $g^{\mathcal{P}_i}$ and D_i, keeping in mind the dependence on ε through the underlying sampling set V_ε.

We also sample the continuous reference state field $\mu : M \to \mathcal{P}$.

Assumption 8 (Discrete reference states)**.** *For each* $x_i \in V_\varepsilon$*, we are given a reference state* $\mu_i^{(\varepsilon)} \in \mathcal{P}_i^{(\varepsilon)}$ *such that, under the identification* $\iota_i^{(\varepsilon)}$*, one has*

$$\iota_i^{(\varepsilon)}\big(\mu_i^{(\varepsilon)}\big) = \mu_{x_i}.$$

Thus, the discrete fibres, metrics, divergences, and reference states are obtained by restriction of the continuous state bundle to the sampling set V_ε.

3.3. Discrete Channels and Local Consistency

We now assign discrete informational channels to the edges of the sampling graph G_ε. For each oriented edge (x_i, x_j) with $\{x_i, x_j\} \in E_\varepsilon$, we are given a channel

$$\Phi_{ij}^{(\varepsilon)} : \mathcal{P}_i^{(\varepsilon)} \to \mathcal{P}_j^{(\varepsilon)}.$$

Again, we often drop the superscript (ε) when no confusion arises.

The channels Φ_{ij} are required to be local and to approximate the parallel transport maps in the continuous state bundle. Let γ_{ij} denote the unique minimising geodesic segment in (M, g) joining x_i to x_j, which lies in a convex normal neighbourhood when $d_g(x_i, x_j)$ is sufficiently small. By Assumption 5, we may (and do) assume that each edge length $d_g(x_i, x_j)$ is bounded by a constant multiple of r_ε.

Assumption 9 (Locality and consistency of channels)**.** *There exist constants* $C > 0$ *and* $\alpha > 1$ *such that, for all sufficiently small* ε *and all oriented edges* (x_i, x_j)*:*

1. *(Locality) The channel* Φ_{ij} *depends only on the geometry of* $(M, g, \mathcal{P})$ *in a neighbourhood of the geodesic segment* γ_{ij} *and satisfies*

$$D_j\big(\Phi_{ij}(s), \Phi_{ij}(t)\big) \le C\, D_i(s, t)$$

for all $s, t \in \mathcal{P}_i$*, i.e.,* Φ_{ij} *is (locally) Lipschitz with respect to the divergences.*
2. *(Consistency with continuous parallel transport) Let* $\mathrm{PT}_{\gamma_{ij}} : \mathcal{P}_{x_i} \to \mathcal{P}_{x_j}$ *be the continuous parallel transport map along the geodesic segment* γ_{ij} *joining* x_i *to* x_j*. There exist constants* $C > 0$ *and* $\alpha > 1$ *such that, for all* $s \in \mathcal{P}_i$*,*

$$d_j^{\mathrm{R}}\Big(\Phi_{ij}(s),\, \mathrm{PT}_{\gamma_{ij}}(s)\Big) \le C\, d_g(x_i, x_j)^{1+\alpha}.$$

The exponent $\alpha > 1$ ensures that the error in approximating continuous parallel transport along an edge is of order strictly higher than the edge length, which will imply that the error in approximating holonomy around small loops is of order strictly higher than the area of the loop.

Remark 6. *In many concrete models, one can take $\alpha = 1$ or $\alpha = 2$, depending on how the channels are constructed from the underlying connection on $\mathcal{P}$. For the purposes of the convergence theorems, any $\alpha > 1$ suffices.*

3.4. Discrete Holonomy Operators

Given the channels on edges, we can define discrete holonomy operators by composition along loops in the graph G_ε.

Definition 7 (Discrete paths and loops)**.** *A discrete path in G_ε of length $m \geq 1$ is a sequence of vertices*

$$\gamma = (x_{i_0}, x_{i_1}, \ldots, x_{i_m})$$

such that $\{x_{i_{k-1}}, x_{i_k}\} \in E_\varepsilon$ for all $k = 1, \ldots, m$. The path is closed or a loop if $x_{i_m} = x_{i_0}$.

We denote by $\mathrm{Paths}_\varepsilon$ the set of all discrete paths and by $\mathrm{Loops}_\varepsilon(x_i)$ the set of all loops based at a given vertex x_i.

To each oriented edge $(x_{i_{k-1}}, x_{i_k})$ along a path, we associate the channel $\Phi_{i_{k-1}i_k}$. The channel associated with a path is the composition of the edge channels in order.

Definition 8 (Discrete transport and holonomy)**.** *Let $\gamma = (x_{i_0}, \ldots, x_{i_m})$ be a discrete path in G_ε. The discrete transport along γ is the map*

$$T_\gamma^{(\varepsilon)} := \Phi_{i_{m-1}i_m} \circ \cdots \circ \Phi_{i_1 i_2} \circ \Phi_{i_0 i_1}.$$

If γ is a loop based at x_{i_0}, i.e., $x_{i_m} = x_{i_0}$, we call $T_\gamma^{(\varepsilon)} : \mathcal{P}_{i_0} \to \mathcal{P}_{i_0}$ the discrete holonomy operator of γ and denote it by

$$\mathrm{Hol}_\gamma^{(\varepsilon)} := T_\gamma^{(\varepsilon)}.$$

The discrete transport operators satisfy the obvious composition rules: if $\gamma_1 = (x_{i_0}, \ldots, x_{i_m})$ and $\gamma_2 = (x_{i_m}, \ldots, x_{i_{m+\ell}})$ are two paths with matching endpoint and starting point, then

$$T_{\gamma_2 \star \gamma_1}^{(\varepsilon)} = T_{\gamma_2}^{(\varepsilon)} \circ T_{\gamma_1}^{(\varepsilon)}.$$

If $\bar{\gamma} = (x_{i_m}, \ldots, x_{i_0})$ denotes the reversed path, then

$$T_{\bar{\gamma}}^{(\varepsilon)} \approx \left(T_\gamma^{(\varepsilon)}\right)^{-1},$$

with the approximation becoming exact if the channels satisfy an exact involutive property. In our setting, we will only need approximate inversion properties at the infinitesimal level, which follow from Assumption 9.

In the sequel, we will focus on loops associated with small discrete triangles.

Definition 9 (Discrete triangles)**.** *A discrete triangle in G_ε is an ordered triple of distinct vertices (x_i, x_j, x_k) such that all three edges $\{x_i, x_j\}$, $\{x_j, x_k\}$ and $\{x_k, x_i\}$ belong to E_ε. The associated oriented loop is*

$$\gamma_{ijk} := (x_i, x_j, x_k, x_i).$$

We denote by $\mathcal{T}_\varepsilon$ the set of all discrete triangles in G_ε.

For each discrete triangle (x_i, x_j, x_k), we define the corresponding holonomy operator

$$\mathrm{Hol}_{ijk}^{(\varepsilon)} := \mathrm{Hol}_{\gamma_{ijk}}^{(\varepsilon)} = \Phi_{ki} \circ \Phi_{jk} \circ \Phi_{ij} : \mathcal{P}_i \to \mathcal{P}_i.$$

3.5. Triangle Geometry and Area Approximation

In order to compare discrete holonomy around triangles with continuous holonomy around geodesic triangles, we need to relate the combinatorial triangles in G_ε to small geodesic triangles in (M, g) and to assign an appropriate area to each discrete triangle.

Given a discrete triangle $(x_i, x_j, x_k) \in \mathcal{T}_\varepsilon$, we consider the unique geodesic triangle $\Delta_g(x_i, x_j, x_k)$ formed by the minimising geodesics between x_i, x_j, x_k. For ε sufficiently small, Assumption 5 ensures that all edge lengths $d_g(x_i, x_j)$, $d_g(x_j, x_k)$, and $d_g(x_k, x_i)$ are bounded by a constant multiple of r_ε, so that $\Delta_g(x_i, x_j, x_k)$ is contained in a convex normal neighbourhood. We denote by

$$A_g(i, j, k) := A_g(x_i, x_j, x_k)$$

the Riemannian area of this geodesic triangle.

In purely informational settings, one may not have direct access to $A_g(i, j, k)$, but for the analytical convergence results we only require that the area used in the discrete curvature is a good approximation of $A_g(i, j, k)$.

Assumption 10 (Area approximation)**.** *For each triangle $(x_i, x_j, x_k) \in \mathcal{T}_\varepsilon$, we are given a non-negative number $A_\varepsilon(i, j, k)$, called its discrete area, such that there exists a sequence $q_\varepsilon \to 0$ with*

$$\left|A_\varepsilon(i, j, k) - A_g(i, j, k)\right| \leq q_\varepsilon \, r_\varepsilon^2$$

for all $(x_i, x_j, x_k) \in \mathcal{T}_\varepsilon$ when ε is sufficiently small.

The factor r_ε^2 reflects the typical area of a small triangle with edge lengths of order r_ε.

3.6. Families of Triangles Approximating Two-Planes

To define discrete sectional curvatures, we will need to average over families of discrete triangles that approximate a given point and two-plane in the manifold. This requires an additional isotropy assumption on the sampling.

Let $K \subset M$ be a fixed compact subset. For each $x \in K$ and each two-dimensional subspace $\Pi \subset T_x M$, we wish to consider families of discrete triangles based at vertices close to x whose geodesic realisations are small and whose edge directions approximate Π in an approximately isotropic fashion.

Assumption 11 (Directional sampling and triangle families)**.** *There exists a sequence $\eta_\varepsilon \to 0$ such that the following holds. For each compact $K \subset M$, there exists $\varepsilon_K > 0$ such that, for all $0 < \varepsilon < \varepsilon_K$, for every $x \in K$, and every two-dimensional subspace $\Pi \subset T_x M$, one can choose:*

- *A vertex $x_i \in V_\varepsilon$ with $d_g(x, x_i) \leq c_2 r_\varepsilon$;*
- *A finite non-empty set of triangles*

$$\mathcal{T}_\varepsilon(x_i, \Pi) \subset \{(x_i, x_j, x_k) \in \mathcal{T}_\varepsilon\}$$

such that:

1. *(Scale and non-degeneracy) For all $(x_i, x_j, x_k) \in \mathcal{T}_\varepsilon(x_i, \Pi)$, all edge lengths are bounded by $C r_\varepsilon$ and bounded below by $c r_\varepsilon$ for some constants $0 < c \leq C < \infty$ independent of*

x, Π, ε. *Moreover, the associated geodesic triangle is uniformly non-degenerate: there exists* $c_A > 0$ *(independent of* x, Π, ε*) such that*

$$A_g(i,j,k) \;\geq\; c_A\, r_\varepsilon^2,$$

and consequently (by Assumption 10) also $A_\varepsilon(i,j,k) \geq \frac{c_A}{2} r_\varepsilon^2$ *for* ε *small.*

2. *(Planarity) Let* $u_{ij}, u_{ik} \in T_xM$ *denote the initial velocity vectors of the geodesics from* x *to* x_j *and* x *to* x_k *(transported back to* T_xM *via parallel transport if necessary). Then, the angle between the plane spanned by* $\{u_{ij}, u_{ik}\}$ *and* Π *is at most* η_ε.
3. *(Directional isotropy) The distribution of directions of the edges incident to* x_i *within* $\mathcal{T}_\varepsilon(x_i, \Pi)$ *is approximately isotropic in* Π*, in the sense that for any continuous function* φ *on the unit circle of* Π *one has*

$$\left| \frac{1}{|\mathcal{T}_\varepsilon(x_i,\Pi)|} \sum_{(x_i,x_j,x_k)\in\mathcal{T}_\varepsilon(x_i,\Pi)} \varphi(\theta_{ij}) - \frac{1}{2\pi}\int_0^{2\pi} \varphi(\theta)\,\mathrm{d}\theta \right| \longrightarrow 0,$$

as $\varepsilon \to 0$*, where* θ_{ij} *denotes the direction of the geodesic from* x *to* x_j *projected onto* Π *and normalized.*

Assumption 11 is a discrete isotropy condition ensuring that the graph carries enough small, uniformly non-degenerate triangles to sample each two-plane uniformly in the limit $\varepsilon \to 0$. It is analogous to assumptions used in the analysis of discrete Laplace–Beltrami operators and curvature estimators on random or quasi-uniform point clouds.

The structures and assumptions introduced in this section—sampling graphs, discrete fibres and divergences, channels, holonomy operators, triangle areas, and isotropic triangle families—provide the discrete environment in which we will define the informational holonomy curvature. In the next section, we use these ingredients to formulate precise discrete and continuous curvature quantities and to state the main convergence results.

4. Informational Holonomy Curvature: Definitions

In this section, we define the informational holonomy defect and the associated curvature, both in the discrete setting of sampling graphs and in the continuous state-bundle setting. A key point, already implicit in Assumption 1, is that the divergence on each fibre induces a Riemannian metric and therefore a natural local distance. The curvature will be defined using this distance, which is linear in the holonomy displacement to first order, rather than the divergence itself, which only captures a quadratic effect.

4.1. Informational Distances and Defects for Discrete Loops

We begin by extracting from each fibre divergence a distance-like function.

Recall that, for each $\varepsilon > 0$ and each vertex $x_i \in V_\varepsilon$, we have a fibre $\mathcal{P}_i$ endowed with a Riemannian metric $g^{\mathcal{P}_i}$ and a divergence $D_i : \mathcal{P}_i \times \mathcal{P}_i \to [0, \infty)$ satisfying Assumption 1 (after transport to the continuous fibre). We define:

Definition 10 (Discrete informational distance)**.** *For each vertex* $x_i \in V_\varepsilon$ *and* $s, t \in \mathcal{P}_i$ *we define the informational distance*

$$d_i(s,t) := \sqrt{2\, D_i(s,t)}.$$

For the purposes of the estimates below, one may replace d_i by the Riemannian distance on $(\mathcal{P}_i, g^{\mathcal{P}_i})$, which is locally equivalent to d_i by Lemma 4. We implicitly make this replacement whenever the triangle inequality is invoked.

In general, d_i is only guaranteed to be a local distance function near the diagonal (we do not assume the triangle inequality). In the Jensen–Shannon setting, d_i coincides (up to the global constant factor $\sqrt{2}$) with the usual Jensen–Shannon distance on the probability simplex.

We now assign informational defects to discrete loops. Recall that, for each loop $\gamma \in \mathrm{Loops}_\varepsilon(x_i)$ based at x_i, we have a holonomy operator

$$\mathrm{Hol}_\gamma^{(\varepsilon)} : \mathcal{P}_i \to \mathcal{P}_i$$

defined by composition of edge channels along γ, and a reference state $\mu_i \in \mathcal{P}_i$.

Remark 7 (Riemannian vs. informational distance)**.** *For later estimates, it will be convenient to work with the genuine Riemannian distance on each fibre $(\mathcal{P}_i, g^{\mathcal{P}_i})$, which we denote by d_i^{R}. By Lemma 4 and compactness, d_i and d_i^{R} are locally equivalent: there exist constants $0 < c_1 \le c_2 < \infty$ such that, whenever $d_i^{\mathrm{R}}(s,t)$ is sufficiently small,*

$$c_1\, d_i^{\mathrm{R}}(s,t) \;\le\; d_i(s,t) \;\le\; c_2\, d_i^{\mathrm{R}}(s,t).$$

In particular, all notions of "defect" and "curvature" defined using d_i are unchanged, up to uniform multiplicative constants, if one replaces d_i by d_i^{R}. From this point on, whenever the triangle inequality is invoked, we implicitly work with d_i^{R}; the symbol d_i may be read as either distance, since they are locally equivalent.

Lemma 3 (From divergence contraction to distance contraction)**.** *Let D be a non-negative divergence and define $d(s,t) := \sqrt{2D(s,t)}$. If a map Φ satisfies*

$$D(\Phi(s), \Phi(t)) \le C\, D(s,t) \qquad \forall\, s,t,$$

then

$$d(\Phi(s), \Phi(t)) \le \sqrt{C}\, d(s,t) \qquad \forall\, s,t.$$

Proof. Immediate from $d = \sqrt{2D}$ and monotonicity of the square root. □

Definition 11 (Discrete informational holonomy defects)**.** *Let $\gamma \in \mathrm{Loops}_\varepsilon(x_i)$ be a discrete loop based at $x_i \in V_\varepsilon$. We define:*

1. *The divergence defect*
$$\delta_\gamma^{(\varepsilon)}(x_i) := D_i\big(\mu_i,\ \mathrm{Hol}_\gamma^{(\varepsilon)}(\mu_i)\big) \ge 0;$$
2. *The distance defect*
$$\Delta_\gamma^{(\varepsilon)}(x_i) := d_i\big(\mu_i,\ \mathrm{Hol}_\gamma^{(\varepsilon)}(\mu_i)\big) = \sqrt{2\,\delta_\gamma^{(\varepsilon)}(x_i)} \ge 0.$$

For curvature purposes, $\delta_\gamma^{(\varepsilon)}$ and $\Delta_\gamma^{(\varepsilon)}$ contain equivalent information, but the distance defect $\Delta_\gamma^{(\varepsilon)}$ scales linearly with the holonomy displacement and is therefore the natural quantity to normalize by the area of small loops.

As in Section 3.4, we now specialise to loops associated with discrete triangles. For a triangle $(x_i, x_j, x_k) \in \mathcal{T}_\varepsilon$, with associated loop $\gamma_{ijk} = (x_i, x_j, x_k, x_i)$ and holonomy operator

$$\mathrm{Hol}_{ijk}^{(\varepsilon)} := \Phi_{ki} \circ \Phi_{jk} \circ \Phi_{ij} : \mathcal{P}_i \to \mathcal{P}_i,$$

we use the shorthand

$$\delta_{ijk}^{(\varepsilon)} := \delta_{\gamma_{ijk}}^{(\varepsilon)}(x_i), \qquad \Delta_{ijk}^{(\varepsilon)} := \Delta_{\gamma_{ijk}}^{(\varepsilon)}(x_i).$$

Definition 12 (Triangle defects)**.** *For a discrete triangle $(x_i, x_j, x_k) \in \mathcal{T}_\varepsilon$, the divergence defect and* distance defect *of the triangle based at x_i are given by*

$$\delta^{(\varepsilon)}_{ijk} := D_i(\mu_i,\ \mathrm{Hol}^{(\varepsilon)}_{ijk}(\mu_i)), \qquad \Delta^{(\varepsilon)}_{ijk} := d_i(\mu_i,\ \mathrm{Hol}^{(\varepsilon)}_{ijk}(\mu_i)).$$

4.2. Discrete Informational Holonomy Curvature of Triangles

To obtain a curvature quantity from the defects, we normalise by the area associated with each triangle. Assumption 10 provides a discrete area $A_\varepsilon(i,j,k)$ which approximates the Riemannian area $A_g(i,j,k)$ of the geodesic triangle with vertices x_i, x_j, x_k.

Definition 13 (Triangle-wise informational holonomy curvature)**.** *Let $(x_i, x_j, x_k) \in \mathcal{T}_\varepsilon$ be a discrete triangle. The informational holonomy curvature of the triangle (x_i, x_j, x_k) is*

$$K^{(\varepsilon)}_{\mathrm{hol}}(i,j,k) := \frac{\Delta^{(\varepsilon)}_{ijk}}{A_\varepsilon(i,j,k)},$$

with the convention that $K^{(\varepsilon)}_{\mathrm{hol}}(i,j,k) = 0$ if $A_\varepsilon(i,j,k) = 0$.

Thus, $K^{(\varepsilon)}_{\mathrm{hol}}(i,j,k)$ measures the informational distance travelled by the reference state per unit area when it is transported around the discrete triangle (x_i, x_j, x_k).

In order to define a discrete sectional curvature at a point and a two-dimensional direction, we now average the triangle-wise curvature over suitable families of triangles.

Let $K \subset M$ be a fixed compact set. For each $x \in K$ and each 2-plane $\Pi \subset T_xM$, Assumption 11 provides, for ε small enough, a vertex $x_i \in V_\varepsilon$ close to x and a finite family of discrete triangles

$$\mathcal{T}_\varepsilon(x_i, \Pi) \subset \mathcal{T}_\varepsilon$$

which are small, non-degenerate, have edge directions close to Π, and are approximately isotropically distributed in Π.

Definition 14 (Discrete informational sectional curvature)**.** *Let $K \subset M$ be compact, $x \in K$, and $\Pi \subset T_xM$ a two-dimensional subspace. For $\varepsilon > 0$ small enough, let $x_i \in V_\varepsilon$ and $\mathcal{T}_\varepsilon(x_i, \Pi)$ be as in Assumption 11. The discrete informational sectional curvature at scale ε associated with (x, Π) is*

$$K^{(\varepsilon)}_{\mathrm{hol}}(x, \Pi) := \frac{\displaystyle\sum_{(x_i,x_j,x_k)\in\mathcal{T}_\varepsilon(x_i,\Pi)} A_\varepsilon(i,j,k)\, K^{(\varepsilon)}_{\mathrm{hol}}(i,j,k)}{\displaystyle\sum_{(x_i,x_j,x_k)\in\mathcal{T}_\varepsilon(x_i,\Pi)} A_\varepsilon(i,j,k)}. \tag{7}$$

If the denominator vanishes, we set $K^{(\varepsilon)}_{\mathrm{hol}}(x, \Pi) := 0$ by convention.

In words, $K^{(\varepsilon)}_{\mathrm{hol}}(x, \Pi)$ is the area-weighted average of the triangle-wise informational holonomy curvature over all triangles in $\mathcal{T}_\varepsilon(x_i, \Pi)$.

Remark 8 (Choice of base vertex and triangle family)**.** *The definition of $K^{(\varepsilon)}_{\mathrm{hol}}(x, \Pi)$ in (7) involves several auxiliary choices: for each $x \in K$ and two-plane $\Pi \subset T_xM$, we pick a nearby vertex x_i with $d_g(x, x_i) \le c_2 r_\varepsilon$ and a finite family of triangles $\mathcal{T}_\varepsilon(x_i, \Pi)$ as in Assumption 11. A priori, different admissible choices $(x_i, \mathcal{T}_\varepsilon(x_i, \Pi))$ could lead to different values of $K^{(\varepsilon)}_{\mathrm{hol}}(x, \Pi)$.*

Under Assumptions 4–11, however, Lemmas 10 and 11 imply that any two such choices produce values that differ by at most $C\,\kappa_\varepsilon$, with κ_ε as in (10). In particular, the limit $K^{(\varepsilon)}_{\mathrm{hol}}(x, \Pi) \to K^{\mathrm{cont}}_{\mathrm{hol}}(x, \Pi)$ in Theorem 2 is independent of these auxiliary choices.

4.3. Continuous Informational Holonomy Curvature Revisited

We now recall the continuous framework and align the definitions with the discrete case by introducing the corresponding informational distance on each fibre.

Let (M, g) be a Riemannian manifold, $\pi : \mathcal{P} \to M$ a state bundle with fibre metrics and divergences as in Definition 2 and Assumption 1, equipped with a connection satisfying Assumption 3. Let $\mu : M \to \mathcal{P}$ be a smooth reference state field.

For each $x \in M$ and $s, t \in \mathcal{P}_x$, we define the continuous informational distance

$$d_x(s,t) := \sqrt{2\, D_x(s,t)}. \tag{8}$$

By Assumption 1, d_x is locally equivalent to the Riemannian distance induced by $g^{\mathcal{P}_x}$ on the fibre (see Lemma 4 in Section 5).

For $x \in M$ and a two-dimensional subspace $\Pi \subset T_xM$, consider pairs of tangent vectors $u, v \in \Pi$ of sufficiently small norm and the associated geodesic triangle with vertices

$$x, \quad y = \exp_x(u), \quad z = \exp_x(v).$$

Let $\gamma_{x,u,v}$ denote the closed loop $x \to y \to z \to x$ obtained by traversing the geodesic segments in order, and let $A_g(x, y, z)$ be its Riemannian area. The continuous holonomy map

$$\mathrm{Hol}_{\gamma_{x,u,v}} : \mathcal{P}_x \to \mathcal{P}_x$$

is defined by parallel transport along $\gamma_{x,u,v}$.

Definition 15 (Continuous informational holonomy defects)**.** *The continuous divergence defect and continuous distance defect of the loop $\gamma_{x,u,v}$ at x are*

$$\delta_{x,u,v} := D_x\big(\mu_x,\ \mathrm{Hol}_{\gamma_{x,u,v}}(\mu_x)\big), \qquad \Delta_{x,u,v} := d_x\big(\mu_x,\ \mathrm{Hol}_{\gamma_{x,u,v}}(\mu_x)\big) = \sqrt{2\,\delta_{x,u,v}}.$$

We are interested in the behaviour of $\Delta_{x,u,v}$ as the triangle shrinks to x within the plane Π.

Definition 16 (Continuous informational sectional curvature)**.** *Let $(M, g, \mathcal{P}, \mu)$ be as above, and fix $x \in M$ and a two-dimensional subspace $\Pi \subset T_xM$. For $u, v \in \Pi$ sufficiently small we denote by $\gamma_{x,u,v}$ the associated piecewise-geodesic loop and by $\Delta_{x,u,v}$ its informational distance defect at x (cf. Definition 15).*

We say that the continuous informational sectional curvature at (x, Π) exists if there are constants $r_0 > 0$ and $c_0 > 0$ and, for each $r \in (0, r_0]$, vectors $u_r, v_r \in \Pi$ such that:

1. *The corresponding geodesic triangle*

$$\Delta_g\big(x, \exp_x(u_r), \exp_x(v_r)\big)$$

is contained in a normal neighbourhood of x and its vertices converge to x as $r \to 0$;

2. *The triangle has uniformly non-degenerate shape in Π, in the sense that*

$$c_0\, r \ \le\ \max\{\|u_r\|_g, \|v_r\|_g\} \ \le\ r, \qquad A_g\big(x, \exp_x(u_r), \exp_x(v_r)\big) \ \ge\ c_0\, r^2,$$

so that the side lengths and the area are uniformly comparable to r and r^2, respectively, independently of r.

Whenever this holds and, for every such admissible family (u_r, v_r), the limit

$$K_{\mathrm{hol}}^{\mathrm{cont}}(x,\Pi) := \lim_{r\to 0} \frac{\Delta_{x,u_r,v_r}}{A_g\big(x, \exp_x(u_r), \exp_x(v_r)\big)} \tag{9}$$

exists and has the same value, we call this common value the continuous informational sectional curvature at (x,Π).

In Section 5 we show, under Assumptions 1 and 3, that this limit exists for all $x \in M$ and all two-planes $\Pi \subset T_xM$, and that it can be expressed explicitly in terms of the curvature of the connection on $\mathcal{P}$ along Π.

4.4. Main Curvature Theorems

We can now formulate the main results of this work, which will be proved in Sections 5 and 6. The first theorem relates the continuous informational sectional curvature to the curvature of the connection on the state bundle. The second theorem shows that the discrete informational sectional curvature converges to this continuous quantity as the sampling becomes dense.

Throughout this subsection, we fix a compact subset $K \subset M$ and tacitly restrict attention to points $x \in K$.

Theorem 1 (Continuous informational holonomy curvature)**.** *Let (M,g) be a smooth Riemannian manifold, $\pi : \mathcal{P} \to M$ a state bundle with fibre metrics and divergences satisfying Assumption 1, endowed with a connection satisfying Assumption 3, and let $\mu : M \to \mathcal{P}$ be a smooth reference state field.*

Then, for every $x \in M$ and every two-dimensional subspace $\Pi \subset T_xM$, the continuous informational sectional curvature $K_{\mathrm{hol}}^{\mathrm{cont}}(x,\Pi)$ exists in the sense of Definition 16 and can be written as

$$K_{\mathrm{hol}}^{\mathrm{cont}}(x,\Pi) = \big\|W_x(\Pi;\mu_x)\big\|_{g^{\mathcal{P}_x}},$$

where $W_x(\Pi;\mu_x) \in T_{\mu_x}\mathcal{P}_x$ is a vector depending linearly on the curvature of the connection on $\mathcal{P}$ along Π. In particular, when the connection on $\mathcal{P}$ is induced by the Levi–Civita connection via a linear isometric representation, $K_{\mathrm{hol}}^{\mathrm{cont}}(x,\Pi)$ is a scalar invariant built from the Riemann curvature tensor R_x^g restricted to Π and is proportional to $|\sec_g(x,\Pi)|$ in spaces of constant sectional curvature.

Remark 9 (Dependence on the informational structure)**.** *The quantity $K_{\mathrm{hol}}^{\mathrm{cont}}(x,\Pi)$ should not be thought of as a curvature of the Riemannian manifold (M,g) alone. It depends on the choice of state bundle $\pi : \mathcal{P} \to M$, on the fibre divergences $\{D_x\}_{x\in M}$ (equivalently, on the fibre metrics $g^{\mathcal{P}_x}$), on the Ehresmann connection used to define parallel transport, and on the reference state field μ. Different choices over the same (M,g) may lead to different informational holonomy curvatures. In particular, two bundles with the same base (M,g) but different fibre geometries or connections need not produce the same values of $K_{\mathrm{hol}}^{\mathrm{cont}}$.*

We now turn to the discrete setting. In addition to the continuous structures above, we assume that the manifold (M,g) is sampled by graphs $(G_\varepsilon)_{\varepsilon>0}$ and that discrete fibres, divergences, reference states, channels, areas, and triangle families are given as in Assumptions 4–11.

For convenience, we introduce a sequence of positive numbers $(\kappa_\varepsilon)_{\varepsilon>0}$ that captures the various discretisation errors. More precisely, we set

$$\kappa_\varepsilon := r_\varepsilon + \eta_\varepsilon + q_\varepsilon + \rho_\varepsilon, \qquad \rho_\varepsilon := r_\varepsilon^{\alpha-1}. \tag{10}$$

Here, r_ε is the sampling radius from Assumption 4, η_ε and q_ε are the anisotropy and area errors from Assumptions 11 and 10, and $\rho_\varepsilon \to 0$ controls the channel consistency scale induced by Assumption 9(2) with exponent $\alpha > 1$.

Theorem 2 (Discrete-to-continuous convergence of informational holonomy curvature)**.** *Let $(M, g, \mathcal{P}, \mu)$ be as in Theorem 1, and let $(G_\varepsilon)_{\varepsilon>0}$, together with discrete fibres, divergences, reference states, channels, areas and triangle families, satisfy Assumptions 4–11. Let $K \subset M$ be compact.*

Then, there exist constants $C > 0$ and $\varepsilon_K > 0$ such that, for all $0 < \varepsilon < \varepsilon_K$, for every $x \in K$ and every two-dimensional subspace $\Pi \subset T_xM$, the discrete informational sectional curvature $K^{(\varepsilon)}_{\mathrm{hol}}(x, \Pi)$ is well defined and satisfies

$$\left|K^{(\varepsilon)}_{\mathrm{hol}}(x, \Pi) - K^{\mathrm{cont}}_{\mathrm{hol}}(x, \Pi)\right| \leq C\,\kappa_\varepsilon, \tag{11}$$

where κ_ε is given by (10). In particular,

$$\left|K^{(\varepsilon)}_{\mathrm{hol}}(x, \Pi) - \left\|W_x(\Pi; \mu_x)\right\|_{g^{\mathcal{P}_x}}\right| \leq C\,\kappa_\varepsilon.$$

Consequently, as $\varepsilon \to 0$, the discrete informational sectional curvatures converge uniformly on compact subsets of M and on the Grassmannian of two-planes to the continuous informational sectional curvature $K^{\mathrm{cont}}_{\mathrm{hol}}$.

The proof of Theorem 2, given in Section 6, proceeds in two steps. First, we compare the discrete holonomy operators $\mathrm{Hol}^{(\varepsilon)}_{ijk}$ around small triangles with the continuous holonomy operators associated with the corresponding geodesic triangles, using Assumption 9 and the local consistency of the sampling. Second, we exploit the isotropy of the triangle families from Assumption 11 to show that the area-weighted average (7) converges to the continuous limit (9), with an error controlled by κ_ε.

Overview of Assumptions

For the reader's convenience, we briefly summarise the role of the hypotheses. Assumption 1 requires each fibre divergence D_x to admit a second-order expansion whose quadratic part induces the fibre Riemannian metric $g^{\mathcal{P}_x}$; together with Assumption 3, this ensures that informational distances behave, locally and under parallel transport, like the Riemannian distances of the fibre metrics. Assumption 2 is a structural regularity assumption guaranteeing that the Ehresmann connection on $\mathcal{P}$ arises from a smooth principal connection.

On the discrete side, Assumptions 4 and 5 encode that the sampling graphs G_ε form quasi-uniform discretisations of (M, g) with bounded degree and uniformly controlled edge lengths, whereas Assumption 6 provides vertex weights approximating the Riemannian volume. Assumption 7 specifies discrete fibres and divergences approximating the continuous ones, and Assumption 8 does the same for the reference states. Assumption 9 postulates channels that are local and Lipschitz and whose first-order behaviour approximates continuous parallel transport, with an error measured by ρ_ε. Finally, Assumption 10 ensures that the discrete triangle areas $A_\varepsilon(i, j, k)$ approximate the Riemannian areas, and Assumption 11 provides, for each (x, Π), families of triangles that sample directions in Π in an almost isotropic way. The quantity κ_ε in (10) summarises the various errors contributed by these assumptions.

Definitions 13, 14 and 16, together with Theorems 1 and 2, provide the conceptual and analytical core of the informational holonomy curvature framework. In the next sections, we make these statements precise by deriving the continuous holonomy–curvature relation and then establishing the discrete-to-continuous convergence.

5. Continuous Holonomy Curvature and Connection Curvature

In this section, we prove Theorem 1. We work in the continuous framework of Section 2: (M, g) is a smooth Riemannian manifold, $\pi : \mathcal{P} \to M$ is a state bundle with fibre metrics and divergences satisfying Assumption 1, endowed with a connection satisfying Assumption 3, and $\mu : M \to \mathcal{P}$ is a smooth reference state field.

The core of the argument consists of two ingredients:

- The second-order expansion of the divergence on each fibre, which implies that the informational distance $d_x(s,t) = \sqrt{2D_x(s,t)}$ is locally equivalent to the Riemannian distance induced by $g^{\mathcal{P}_x}$;
- The first-order (in area) expansion of the holonomy map of the connection on $\mathcal{P}$ around small geodesic triangles, controlled by the curvature of the connection.

Combining these, we obtain a linear-in-area behaviour for the distance defect $\Delta_{x,u,v}$ and hence the existence and explicit form of the continuous informational sectional curvature.

Throughout this section, we fix a compact set $K \subset M$, and all constants will be uniform over $x \in K$ and over two-planes $\Pi \subset T_xM$.

5.1. Local Expansion of the Informational Distance

We recall that, for each $x \in M$ and $s, t \in \mathcal{P}_x$, the informational distance is defined by

$$d_x(s,t) := \sqrt{2\, D_x(s,t)},$$

where D_x is the divergence on the fibre $\mathcal{P}_x$ (see (8)). The following lemma makes precise the local behaviour of d_x in terms of the Riemannian metric $g^{\mathcal{P}_x}$ on the fibre.

Lemma 4 (Local expansion of the informational distance)**.** *Let $K \subset M$ be compact. There exist constants $C > 0$ and $r > 0$ such that for every $x \in K$, for every $s \in \mathcal{P}_x$, and for every $v \in T_s\mathcal{P}_x$ with $\|v\|_{g^{\mathcal{P}_x}} \le r$, the following holds. Let*

$$t := \exp_s^{\mathcal{P}_x}(v),$$

where $\exp_s^{\mathcal{P}_x}$ is the Riemannian exponential map on $(\mathcal{P}_x, g^{\mathcal{P}_x})$. Then,

$$\left| d_x(s,t) - \|v\|_{g^{\mathcal{P}_x}} \right| \le C\, \|v\|^2_{g^{\mathcal{P}_x}}. \tag{12}$$

In particular, there exist constants $c_1, c_2 > 0$ such that, for all such v,

$$c_1\, \|v\|_{g^{\mathcal{P}_x}} \le d_x(s,t) \le c_2\, \|v\|_{g^{\mathcal{P}_x}}.$$

Proof. Fix $x \in K$ and $s \in \mathcal{P}_x$. By Assumption 1, in a normal coordinate chart for $(\mathcal{P}_x, g^{\mathcal{P}_x})$ centred at s, the divergence D_x satisfies

$$D_x(s,t) = \frac{1}{2}\, g_s^{\mathcal{P}_x}(v,v) + R_s(v),$$

where $t = \exp_s^{\mathcal{P}_x}(v)$ and $R_s(v)$ is a remainder term with

$$|R_s(v)| \le C_0\, \|v\|^3_{g^{\mathcal{P}_x}}$$

for $\|v\|_{g^{\mathcal{P}_x}} \le r_0$, with $C_0, r_0 > 0$ independent of x and s in compact sets (smoothness and compactness of K).

By definition,

$$d_x(s,t) = \sqrt{2D_x(s,t)} = \sqrt{g_s^{\mathcal{P}_x}(v,v) + 2R_s(v)}.$$

Let $L := \|v\|_{g^{\mathcal{P}_x}}$. Then, $g_s^{\mathcal{P}_x}(v,v) = L^2$ and

$$d_x(s,t) = L\sqrt{1 + \frac{2R_s(v)}{L^2}}.$$

For $0 < L \le r_0$, we have

$$\left|\frac{2R_s(v)}{L^2}\right| \le 2C_0\, L.$$

Choose $r \le r_0$ such that $2C_0 r \le 1/2$. Then, for $L \le r$, we have

$$\left|\frac{2R_s(v)}{L^2}\right| \le \frac{1}{2}.$$

For $|z| \le 1/2$, the Taylor expansion of $\sqrt{1+z}$ yields

$$\sqrt{1+z} = 1 + \frac{z}{2} + \theta(z)\, z^2,$$

where $|\theta(z)| \le C_1$ for some universal constant C_1. Taking $z = 2R_s(v)/L^2$, we obtain

$$d_x(s,t) = L\left(1 + \frac{R_s(v)}{L^2} + \theta\Big(\frac{2R_s(v)}{L^2}\Big)\frac{4R_s(v)^2}{L^4}\right).$$

Using $|R_s(v)| \le C_0 L^3$, we obtain

$$\left|\frac{R_s(v)}{L^2}\right| \le C_0 L, \qquad \left|\frac{4R_s(v)^2}{L^4}\right| \le 4C_0^2 L^2.$$

Therefore,

$$\left|d_x(s,t) - L\right| \le L\Big(C_0 L + C_1 \cdot 4C_0^2 L^2\Big) \le C\, L^2,$$

for some constant $C > 0$ depending only on C_0 and C_1. This proves (12).

The two-sided inequality follows from (12) by taking L sufficiently small and absorbing the quadratic term into the linear one. □

5.2. Curvature of the Connection and Small-Loop Holonomy

We now recall the curvature of an Ehresmann connection on $\pi : \mathcal{P} \to M$ and its relation with holonomy around small loops. We adopt a local viewpoint sufficient for our application to small geodesic triangles.

Let $\mathcal{H}_p \subset T_p\mathcal{P}$ be the horizontal subspace at $p \in \mathcal{P}$ and $\mathcal{V}_p = \ker(d\pi_p)$ the vertical subspace. For each vector field X on M, there exists a unique horizontal lift X^H on $\mathcal{P}$ such that $d\pi(X_p^H) = X_{\pi(p)}$ for all $p \in \mathcal{P}$.

The curvature of the connection is the vertical-valued 2-form Ω on $\mathcal{P}$ defined by

$$\Omega_p(X,Y) := [X^H, Y^H]_p^{\mathrm{vert}},$$

where X, Y are vector fields on M and the superscript vert denotes projection onto $\mathcal{V}_p$ in the decomposition $T_p\mathcal{P} = \mathcal{H}_p \oplus \mathcal{V}_p$. This definition is independent of the choice of extensions of X, Y.

For each $x \in M$ and $p \in \mathcal{P}_x := \pi^{-1}(x)$, the restriction of Ω gives a bilinear alternating map

$$\Omega_{x,p} : \Lambda^2 T_x M \to T_p \mathcal{P}_x,$$

by setting $\Omega_{x,p}(u, v) := \Omega_p(\tilde{u}, \tilde{v})$, where $\tilde{u}, \tilde{v}$ are any vector fields extending u, v in a neighbourhood of x. This is well defined and smooth in (x, p).

We now state the holonomy expansion around small geodesic triangles. The result is a specialization of standard holonomy–curvature relations (see, for example, Chapter II, Sections 3 and 4 of [7]), adapted to our two-dimensional situation and with explicit attention to the dependence on the area of the triangle.

Let $x \in M$ and $\Pi \subset T_x M$ be a two-dimensional subspace. For $u, v \in \Pi$ sufficiently small, we set

$$y = \exp_x(u), \qquad z = \exp_x(v),$$

and consider the geodesic triangle with vertices x, y, z. Let $\gamma_{x,u,v}$ denote the closed loop obtained by traversing the geodesic segments $x \to y \to z \to x$ in order, and let $A_g(x, y, z)$ be the Riemannian area of the geodesic triangle.

Fix $p_x \in \mathcal{P}_x$. Parallel transport along $\gamma_{x,u,v}$ yields a holonomy map

$$\mathrm{Hol}_{\gamma_{x,u,v}} : \mathcal{P}_x \to \mathcal{P}_x$$

sending p_x to a point $p_{x,u,v} := \mathrm{Hol}_{\gamma_{x,u,v}}(p_x) \in \mathcal{P}_x$.

Lemma 5 (Small-loop holonomy expansion)**.** *Assume the associated-bundle framework of Assumption 2. Let $K \subset M$ be compact. Then, there exist constants $C > 0$ and $r > 0$ such that for every $x \in K$, every 2-plane $\Pi \subset T_x M$, every pair $u, v \in \Pi$ with $\|u\|_g, \|v\|_g \le r$, and every $p_x \in \mathcal{P}_x$, the following holds.*

Let $y = \exp_x(u)$, $z = \exp_x(v)$, let $A := A_g(x, y, z)$ be the Riemannian area of the geodesic triangle $\Delta_g(x, y, z)$, and let $\gamma_{x,u,v}$ be the piecewise-geodesic loop $x \to y \to z \to x$. Denote $p_{x,u,v} := \mathrm{Hol}_{\gamma_{x,u,v}}(p_x) \in \mathcal{P}_x$ and $\ell := \max\{\|u\|_g, \|v\|_g\}$. Then,

$$\exp_{p_x}^{\mathcal{P}_x\,-1}(p_{x,u,v}) = A\, W_x(\Pi; p_x) + R_x(u, v; p_x), \tag{13}$$

where $W_x(\Pi; p_x) \in T_{p_x}\mathcal{P}_x$ depends smoothly on (x, Π, p_x) and linearly on the curvature of the inducing principal connection, and the remainder satisfies

$$\big\|R_x(u, v; p_x)\big\|_{g^{\mathcal{P}_x}} \le C\, \ell^3. \tag{14}$$

More explicitly: fix an orientation of Π and let (e_1, e_2) be any oriented g-orthonormal basis of Π. Let $X_x(\Pi) \in \mathfrak{g}$ be the curvature element determined by the principal curvature F_ω (in any local gauge) evaluated on (e_1, e_2) at x. Let $\xi : \mathfrak{g} \to \Gamma(TS)$ be the infinitesimal action (fundamental vector fields) of G on S. Then, identifying $\mathcal{P}_x \simeq S$ via any choice of $q \in \mathcal{Q}_x$,

$$W_x(\Pi; p_x) = \xi\big(X_x(\Pi)\big)\big|_{p_x} \in T_{p_x}\mathcal{P}_x,$$

and this definition is independent of the chosen gauge because F_ω is Ad*-equivariant and the action of G on S is by isometries.*

Proof. We work on a convex normal neighbourhood $U \subset M$ of x contained in a compact set K and choose a C^3 local section $\sigma : U \to \mathcal{Q}$. In this gauge, the principal connection is represented by a $\mathfrak{g}$-valued 1-form $A := \sigma^* \omega$ on U, with curvature $F := \sigma^* F_\omega = dA + \frac{1}{2}[A \wedge A]$, whose coefficients are C^1 and uniformly bounded on K (by Assumption 2).

Let $\Sigma_{x,u,v} \subset U$ be the geodesic triangle surface with boundary $\partial\Sigma_{x,u,v} = \gamma_{x,u,v}$. The holonomy of the principal connection around $\gamma_{x,u,v}$ is an element $g_{x,u,v} \in G$ obtained by the path-ordered exponential of A along $\gamma_{x,u,v}$. In this standard principal-connection setting, the non-abelian Stokes/holonomy–curvature expansion (see, e.g., Chapter II, §3–§4 of [7]) yields

$$g_{x,u,v} = \exp\Big(\int_{\Sigma_{x,u,v}} F\Big)\, \exp(E_{x,u,v}), \qquad \|E_{x,u,v}\|_{\mathfrak{g}} \le C_0\, \ell^3, \tag{15}$$

for ℓ small, with constants uniform for $x \in K$ and Π varying in the Grassmannian. Moreover, since F is C^1, Taylor expansion on $\Sigma_{x,u,v}$ gives

$$\int_{\Sigma_{x,u,v}} F = A\, F_x(e_1, e_2) + O(\ell^3) \quad \text{in } \mathfrak{g},$$

where $F_x(e_1, e_2)$ denotes the curvature evaluated at x on an oriented orthonormal basis (e_1, e_2) of Π (and the $O(\ell^3)$ term is uniform on K). Combining with (15) yields

$$g_{x,u,v} = \exp\big(A\, X_x(\Pi) + \widetilde{E}_{x,u,v}\big), \qquad \|\widetilde{E}_{x,u,v}\|_{\mathfrak{g}} \le C_1\, \ell^3. \tag{16}$$

Now pass to the associated bundle. Choose $q \in \mathcal{Q}_x$ and identify $\mathcal{P}_x \simeq S$ by $[q, s] \leftrightarrow s$. Under this identification, the holonomy acts by the G-action:

$$p_{x,u,v} = \mathrm{Hol}_{\gamma_{x,u,v}}(p_x) = g_{x,u,v} \cdot p_x.$$

Consider the smooth map

$$H_{p_x} : \mathfrak{g} \supset B(0, \delta) \to T_{p_x}S, \qquad H_{p_x}(Z) := \exp_{p_x}^{S\ -1}\big(\exp(Z) \cdot p_x\big),$$

defined for $\delta > 0$ small. Since the G-action is smooth, H_{p_x} is C^2 and satisfies $DH_{p_x}(0)[Z] = \xi(Z)|_{p_x}$ (the fundamental vector field). Therefore, Taylor expansion at 0 gives

$$H_{p_x}(Z) = \xi(Z)|_{p_x} + O(\|Z\|_{\mathfrak{g}}^2) \quad \text{in } T_{p_x}S, \tag{17}$$

with a uniform constant for $x \in K$ and p_x ranging in compact subsets of the fibres.

Apply (17) to $Z := A\, X_x(\Pi) + \widetilde{E}_{x,u,v}$ from (16). Since $A = O(\ell^2)$ and $\|\widetilde{E}_{x,u,v}\| = O(\ell^3)$, we obtain

$$\exp_{p_x}^{\mathcal{P}_x\ -1}(p_{x,u,v}) = H_{p_x}(Z) = A\, \xi\big(X_x(\Pi)\big)\big|_{p_x} + O(\ell^3),$$

because the quadratic term $O(\|Z\|^2)$ is $O(\ell^4)$ and hence dominated by $O(\ell^3)$. This yields (13) and (14) with $W_x(\Pi; p_x) = \xi(X_x(\Pi))|_{p_x}$. □

Remark 10. *The vector $W_x(\Pi; p_x)$ depends linearly on the curvature Ω_{x,p_x} restricted to Π. In particular, when $\mathcal{P}$ is an associated bundle to a principal bundle with a connection induced from the Levi–Civita connection via an isometric representation, $W_x(\Pi; p_x)$ is obtained by applying the differential of the representation to the Riemann curvature tensor R^g_x restricted to Π.*

5.3. Proof of Theorem 1

Fix $x \in M$ and a two-dimensional subspace $\Pi \subset T_xM$. Let $(u_r, v_r) \in \Pi \times \Pi$ be a family as in Definition 16, with $\|u_r\|_g, \|v_r\|_g \to 0$ and with uniformly non-degenerate shape in Π. Set

$$y_r := \exp_x(u_r), \qquad z_r := \exp_x(v_r), \qquad A_r := A_g(x, y_r, z_r),$$

and let $\gamma_r := \gamma_{x,u_r,v_r}$ be the boundary loop of the geodesic triangle $\Delta_g(x, y_r, z_r)$. Finally, set

$$p_x := \mu_x, \qquad p_r := \mathrm{Hol}_{\gamma_r}(p_x) \in \mathcal{P}_x.$$

By Lemma 5 (applied with $p_x = \mu_x$) we have, for r small,

$$v_r := \exp_{p_x}^{\mathcal{P}_x\,-1}(p_r) = A_r\, W_x(\Pi; \mu_x) + R_r, \tag{18}$$

where $\|R_r\|_{g^{\mathcal{P}_x}} \le C\,\ell_r^3$ and $\ell_r := \max\{\|u_r\|_g, \|v_r\|_g\}$. Since the triangles have uniformly non-degenerate shape in Π, there exists $c > 0$ such that $A_r \ge c\,\ell_r^2$ for r small. Consequently,

$$\|v_r\|_{g^{\mathcal{P}_x}} = A_r\, \|W_x(\Pi; \mu_x)\|_{g^{\mathcal{P}_x}} + O(A_r^{3/2}), \tag{19}$$

where the implicit constant is uniform for x in compact sets and for Π.

Now, the continuous distance defect satisfies

$$\Delta_{x,u_r,v_r} = d_x(p_x, p_r) = d_x\big(p_x, \exp_{p_x}^{\mathcal{P}_x}(v_r)\big).$$

Applying Lemma 4 with $s = p_x$ and $v = v_r$ yields

$$\Delta_{x,u_r,v_r} = \|v_r\|_{g^{\mathcal{P}_x}} + O(\|v_r\|_{g^{\mathcal{P}_x}}^2) = A_r\, \|W_x(\Pi; \mu_x)\|_{g^{\mathcal{P}_x}} + O(A_r^{3/2}). \tag{20}$$

Dividing by A_r and letting $r \to 0$, we obtain

$$\lim_{r\to 0} \frac{\Delta_{x,u_r,v_r}}{A_r} = \|W_x(\Pi; \mu_x)\|_{g^{\mathcal{P}_x}}.$$

This limit is independent of the chosen shrinking family (u_r, v_r) (subject to the non-degeneracy condition), and therefore $K_{\mathrm{hol}}^{\mathrm{cont}}(x, \Pi)$ exists and equals $\|W_x(\Pi; \mu_x)\|_{g^{\mathcal{P}_x}}$.

5.4. Geometric Models Induced from the Levi–Civita Connection

We briefly justify the last assertion of Theorem 1. Assume that $\mathcal{P}$ is an associated bundle to the orthonormal frame bundle $\mathrm{Fr}(M) \to M$, with structure group $\mathrm{O}(n)$ acting on the model fibre S by isometries of (S, g^S). Let $\rho : \mathrm{O}(n) \to \mathrm{Isom}(S, g^S)$ denote this representation and let ∇ be the Levi–Civita connection on $\mathrm{Fr}(M)$; it induces an Ehresmann connection on $\mathcal{P}$.

Let $x \in M$ and $\Pi \subset T_xM$ be a two-plane. Denote by $\mathcal{R}_x^g(\Pi) \in \mathfrak{so}(T_xM, g_x)$ the curvature endomorphism of ∇ restricted to Π (equivalently, the image of R_x^g under the identification $\Lambda^2 T_xM \simeq \mathfrak{so}(T_xM, g_x)$). The curvature of the induced connection on $\mathcal{P}$ is obtained by applying the differential $d\rho$ fibrewise, and therefore there exists a smooth linear map

$$L_{x,p} : \mathfrak{so}(T_xM, g_x) \longrightarrow T_p\mathcal{P}_x, \qquad p \in \mathcal{P}_x,$$

such that

$$W_x(\Pi; p) = L_{x,p}\big(\mathcal{R}_x^g(\Pi)\big). \tag{21}$$

In particular, $K_{\mathrm{hol}}^{\mathrm{cont}}(x, \Pi) = \|L_{x,\mu_x}(\mathcal{R}_x^g(\Pi))\|$ is a scalar invariant determined by the restriction of R_x^g to Π.

If (M, g) has constant sectional curvature κ, then for every x and Π one has $\mathcal{R}_x^g(\Pi) = \kappa\, J_\Pi$, where J_Π is the infinitesimal generator of the g_x-rotation in Π (normalized so that $\|J_\Pi\|_{\mathrm{HS}} = 1$). Combining this with (21) yields

$$K_{\mathrm{hol}}^{\mathrm{cont}}(x, \Pi) = |\kappa|\, \big\|L_{x,\mu_x}(J_\Pi)\big\|_{g^{\mathcal{P}_x}}.$$

Under the additional natural hypothesis that the model is $\mathrm{O}(n)$-equivariant and the reference field μ is chosen compatibly with the symmetry, the factor $\|L_{x,\mu_x}(J_\Pi)\|$ is independent of x and Π, so that $K_{\mathrm{hol}}^{\mathrm{cont}}$ reduces to a constant multiple of $|\kappa| = |\mathrm{sec}_g|$, as claimed.

6. Discrete-to-Continuous Convergence of Informational Holonomy Curvature

In this section, we prove Theorem 2. We work under Assumptions 4–11 and the continuous hypotheses of Theorem 1. Throughout, $K \subset M$ is a fixed compact set and all constants are uniform over $x \in K$ and over two-planes $\Pi \subset T_xM$.

The proof proceeds in three steps:

1. We compare discrete and continuous holonomy on individual small triangles (Section 6.1);
2. We convert this comparison into a bound between discrete and continuous *triangle-wise* informational holonomy curvature (Section 6.2);
3. We pass to the averaged, *sectional* quantity by exploiting the isotropic triangle families from Assumption 11 and the area approximation from Assumption 10 (Section 6.3).

6.1. Comparison of Discrete and Continuous Holonomy on Small Triangles

Fix $\varepsilon > 0$ small, a vertex $x_i \in V_\varepsilon$ and a triangle $(x_i, x_j, x_k) \in \mathcal{T}_\varepsilon$. Let $\gamma_{ijk} = (x_i, x_j, x_k, x_i)$ be the associated discrete loop and

$$\mathrm{Hol}_{ijk}^{(\varepsilon)} := \Phi_{ki} \circ \Phi_{jk} \circ \Phi_{ij} : \mathcal{P}_i \to \mathcal{P}_i$$

the discrete holonomy operator based at x_i.

On the continuous side, consider the geodesic triangle $\Delta_g(x_i, x_j, x_k)$ in (M, g) and the loop $\tilde{\gamma}_{ijk}$ obtained by traversing the minimizing geodesic segments

$$x_i \to x_j \to x_k \to x_i$$

in order. Let

$$\mathrm{Hol}_{\tilde{\gamma}_{ijk}} : \mathcal{P}_{x_i} \to \mathcal{P}_{x_i}$$

be the corresponding holonomy map induced by the connection on $\mathcal{P}$. Via the identification $\iota_i^{(\varepsilon)} : \mathcal{P}_i \to \mathcal{P}_{x_i}$ from Assumption 7, we may view both $\mathrm{Hol}_{ijk}^{(\varepsilon)}$ and $\mathrm{Hol}_{\tilde{\gamma}_{ijk}}$ as acting on the same fibre.

For notational simplicity, we replace $\iota_i^{(\varepsilon)}$ by the identity and treat $\mathcal{P}_i$ as $\mathcal{P}_{x_i}$, understanding that all maps and distances are transported accordingly. Thus, in what follows, μ_i denotes μ_{x_i} and d_i denotes the informational distance d_{x_i}.

We first estimate how well the discrete holonomy operator approximates the continuous one near the reference state.

Lemma 6 (Staying in the local regime along short compositions)**.** *Let $K \subset M$ be compact. Assume:*

(i) The fibre divergences satisfy Assumption 1 uniformly on K.
(ii) The continuous model is in the associated-bundle setting Assumption 2, so that parallel transport preserves fibre Riemannian distances (Lemma 2).
(iii) The discrete channels satisfy Assumption 9(2) in the fibre Riemannian distance d^{R} with exponent $\alpha > 1$.
(iv) The channel locality Assumption 9(1) holds.

Then, there exist constants $r_0 > 0$, $C > 0$ and $\varepsilon_K > 0$ such that for all $0 < \varepsilon < \varepsilon_K$ the following holds.

Let (x_i, x_j) be any oriented edge with $x_i, x_j \in K$, and let γ_{ij} be the minimizing geodesic from x_i to x_j. For any $s, t \in \mathcal{P}_i$ with $d_i^{\mathrm{R}}(s,t) \le r_0$ one has the local Lipschitz bound

$$d_j^{\mathrm{R}}\big(\Phi_{ij}(s), \Phi_{ij}(t)\big) \le C\, d_i^{\mathrm{R}}(s,t). \tag{22}$$

Moreover, for any (nondegenerate) discrete triangle (x_i, x_j, x_k) in K, define the continuous and discrete intermediate states starting from μ_i by

$$\tilde{s}_0 := \mu_i, \quad \tilde{s}_1 := \mathrm{PT}_{\gamma_{ij}}(\tilde{s}_0), \quad \tilde{s}_2 := \mathrm{PT}_{\gamma_{jk}}(\tilde{s}_1), \quad \tilde{s}_3 := \mathrm{PT}_{\gamma_{ki}}(\tilde{s}_2),$$

$$s_0 := \mu_i, \quad s_1 := \Phi_{ij}(s_0), \quad s_2 := \Phi_{jk}(s_1), \quad s_3 := \Phi_{ki}(s_2).$$

Then, each pair $(s_m, \tilde{s}_m)$ remains in the local regime:

$$d^{\mathrm{R}}(s_m, \tilde{s}_m) \le r_0 \quad \textit{for } m = 1, 2, 3,$$

and, quantitatively,

$$d^{\mathrm{R}}(s_m, \tilde{s}_m) \le C\, r_\varepsilon^{1+\alpha} \qquad (m = 1, 2, 3), \tag{23}$$

where r_ε is the sampling radius.

Proof. *Step 1: uniform local equivalence and choice of r_0.* Since K is compact and μ is smooth, $\mu(K)$ is compact in $\mathcal{P}$. By Assumption 1 and smooth dependence of the fibre metrics/divergences, there exists $r_0 > 0$ and constants $0 < c_1 \le c_2 < \infty$ such that for any $x \in K$ and any $s, t \in \mathcal{P}_x$ with $d_x^{\mathrm{R}}(s,t) \le r_0$,

$$c_1\, d_x^{\mathrm{R}}(s,t) \le d_x(s,t) \le c_2\, d_x^{\mathrm{R}}(s,t). \tag{24}$$

Step 2: local Lipschitz of Φ_{ij} in d^{R}. Assumption 9(1) gives $D_j(\Phi_{ij}(s), \Phi_{ij}(t)) \le C_D\, D_i(s,t)$. By Lemma 3, this implies $d_j(\Phi_{ij}(s), \Phi_{ij}(t)) \le \sqrt{C_D}\, d_i(s,t)$. If $d_i^{\mathrm{R}}(s,t) \le r_0$ then (24) yields

$$d_j^{\mathrm{R}}(\Phi_{ij}(s), \Phi_{ij}(t)) \le c_1^{-1} d_j(\Phi_{ij}(s), \Phi_{ij}(t)) \le c_1^{-1} \sqrt{C_D}\, d_i(s,t) \le c_1^{-1} \sqrt{C_D}\, c_2\, d_i^{\mathrm{R}}(s,t),$$

which is (22) with $C = c_1^{-1} \sqrt{C_D}\, c_2$.

Step 3: staying in the local regime and the bound (23). For $m = 1$, Assumption 9(2) (applied at $s = \mu_i$) gives

$$d^{\mathrm{R}}(s_1, \tilde{s}_1) = d^{\mathrm{R}}\big(\Phi_{ij}(\mu_i), \mathrm{PT}_{\gamma_{ij}}(\mu_i)\big) \le C_0\, d_g(x_i, x_j)^{1+\alpha} \le C\, r_\varepsilon^{1+\alpha}.$$

Choose ε_K so that $C r_\varepsilon^{1+\alpha} \le r_0$ for all $\varepsilon < \varepsilon_K$. Then, $(s_1, \tilde{s}_1)$ lies in the local regime.

Assume inductively that $d^{\mathrm{R}}(s_m, \tilde{s}_m) \le r_0$ and $d^{\mathrm{R}}(s_m, \tilde{s}_m) \le C r_\varepsilon^{1+\alpha}$ for $m = 1, 2$. Using the triangle inequality in d^{R},

$$\begin{aligned} d^{\mathrm{R}}(s_{m+1}, \tilde{s}_{m+1}) &= d^{\mathrm{R}}\big(\Phi(s_m), \mathrm{PT}(\tilde{s}_m)\big) \\ &\le d^{\mathrm{R}}\big(\Phi(s_m), \mathrm{PT}(s_m)\big) + d^{\mathrm{R}}\big(\mathrm{PT}(s_m), \mathrm{PT}(\tilde{s}_m)\big). \end{aligned}$$

The first term is bounded by Assumption 9(2): $\le C\, r_\varepsilon^{1+\alpha}$. The second term equals $d^{\mathrm{R}}(s_m, \tilde{s}_m)$ by Lemma 2. Therefore,

$$d^{\mathrm{R}}(s_{m+1}, \tilde{s}_{m+1}) \le C\, r_\varepsilon^{1+\alpha} + d^{\mathrm{R}}(s_m, \tilde{s}_m) \le C'\, r_\varepsilon^{1+\alpha}.$$

For $\varepsilon < \varepsilon_K$ this is $\le r_0$, closing the induction. □

Lemma 7 (Discrete vs continuous holonomy)**.** *Let $K \subset M$ be compact. There exist constants $C > 0$ and $\varepsilon_K > 0$ such that, for all $0 < \varepsilon < \varepsilon_K$, for every vertex $x_i \in V_\varepsilon \cap K$ and every discrete triangle $(x_i, x_j, x_k) \in \mathcal{T}_\varepsilon$ with all vertices in K and satisfying the non-degeneracy scale condition of Assumption 11(1), the following holds. Let $A_g(i,j,k) := A_g(x_i, x_j, x_k)$ and let $\ell_{ijk} := \max\{d_g(x_i, x_j), d_g(x_j, x_k), d_g(x_k, x_i)\}$. Then:*

1. $\ell_{ijk} \leq C\, r_\varepsilon$.
2. *In the fibre $\mathcal{P}_i$,*

$$d_i\Big(\mathrm{Hol}^{(\varepsilon)}_{ijk}(\mu_i),\ \mathrm{Hol}_{\tilde{\gamma}_{ijk}}(\mu_i)\Big) \leq C\,\rho_\varepsilon\, A_g(i,j,k), \qquad \rho_\varepsilon = r_\varepsilon^{\alpha-1}. \tag{25}$$

Proof. The first assertion is immediate from Assumption 5: each edge $\{x_i, x_j\}$ has length $d_g(x_i, x_j) \leq c_3 r_\varepsilon$, so any triangle with vertices connected by edges has all side lengths bounded by a constant multiple of r_ε. Adjusting constants yields $\ell_{ijk} \leq C r_\varepsilon$.

For the second assertion, write the continuous holonomy map as a composition of continuous parallel transport along the three geodesic edges:

$$\mathrm{Hol}_{\tilde{\gamma}_{ijk}} = \mathrm{PT}_{\gamma_{ki}} \circ \mathrm{PT}_{\gamma_{jk}} \circ \mathrm{PT}_{\gamma_{ij}},$$

where γ_{ij} denotes the minimizing geodesic segment from x_i to x_j.

We use a telescoping argument. Set

$$\begin{aligned} s_0 &:= \mu_i, & \tilde{s}_0 &:= \mu_i, \\ s_1 &:= \Phi_{ij}(s_0), & \tilde{s}_1 &:= \mathrm{PT}_{\gamma_{ij}}(\tilde{s}_0), \\ s_2 &:= \Phi_{jk}(s_1), & \tilde{s}_2 &:= \mathrm{PT}_{\gamma_{jk}}(\tilde{s}_1), \\ s_3 &:= \Phi_{ki}(s_2), & \tilde{s}_3 &:= \mathrm{PT}_{\gamma_{ki}}(\tilde{s}_2). \end{aligned}$$

Then, $\mathrm{Hol}^{(\varepsilon)}_{ijk}(\mu_i) = s_3$ and $\mathrm{Hol}_{\tilde{\gamma}_{ijk}}(\mu_i) = \tilde{s}_3$.

By Assumption 9(2), the edge-wise channel error is controlled in the *fibre Riemannian distance* d^{R}: after adjusting constants, for every oriented edge (x_a, x_b) and every $s \in \mathcal{P}_a$,

$$d^{\mathrm{R}}_b\Big(\Phi_{ab}(s),\ \mathrm{PT}_{\gamma_{ab}}(s)\Big) \leq C_0\, d_g(x_a, x_b)^{1+\alpha}.$$

We therefore perform the telescoping estimate entirely in the distances d^{R} (so that the triangle inequality holds globally), and only at the end pass back to the informational distance $d = \sqrt{2D}$ using the local equivalence of Remark 7 (which applies since the final displacement is $O(r_\varepsilon^{1+\alpha}) \to 0$).

Moreover, by Assumption 3(1), parallel transport is a fibrewise Riemannian isometry; in particular, it preserves fibre distances:

$$d^{\mathrm{R}}_b\big(\mathrm{PT}_{\gamma_{ab}}(s),\ \mathrm{PT}_{\gamma_{ab}}(t)\big) = d^{\mathrm{R}}_a(s,t).$$

We estimate recursively:

$$d^{\mathrm{R}}_j(s_1, \tilde{s}_1) = d^{\mathrm{R}}_j\big(\Phi_{ij}(s_0), \mathrm{PT}_{\gamma_{ij}}(\tilde{s}_0)\big) \leq C_0\, d_g(x_i, x_j)^{1+\alpha}.$$

Next,

$$\begin{aligned} d_k^{\mathrm{R}}(s_2, \tilde{s}_2) &= d_k^{\mathrm{R}}\big(\Phi_{jk}(s_1), \mathrm{PT}_{\gamma_{jk}}(\tilde{s}_1)\big) \\ &\le d_k^{\mathrm{R}}\big(\Phi_{jk}(s_1), \mathrm{PT}_{\gamma_{jk}}(s_1)\big) + d_k^{\mathrm{R}}\big(\mathrm{PT}_{\gamma_{jk}}(s_1), \mathrm{PT}_{\gamma_{jk}}(\tilde{s}_1)\big) \\ &\le C_0\, d_g(x_j, x_k)^{1+\alpha} + d_j^{\mathrm{R}}(s_1, \tilde{s}_1), \end{aligned}$$

and similarly,

$$\begin{aligned} d_i^{\mathrm{R}}(s_3, \tilde{s}_3) &= d_i^{\mathrm{R}}\big(\Phi_{ki}(s_2), \mathrm{PT}_{\gamma_{ki}}(\tilde{s}_2)\big) \\ &\le d_i^{\mathrm{R}}\big(\Phi_{ki}(s_2), \mathrm{PT}_{\gamma_{ki}}(s_2)\big) + d_i^{\mathrm{R}}\big(\mathrm{PT}_{\gamma_{ki}}(s_2), \mathrm{PT}_{\gamma_{ki}}(\tilde{s}_2)\big) \\ &\le C_0\, d_g(x_k, x_i)^{1+\alpha} + d_k^{\mathrm{R}}(s_2, \tilde{s}_2). \end{aligned}$$

Combining the three bounds yields

$$d_i^{\mathrm{R}}(s_3, \tilde{s}_3) \le C_1 \Big(d_g(x_i, x_j)^{1+\alpha} + d_g(x_j, x_k)^{1+\alpha} + d_g(x_k, x_i)^{1+\alpha} \Big) \le C_2\, \ell_{ijk}^{1+\alpha}.$$

Finally, since $d_i^{\mathrm{R}}(s_3, \tilde{s}_3) = O(\ell_{ijk}^{1+\alpha}) \to 0$, Remark 7 applies for ε small, and after adjusting constants we obtain the same bound in the informational distance:

$$d_i(s_3, \tilde{s}_3) \le C_2\, \ell_{ijk}^{1+\alpha}.$$

Now write

$$\ell_{ijk}^{1+\alpha} = \ell_{ijk}^{\alpha-1}\, \ell_{ijk}^2.$$

By the uniform non-degeneracy in Assumption 11(1), the triangle area satisfies $A_g(i,j,k) \ge c_A r_\varepsilon^2$, and since $\ell_{ijk} \le C r_\varepsilon$ we also have $\ell_{ijk}^2 \le C' A_g(i,j,k)$ (after adjusting constants). Hence,

$$d_i(s_3, \tilde{s}_3) \le C_3\, \ell_{ijk}^{\alpha-1}\, A_g(i,j,k) \le C_4\, r_\varepsilon^{\alpha-1}\, A_g(i,j,k) = C_4\, \rho_\varepsilon\, A_g(i,j,k),$$

which is exactly (25). □

6.2. A Triangle-Wise Discrete-to-Continuum Bound

Fix $\varepsilon > 0$ small, $x_i \in V_\varepsilon \cap K$, and a discrete triangle $(x_i, x_j, x_k) \in \mathcal{T}_\varepsilon$ with vertices in K satisfying the non-degeneracy condition of Assumption 11(1). Recall the triangle-wise curvature

$$K_{\mathrm{hol}}^{(\varepsilon)}(i,j,k) = \frac{\Delta_{ijk}^{(\varepsilon)}}{A_\varepsilon(i,j,k)}, \qquad \Delta_{ijk}^{(\varepsilon)} = d_i\big(\mu_i, \mathrm{Hol}_{ijk}^{(\varepsilon)}(\mu_i)\big),$$

and the corresponding continuous quantity

$$K_{\mathrm{hol}}^{\mathrm{cont}}(x_i, \Pi_{ijk}) = \lim_{A \to 0} \frac{d_i\big(\mu_i, \mathrm{Hol}_\gamma(\mu_i)\big)}{A},$$

where $\Pi_{ijk} \subset T_{x_i} M$ denotes the plane spanned by the geodesic initial directions from x_i to x_j and x_i to x_k (and γ is any shrinking family of geodesic triangles tangent to Π_{ijk}). In practice, we will compare the discrete loop γ_{ijk} with the continuous holonomy around the *geodesic* triangle $\Delta_g(x_i, x_j, x_k)$.

Let

$$\Delta_{ijk}^{\mathrm{cont}} := d_i\big(\mu_i, \mathrm{Hol}_{\tilde{\gamma}_{ijk}}(\mu_i)\big),$$

where $\tilde{\gamma}_{ijk}$ is the loop traversing the geodesic edges $i \to j \to k \to i$.

Lemma 8 (Triangle-wise defect comparison)**.** *There exist constants $C > 0$ and $\varepsilon_K > 0$ such that, for all $0 < \varepsilon < \varepsilon_K$ and all triangles $(x_i, x_j, x_k) \in \mathcal{T}_\varepsilon$ with vertices in K satisfying Assumption 11(1), one has*

$$\left|\Delta^{(\varepsilon)}_{ijk} - \Delta^{\mathrm{cont}}_{ijk}\right| \le C\,\rho_\varepsilon\,A_g(i,j,k), \tag{26}$$

where $\rho_\varepsilon = r_\varepsilon^{\alpha-1}$.

Proof. By the triangle inequality for the fibre Riemannian distance (Remark 7) and the local equivalence with d_i, we have

$$\left|\Delta^{(\varepsilon)}_{ijk} - \Delta^{\mathrm{cont}}_{ijk}\right| \le d_i\Big(\mathrm{Hol}^{(\varepsilon)}_{ijk}(\mu_i),\ \mathrm{Hol}_{\tilde{\gamma}_{ijk}}(\mu_i)\Big).$$

The right-hand side is bounded by Lemma 7 (25), which gives (26). □

We now incorporate the area approximation. Recall from Assumption 10 that $A_\varepsilon(i,j,k)$ satisfies

$$\left|A_\varepsilon(i,j,k) - A_g(i,j,k)\right| \le q_\varepsilon\,r_\varepsilon^2.$$

Lemma 9 (Triangle-wise curvature comparison)**.** *There exist constants $C > 0$ and $\varepsilon_K > 0$ such that, for all $0 < \varepsilon < \varepsilon_K$ and all triangles $(x_i, x_j, x_k) \in \mathcal{T}_\varepsilon$ with vertices in K satisfying Assumption 11(1), one has*

$$\left|K^{(\varepsilon)}_{\mathrm{hol}}(i,j,k) - \frac{\Delta^{\mathrm{cont}}_{ijk}}{A_g(i,j,k)}\right| \le C\,(\rho_\varepsilon + q_\varepsilon). \tag{27}$$

Proof. Write

$$K^{(\varepsilon)}_{\mathrm{hol}}(i,j,k) - \frac{\Delta^{\mathrm{cont}}_{ijk}}{A_g(i,j,k)} = \frac{\Delta^{(\varepsilon)}_{ijk}}{A_\varepsilon(i,j,k)} - \frac{\Delta^{\mathrm{cont}}_{ijk}}{A_g(i,j,k)}.$$

Add and subtract $\Delta^{\mathrm{cont}}_{ijk}/A_\varepsilon(i,j,k)$:

$$\frac{\Delta^{(\varepsilon)}_{ijk} - \Delta^{\mathrm{cont}}_{ijk}}{A_\varepsilon(i,j,k)} + \Delta^{\mathrm{cont}}_{ijk}\left(\frac{1}{A_\varepsilon(i,j,k)} - \frac{1}{A_g(i,j,k)}\right).$$

For the first term, Lemma 8 gives

$$\frac{\left|\Delta^{(\varepsilon)}_{ijk} - \Delta^{\mathrm{cont}}_{ijk}\right|}{A_\varepsilon(i,j,k)} \le C\,\rho_\varepsilon\frac{A_g(i,j,k)}{A_\varepsilon(i,j,k)}.$$

By Assumption 11(1) and Assumption 10, $A_\varepsilon(i,j,k)$ is uniformly comparable to $A_g(i,j,k)$ from below and above (for ε small), hence the ratio is bounded and the first term is $O(\rho_\varepsilon)$.

For the second term, note that by the holonomy expansion (Lemma 5) and the distance expansion (Lemma 4), there is a uniform constant C such that

$$\Delta^{\mathrm{cont}}_{ijk} = d_i\big(\mu_i, \mathrm{Hol}_{\tilde{\gamma}_{ijk}}(\mu_i)\big) \le C\,A_g(i,j,k),$$

for triangles in K sufficiently small. Therefore,

$$\left|\Delta^{\mathrm{cont}}_{ijk}\left(\frac{1}{A_\varepsilon} - \frac{1}{A_g}\right)\right| \le C\,A_g(i,j,k)\,\frac{\left|A_\varepsilon - A_g\right|}{A_\varepsilon A_g}.$$

Using Assumption 10 and the uniform lower bound $A_g(i,j,k) \ge c_A r_\varepsilon^2$ (Assumption 11(1)), together with the comparability $A_\varepsilon \ge \frac{c_A}{2} r_\varepsilon^2$, we obtain that this term is $O(q_\varepsilon)$. Combining the bounds yields (27). □

6.3. From Triangle-Wise to Sectional Curvature by Averaging

We now pass from the triangle-wise comparison to the averaged sectional quantity $K^{(\varepsilon)}_{\mathrm{hol}}(x,\Pi)$ defined in (7).

Fix $x \in K$ and a two-plane $\Pi \subset T_xM$. By Assumption 11, for ε sufficiently small we can choose a vertex $x_i \in V_\varepsilon$ with $d_g(x,x_i) \le c_2 r_\varepsilon$ and a finite non-empty family $\mathcal{T}_\varepsilon(x_i,\Pi)$ of triangles based at x_i satisfying:

- Uniform scale and non-degeneracy at scale r_ε;
- Planarity up to η_ε with respect to Π;
- Approximate directional isotropy in Π.

We write the discrete sectional curvature as an area-weighted average:

$$K^{(\varepsilon)}_{\mathrm{hol}}(x,\Pi) = \frac{\sum_{(x_i,x_j,x_k)\in\mathcal{T}_\varepsilon(x_i,\Pi)} \Delta^{(\varepsilon)}_{ijk}}{\sum_{(x_i,x_j,x_k)\in\mathcal{T}_\varepsilon(x_i,\Pi)} A_\varepsilon(i,j,k)}.$$

Similarly, define the corresponding continuous average over the same *geodesic* triangles:

$$\widetilde{K}^{\mathrm{cont}}_{\mathrm{hol}}(x_i,\Pi) := \frac{\sum_{(x_i,x_j,x_k)\in\mathcal{T}_\varepsilon(x_i,\Pi)} \Delta^{\mathrm{cont}}_{ijk}}{\sum_{(x_i,x_j,x_k)\in\mathcal{T}_\varepsilon(x_i,\Pi)} A_g(i,j,k)}.$$

Lemma 10 (Averaging stability)**.** *There exist constants $C > 0$ and $\varepsilon_K > 0$ such that, for all $0 < \varepsilon < \varepsilon_K$, for every $x \in K$ and every two-plane $\Pi \subset T_xM$,*

$$\left|K^{(\varepsilon)}_{\mathrm{hol}}(x,\Pi) - \widetilde{K}^{\mathrm{cont}}_{\mathrm{hol}}(x_i,\Pi)\right| \le C\,(\rho_\varepsilon + q_\varepsilon). \tag{28}$$

Proof. By Lemma 9, for each triangle in the family we have

$$\left|\frac{\Delta^{(\varepsilon)}_{ijk}}{A_\varepsilon(i,j,k)} - \frac{\Delta^{\mathrm{cont}}_{ijk}}{A_g(i,j,k)}\right| \le C(\rho_\varepsilon + q_\varepsilon).$$

Multiplying by $A_\varepsilon(i,j,k)$ and summing over the family gives

$$\left|\sum \Delta^{(\varepsilon)}_{ijk} - \sum \Delta^{\mathrm{cont}}_{ijk}\frac{A_\varepsilon(i,j,k)}{A_g(i,j,k)}\right| \le C(\rho_\varepsilon + q_\varepsilon)\sum A_\varepsilon(i,j,k).$$

Using uniform comparability $A_\varepsilon \simeq A_g$ on the family (Assumption 11(1) and Assumption 10), we may replace A_ε/A_g by 1 at the cost of an additional $O(q_\varepsilon)$ relative error. Dividing by the denominators and using again that $\sum A_\varepsilon \simeq \sum A_g$ yields (28). □

We now compare $\widetilde{K}^{\mathrm{cont}}_{\mathrm{hol}}(x_i,\Pi)$ with the intrinsic continuous quantity $K^{\mathrm{cont}}_{\mathrm{hol}}(x,\Pi)$.

Lemma 11 (Planarity and base-point stability)**.** *There exist constants $C > 0$ and $\varepsilon_K > 0$ such that, for all $0 < \varepsilon < \varepsilon_K$, for every $x \in K$, and every two-plane $\Pi \subset T_xM$,*

$$\left|\widetilde{K}^{\mathrm{cont}}_{\mathrm{hol}}(x_i,\Pi) - K^{\mathrm{cont}}_{\mathrm{hol}}(x,\Pi)\right| \le C\,(r_\varepsilon + \eta_\varepsilon). \tag{29}$$

Proof. First, by smoothness of $K^{\mathrm{cont}}_{\mathrm{hol}}$ in (x,Π) (Theorem 1 and smooth dependence of $W_x(\Pi;\mu_x)$), moving the base point from x to x_i induces an error bounded by $C\,d_g(x,x_i) \le Cr_\varepsilon$.

Second, by Assumption 11(2), the planes Π_{ijk} spanned by the geodesic directions of each triangle in the family are within angle η_ε of Π. Again by smoothness in Π (uniform on K), replacing Π_{ijk} by Π induces an additional error $O(\eta_\varepsilon)$.

Finally, the quantity $\widetilde{K}^{\mathrm{cont}}_{\mathrm{hol}}(x_i, \Pi)$ is an average over finitely many triangles of uniformly comparable shape and size $O(r_\varepsilon)$, hence the above pointwise stability bounds propagate to the average with the same order. □

6.4. Conclusion of the Proof of Theorem 2

Combining Lemmas 10 and 11, we obtain

$$\left|K^{(\varepsilon)}_{\mathrm{hol}}(x, \Pi) - K^{\mathrm{cont}}_{\mathrm{hol}}(x, \Pi)\right| \leq C\big(\rho_\varepsilon + q_\varepsilon + r_\varepsilon + \eta_\varepsilon\big),$$

uniformly for $x \in K$ and $\Pi \subset T_xM$. Recalling the definition $\kappa_\varepsilon = r_\varepsilon + \eta_\varepsilon + q_\varepsilon + \rho_\varepsilon$ from (10), this is exactly (11). This completes the proof of Theorem 2.

7. Examples and Model Constructions

In this section, we discuss several model constructions that illustrate the notion of informational holonomy curvature and the hypotheses of our convergence theorem. We first describe a basic classical choice of state space and divergence (the Jensen–Shannon model), then introduce a natural geometric state bundle built from tangent distributions and parallel transport, and finally discuss spaces of constant curvature and discrete sampling schemes.

7.1. The Classical Jensen–Shannon Model

We begin with a simple and concrete choice of state space and divergence, which fits into the general framework of Sections 2 and 4.

Fix a finite set

$$\Omega = \{1, \ldots, m\}, \qquad m \geq 2,$$

and let

$$S := \Delta^\circ(\Omega) = \Big\{p \in \mathbb{R}^m : \; p_i > 0, \; \sum_{i=1}^m p_i = 1\Big\}$$

denote the open probability simplex. We endow S with:

- The Fisher information metric g^{Fisher} (restricted to S),
- The Jensen–Shannon divergence

$$D_{\mathrm{JS}}(p, q) := H\Big(\frac{p+q}{2}\Big) - \frac{1}{2}H(p) - \frac{1}{2}H(q), \qquad H(p) = -\sum_{i=1}^m p_i \log p_i.$$

It is well known that D_{JS} is symmetric and non-negative and that $\sqrt{D_{\mathrm{JS}}}$ defines a genuine metric on S [9]. With our normalization $d = \sqrt{2D}$, we set

$$d_{\mathrm{JS}}(p, q) := \sqrt{2\, D_{\mathrm{JS}}(p, q)}$$

which is also a genuine metric on S. Moreover, the second-order expansion of D_{JS} at the diagonal yields a constant multiple of the Fisher metric:

$$D_{\mathrm{JS}}(p, q) = \frac{c_{\mathrm{JS}}}{2}\, g^{\mathrm{Fisher}}_p(v, v) + O\big(\|v\|^3_{g^{\mathrm{Fisher}}}\big),$$

where $v \in T_pS$ is the tangent vector such that $q = \exp^S_p(v)$ and $c_{\mathrm{JS}} > 0$ is a constant.

Thus Assumption 1 is satisfied with g^S proportional to the Fisher metric and $D = D_{\mathrm{JS}}$. In particular, the informational distance $d(s,t) = \sqrt{2D(s,t)}$ is locally equivalent to the Riemannian distance induced by g^{Fisher} on S.

Given a Riemannian manifold (M,g), the simplest associated state bundle is the trivial bundle

$$\pi : \mathcal{P} = M \times S \to M,$$

with fibre $(S, g^S, D_{\mathrm{JS}})$ independent of $x \in M$. To obtain a non-trivial informational holonomy curvature, however, one needs a connection on $\mathcal{P}$ whose curvature reflects the geometry of (M,g). The trivial product connection on $M \times S$ has zero curvature and yields vanishing holonomy and hence vanishing informational holonomy curvature. Thus, in interesting examples, the state bundle and its connection must be constructed from the Levi–Civita connection in a non-trivial manner. In this trivial product situation, the connection on $\mathcal{P}$ is taken to be the product of the Levi–Civita connection on (M,g) with the trivial connection on S, so that parallel transport acts as the identity on the fibre. Consequently, the Jensen–Shannon divergence is exactly preserved under parallel transport, i.e.,

$$D_{\mathrm{JS}}\big(\mathrm{PT}_\gamma(s), \mathrm{PT}_\gamma(t)\big) = D_{\mathrm{JS}}(s,t)$$

for every curve γ in M and every $s,t \in S$. Thus, Assumption 3 holds with $C_K = 0$, and Assumption 9(2) may be realised with vanishing channel-consistency error: in this example, the term ρ_ε in (10) does not contribute to the bound of Theorem 2 if the discrete channels Φ_{ij} are chosen to coincide with the exact parallel transport on $\mathcal{P}$.

7.2. A Geometric State Bundle from Tangent Distributions

We now describe a natural geometric construction in which the state bundle is built from probability distributions on tangent spaces and the connection is induced by parallel transport in (M,g).

Let (M,g) be a smooth Riemannian manifold of dimension n. For each $x \in M$, consider the tangent space T_xM with its Euclidean inner product g_x. Let

$$S_x := \mathcal{P}(T_xM)$$

be a chosen smooth manifold of probability measures on T_xM, for instance:

- The manifold of non-degenerate Gaussian measures on T_xM;
- Or a finite-dimensional exponential family of probability measures with smooth densities with respect to Lebesgue measure on T_xM.

To fix ideas, one may take S_x to be the set of Gaussian measures $\mathcal{N}(m,\Sigma)$ on T_xM, with mean $m \in T_xM$ and covariance matrix Σ in some fixed compact subset of the positive definite cone.

We define the state bundle

$$\pi : \mathcal{P} \to M, \qquad \mathcal{P}_x = S_x,$$

by realizing $\mathcal{P}$ as an associated bundle to the orthonormal frame bundle of (M,g): fix a typical fibre S_0 (e.g., a family of measures on $\mathbb{R}^n$), let $\mathrm{O}(n)$ act on S_0 by push-forward, and set $\mathcal{P} := \mathcal{F}_{\mathrm{O}}(M) \times_{\mathrm{O}(n)} S_0$; local trivializations are given by local frames (the bundle need not be globally trivial). The Riemannian metric $g^{\mathcal{P}_x}$ on each fibre is taken to be the Fisher information metric associated with the chosen statistical model on T_xM, and the divergence D_x is taken to be the Jensen–Shannon divergence between distributions in S_x.

To define the connection on $\mathcal{P}$, we use parallel transport in (M, g). Let $\gamma : [0, 1] \to M$ be a smooth curve with $\gamma(0) = x$ and $\gamma(1) = y$, and let

$$P_\gamma^g : T_xM \to T_yM$$

denote the parallel transport map associated with the Levi–Civita connection of g. We define the parallel transport on $\mathcal{P}$ along γ by pushing forward measures under P_γ^g:

$$\mathrm{PT}_\gamma : \mathcal{P}_x \to \mathcal{P}_y, \qquad \mathrm{PT}_\gamma(\nu) := (P_\gamma^g)_*\nu.$$

In particular, if $\nu = \mathcal{N}(m, \Sigma) \in S_x$ is Gaussian, then

$$\mathrm{PT}_\gamma(\nu) = \mathcal{N}(P_\gamma^g m,\ P_\gamma^g \circ \Sigma \circ (P_\gamma^g)^{-1}) \in S_y,$$

so that the family $(S_x)_{x\in M}$ is preserved by parallel transport. The resulting parallel transport maps satisfy the functoriality conditions (4), and thus define an Ehresmann connection on $\pi : \mathcal{P} \to M$.

The curvature of this connection is induced by the curvature of (M, g): the curvature of the Levi–Civita connection acts on T_xM via the Riemann curvature tensor R_x^g, and this, in turn, induces a curvature 2-form Ω on $\mathcal{P}$ by differentiation of the pushforward action on S_x. In particular, if the base manifold (M, g) has zero curvature, then the induced connection on $\mathcal{P}$ is flat and the informational holonomy curvature vanishes.

The reference state field $\mu : M \to \mathcal{P}$ can be chosen, for instance, as the isotropic Gaussian with mean $0 \in T_xM$ and covariance $\sigma^2\mathrm{Id}$ at each $x \in M$, for some fixed $\sigma > 0$. This is invariant under orthogonal transformations of T_xM, which simplifies the structure of $W_x(\Pi; \mu_x)$.

Under this construction, Assumptions 1 and 3 are satisfied: the Jensen–Shannon divergence on each S_x induces the Fisher metric, the pushforward by isometries preserves the Fisher metric and the second-order expansion of D_x, and the connection on $\mathcal{P}$ is metric along the fibres. Thus, the continuous informational holonomy curvature $K_{\mathrm{hol}}^{\mathrm{cont}}(x, \Pi)$ is well defined and determined by the curvature tensor R_x^g.

7.3. Spaces of Constant Curvature

We now consider the case where (M, g) has constant sectional curvature and specialise the construction of the previous subsection. The aim is to illustrate how the informational holonomy curvature reflects the constant curvature of the base manifold.

Let (M, g) be complete, simply connected, and of constant sectional curvature $\kappa \in \mathbb{R}$. Thus, (M, g) is isometric to the Euclidean space $\mathbb{R}^n$ (if $\kappa = 0$), the round sphere S^n (if $\kappa > 0$), or the hyperbolic space $\mathbb{H}^n$ (if $\kappa < 0$). The Riemann curvature tensor satisfies

$$R_x^g(u, v)w = \kappa\big(\langle v, w\rangle u - \langle u, w\rangle v\big),$$

for all $x \in M$ and $u, v, w \in T_xM$, where $\langle\cdot,\cdot\rangle$ is the Riemannian inner product.

We equip M with the geometric state bundle $\mathcal{P}$ of tangent distributions described in Section 7.2, with fibres consisting of Gaussian measures on T_xM and divergence given by the Jensen–Shannon divergence. We choose the reference state field μ_x to be the isotropic Gaussian $\mathcal{N}(0, \sigma^2\mathrm{Id})$ on T_xM, with fixed $\sigma > 0$ independent of x.

The isotropy of (M, g) implies that, for any $x \in M$ and any two 2-planes $\Pi, \Pi' \subset T_xM$, there exists an isometry of (M, g) mapping (x, Π) to (x, Π'). The induced action on the

state bundle $\mathcal{P}$ preserves the connection, the fibre metric and the divergence, and sends μ_x to itself. Thus, for each fixed x, the map

$$\Pi \longmapsto W_x(\Pi; \mu_x)$$

must have constant norm on the Grassmannian of 2-planes at x, and this norm can depend only on κ and on the parameters of the state bundle (e.g., σ and the choice of divergence). In particular, there exists a constant $c_{\mathrm{hol}}(\kappa, \sigma) > 0$ such that

$$\left\| W_x(\Pi; \mu_x) \right\|_{g^{\mathcal{P}_x}} = c_{\mathrm{hol}}(\kappa, \sigma)$$

for all $x \in M$ and all 2-planes $\Pi \subset T_x M$.

When $\kappa = 0$, the Levi–Civita connection is flat and the parallel transport maps P^g_γ along closed loops are the identity. Hence, the induced connection on $\mathcal{P}$ is flat, the holonomy maps on $\mathcal{P}$ are trivial, and $W_x(\Pi; \mu_x) = 0$, so

$$c_{\mathrm{hol}}(0, \sigma) = 0.$$

For $\kappa \neq 0$, the curvature of the Levi–Civita connection is non-zero and so is the curvature of the induced connection on $\mathcal{P}$; therefore, $W_x(\Pi; \mu_x)$ is non-zero and $c_{\mathrm{hol}}(\kappa, \sigma) > 0$.

By Theorem 1, the continuous informational sectional curvature is

$$K^{\mathrm{cont}}_{\mathrm{hol}}(x, \Pi) = \left\| W_x(\Pi; \mu_x) \right\|_{g^{\mathcal{P}_x}} = c_{\mathrm{hol}}(\kappa, \sigma),$$

which is constant in (x, Π). In particular, the informational holonomy curvature detects the constant curvature of (M, g) modulo the scale factor $c_{\mathrm{hol}}(\kappa, \sigma)$ coming from the choice of state bundle and divergence. In spaces of constant curvature, $K^{\mathrm{cont}}_{\mathrm{hol}}$ is thus a constant multiple of $|\kappa|$.

7.4. Discrete Sampling Schemes

Finally, we discuss concrete choices of sampling graphs, areas, and triangle families that satisfy the assumptions of Section 3 and enable the application of Theorem 2.

7.4.1. Quasi-Uniform Point Clouds

Let (M, g) be a compact Riemannian manifold. For each small parameter $\varepsilon > 0$, consider a finite set of points $V_\varepsilon \subset M$ obtained, for example, by:

- A deterministic quasi-uniform mesh (e.g., a geodesic triangulation or a regular grid in local charts);
- An i.i.d. sample of points with respect to the Riemannian volume measure, followed by a thinning procedure to enforce minimal separation.

Under natural mesh regularity conditions or with high probability under suitable random sampling, one can ensure that V_ε satisfies the separation and covering properties of Assumption 4 with $r_\varepsilon \sim |V_\varepsilon|^{-1/n}$. For example, in the random setting, results from geometric probability show that the typical spacing between neighbouring points is of order $|V_\varepsilon|^{-1/n}$ and that an appropriate choice of thresholds yields a Delone set.

7.4.2. Neighbour Graphs and Edges

Given V_ε, one natural choice of graph is the ρ_ε-neighbourhood graph: for a suitable radius R_ε satisfying

$$c_- r_\varepsilon \leq R_\varepsilon \leq c_+ r_\varepsilon,$$

one sets

$$E_\varepsilon := \{\{x_i, x_j\} : d_g(x_i, x_j) \le R_\varepsilon\}.$$

Alternatively, one can consider k-nearest neighbour graphs with k fixed or slowly increasing as $\varepsilon \to 0$. Under standard conditions, these constructions satisfy Assumption 5: edges connect points at distance of order r_ε, and the vertex degrees are uniformly bounded.

7.4.3. Discrete Areas and Triangle Families

Given a sampling graph $G_\varepsilon = (V_\varepsilon, E_\varepsilon)$ embedded in M, one can define the set of discrete triangles $\mathcal{T}_\varepsilon$ as in Definition 9. For each triangle $(x_i, x_j, x_k) \in \mathcal{T}_\varepsilon$, the discrete area $A_\varepsilon(i, j, k)$ can be chosen, for example, as:

- The Euclidean area of the triangle formed by the images of x_i, x_j, x_k in a normal coordinate chart centred at x_i;
- Or the area of a piecewise flat triangle obtained by approximating the metric g locally by its value at x_i.

In both cases, standard Taylor expansions in normal coordinates show that

$$|A_\varepsilon(i, j, k) - A_g(i, j, k)| \le C\, \ell_{ijk}^3 \le C'\, r_\varepsilon^3,$$

so Assumption 10 is satisfied with $q_\varepsilon \sim r_\varepsilon$.

Families of triangles $\mathcal{T}_\varepsilon(x_i, \Pi)$ satisfying Assumption 11 can be constructed by selecting, for each vertex x_i and each approximate direction in $T_{x_i}M$, a finite number of neighbouring vertices whose geodesic directions approximate a given 2-plane $\Pi \subset T_{x_i}M$ in an approximately isotropic fashion. In random sampling models, the law of large numbers ensures that the empirical distribution of edge directions becomes asymptotically isotropic, with deviations captured by a parameter $\eta_\varepsilon \to 0$. In deterministic constructions, Assumption 11 is ensured by imposing an additional directional regularity (isotropy) condition on the chosen local triangle stencils.

7.4.4. Discrete Channels from Continuous Transport

Finally, the discrete channels $\Phi_{ij} : \mathcal{P}_i \to \mathcal{P}_j$ on edges (x_i, x_j) can be defined by approximating the continuous parallel transport and local equivalence of transport maps $\mathrm{PT}_{\gamma_{ij}}$ along the minimizing geodesics $\gamma_{ij} : [0, 1] \to M$. In the geometric state bundle of Section 7.2, this amounts to approximating the pushforward of tangent distributions by the parallel transport $P^g_{\gamma_{ij}} : T_{x_i}M \to T_{x_j}M$.

For instance, one can set

$$\Phi_{ij} := \mathrm{PT}_{\gamma_{ij}},$$

whenever γ_{ij} is uniquely defined and computable, in which case Assumption 9(2) holds with $\rho_\varepsilon = 0$. In numerical settings where γ_{ij} and $P^g_{\gamma_{ij}}$ are approximated by finite-difference schemes or local polynomial approximations of g, consistency estimates of the form

$$d_j\big(\Phi_{ij}(s),\ \mathrm{PT}_{\gamma_{ij}}(s)\big) \le C\, d_g(x_i, x_j)^{1+\alpha}$$

can be obtained for suitable $\alpha > 1$, leading to a non-zero but convergent ρ_ε.

Under these constructions, all the assumptions of Sections 3 and 4 are satisfied, and Theorem 2 applies. Thus, the discrete informational sectional curvature $K^{(\varepsilon)}_{\mathrm{hol}}(x, \Pi)$ computed from sampling graphs, approximate geodesic triangles, and discrete channels converges, as $\varepsilon \to 0$, to the continuous informational sectional curvature $K^{\mathrm{cont}}_{\mathrm{hol}}(x, \Pi)$ determined by the geometric data $(M, g, \mathcal{P}, \mu)$.

7.5. A Worked Geometric Model and Verification of Assumptions

We briefly indicate a standard geometric setting in which the discrete assumptions of Section 3 can be verified by construction. Let (M, g) be closed (or, more generally, let us work on a fixed compact subset $K \subset M$ with bounded geometry: positive injectivity radius and uniform bounds on curvature). Choose $V_\varepsilon \subset K$ to be a maximal r_ε-separated set; then, V_ε is an r_ε-net and, in particular, a Delone set, yielding Assumption 4. Define the sampling graph by a radius rule, i.e., $(x, y) \in E_\varepsilon$ whenever $d_g(x, y) \leq R\, r_\varepsilon$ for a fixed $R > 1$; this gives local connectivity and uniform degree bounds as in Assumption 5. Triangle families approximating a plane $\Pi \subset T_x M$ can be obtained, for ε small, by geodesic triangles inside a convex normal neighbourhood of x (e.g., via local Delaunay/Voronoi constructions), so that their areas satisfy the required approximation estimates in Assumptions 10 and 11, … with errors of order $O(r_\varepsilon^3)$ (hence, Assumption 10 holds with $q_\varepsilon = O(r_\varepsilon)$) by standard normal-coordinate expansions. Finally, given any smooth state bundle $\mathcal{P} \to M$ with a smooth connection-like transport $y \mapsto \Phi_{x\to y}$, define edge channels by restricting $\Phi_{x\to y}$ to neighbours $(x, y) \in E_\varepsilon$; smoothness implies the channel-consistency estimates in Assumption 9 (2), with $\rho_\varepsilon \to 0$ as $r_\varepsilon \to 0$. For the divergence, choosing the Jensen–Shannon divergence on finite probability simplices (Section 7.1) provides a canonical instance satisfying Assumption 1; moreover, it supports finite-scale defects as discussed in Remark 4.

8. Discussion and Outlook

The constructions developed in this work provide a framework for defining and estimating curvature from *informational holonomy*. Starting from a Riemannian manifold (M, g) and a state bundle $\pi : \mathcal{P} \to M$ endowed with fibrewise divergences and a compatible connection, we defined a continuous informational holonomy curvature $K_{\mathrm{hol}}^{\mathrm{cont}}(x, \Pi)$ associated with a point $x \in M$ and a two-plane $\Pi \subset T_x M$ by measuring, via the informational distance induced by the divergence, the leading (area–linear) effect of transporting a reference state around small geodesic triangles. We then showed that, under explicit sampling, area-approximation, and channel-consistency assumptions, a purely discrete estimator $K_{\mathrm{hol}}^{(\varepsilon)}(x, \Pi)$ constructed on graphs embedded in M converges to $K_{\mathrm{hol}}^{\mathrm{cont}}(x, \Pi)$ as the sampling scale $\varepsilon \to 0$, with a quantitative error bound controlled by the discretization scale.

8.1. Summary of the Framework

The continuous construction hinges on three ingredients:

- A state bundle $\pi : \mathcal{P} \to M$ whose fibres represent informational states (classical or quantum), equipped with a fibre Riemannian metric and a divergence whose second-order expansion induces this metric;
- An Ehresmann connection on $\mathcal{P}$ compatible with the fibre metrics and divergences, so that parallel transport acts as a fibrewise isometry to first order and preserves the informational structure infinitesimally;
- A reference state field $\mu : M \to \mathcal{P}$ serving as a basepoint for measuring informational defects.

Holonomy of the connection on $\mathcal{P}$ along small geodesic triangles based at x and tangent to Π produces a displacement of μ_x in $\mathcal{P}_x$ which, by holonomy–curvature expansions, is proportional to the triangle area to first order. The informational distance associated with the fibre divergence then yields a scalar quantity per unit area, identified as the continuous informational sectional curvature $K_{\mathrm{hol}}^{\mathrm{cont}}(x, \Pi)$ (Theorem 1).

On the discrete side, we considered quasi-uniform sampling graphs $(G_\varepsilon)_{\varepsilon>0}$ on M, endowed with:

- discrete fibres $\mathcal{P}_i$ and divergences D_i at vertices $x_i \in V_\varepsilon$ approximating the continuous fibres and divergences;
- edge channels $\Phi_{ij} : \mathcal{P}_i \to \mathcal{P}_j$ approximating continuous parallel transport along short geodesic segments;
- discrete areas $A_\varepsilon(i,j,k)$ for triangles and triangle families $\mathcal{T}_\varepsilon(x_i, \Pi)$ that are asymptotically isotropic in prescribed directions.

The resulting discrete holonomy operators $\mathrm{Hol}_{ijk}^{(\varepsilon)}$ yield distance defects $\Delta_{ijk}^{(\varepsilon)}$, and the triangle-wise curvatures $K_{\mathrm{hol}}^{(\varepsilon)}(i,j,k)$ are defined by normalising by $A_\varepsilon(i,j,k)$. Averaging over $\mathcal{T}_\varepsilon(x_i, \Pi)$ produces a discrete sectional curvature $K_{\mathrm{hol}}^{(\varepsilon)}(x, \Pi)$ which converges, at a rate governed by κ_ε, to the continuous quantity $K_{\mathrm{hol}}^{\mathrm{cont}}(x, \Pi)$.

8.2. Relation to Classical and Discrete Curvature Notions

The informational holonomy curvature sits at the intersection of several strands of work on curvature:

- *Riemannian sectional curvature.* In classical Riemannian geometry, sectional curvature can be characterised in terms of angle defects, Jacobi fields, or holonomy of the Levi–Civita connection. Our framework replaces the linear tangent bundle by a (generally non-linear) state bundle and linear norms by informational distances induced by divergences. When the connection on $\mathcal{P}$ is induced by the Levi–Civita connection via a linear isometric representation, the vector $W_x(\Pi; \mu_x)$ in Theorem 1 is obtained as a linear image of the restriction of R_x^g to Π, and $K_{\mathrm{hol}}^{\mathrm{cont}}(x, \Pi) = \|W_x(\Pi; \mu_x)\|$ becomes an invariant of this restriction. In spaces of constant sectional curvature, this reduces to a constant multiple of $|\mathrm{sec}_g(x, \Pi)|$ (equivalently, of $|\kappa|$ when $\mathrm{sec}_g \equiv \kappa$).
- *Discrete and combinatorial curvature.* Various notions of curvature for graphs and discrete spaces have been proposed, including Ollivier–Ricci curvature, Forman curvature, and Regge-type discretisations. The discrete informational holonomy curvature differs from these in two key aspects: it is based on holonomy of a bundle connection (rather than on pairwise comparisons of neighbourhood measures or purely combinatorial angle/defect data), and it uses divergences on state spaces attached to vertices (rather than solely distances in the ambient manifold or graph). In particular, it blends geometric information about (M, g) with an informational structure in the fibres.
- *Curvature in information geometry.* In information geometry, Fisher metrics and α-connections yield Riemannian and affine structures on statistical manifolds, and their curvature encodes statistical properties of models. The present construction can be viewed as a "mixed" curvature: it is controlled by the curvature of a connection on a state bundle over a geometric base, while the informational structure enters via the choice of fibre divergence and reference state. The Jensen–Shannon model of Section 7.1 provides a particularly transparent example where the divergence has a direct information-theoretic meaning.

8.3. Limitations and Choices of State Bundle

The informational holonomy curvature is not a curvature of (M, g) alone: it depends on the choice of state bundle, divergence, connection, and reference state field. Different choices can therefore produce different curvature functionals over the same base manifold. This flexibility is both a strength and a limitation.

On the one hand, it allows the notion of curvature to be adapted to an application: classical probability distributions on finite sets with Jensen–Shannon divergence, Gaussian distributions on tangent spaces, or quantum density matrices with quantum Jensen–Shannon or other quantum divergences all fit naturally into the framework. On the

other hand, it raises the question of which choices are canonical, or geometrically natural, for a given problem.

A natural option in a purely geometric setting is the geometric state bundle built from tangent distributions (Section 7.2), whose connection is induced canonically by parallel transport in (M, g) and whose fibres behave naturally under base isometries. In data-driven or statistical settings, other choices may be more appropriate; for instance, state spaces encoding empirical distributions of local observations, feature vectors, or structured data attached to points in M.

From a foundational perspective, the present work treats the state bundle and its connection as given. Understanding how to construct such bundles in a canonical or data-driven way, and how the resulting $K_{\mathrm{hol}}^{\mathrm{cont}}$ varies across different constructions, remain interesting open questions.

8.4. Potential Applications and Further Directions

We conclude by mentioning several directions in which the informational holonomy curvature framework may be developed further.

8.4.1. Data Analysis and Manifold Learning

In applications where only a point cloud in M is observed, possibly together with empirical distributions or feature states at each point, the discrete curvature $K_{\mathrm{hol}}^{(\varepsilon)}(x, \Pi)$ provides a way to estimate curvature-like quantities that combine geometric and informational structure. Compared to purely metric estimators based on distances or angles, informational holonomy curvature incorporates how local states are transported along the graph via channels, which may reflect dynamics, diffusion, or parallel transport in latent spaces. Analysing statistical properties and robustness of such estimators in the presence of noise and finite-sample effects is a natural next step.

8.4.2. Other Divergences and Connections

While we focused on divergences whose second-order expansion induces a Riemannian metric (e.g., Jensen–Shannon), one could consider more general f-divergences or Bregman divergences, potentially leading to non-Riemannian local geometry on fibres. Extending the holonomy curvature construction to such settings would require an appropriate notion of distance defect and a careful analysis of higher-order terms. Similarly, one may study families of connections on $\mathcal{P}$ (for example, α-connections in information geometry) and compare the corresponding informational holonomy curvatures.

8.4.3. Ricci-Type and Scalar Informational Curvatures

The present work is focused on sectional-type curvature attached to two-planes. In analogy with Riemannian geometry, one may seek informational versions of Ricci and scalar curvature. One possibility is to average $K_{\mathrm{hol}}^{\mathrm{cont}}(x, \Pi)$ over the Grassmannian of 2-planes at x with respect to a suitable measure, obtaining a scalar quantity $K_{\mathrm{hol}}^{\mathrm{scal}}(x)$, and relating it to classical scalar curvature or to information-theoretic quantities such as entropy production or functional inequalities. On the discrete side, different averaging schemes over triangle families may yield Ricci-type informational curvatures along edges or preferred directions in the graph.

8.4.4. Quantum and Non-Commutative Models

The state-bundle viewpoint naturally accommodates quantum state spaces, where fibres consist of density matrices on finite-dimensional Hilbert spaces and divergences are given by quantum generalisations of Jensen–Shannon or relative entropy. In such settings, the connection on $\mathcal{P}$ may encode both geometric parallel transport and quantum channels

acting along paths in M. Extending the convergence analysis to non-commutative state bundles, and understanding how informational holonomy curvature reflects underlying quantum geometric structure, are promising directions.

8.4.5. Algorithmic and Numerical Aspects

From a practical perspective, computing $K_{\mathrm{hol}}^{(\varepsilon)}(x, \Pi)$ requires:

1. Constructing a sampling graph and identifying triangle families $\mathcal{T}_\varepsilon(x_i, \Pi)$;
2. Specifying discrete channels Φ_{ij} and evaluating their composition along loops;
3. Computing divergences and distances in the fibre state spaces.

Each of these steps has algorithmic consequences, and different applications may favour different trade-offs between accuracy and complexity. Designing efficient algorithms for informational holonomy curvature in high-dimensional state spaces, and testing them on simulated and real data, would help assess the practical relevance of the notion.

In summary, the informational holonomy curvature introduced here provides a bridge between classical Riemannian curvature, graph-based approximations, and information-theoretic structures on state spaces. It offers a geometrically grounded way of measuring how "information" twists under transport around small loops. The results of this paper establish its mathematical foundation and discrete-to-continuous consistency; its full potential will likely emerge in concrete applications and in further theoretical developments linking geometry, probability, and information.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Amari, S.; Nagaoka, H. *Methods of Information Geometry*; Translations of Mathematical Monographs; American Mathematical Society: Providence, RI, USA, 2000; Volume 191.
2. Ollivier, Y. Ricci curvature of Markov chains on metric spaces. *J. Funct. Anal.* **2009**, *256*, 810–864. [CrossRef]
3. Forman, R. Bochner's method for cell complexes and combinatorial Ricci curvature. *Discret. Comput. Geom.* **2003**, *29*, 323–374. [CrossRef]
4. Regge, T. General relativity without coordinates. *Il Nuovo C.* **1961**, *19*, 558–571. [CrossRef]
5. Gutierrez-Ule, D. Discrete Informational Curvature from the Jensen–Shannon Metric: A Conditional Regge-Type Convergence to the Scalar Curvature. *Zenodo* **2026**, *preprint*. [CrossRef]
6. do Carmo, M.P. *Riemannian Geometry*; Birkhäuser: Boston, MA, USA, 1992.
7. Kobayashi, S.; Nomizu, K. *Foundations of Differential Geometry*; Interscience Publishers: New York, NY, USA, 1963; Volume 1.
8. Lin, J. Divergence measures based on the Shannon entropy. *IEEE Trans. Inf. Theory* **1991**, *37*, 145–151. [CrossRef]
9. Endres, D.M.; Schindelin, J.E. A new metric for probability distributions. *IEEE Trans. Inf. Theory* **2003**, *49*, 1858–1860. [CrossRef]
10. Lagarias, J.C. Geometric models for quasicrystals I. Delone sets of finite type. *Discret. Comput. Geom.* **1992**, *7*, 331–406.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.