# SMALL REGIONS, PERSISTENT HOMOLOGY: AN ILLUSTRATED COMPANION TO THE CONSTRUCTION

ANASS NIFA

Abstract. This companion explains the construction in *Homological obstructions to local contractibility in noncollapsed Ricci limits*. Boundary-collapse maps detect independent second-homology classes in the ambient limit. Local metric control and a perturbation argument preserve these classes after deletion of an arbitrary set below the stated Hausdorff-dimension threshold. We describe the regularized circle profile, the compensation of spatial Ricci terms, the flat parameters and positive floors, and the annuli and fixed collars that retain each inserted domain. For the isolated example, localized distance estimates identify a unique product-cone tangent without imposing a topologically conical neighborhood. The figures distinguish geometric models, algebraic estimates, and schematic decompositions.

## 1. Small for distance, persistent in homology

The construction is organized around two ways of measuring a small region. Its intrinsic diameter controls the error made when paths pass through it. A continuous map from the whole ambient space onto a quotient of that region can nevertheless detect a homology class inside it. These statements remain compatible when infinitely many disjoint regions accumulate. The pieces are

$$P_4 = \mathbb{CP}^2 \setminus \operatorname{Int} D^4, \qquad P_5 = (S^2 \times S^3) \setminus \operatorname{Int} D^5.$$

Each has second homology $\mathbb{Z}$. Scaled copies, transition annuli, and collars are inserted into a sphere with regularized codimension-two geometry. Once one domain has been inserted, its Riemannian metric is retained exactly in all later approximants. Its boundary-collapse map then passes to the limit. A class detected by a map defined on the whole limit cannot be killed by allowing a bounding chain to leave the inserted domain. [M, §§1.5, 2].

For $n = 4, 5$, the accumulation set is a nonempty compact set $F$ with empty interior in a small coordinate ball of $S^{n-2}$. The insertion centers approach every point of $F$. A positive-measure choice gives an obstruction to deleting an arbitrary set of smaller dimension. A singleton choice permits a separate, precise calculation of a unique tangent.

The obstruction here is *absolute neighborhood homology detected in the ambient space*. It is not the relative local-homology calculation for a cone over a fixed compact link. In particular, the isolated example below is not locally contractible and has no open topologically conical neighborhood. The distinct construction discussed in [M, §1.4] has contractible conical neighborhoods; its fixed-link relative-homology formula is not used here.

## 2. The two conclusions

Homology means singular homology with integer coefficients, and $\operatorname{rank} A = \dim_{\mathbb{Q}}(A \otimes_{\mathbb{Z}} \mathbb{Q})$. Write $S_{\mathrm{top}}(X)$ for the points without an ambient Euclidean neighborhood. The dimension $n \in \{4, 5\}$ is that of the basic construction; $N = n + k$ is the dimension after taking a product with a fixed flat torus $T^k$. The space $T^0$ is a point. A closed manifold is compact and without boundary. All second fundamental forms use the outward normal: $\mathrm{II}(V, W) = \langle \nabla_V \nu_{\mathrm{out}}, W \rangle$. Balls are open unless a bar is displayed; closed balls in metric comparisons carry restricted ambient distances, not their intrinsic length distances.

**An obstruction along a large set.** For each integer $N \geq 4$, [M] constructs a compact noncollapsed Ricci limit $X_N$ and a compact subset $Z_N \subset S_{\mathrm{top}}(X_N)$ such that

$$0 < \mathcal{H}^{N-2}(Z_N) \leq \mathcal{H}^{N-2}(S_{\mathrm{top}}(X_N)) < \infty. \tag{2.1}$$

For *every* subset $E \subset X_N$ with $\dim_{\mathcal{H}} E < N-2$, every $z \in Z_N \setminus E$, and every relative open neighborhood $U$ of $z$ in $X_N \setminus E$,

$$\operatorname{rank} \operatorname{im}[H_2(U;\mathbb{Z}) \longrightarrow H_2(X_N;\mathbb{Z})] = \infty. \tag{2.2}$$

The arrow is induced by inclusion into the *original ambient space*. The set $E$ need not be closed or measurable; Hausdorff dimension uses outer measures. Since $\mathcal{H}^{N-2}(E) = 0$, some point of $Z_N$ survives. The complement fails to be locally contractible at every surviving distinguished point and is not a topological manifold. [M, Theorem 1.1; Corollary 1.2].

**An isolated obstruction with an explicitly identified tangent.** A separate compact four-dimensional example satisfies

$$S_{\mathrm{top}}(Y) = \{p\}, \qquad \mathrm{Tan}(Y,p) = \{\mathbb{R}^2 \times \mathrm{C}(S^1_{2\pi q})\}, \qquad 0 < q < 1. \tag{2.3}$$

Here $q \in (0,1)$ is a suitable fixed choice supplied by the construction, and $S^1_{2\pi q}$ is the intrinsic circle of circumference $2\pi q$. Every open neighborhood $U$ of $p$ has infinite-rank image in $H_2(Y;\mathbb{Z})$. The puncture $Y \setminus \{p\}$ is locally isometric to a smooth Riemannian four-manifold. The tangent link is $S^1 * S^1_{2\pi q}$, homeomorphic to $S^3$, but it is not a topological link for a neighborhood of $p$. In fact, no such neighborhood is an open cone over any compact space. [M, Theorem 1.3].

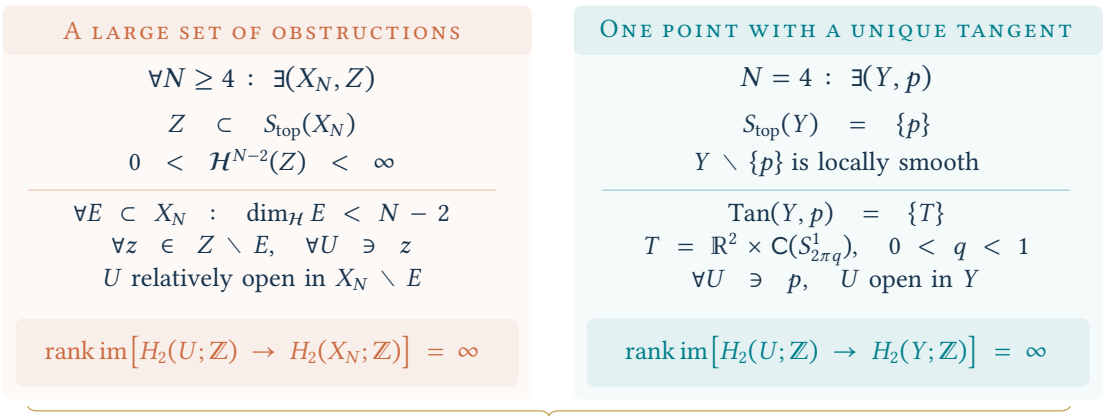



FIGURE 1. For every integer $N \geq 4$, the first construction has a positive-measure distinguished set and quantifies over arbitrary smaller-dimensional deletions and relative open neighborhoods. The separate four-dimensional construction has one nonmanifold point and a unique metric tangent. Both homology arrows have the original ambient limit as target.

In both constructions the realizing manifolds are closed, with a common positive lower bound for based unit-ball volumes and

$$\sup_j \int_{M_j} \left(\mathrm{Scal}_{g_j} + |\mathrm{Ric}_{g_j}|_{g_j}\right) dv_{g_j} < \infty. \tag{2.4}$$

Their Ricci curvature is nonnegative. In the basic dimensions four and five there is a common strictly positive lower bound. The Hilbert–Schmidt norm is used in (2.4). These are estimates on the constructed sequences, not a general scalar-integral estimate from a lower Ricci bound.

## 3. Three tasks in the construction

First construct the maps that recognize one inserted class. Keeping both the inclusion and the boundary-collapse map through metric convergence proves independence. Local distance preservation then transfers a thin deleted set to a smooth domain, where representative spheres can be perturbed away from it.

Next supply the geometry required by these maps: regularize the circle, compensate for spatial derivatives, build flat parameters with exact plateaux, and join each core to the background through an annulus with strict boundary margins. The collars must remain fixed, not merely become progressively closer.

Finally compare distances. Noncollapse and a diameter bound supply a compact limit; summable replacement costs supply the integral bound. For the isolated example a *localized* tail estimate, stronger than global summability, shows that replacements do not alter the background tangent.

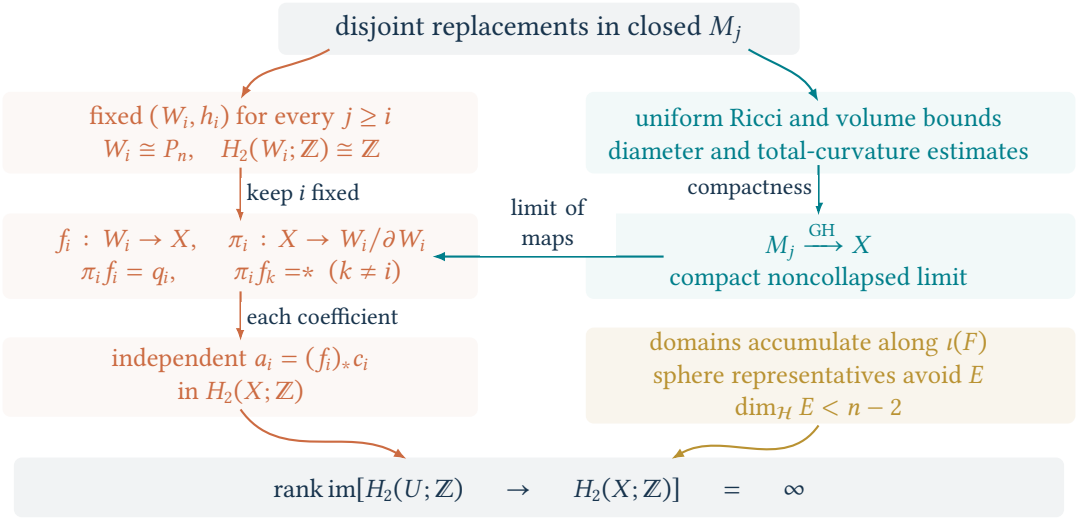



Figure 2. The topology is carried by fixed domains and limiting maps; compactness is supplied by the metric estimates. Localizing the domains and moving sphere representatives give the neighborhood assertion. No continuity of singular homology under Gromov–Hausdorff convergence is assumed.

**Reading routes.** For the topological argument, read Sections 4–9, using Visuals F–H to follow avoidance and the neighborhood obstruction. For the geometric construction, read Sections 10–16, especially the flat parameters in Section 11 and the boundary identifications in Sections 13–15. Visuals K and M explain exact parameter stabilization; Visual L is an optional reminder of exponential flatness. For the isolated tangent, read Sections 17–19, following Visuals V–X from quotient chains to same-radius comparisons. The source table in Section 20 records the corresponding results of [M]. Each inserted domain has dimension $n$; its boundary and the annular links have dimension $n-1$; the parameter sphere has dimension $n-2$.

The numbered figures provide the main sequence of geometric models and estimates. The lettered visuals isolate individual proof mechanisms, such as an unavoidable path cost, exact stabilization, or a localized sum. Use them alongside the argument identified in their captions. A drawn outline is not an embedding of a high-dimensional core. A plotted algebraic bound is not a computed field of Ricci eigenvalues. Captions specify these distinctions.

## 4. Collapsing a boundary gives a detector

Let $W$ be a compact connected Riemannian domain with nonempty boundary, with intrinsic distance $d_W$, and put $r_W(u) = d_W(u, \partial W)$. Collapsing the entire boundary gives the exact metric

$$d_{W/\partial W}([u],[v]) = \min\{d_W(u,v), r_W(u) + r_W(v)\}. \tag{4.1}$$

The two alternatives mean staying in $W$, or reaching the collapsed boundary and returning from it. The boundary-distance function is 1-Lipschitz. This and the triangle inequality for $d_W$ verify the triangle inequality in each combination of the alternatives. The only zero classes are the entire boundary and the individual interior points. The quotient map is continuous; the resulting metric topology equals the compact topological quotient topology. [M, Lemma 2.2].

If $W \subset M$ is a smooth domain, extend its quotient map by sending $M \setminus \operatorname{Int} W$ to the boundary point. Denote the result by $\pi_W : M \to W/\partial W$. A curve between interior endpoints either stays inside $W$, or pays at least $r_W(u) + r_W(v)$ before its first exit and after its last entrance. If only one endpoint is interior, its first boundary encounter costs at least its boundary distance; two exterior endpoints have the same image. Taking infima proves

$$d_{W/\partial W}(\pi_W(u), \pi_W(v)) \le d_M(u,v).$$

Thus shortcuts outside $W$ do not invalidate the detector. In contrast, the inclusion $(W, d_W) \to (M, d_M)$ need not be an isometry: it is only 1-Lipschitz at this stage.

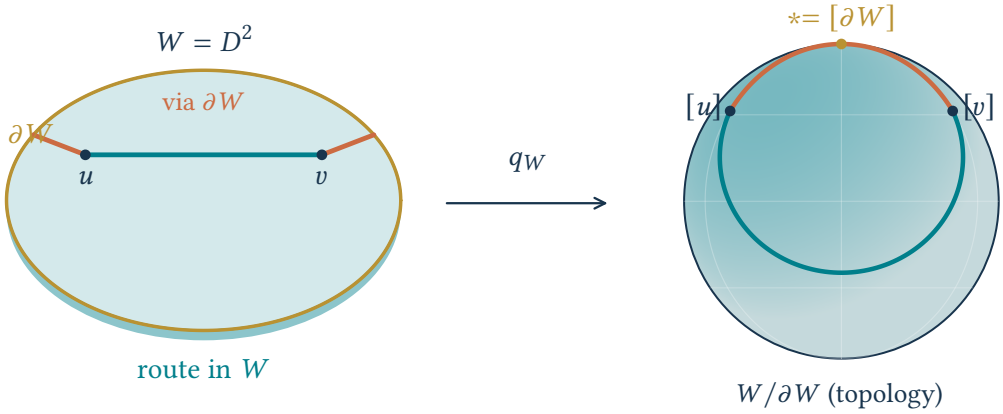


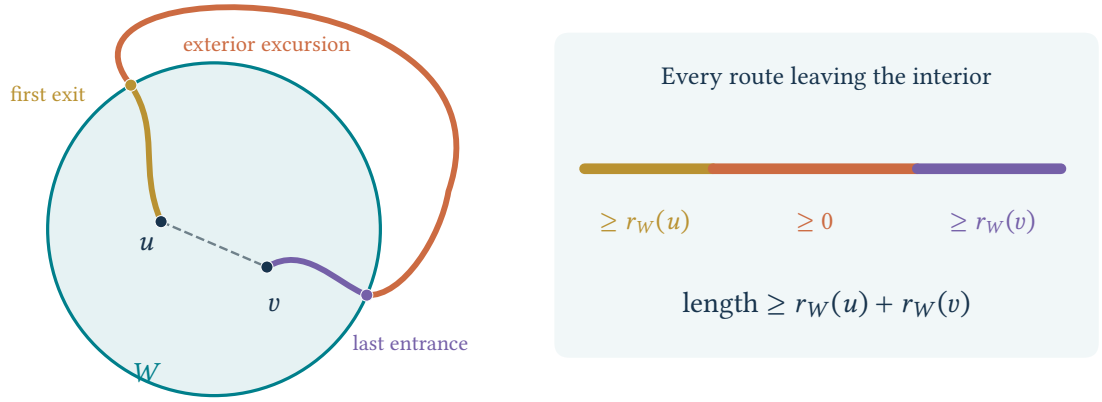

Figure 3. The unit-disc illustration shows the two routes in the quotient metric. The sphere represents the topology of $D^2/\partial D^2$, not a round metric on that quotient. Formula (4.1) applies to every compact connected Riemannian domain with nonempty boundary.

## 5. Which homology class does the quotient see?

Write $Q = \mathbb{CP}^2$ or $S^3 \times S^2$, with $P_n = Q \setminus \operatorname{Int} D^n$. The cellular decomposition of $\mathbb{CP}^2$ and the Künneth formula for $S^3 \times S^2$ give $H_2(Q;\mathbb{Z}) = \mathbb{Z}$. A boundary collar makes $(Q, P_n)$ a good pair, and $Q/P_n \cong S^n$. Since $n \ge 4$, $H_3(Q, P_n) = H_2(Q, P_n) = 0$; hence $H_2(P_n) \cong H_2(Q)$. The boundary-pair sequence contains

$$0 = H_2(S^{n-1}) \longrightarrow H_2(P_n) \longrightarrow H_2(P_n, \partial P_n) \longrightarrow H_1(S^{n-1}) = 0.$$

An exterior shortcut still pays twice at the boundary




$$d_{W/\partial W}(\pi_W u, \pi_W v) \le d_M(u,v)$$

The detector remains Lipschitz even when the inclusion is not isometric.

VISUAL A. The first exit and last entrance isolate two compulsory pieces of any exterior route. Their lengths dominate the distances of the endpoints to the boundary, regardless of the shortcut between them. Interior routes are controlled by $d_W$. Taking the infimum over both types proves the detector inequality. The disc is a geometric model for this length argument, not the topology of an inserted core.

The collar also identifies the relative group with the reduced homology of the quotient. Thus the actual quotient map induces

$$(q_n)_* : H_2(P_n;\mathbb{Z}) \xrightarrow{\cong} \widetilde{H}_2(P_n/\partial P_n;\mathbb{Z}) \cong \mathbb{Z}. \tag{5.1}$$

This is an infinite-order class, not merely a nonzero class. [M, Lemma 2.1].

The geometric input is different from this homology calculation. The core metrics supplied in [M, Lemma 2.1] have unit round boundary and constants $\mu, \alpha > 0$ with

$$\mathrm{Ric}_k \ge \mu k, \qquad \mathrm{II}_{\partial P,k} \ge \alpha g_{S^{n-1}}.$$

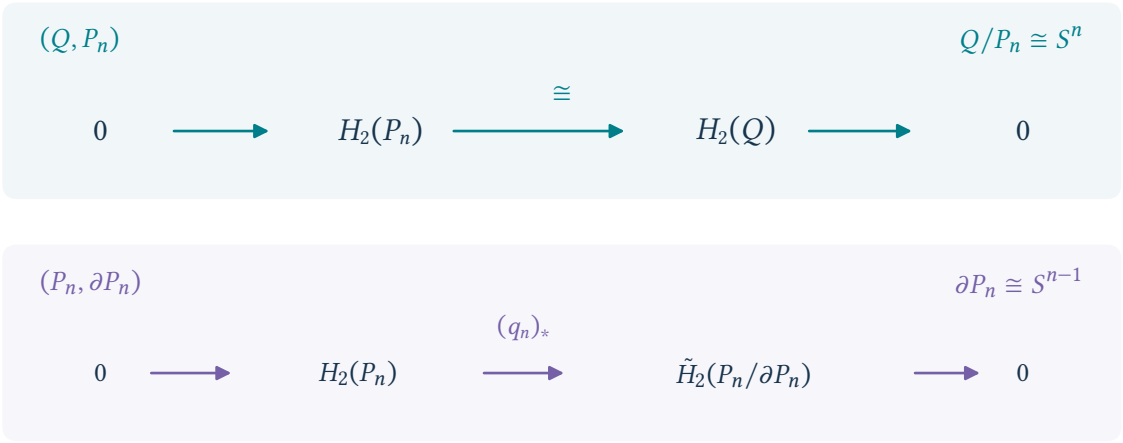
Their existence uses the core-metric results cited there. Rescaling a round boundary to unit radius preserves positive Ricci curvature and strict convexity; compactness then gives the displayed positive minima. The companion uses this geometric input and does not replace it by a topological assertion.

## 6. KEEP THE MAPS WHILE THE SPACES CONVERGE

Suppose compact connected Riemannian manifolds $M_j$ converge to $X$ in Gromov–Hausdorff distance. For each fixed $i$, assume that one smooth domain $(W_i, h_i^W)$, diffeomorphic to $P_n$, occurs with unchanged Riemannian metric in every $M_j$, $j \ge i$. At each finite stage these domains are pairwise disjoint. Their ambient distances may change because of shortcuts elsewhere. Their intrinsic metrics and the quotient targets $Q_i = W_i/\partial W_i$, however, are fixed.

**Why a limiting detector exists.** Realize the spaces isometrically in a common compact metric space $Z$. The inclusions $e_{i,j} : W_i \to Z$ are 1-Lipschitz; Arzelà–Ascoli gives uniform limits $f_i : W_i \to X$ after a diagonal subsequence. For the maps $\pi_{i,j} : M_j \to Q_i$, take Hausdorff limits of their compact graphs in $Z \times Q_i$. Every point $x \in X$ is the limit of points $x_j \in M_j$, so the limiting graph projects onto $X$. If $(x,u),(y,v)$ belong to it, the original Lipschitz inequalities give $d_{Q_i}(u,v) \le d_X(x,y)$. Taking $y = x$ proves single-valuedness. It is therefore the graph of a 1-Lipschitz map $\pi_i : X \to Q_i$.

Removing a ball and collapsing its boundary preserve degree two




$$n \geq 4 : \quad H_3(S^n) = H_2(S^n) = H_2(S^{n-1}) = H_1(S^{n-1}) = 0$$

VISUAL B. Two exact-sequence segments identify the same infinite-order class. Removing the open ball preserves $H_2$ because the relevant relative groups vanish. Collapsing the boundary preserves it because the neighboring boundary groups vanish. The collar identifies relative homology with quotient homology, so the lower arrow is induced by the actual quotient map $q_n$, which is the map later used by the detector.

There are only countably many domains. One diagonal subsequence retains both kinds of maps for every fixed index. Passing the finite-stage identities to the closed limiting graphs gives

$$\pi_i f_i = q_i, \qquad \pi_i f_k = * \quad (k \neq i), \qquad f_i : W_i \to X, \quad \pi_i : X \to Q_i. \tag{6.1}$$

Here $* = [\partial W_i]$. These are identities between continuous maps, not consequences of continuity of singular homology under metric convergence. [M, Proposition 2.3].

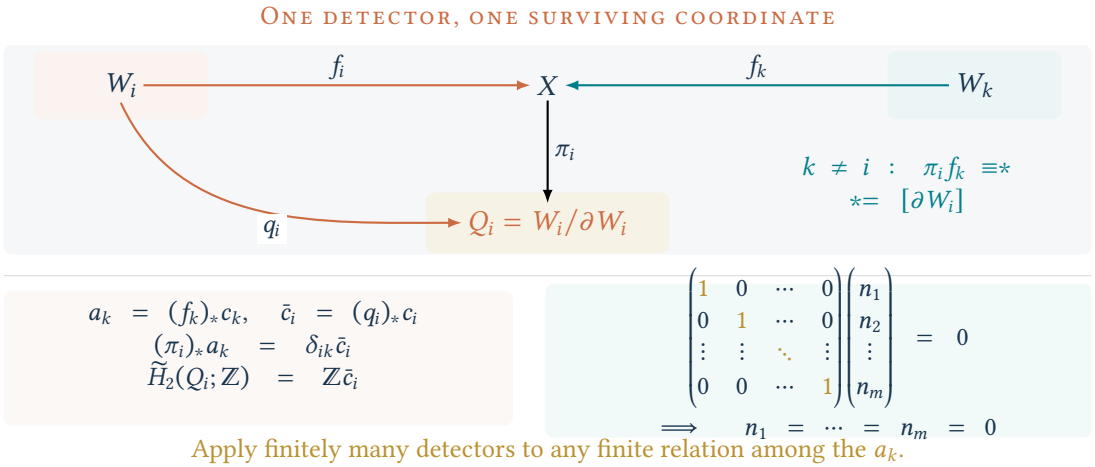
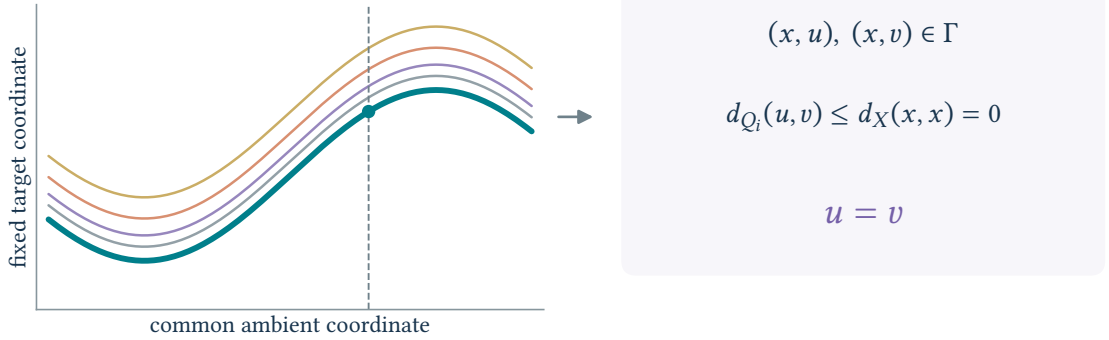



FIGURE 4. The limiting triangle commutes. The quotient of $W_i$ detects its own class and annihilates classes from the other domains. Applying finitely many detectors to a finite relation gives an identity matrix, both integrally and after tensoring with $\mathbb{Q}$.

A limiting Lipschitz graph remains single-valued

Scalar 1-Lipschitz graph prototype




$$\pi_i f_i = q_i \qquad \pi_i f_k = * \quad (k\neq i)$$

Compact graphs retain the maps and their identities.

Visual C. Uniform Lipschitz control prevents branching in the limiting graph: two target values above the same point have distance zero. Surjectivity of the graph projection therefore produces a map on all of $X$. The plotted scalar graphs illustrate this mechanism only; the proof uses compact graphs in $Z\times Q_i$ and passes both detector identities through their closed limits.

Choose a generator $c_i \in H_2(W_i;\mathbb{Z})$, and set $a_i = (f_i)_* c_i$. By (5.1) and (6.1),

$$(\pi_i)_* a_i = (q_i)_* c_i \text{ generates } \mathbb{Z}, \qquad (\pi_i)_* a_k = 0 \quad (k\neq i).$$

If $\sum_{i\in I} m_i a_i = 0$, with $I$ finite, applying $(\pi_k)_*$ gives $m_k (q_k)_* c_k = 0$, so $m_k = 0$. Applying $(\pi_k)_* \otimes \mathrm{Id}_{\mathbb{Q}}$ proves the same statement for rational relations. Any subgroup containing infinitely many $a_i$ has infinite rank. No infinite singular chain or infinite sum of classes is required. [M, Proposition 2.4].

The geometry supplies markings $z_i \in F$, a bi-Lipschitz embedding $\iota : F \to X$, and $\eta_i \downarrow 0$ with

$$\sup_{w\in W_i} d_X(f_i(w), \iota(z_i)) \le \eta_i.$$

For each $z \in F$, a strictly increasing subsequence has $z_i \to z$. Along it, the *whole* images $f_i(W_i)$ eventually lie in any open neighborhood $U$ of $\iota(z)$. Thus these maps factor through $U$, and

$$\operatorname{rank}\operatorname{im}[H_2(U;\mathbb{Z}) \longrightarrow H_2(X;\mathbb{Z})] = \infty. \tag{6.2}$$

Section 15 explains the metric marking that supplies this estimate. The ambient target in (6.2) is essential: all of $X$ is already available for bounding chains.

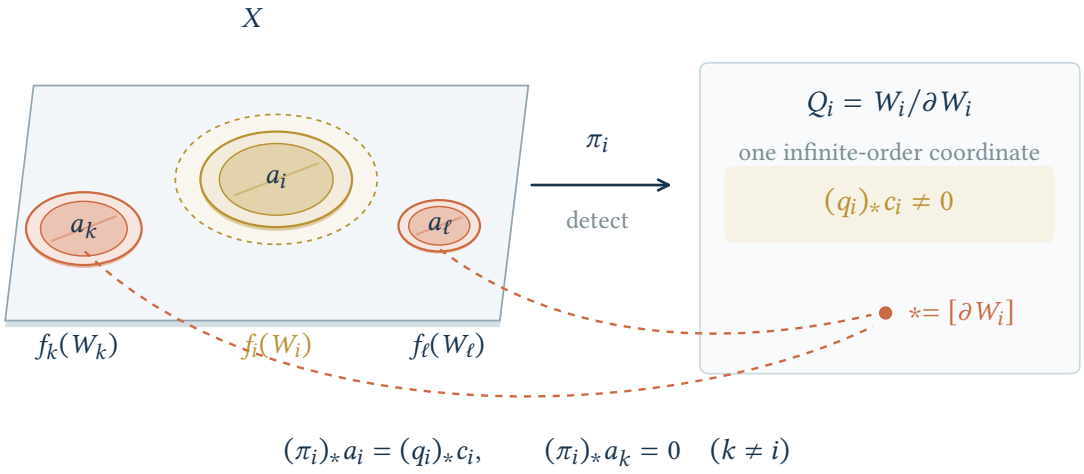
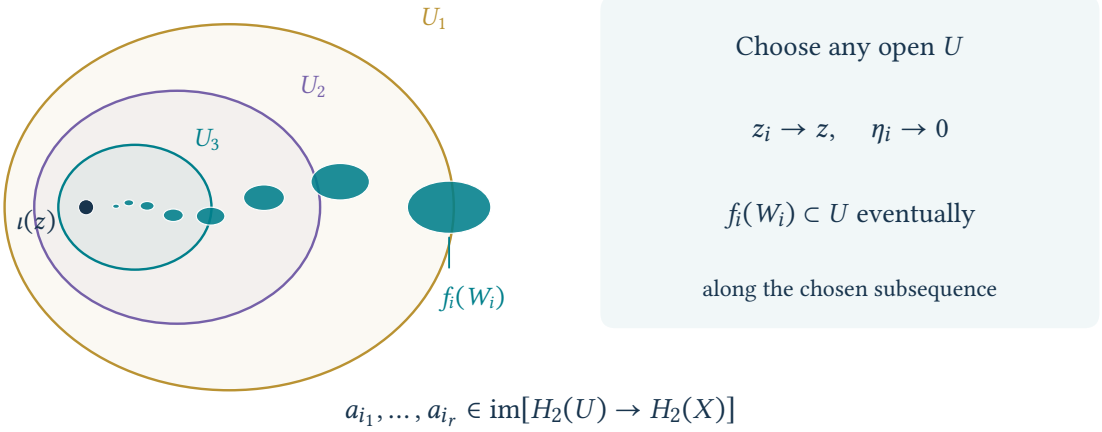



FIGURE 5. The detector has the whole ambient limit as its domain. It retains the coordinate contributed by $W_i$ and sends every other fixed domain to the collapsed boundary. The drawn bodies indicate positions, not core topology or embedded representing spheres.

Whole domains eventually fit inside every neighborhood




VISUAL D. The estimate controls every point of each domain, not just its center. Along a subsequence with $z_i \to z$, the shrinking error places the whole image $f_i(W_i)$ inside any prescribed neighborhood. Arbitrarily long finite families of their detected classes then come from that neighborhood. Nested outlines and domain silhouettes indicate containment and accumulation, not embeddings or core topology.

## 7. THE INTERIOR RETAINS ITS METRIC

The maps in (6.1) give

$$\min\{d_{W_i}(u,v), r_{W_i}(u) + r_{W_i}(v)\} \le d_X(f_i(u), f_i(v)) \le d_{W_i}(u,v). \tag{7.1}$$

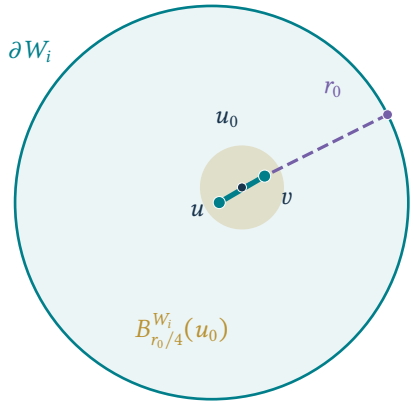
For distinct interior points the left side is positive, proving injectivity on Int $W_i$. At $u_0 \in \operatorname{Int} W_i$, put $r_0 = r_{W_i}(u_0) > 0$. If $u, v \in B^{W_i}_{r_0/4}(u_0)$, then $d_{W_i}(u, v) < r_0/2$, whereas $r_{W_i}(u) + r_{W_i}(v) > 3r_0/2$. Both inequalities in (7.1) are therefore equalities on this ball.

This does not assert that $f_i(\operatorname{Int} W_i)$ is open in $X$. The inverse on each small image is an isometry for restricted distances. A countable family of such neighborhoods covers Int $W_i$. Applying the covering definition of Hausdorff dimension to their inverse maps gives, for every subset $E \subset X$,

$$\dim_{\mathcal{H}}(f_i^{-1}(E) \cap \operatorname{Int} W_i) \le \dim_{\mathcal{H}} E. \tag{7.2}$$

This is the precise passage from a deleted subset of the metric limit to a thin subset of a smooth domain. [M, Lemma 2.5].

An interior ball is too deep for a boundary shortcut




The two competing route lengths

$$d_{W_i}(u, v) < r_0/2$$

$$r_{W_i}(u) + r_{W_i}(v) > 3r_0/2$$

$$d_X(f_i(u), f_i(v)) = d_{W_i}(u, v)$$

Restricted distances agree on the image; openness in $X$ is unnecessary.

Visual E. In the radius-$r_0/4$ ball, the direct intrinsic distance is less than $r_0/2$, while a route through the collapsed boundary costs more than $3r_0/2$. The quotient lower bound and intrinsic upper bound therefore coincide. The Euclidean disc displays these inequalities geometrically; their proof uses only boundary distance and the triangle inequality, so it applies to every retained domain.

## 8. A sphere can avoid a sufficiently thin set

A generator of $H_2(P_n; \mathbb{Z})$ has a smooth representative $\sigma : S^2 \to \operatorname{Int} P_n$. Indeed, van Kampen gives $\pi_1(P_n) = 0$: both $Q$ and the boundary sphere are simply connected. Pushing the boundary inward through its collar makes Int $P_n \hookrightarrow P_n$ a homotopy equivalence. The degree-two Hurewicz theorem then provides a sphere map, and smooth approximation preserves its homology class. An embedded sphere is not required. [M, Lemma 2.6].

Let $V$ be a smooth Riemannian $d$-manifold without boundary, let $A \subset V$ satisfy $\dim_{\mathcal{H}} A < d - 2$, and let $\sigma : S^2 \to V$ be smooth. Inside any open neighborhood $\Omega$ of its image choose finitely many compactly supported vector fields $Y_1, \dots, Y_L$ spanning the tangent spaces near that image. Composing their flows defines $\Phi_t$, for $t$ in a small cube $B \subset \mathbb{R}^L$. For

$$\mathcal{E} : S^2 \times B \to V, \qquad \mathcal{E}(z, t) = \Phi_t(\sigma(z)),$$

the parameter differential at $t = 0$ is $v \mapsto \sum_\ell v_\ell Y_\ell(\sigma(z))$, which is onto. Compactness of $S^2$ and continuity give one cube on which $\mathcal{E}$ is a submersion everywhere. All these diffeomorphisms are supported in a fixed compact subset of $\Omega$.

Submersion coordinates identify $\mathcal{E}$ locally with the projection $\mathbb{R}^d \times \mathbb{R}^{L+2-d} \to \mathbb{R}^d$. On relatively compact boxes the coordinate changes are bi-Lipschitz. For an arbitrary set $A_0 \subset \mathbb{R}^d$, the outer cover proof of [M, Lemma 2.7] gives $\dim_{\mathcal{H}}(A_0 \times \mathbb{R}^b) \le \dim_{\mathcal{H}} A_0 + b$. Indeed, over a covering ball of radius $r_\ell$, a bounded Euclidean $b$-box needs at most $Cr_\ell^{-b}$ boxes of diameter $Cr_\ell$. Summing their $(s+b)$-powers for $s > \dim_{\mathcal{H}} A_0$ and then exhausting by bounded boxes proves the estimate.

A countable submersion-chart cover consequently gives

$$\dim_{\mathcal{H}} \mathcal{E}^{-1}(A) \le \dim_{\mathcal{H}} A + L + 2 - d < L. \tag{8.1}$$

Projection onto $B$ is Lipschitz. Thus the bad parameters whose spheres meet $A$ cannot fill $B$, which has dimension $L$. Choose $t_*$ outside that projection. The isotopy $s \mapsto \Phi_{st_*}$ keeps the sphere in $\Omega$, preserves its homology class, and moves it off $A$. Smaller parameter cubes make the motion arbitrarily small. Only outer covers have been used; $A$ need not be measurable or closed. [M, Lemma 2.8].

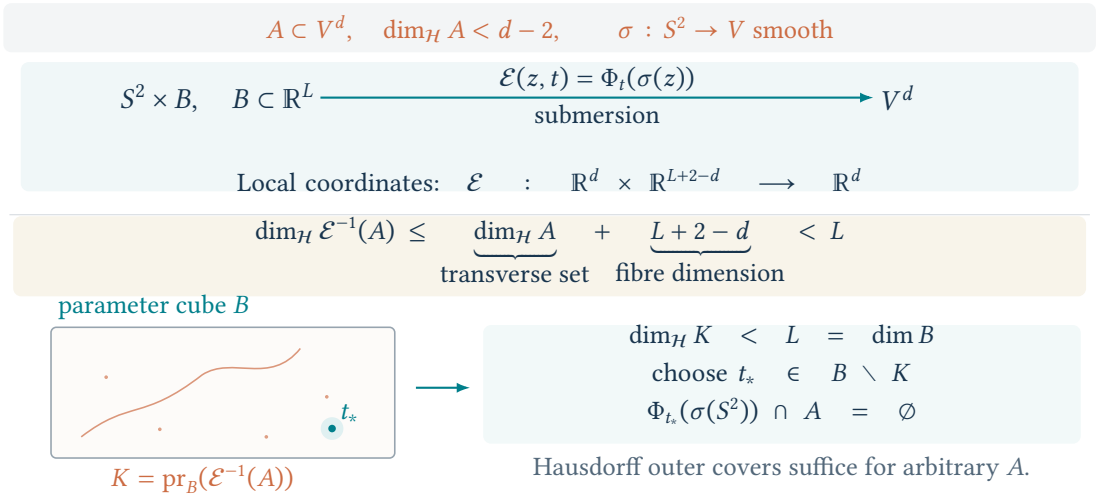



FIGURE 6. The evaluation map is a submersion on a small parameter cube. Its inverse image of the deleted set has dimension less than that cube; projecting it leaves a good parameter. The pictured bad set is schematic and is not assumed to be a submanifold or measurable.

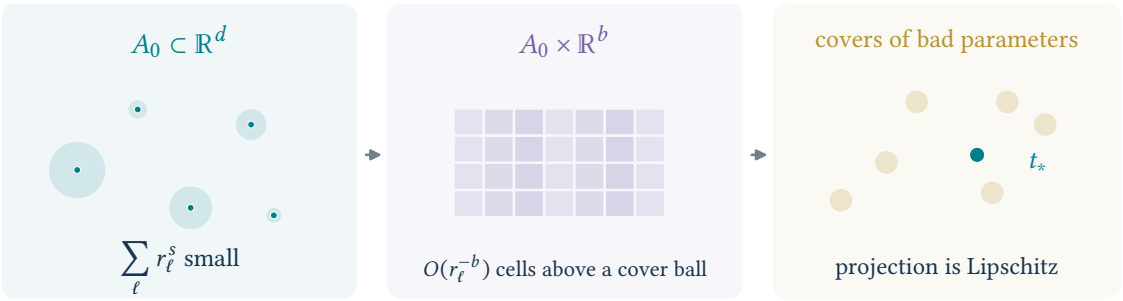
## 9. DELETION IN THE FULL PRODUCT DIMENSION

Take $X \times T^k$, put $N = n + k$, and suppose $\dim_{\mathcal{H}} E < N - 2$. Let $(x, y) \in (\iota(F) \times T^k) \setminus E$ and write a relative open neighborhood as $U = O \setminus E$, with $O$ ambient open. Choose $O_x \times B \subset O$, where $B$ is a small torus ball. Infinitely many $f_i(W_i)$ lie in $O_x$.

On the smooth $N$-manifold Int $W_i \times B$, the map $F_i = f_i \times \mathrm{Id}$ is locally distance-preserving by the product-distance formula. Therefore $A_i = F_i^{-1}(E)$ has dimension less than $N - 2$. Apply Section 8 to $z \mapsto (\sigma_i(z), y)$ in this *entire product domain*. The perturbed sphere avoids $A_i$, and its image under $F_i$ lies in $U$. Motion in torus directions is allowed; constraining it to a fixed slice would not justify the threshold $N - 2$.

If $b_i \in H_2(X \times T^k; \mathbb{Z})$ is its ambient class, homotopy invariance gives $(\mathrm{pr}_X)_* b_i = a_i$. Applying this projection to a finite relation proves independence of the $b_i$. Each comes from $H_2(U; \mathbb{Z})$, so the image has infinite rank. The proof also includes $k = 0$, when $B = T^0$ is a point. [M, Proposition 2.9].

Thinness survives the submersion chart and projection



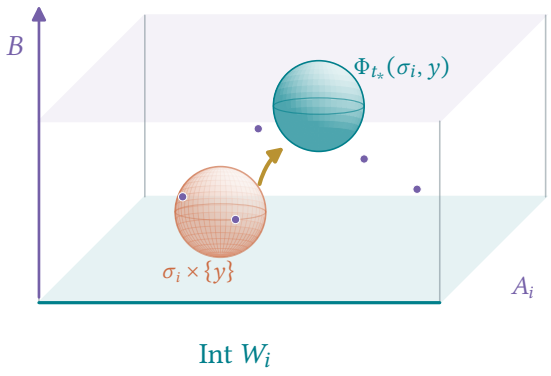

$$b = L + 2 - d, \qquad \dim_{\mathrm{H}} A < d - 2$$

$$\sum O(r_\ell^{-b})(Cr_\ell)^{s+b} \le C_s \sum r_\ell^s$$

$$\dim_{\mathrm{H}} \mathcal{E}^{-1}(A) \le \dim_{\mathrm{H}} A + b < L$$

Visual F. For $s > \dim_{\mathrm{H}} A$, over a radius-$r_\ell$ cover ball, a bounded Euclidean $b$-box requires only $O(r_\ell^{-b})$ cells. Their $(s+b)$-dimensional cover cost is controlled by the original $s$-dimensional cost. With $b = L + 2 - d$, submersion charts give a bad incidence set of dimension below $L$, and its parameter projection cannot fill $B$. The finite marks illustrate covers, not regularity of the arbitrary deleted set.

The perturbation uses the entire product dimension




$$N = n + k$$

$$\dim_{\mathrm{H}} A_i < N - 2$$

Motion in $B$ is allowed.

Support stays in the product domain.

$$(\mathrm{pr}_X)_* b_i = a_i$$

Avoid the deletion, preserve the class, then project to detect independence.

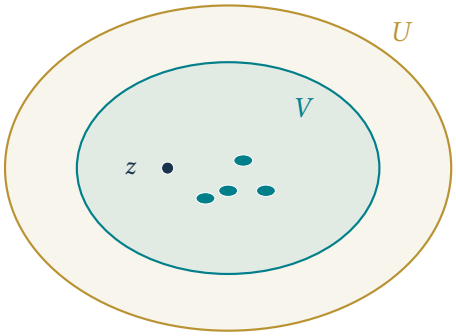
Visual G. The sphere is perturbed inside the smooth $N$-dimensional domain Int $W_i \times B$, including its torus directions. This is what permits the threshold $\dim_{\mathrm{H}} E < N - 2$. The sphere glyphs represent maps, not necessarily embedded spheres; the prism displays factors, not a metric embedding. After applying $F_i$, projection to $X$ recovers the original detected class and proves independence.

If smaller relative open neighborhoods $V \subset U$ admitted a null-homotopic inclusion, the composite

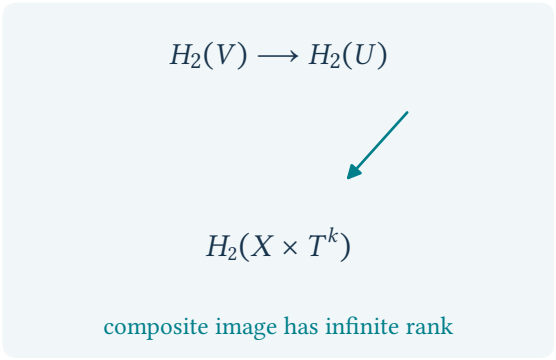
$$H_2(V;\mathbb{Z}) \longrightarrow H_2(U;\mathbb{Z}) \longrightarrow H_2(X \times T^k;\mathbb{Z})$$

would be zero. This contradicts the infinite-rank assertion for $V$. Thus local contractibility fails at each surviving distinguished point. This is stronger than showing that one neighborhood has nonzero homology. A manifold chart would contain a relative Euclidean-ball neighborhood with zero second homology, and is likewise excluded.

Persistence in every smaller neighborhood rules out contraction




Every relative open $V \ni z$




A null-homotopic inclusion $V \hookrightarrow U$ would make this composite zero.

> Visual H. The obstruction persists after passing to every smaller relative open neighborhood of the surviving distinguished point. If one inclusion $V \hookrightarrow U$ were null-homotopic, its map on second homology, and hence the displayed composite, would vanish. The construction gives infinite-rank image for that same $V$. This neighborhood-level contradiction establishes failure of local contractibility, beyond nonzero homology of a single neighborhood.

**Why some distinguished point survives every deletion.** Start with $[0,1]$. At stage $h \geq 1$, remove from the middle of each of its $2^{h-1}$ remaining intervals an interval of length $4^{-h}$. The total removed length is $\sum_{h\geq 1} 2^{h-1}4^{-h} = 1/2$. The lengths of the surviving intervals tend to zero, so the compact limit $K$ has empty interior and measure $1/2$. For a small Euclidean similarity $T$ of ratio $\rho$, the set $F = T(K^{n-2})$ fits the parameter chart and satisfies

$$\mathcal{H}^{n-2}_{|\cdot|}(F) = \rho^{n-2}2^{-(n-2)} > 0.$$

Here the subscript denotes the restricted Euclidean distance in the chart. The identification with the round parameter sphere is bi-Lipschitz on the fixed small chart, so the spherical measure is also positive and finite; its numerical value need not equal the displayed Euclidean value.

Use $n = 4, k = 0$ for $N = 4$, and $n = 5, k = N - 5$ for $N \geq 5$. Small Euclidean cubes in the torus and the bi-Lipschitz map $\iota$ give

$$0 < \mathcal{H}^{N-2}(Z_N) < \infty, \qquad Z_N = \iota(F) \times T^k \subset S_{\mathrm{top}}(X \times T^k).$$

Since $\mathcal{H}^{N-2}(E) = 0$, the set $Z_N \setminus E$ is nonempty. [M, Lemma 11.1; proof of Corollary 1.2].

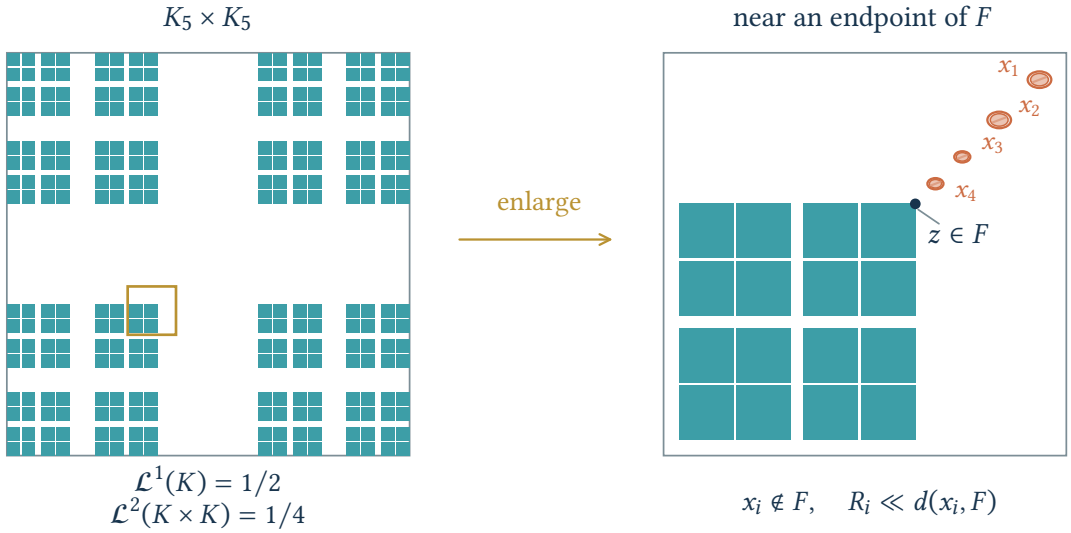



FIGURE 7. A finite stage of the positive-measure Cantor product. The limiting one-dimensional measure is $1/2$, and that of $K \times K$ is $1/4$. The insertion centers lie outside $F$ and approach it; their drawn sizes are enlarged. A similarity places the construction in the fixed small chart.

**The upper measure bound has a separate source.** The positive-measure construction gives a lower bound, not the finiteness of $\mathcal{H}^{N-2}(S_{\mathrm{top}}(X_N))$. For the latter, [M, Lemma 11.2] uses the manifold-structure theorem of Cheeger–Jiang–Naber and its proof. Rebasing the realizing sequence preserves a uniform unit-ball noncollapse bound, because its diameters are bounded. On each unit ball the cited result supplies an exceptional set of finite $(N-2)$-measure and *ambient open* manifold charts outside it. Hence the nonmanifold locus in that ball is contained in the exceptional set. A finite cover of the compact limit and subadditivity give the required upper bound.

A statement about the relative topology of the complement alone would not supply that inclusion. Nor does the construction identify all of $S_{\mathrm{top}}(X_N)$ with $Z_N$. The middle-third variants in [M, §11.1] are different from the positive-measure choice: their distinguished sets have the dimensions computed there, and no uniqueness of tangents along a general $\iota(F)$ is asserted.

## 10. The circle profile has two different scales

The round metric around the equatorial $S^{n-2} \subset S^n$ is

$$dt^2 + \cos^2 t \, g_{S^{n-2}} + \sin^2 t \, d\theta^2, \quad 0 < t < \pi/2, \quad \theta \in \mathbb{R}/(2\pi\mathbb{Z}).$$

At $t = 0$ the circle collapses. Replacing $\sin t$ by $q \sin t$ gives transverse cone angle $2\pi q$, whereas a smooth normal plane requires initial slope one. Fix $0 < q < 1$, put $c = -\log q$, and define

$$e_\delta = e^{-c/\delta}, \qquad b_\delta(t) = \sin t \left(1 + \frac{\sin^2 t}{e_\delta^2}\right)^{-\delta/2}. \tag{10.1}$$

For every fixed $\delta > 0$, this is smooth and odd near zero, with $b_\delta(t) = t + O(t^3)$. Thus the normal metric has a smooth pole. For fixed $t > 0$, the identity

$$P := \log(b_\delta / \sin t) = -c - \frac{\delta}{2} \log(\sin^2 t + e_\delta^2)$$

gives $b_\delta(t) \to q \sin t$ as $\delta \downarrow 0$. [M, §4].

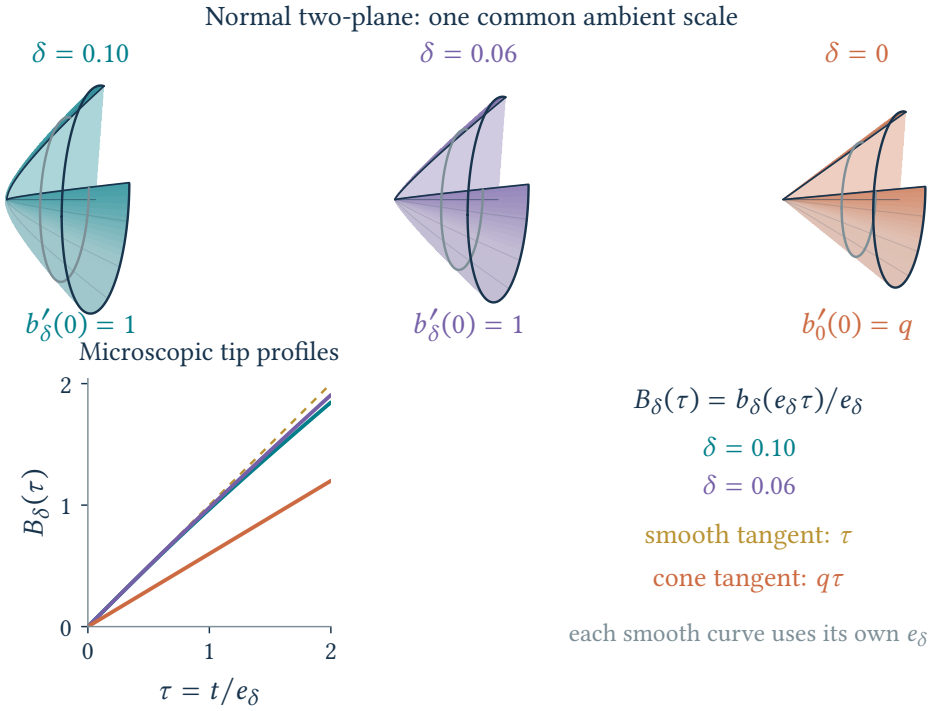



FIGURE 8. The normal surfaces use the exact profile $b_\delta$, with $q = 0.60$ and $0 \le t \le 0.075$. Their heights satisfy $z'_\delta = \sqrt{1-(b'_\delta)^2}$. The upper panels share ambient units; each lower profile is magnified by its own $e_\delta$. Smooth microscopic slopes remain one although the limiting slope is $q$. These numerical parameters illustrate the formula, not the small-$q$ threshold for the complete construction.

Set

$$Z = \frac{\sin^2 t}{\sin^2 t + e_\delta^2}, \qquad L = 1 - \tfrac{1}{2}\log(\sin^2 t + e_\delta^2).$$

Then $b_t/b = (1-\delta Z)\cot t$. The exponentially small $e_\delta$ locates the change of opening, whereas $\delta$ controls the logarithmic derivative defect. The logarithmic coordinate $x = -\log(\sin t)$ places the transition at $x = c/\delta$. Relative convergence cannot be uniform down to the pole: $b_\delta(t)/\sin t \to 1$ as $t \downarrow 0$ for each positive $\delta$, but its limit for fixed positive $t$ is $q$.

The derivatives that enter curvature are

$$-\frac{b_{tt}}{b} = 1 - \delta Z + \delta Z[3-(2+\delta)Z]\cot^2 t, \tag{10.2}$$

$$P_{\delta\delta} = -\frac{2c^2}{\delta^3} Z(1-Z) \le 0, \tag{10.3}$$

$$P_{\delta t} = Z\left(-1 + \frac{2c}{\delta}(1-Z)\right)\cot t. \tag{10.4}$$

The estimates in [M] first fix a domain with $\bar t < 1/10$, $\bar\delta < \min\{1/4, c/4\}$, and $\sin^2 \bar t + e_{\bar\delta}^2 < e^{-2}$. On that domain $L \ge 2$, $0 \le P_\delta \le ZL$ and $|P_{\delta t}| \le 3ZL\cot t$. The curvature may be large near the shrinking core; no uniform upper curvature bound is asserted. [M, Lemma 4.1].

## 11. Compensating for a parameter that varies along the sphere

The parameter must vary on $(S^m, \gamma)$, where $m = n-2 \in \{2,3\}$. An ordinary doubly warped-product calculation is no longer sufficient. At fixed $t$,

$$b^{-1}\,\mathrm{Hess}_\gamma\, b = P_\delta\,\mathrm{Hess}_\gamma\,\delta + (P_{\delta\delta} + P_\delta^2)\,d\delta\otimes d\delta.$$

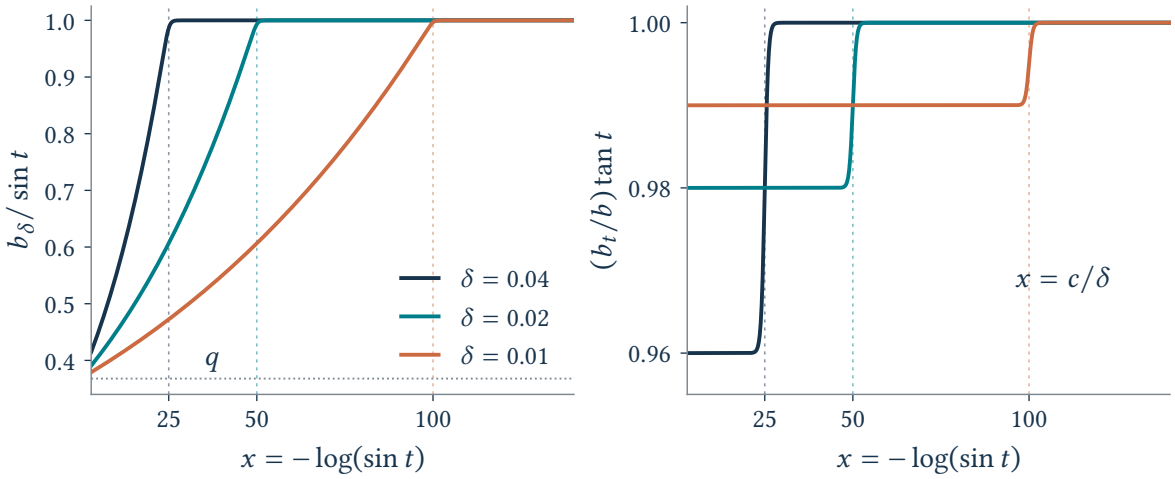



FIGURE 9. Exact profiles for $c = 1$, $q = e^{-1}$ and $\delta = 0.04, 0.02, 0.01$. Moving right approaches the circle collapse. The core is exponentially small, but the logarithmic derivative defect is $\delta Z$. The numerical choices concern the profile only.

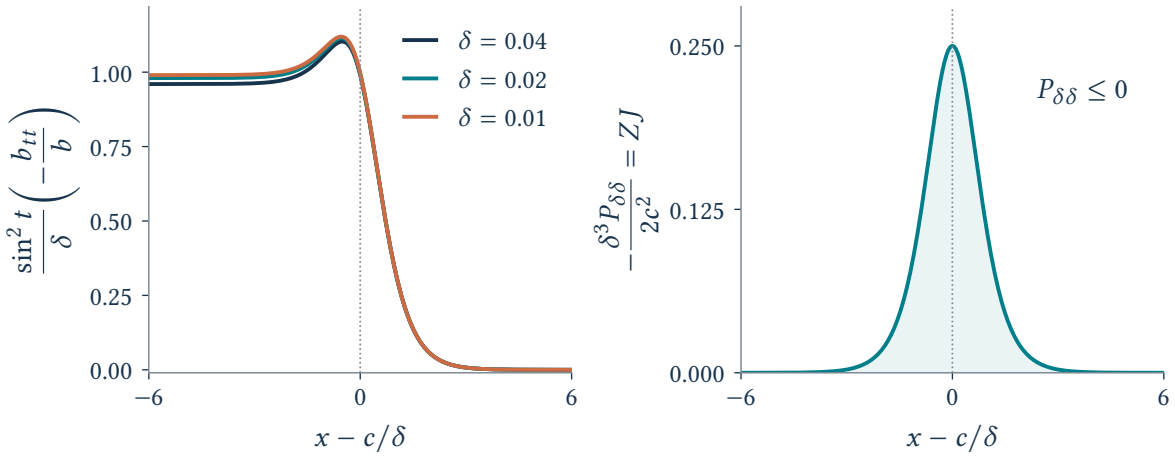



FIGURE 10. The radial curvature term is rescaled by $\sin^2 t/\delta$, and the parameter-concavity term satisfies $-\delta^3 P_{\delta\delta}/(2c^2) = Z(1-Z)$. The panels display exact derivative expressions rather than uniform upper curvature estimates.

This Hessian enters Ricci with a minus sign. An upper bound is therefore needed, and $P_{\delta\delta} \le 0$ is favorable. Taking absolute values of all terms would discard this sign. The full Ricci formulas also account for the change from $\gamma$ to the slice metric $a^2\gamma$. [M, Lemma 5.1; Proposition 5.2].

Define

$$\Psi_\delta(t) = \int_0^t (t-u)L_\delta(u)^2\,du, \quad U = A(x)\Psi_{\delta(x)}(t), \quad a(t,x) = \cos t\, e^{-U},$$

and use $G = dt^2 + a^2\gamma + b_{\delta(x)}^2 d\theta^2$. The double integral has the prescribed second derivative $U_{tt} = AL^2$, so

$$-a_{tt}/a = 1 + AL^2 - 2\tan t\, U_t - U_t^2. \tag{11.1}$$

The positive term has the logarithmic weight of the spatial errors. The remaining terms are controlled by the preliminary radial interval and small $C^2$-norms. For positive smooth parameters, $\Psi_\delta(0) = \Psi_\delta'(0) = 0$ and $\Psi_\delta$ is even in the radial variable. Thus $a(0,x) = 1$ and the induced metric on the

collapsed stratum remains $\gamma$. Large second derivatives can supply curvature while the metric itself changes little.

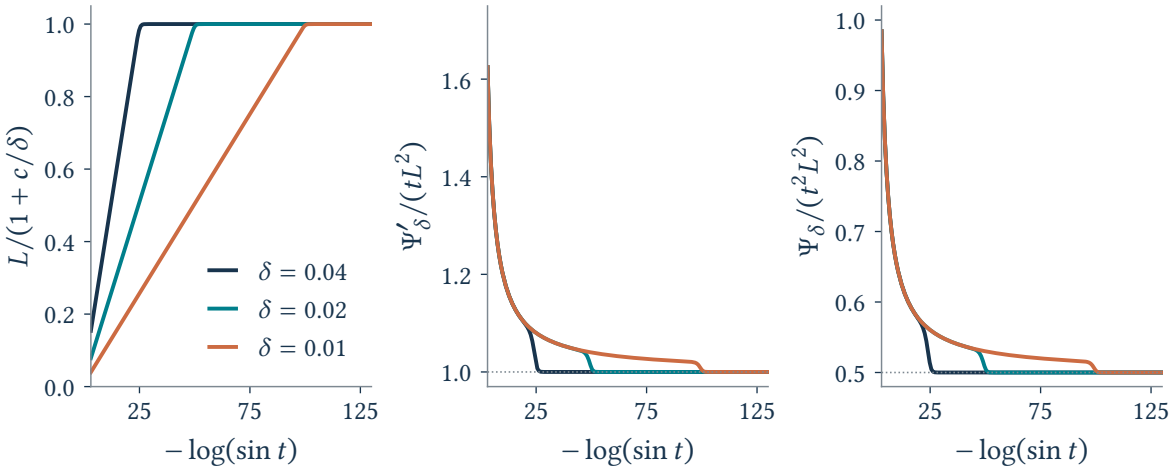
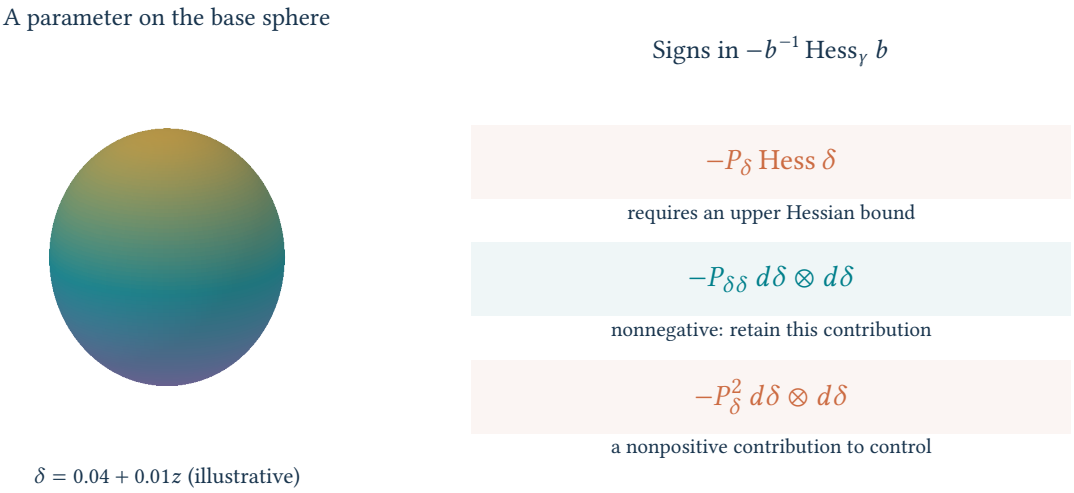



FIGURE 11. The compensation kernel satisfies $\Psi_{tt} = L^2$, while its first integral and double integral are controlled by $tL^2$ and $t^2L^2$. The endpoint limits for each fixed positive $\delta$ are not, by themselves, uniform estimates in $\delta$.

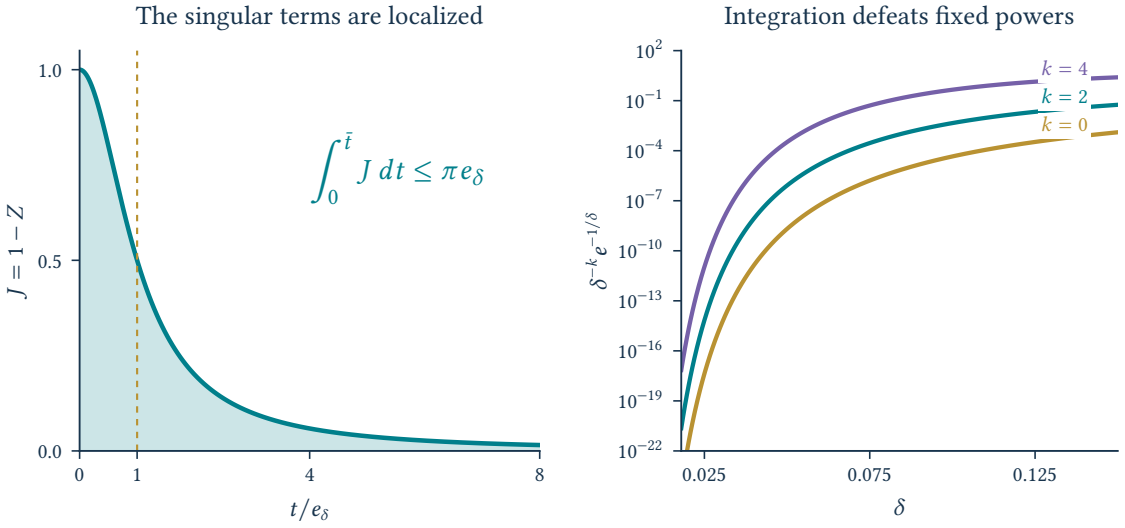



VISUAL I. A parameter varying along the base sphere creates spatial Hessian terms. Their signs matter: $P_{\delta\delta} \le 0$ contributes a nonnegative tensor after the minus sign in Ricci, while the remaining terms require compensation. The sphere displays an illustrative scalar field, not the constructed flat parameter. The construction instead chooses $A$ from its derivative bounds and uses $U_{tt} = AL^2$ to supply the required positive contribution.

**Uniformity does not come from a curvature maximum.** The kernel estimate gives

$$tL^2 \le \Psi_t \le 4tL^2, \qquad 0 \le \Psi \le 4t^2L^2.$$

The singular parameter terms are localized by $J = 1 - Z$. In particular, $\int_0^{\bar{t}} J\,dt \le \pi e^{-c/\delta}$; derivatives of $L^2$ are bounded by fixed powers of $\delta^{-1}$ times $J$. After integration these powers are accompanied by $e^{-c/\delta}$. This controls the first two parameter derivatives of $\Psi$ and $\Psi_t$ uniformly, together with the weighted quantities needed in the Ricci estimate. It does not bound every second radial derivative uniformly. [M, Lemma 4.2].




$c = 1$, $\delta = 0.1$ in the left panel; exact $\sin t$ is used.

VISUAL J. The integral mass of $J = 1 - Z$ is controlled by the scale $e_\delta$ near the pole. The left graph uses the exact sine profile, with $c = 1$ and $\delta = 0.1$. Integrating supplies the exponentially small factor; the right graph shows how it dominates fixed inverse powers of $\delta$. This mechanism controls the integrated parameter derivatives, without asserting a uniform upper bound for all radial second derivatives.

For constants fixed in [M, Proposition 5.2], the required hypothesis is

$$0 < \delta \le \delta_0, \qquad A \ge M\left(|\operatorname{Hess}_\gamma \delta| + \frac{|d\delta|_\gamma^2}{\delta}\right), \qquad \|A\|_{C^2} + \|\delta\|_{C^2} \le \eta. \tag{11.2}$$

It yields $\operatorname{Ric}_G \ge G$ on $0 < t \le t_0$. The constants depend on the fixed $q$ and preceding data, but not on $\min \delta$. The function $A$ is nonnegative, not necessarily positive: it will vanish on the constant-parameter balls used for insertion.

**Choose the regions before the flat parameters.** After fixing the core's convexity constant $\alpha$, choose

$$0 < \beta < \min\{\alpha/4, 1/8\}, \qquad 0 < q < \beta^{n-1}/1000.$$

The background and annular constants are then fixed, followed by a sufficiently small $R_*$. Only afterwards choose centers $x_i \notin F$ and radii $R_i$ so that the enlarged balls $B_{10^6 R_i}(x_i)$ are disjoint and avoid $F$, and

$$R_i < 2^{-i} \min\{R_*, \operatorname{dist}_\gamma(x_i, F)^3\}. \tag{11.3}$$

Finite increasingly fine nets in $F$ supply associated $z_i \in F$ approaching every point of $F$. Empty interior permits the centers to be chosen in the complement; shrinking each new radius enforces the finite separation conditions.

[M, Lemma 7.2] constructs a smooth positive scale function $r(x)$ off $F$ with

$$r(x) \le C \operatorname{dist}_\gamma(x, F), \quad r = R_i \text{ on } B_{60R_i}(x_i), \quad |\nabla^h r| \le C_h r^{1-h}.$$

Its construction smooths a Lipschitz scale on locally finite bounded-overlap balls and corrects it on the disjoint prescribed plateaux. One does not need a smooth distance function to $F$. The prescribed plateaux and derivative estimates, rather than just positivity of $r$, are the properties used next.

**Flatness and compensation from one function.** Put $f = e^{-1/r^2}$ off $F$, extended by zero on $F$. The scale estimates imply that each derivative is a polynomially weighted exponential; hence all derivatives vanish on $F$. Choose $\xi \in C^\infty(S^{n-2} \setminus F, [0,1])$ that is zero on $B_{20R_i}(x_i)$, one outside the union of $B_{40R_i}(x_i)$, with its transitions inside the constant-scale regions and $|\nabla^m \xi| \le C_m R_i^{-m}$ on the $i$-th transition. The cutoff itself need not extend continuously over $F$, but on each transition $f = e^{-1/R_i^2}$, so $|\nabla^m(\xi f^2)| \le C_m R_i^{-m} e^{-2/R_i^2}$. Separation from $F$ bounds each derivative by $C_{m,\ell}\,\mathrm{dist}_\gamma(\cdot, F)^\ell$ for every $\ell \ge 1$; outside the transitions the product is zero or $f^2$, so extending $\xi f^2$ by zero gives a smooth function flat on $F$. Set

$$\delta = \lambda f^2, \qquad A = M_1 \lambda \big(\xi f^2 + |\mathrm{Hess}_\gamma f|^2 + 6|df|_\gamma^2\big). \tag{11.4}$$

Both functions are smooth and flat on $F$, and $\delta^{-1}(0) = F$. Direct differentiation gives

$$\mathrm{Hess}\,\delta = 2\lambda(f\,\mathrm{Hess}\,f + df \otimes df), \qquad \frac{|d\delta|^2}{\delta} = 4\lambda|df|^2$$

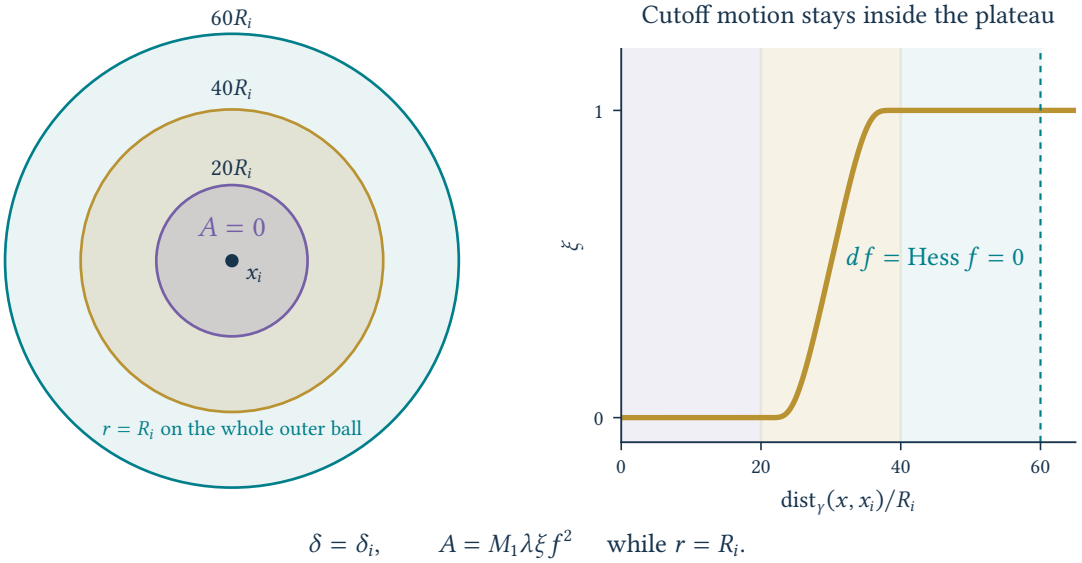
off $F$; the latter quotient extends smoothly across $F$. Thus

$$|\mathrm{Hess}\,\delta| + \frac{|d\delta|^2}{\delta} \le \lambda(f^2 + |\mathrm{Hess}\,f|^2 + 6|df|^2).$$

Where $\xi \ne 1$, the function $f$ is constant, so both derivatives on the left vanish. Consequently the $f^2$ on the right can be replaced by $\xi f^2$. On each insertion ball,

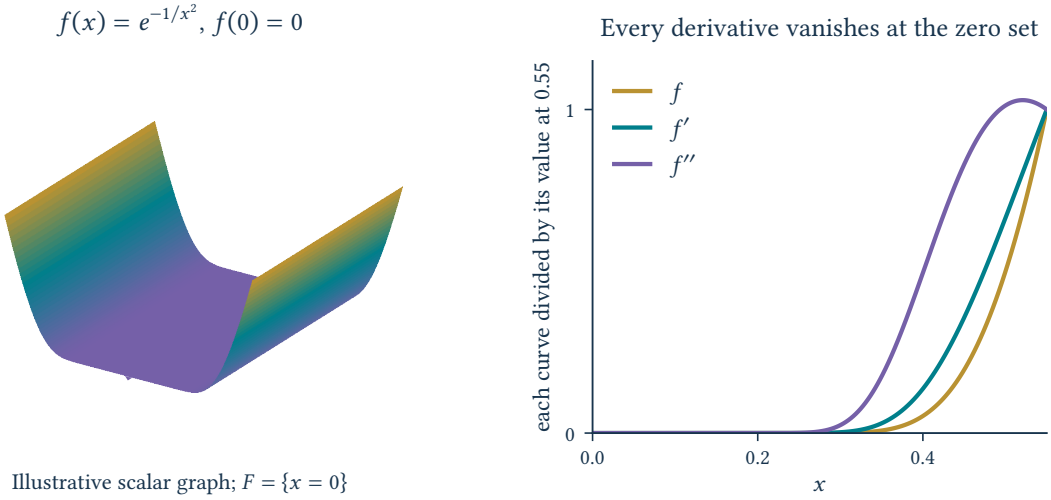
$$\delta = \delta_i := \lambda e^{-2/R_i^2}, \qquad A = 0.$$

Choose $M_1$ large enough for the uniform floor estimate below, and then $\lambda > 0$ small enough for the required $C^2$-bounds. [M, Lemma 7.3].




$$\delta = \delta_i, \qquad A = M_1 \lambda \xi f^2 \quad \text{while } r = R_i.$$

Visual K. The three prescribed radii separate the insertion region, the cutoff transition, and the larger constant-scale region. Since $r = R_i$ throughout the latter, $f$ is constant wherever $\xi$ changes. On the inner ball both derivatives of $f$ and $\xi$ vanish, giving exactly $\delta = \delta_i$ and $A = 0$. The concentric circles are a chart schematic of metric balls, not an isometric embedding.




$$\delta = \lambda f^2 \quad \Longrightarrow \quad |d\delta|^2/\delta = 4\lambda |df|^2$$

VISUAL L. The scalar model $f(x) = e^{-1/x^2}$, extended by zero, illustrates flatness along a zero set. The derivative curves are individually normalized only for visibility. In the construction the controlled scale function replaces $|x|$, with the same exponential domination of polynomial derivative factors. Writing $\delta = \lambda f^2$ also removes the apparent singularity in $|d\delta|^2/\delta$, which extends smoothly across $F$.

**Positive floors preserve earlier insertions exactly.** Choose a smooth step $H_0$ from zero to one, constant outside $(0, 1)$, and put

$$\vartheta(u) = 1 + (u-1)H_0(u-1), \qquad \delta_\varepsilon = \varepsilon\,\vartheta(\delta/\varepsilon).$$

Then $\delta_\varepsilon = \varepsilon$ where $\delta \le \varepsilon$, $\delta_\varepsilon = \delta$ where $\delta \ge 2\varepsilon$, and $\delta_\varepsilon \ge \frac{1}{2}\max\{\delta, \varepsilon\}$. On the transition,

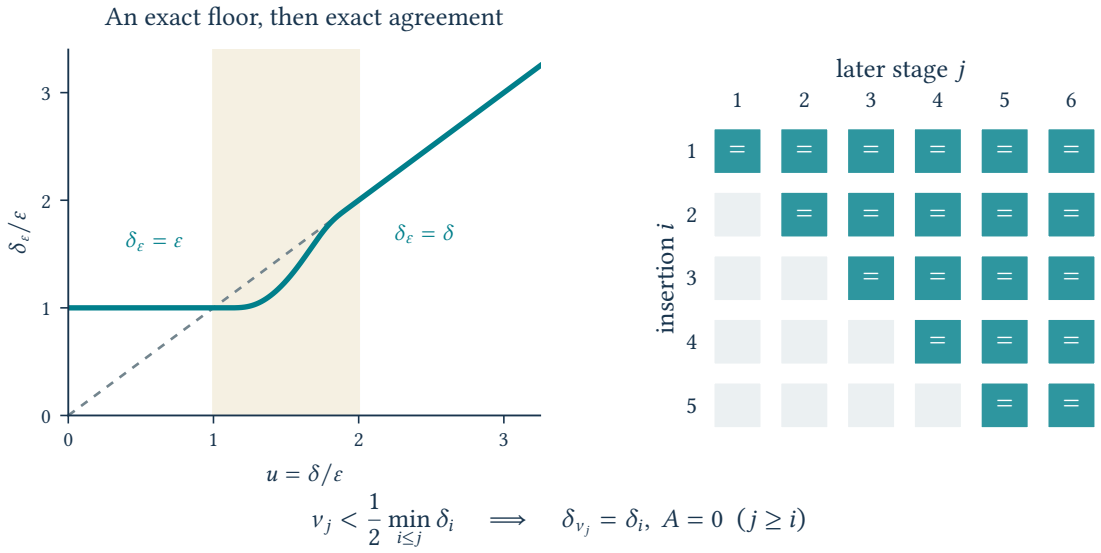
$$\operatorname{Hess}\delta_\varepsilon = \vartheta' \operatorname{Hess}\delta + \varepsilon^{-1}\vartheta''\, d\delta \otimes d\delta, \qquad d\delta_\varepsilon = \vartheta'\, d\delta.$$

Since $\varepsilon^{-1} \le 2/\delta$ there, the defect in (11.2) for $\delta_\varepsilon$ is bounded by a fixed multiple of the defect for $\delta$. The constant depends on $\vartheta$, not on $\varepsilon$. Choosing $M_1$ first and $\lambda, \varepsilon$ sufficiently small gives the same compensation and norm bounds with the *same* $A$.

Choose a decreasing sequence $v_j \downarrow 0$ satisfying

$$v_j < \min\left\{\varepsilon_{\mathrm{adm}}, 2^{-j}, \tfrac{1}{2}\min_{i\le j}\delta_i, \tfrac{1}{2}v_{j-1}\right\} \tag{11.5}$$

(the last entry is omitted at the first index). Then, on the $i$-th prescribed ball, $\delta_{v_j} = \delta_i$ and $A = 0$ for *every* $j \ge i$. There is no approximation error in this assertion. In Visual M, fix an insertion row and follow the later-stage columns: the entries record equality of the parameters at every subsequent stage. These $v_j$ are positive parameter floors, not widths of the collar mollification.




The positive floors change; every previously inserted region stays exactly fixed.

Visual M. The smooth floor equals $\varepsilon$ below the threshold and equals the original parameter above $2\varepsilon$. The highlighted transition is the only place where new parameter derivatives occur. The triangular array records exact stabilization: at stage $j$, all insertion regions indexed by $i \le j$ retain $\delta_i$ and $A = 0$. These positive floors are unrelated to the later collar-mollification widths.

The same small $\lambda$ can ensure

$$\delta_i \le R_i^4/100, \qquad \delta_i|\log\sin R_i| \le \log 2 \tag{11.6}$$

for every $i$: the functions $e^{-2/R^2}/R^4$ and $e^{-2/R^2}|\log\sin R|$ are bounded on $(0, R_*]$ and tend to zero at zero. Thus the already fixed $R_*$ and insertion radii do not have to be changed later. The order of choices is *core and angle, uniform estimates, radii and scale function, compensator constant, small amplitude, positive floors.* [M, §7].

**From the local formula to a closed sphere.** Choose fixed radial cutoffs $\chi_a, \chi_b$, with $\chi_a = 1$ on $[0, t_0]$, zero for $t \ge 2t_0$, and $\chi_b = 1$ on $[0, 2t_0]$, zero for $t \ge 3t_0$. Replace the local coefficients by

$$U = \chi_a A \Psi_\delta, \quad a = \cos t\, e^{-U}, \qquad b = \sin t \exp\{\chi_b P + (1-\chi_b)\log q\}.$$

On the fixed transition strip the derivatives are controlled in $C^2$, and the metric is a small $C^2$-perturbation of $G_q = dt^2 + \cos^2 t\,\gamma + q^2 \sin^2 t\, d\theta^2$, whose Ricci coefficient on its regular part is $n-1$. A further fixed smallness choice retains $\mathrm{Ric} \ge G$. Near $t = 0$, positive smooth floors give the smooth odd circle opening and even other coefficients. Near $\pi/2$ the sphere factor collapses with its ordinary smooth polar behavior. Thus each positive-floor metric is a smooth metric on $S^n$.

[M, Lemma 5.3] gives, with $c_0 = 19/20$,

$$c_0^2 G_q \le G[\delta_{\nu_j}, A] \le g_{S^n}, \qquad c_0 q\, d_{S^n} \le d_{G[\delta_{\nu_j}, A]} \le d_{S^n}. \tag{11.7}$$

The same fixed cutoff formula is used at every stage. Exact agreement of the parameters on an insertion region therefore gives exact agreement of its background tensors and boundary data.

## 12. The mixed Ricci components must also be controlled

Positive diagonal entries do not prove positivity of a quadratic form. Write $h = a^2\gamma$ and $k = a_t/a$. In the radial, spherical, and unit circle directions, [M, Lemma 5.1] gives

$$\begin{aligned} D &= -ma_{tt}/a - b_{tt}/b, \\ B_i &= -(m-1)\partial_i k - (b_{ti} - kb_i)/b, \\ T &= \mathrm{Ric}_h - \big(a_{tt}/a + (m-1)k^2 + kb_t/b\big)h - b^{-1}\,\mathrm{Hess}_h\, b, \\ V &= -b_{tt}/b - mkb_t/b - b^{-1}\Delta_h b. \end{aligned} \tag{12.1}$$

Here $D = \mathrm{Ric}(\partial_t, \partial_t)$, $B$ is the radial–spherical covector, $T$ the spherical block, and $V = \mathrm{Ric}(b^{-1}\partial_\theta, b^{-1}\partial_\theta)$. All remaining circle-mixed entries vanish. The spatial Hessian, radial–spatial derivatives, and change of slice metric are all present in these identities.

Within this calculation only, write $Q = AL^2$ and $X = \delta Z \cot^2 t$; these are nonnegative scalar quantities, not the ambient space $X$. The compensation supplies $Q$; the radial circle curvature supplies $X$. Under (11.2), [M, Proposition 5.2] obtains

$$\begin{aligned} D &\geq \tfrac{5}{2} + Q + \tfrac{1}{2}X, \quad & T &\geq (\tfrac{5}{2} + \tfrac{1}{2}Q)h, \\ |B|_h^2 &\leq \tfrac{1}{4} + \tfrac{1}{64}QX, & V &\geq 2. \end{aligned} \tag{12.2}$$

The mixed bound is allowed to grow with $QX$, because the product of the diagonal lower bounds grows with the same quantity and a larger coefficient.

To prove $\mathrm{Ric}_G \geq G$, subtract the metric before applying the Schur test. If $\lambda_T$ is the least $h$-eigenvalue of $T$, then

$$(D-1)(\lambda_T - 1) - |B|_h^2 \geq 2 + \tfrac{9}{4}Q + \tfrac{1}{2}Q^2 + \tfrac{3}{4}X + \tfrac{15}{64}QX > 0. \tag{12.3}$$

The quadratic form on $s\partial_t + Y$ is

$$(D-1)\left(s + \frac{B(Y)}{D-1}\right)^2 + (T-h)(Y,Y) - \frac{B(Y)^2}{D-1}.$$

Cauchy–Schwarz and (12.3) control the last two terms for every spherical vector $Y$. The unit circle coefficient after subtraction is $V - 1 \geq 1$. This proves the full tensor inequality, not merely positivity on the coordinate axes. Figure 12 displays the mixed-cost bounds as $Q, X$ vary; Visual N shows how completing the square uses those bounds to control the quadratic form.

## 13. A fixed volume form makes slow interpolation possible

**The cutting sphere.** On the protected region about $x_i$, all backgrounds with $j \geq i$ have $A = 0$, $\delta = \delta_i$. Polar coordinates in the base hemisphere give

$$\sin t = \sin r \sin u, \qquad dt^2 + \cos^2 t\,\gamma = dr^2 + \sin^2 r\big(du^2 + \cos^2 u\, g_{S^{n-3}}\big).$$

In this region $|dr| = 1$, and curves with fixed angular coordinates have length equal to their change in $r$. A curve trying to leave the protected region first pays at least $10c_0R_i > R_i$, by the Lipschitz projection to the base hemisphere. Therefore it cannot create a shorter route to a point with $r \leq R_i$. The cutting sphere $r = R_i$ is the ambient distance sphere of radius $R_i$, not just an arbitrarily named coordinate sphere. [M, Lemma 8.1].

Its induced metric is $\sin^2 R_i\,\widetilde{h}_i$, where, with $e_i = e_{\delta_i}/\sin R_i$,

$$\widetilde{h}_i = du^2 + \cos^2 u\, g_{S^{n-3}} + \sin^2 u\left(1 + \frac{\sin^2 u}{e_i^2}\right)^{-\delta_i} d\theta^2. \tag{13.1}$$

This is the boundary metric that the annulus must match.

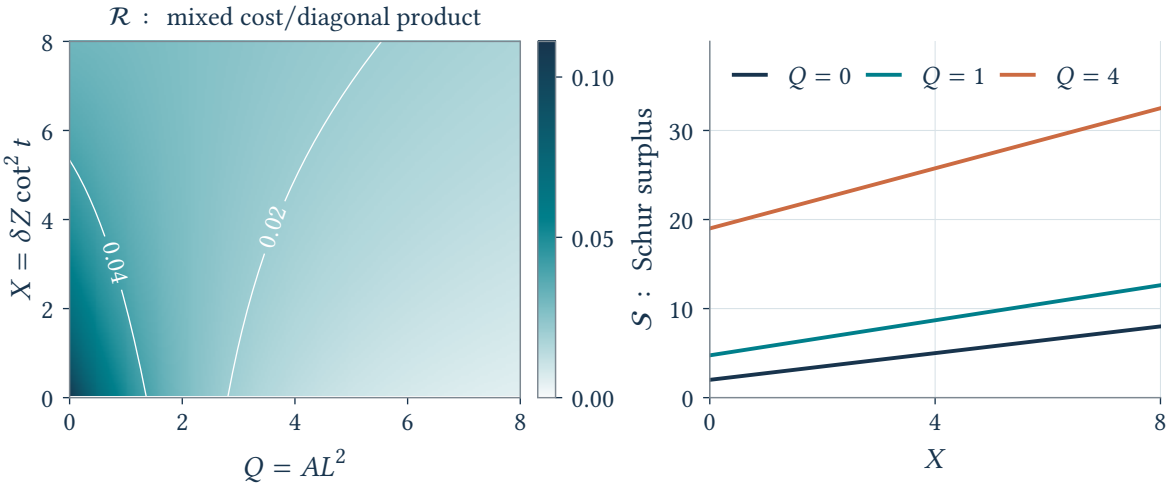
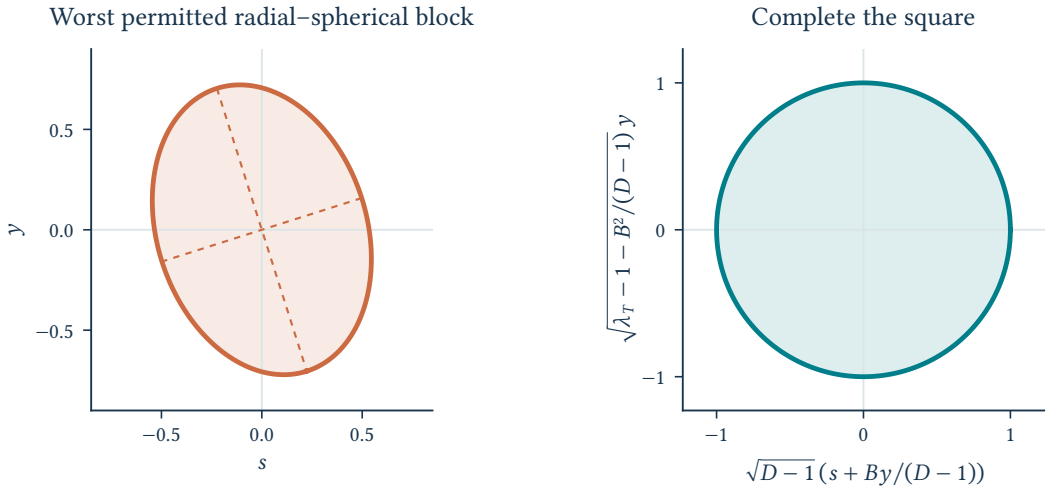



FIGURE 12. The graph displays the algebraic bounds for $\mathrm{Ric}_G - G$. With $Q, X \geq 0$, the mixed-cost ratio $(1/4 + QX/64)/[(3/2 + Q + X/2)(3/2 + Q/2)]$ is at most $1/9$. The curves are sufficient bounds, not computed Ricci eigenvalues or a spatial field of independently prescribed $Q, X$.



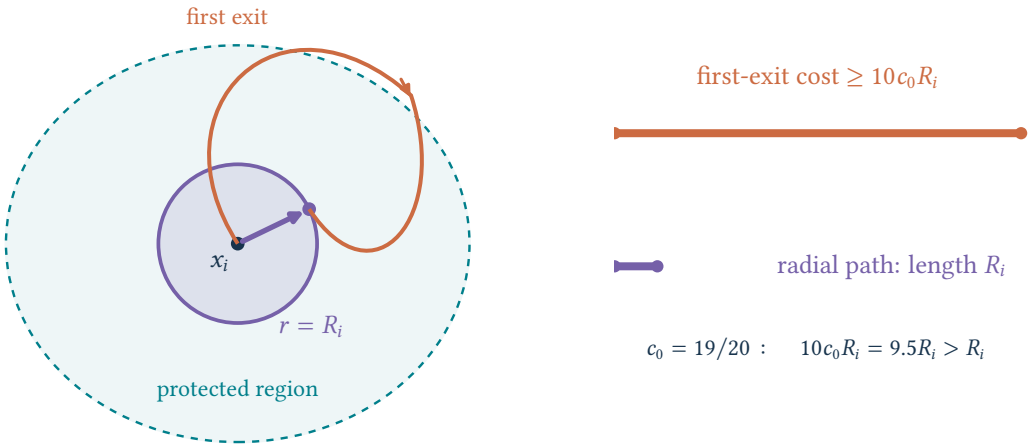

$$(D-1)(\lambda_T - 1) - |B|^2 \geq 2 + \frac{9}{4}Q + \frac{1}{2}Q^2 + \frac{3}{4}X + \frac{15}{64}QX > 0$$

$Q = 1$, $X = 2$ illustrate the bound; the circle direction separately has $V - 1 \geq 1$.

VISUAL N. The Schur test is applied to $\mathrm{Ric}_G - G$. The left ellipse is the unit level set of the least diagonal bounds and largest allowed mixed term at the illustrative values $Q = 1$, $X = 2$. Completing the square transforms this positive quadratic form into the circle on the right. These are algebraic bounds in one spherical direction, not computed Ricci eigenvalues of an independently prescribed spatial field.

Keep $e_i$ fixed and replace the exponent $\delta_i$ in (13.1) by $p \in [0, \delta_i]$; call the result $h(p)$. Then $h(0)$ is round and $h(\delta_i) = \widetilde{h}_i$. The link calculation gives $\mathrm{Ric}_{h(p)} \geq (n-2)(1-\delta_i)h(p)$. Its volume decreases with $p$. Put

$$V_i = (1-\delta_i)^{(n-1)/2}\,\mathrm{Vol}(\widetilde{h}_i), \qquad \widehat{h}(p) = \left(\frac{V_i}{\mathrm{Vol}(h(p))}\right)^{2/(n-1)} h(p). \tag{13.2}$$




$|dr| = 1$ inside the protected region; a first exit already costs too much.
Consequently, the cutting sphere $r = R_i$ is the ambient distance sphere.

VISUAL O. Radial segments reach the cutting sphere with length $R_i$. Inside the protected region, $|dr| = 1$ prevents a shorter path. A path leaving that region first pays at least $10c_0R_i > R_i$, as detected by the projection to the base hemisphere. Thus outside shortcuts cannot change the radius. The left region diagram is schematic; the right length bars use the exact comparison ratio.

The metric scaling factor is at most $1 - \delta_i$, so $\mathrm{Ric}_{\widehat{h}(p)} \geq (n-2)\widehat{h}(p)$, and every slice has volume $V_i$. The endpoints are $\tau_i^2 g_{S^{n-1}}$ and $h_i = (1-\delta_i)\widetilde{h}_i$, with $\tau_i^{n-1} = V_i/\sigma_{n-1}$. Holding $e_i$ fixed is essential: this path is not obtained by replacing $\delta_i$ everywhere by $p$ in the original two-scale profile. [M, Lemma 8.2].

**Fix the form, not just its integral.** Reparametrize the path to be stationary near both ends, and let $\omega_u = dv_{\widehat{h}_u}$. On the oriented closed sphere choose one background metric, with Green operator $\mathscr{G}$ and codifferential $d^*$. Since $\int \partial_u \omega_u = 0$, the top-degree Hodge identity gives

$$\eta_u = d^*\mathscr{G}(\partial_u \omega_u), \qquad d\eta_u = \partial_u \omega_u.$$

Define $Y_u$ by $\iota_{Y_u}\omega_u = -\eta_u$, and let $\phi_u$ be its flow starting at the identity at the left endpoint. Compactness gives the flow for the whole finite parameter interval. Cartan's formula yields

$$\frac{d}{du}(\phi_u^*\omega_u) = \phi_u^*(\partial_u\omega_u + d\iota_{Y_u}\omega_u) = 0.$$

Thus $g_u = \phi_u^*\widehat{h}_u$ has one fixed volume form and the same Ricci lower bound. Where the path is stationary, $Y_u = 0$, so the isotopy is constant; at the right endpoint that constant need not be the identity. The terminal map $\phi_i$ is retained as the attaching map from the annular boundary to the background boundary. [M, Lemma 6.1; Proposition 8.3].

For a smooth radial family $g(r)$ with fixed volume form,

$$\mathrm{tr}_g\, g' = 0, \qquad \mathrm{tr}_g\, g'' = |g'|_g^2. \tag{13.3}$$

The second identity differentiates the first and uses $(g^{-1})' = -g^{-1}g'g^{-1}$; its trace is not zero. For $dr^2 + w^2 g(r)$, the radial Ricci component is

$$\mathrm{Ric}_{rr} = -(n-1)w''/w - \tfrac{1}{4}|g'|_g^2. \tag{13.4}$$

Fixing the volume form removes trace terms but does not remove the curvature cost of changing the shape of a slice.

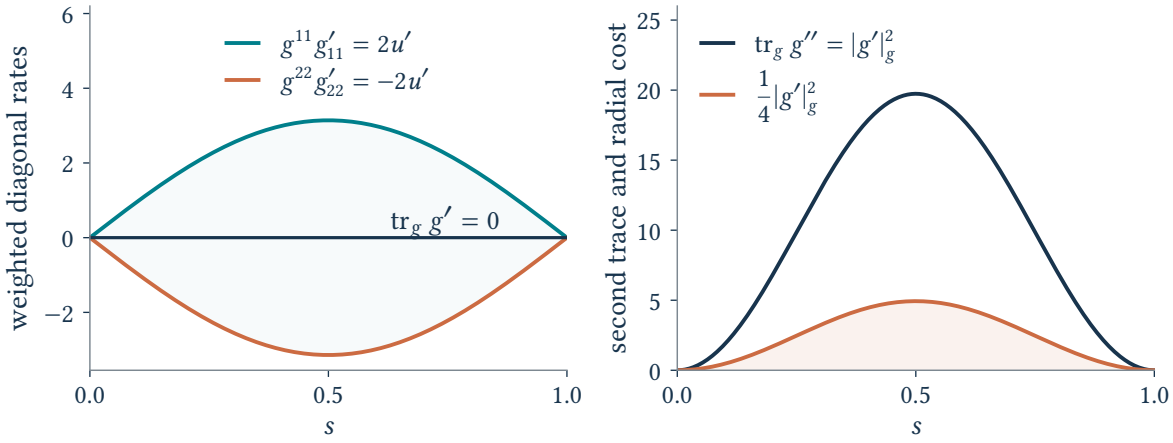



FIGURE 13. The pointwise model $g(s) = \mathrm{diag}(e^{2u(s)}, e^{-2u(s)}, 1)$, $u(s) = (1 - \cos \pi s)/2$, has determinant one and $\mathrm{tr}_g\, g'' = |g'|^2_g = 8(u')^2$. It illustrates the trace identities, not a positive-Ricci path on a sphere.

**Slow the path before inserting the radial parameter.** For each fixed path, its first and second parameter derivatives and the spatial covariant derivative of its first parameter derivative are bounded. Let $\rho_0$ be a fixed smooth nondecreasing map from $\mathbb{R}$ to $[-2, 2]$, with $\rho_0(v) = -2$ for $v \le -1$ and $\rho_0(v) = 2$ for $v \ge 1$. Choose $L_*$ large enough that $k_v = g_{\rho_0(v/L_*)}$ satisfies

$$|\partial_v k_v|_{k_v} + |\partial_v^2 k_v|_{k_v} + |\nabla^{k_v} \partial_v k_v|_{k_v} \le 1.$$

The chain rule gives factors $L_*^{-1}$ for first derivatives and $L_*^{-2}$ for second derivatives. No common bound on the original path derivatives over all insertions is assumed. The chosen $L_*$ may depend on the individual path; the resulting annular size estimates do not. [M, Lemma 6.2].

For the small annular parameter $0 < \theta < 1/4$, set

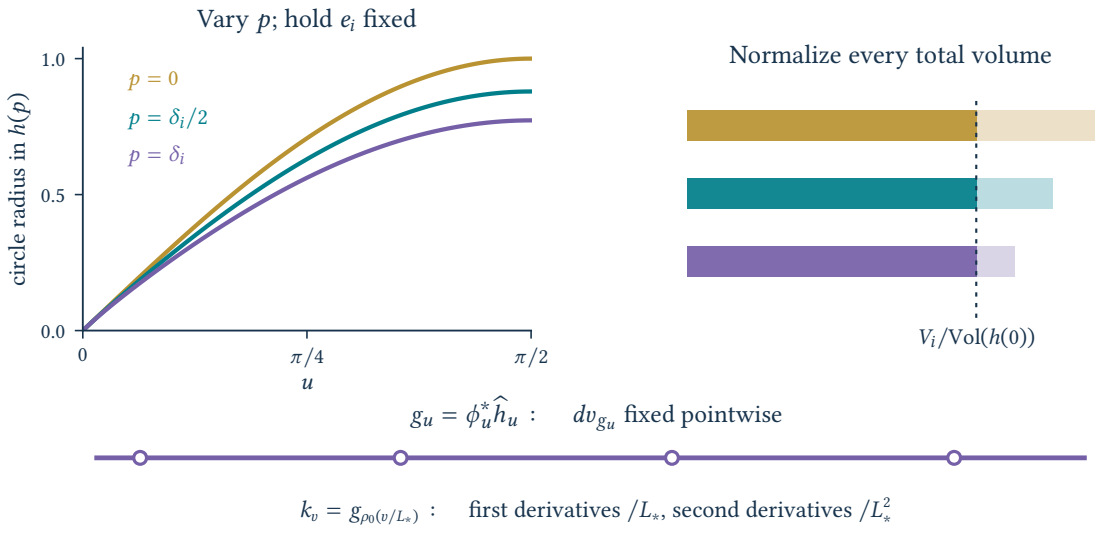
$$E = \theta/8, \quad F_0 = \theta/320, \quad L_0 = 320\theta^{-2}, \quad T = e^{L_0} - \log r, \quad L = \log T,$$
$$\chi = 1 - E/L, \qquad w = r\chi, \qquad f(r) = L_* + 1 - F_0 \log(L/L_0), \qquad g(r) = k_{f(r)}.$$

Here $\theta$ is the small interpolation parameter, not the circle coordinate. Since $f(1) = L_* + 1$, the outer slice is exactly the terminal metric near $r = 1$; as $r \downarrow 0$, $f \to -\infty$, so the inner slice is exactly the initial round metric for all sufficiently small $r$. Also $rf' = F_0/(TL) \to 0$.

The coordinated concavity and derivative estimates in [M, Lemma 6.2] give

$$\mathrm{Ric}_{rr} \ge \frac{E}{2r^2 TL^2}, \quad \lambda_{\tan} \ge \frac{E}{r^2 L}, \quad |B| \le \frac{2F_0}{r^2 TL}.$$

The squared mixed bound divided by the diagonal product equals $8F_0^2 L/(E^2 T) < 1/800$. Hence the whole Ricci tensor is positive. The stationary portions have a separate purpose: their exact warped-product form makes the boundary tensors available for the next gluing step.




VISUAL P. The path varies only the exponent, retaining its microscopic scale $e_i$. The volume bars, in units of $\mathrm{Vol}(h(0))$, have the common darker target $V_i$ after homothety. Moser pullback then fixes the volume form pointwise; its terminal map remains the attaching map. Finally, slowing this fixed-form path controls its derivatives. The displayed parameters illustrate the path formulas, without asserting the smallness conditions of an actual insertion.

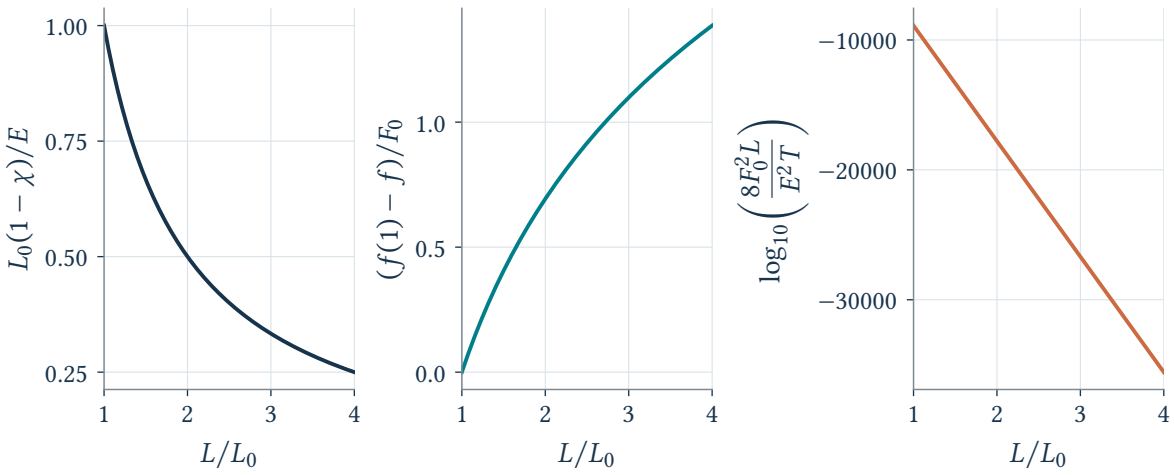



FIGURE 14. The exact annular schedule for $\theta = 1/8$. The horizontal variable approaches the inner end; the parameter continues toward $-\infty$, while the warping defect and mixed-cost fraction decrease. The curve need not contain the entire interpolation. The large logarithms are displayed without numerically evaluating $T = e^L$.

## 14. The annulus carries a quantitative surplus at the seam

**Why the spherical modification is needed.** A nonnegative-Ricci logarithmic annulus alone does not give the boundary convexity required for attachment. Put $\theta = \varepsilon^4$, replace $r\chi(r)$ by $a_\kappa(r) = \chi(r)\sin(\kappa r)/\kappa$, and write $G_\kappa = dr^2 + a_\kappa^2 g(r)$. The expansions

$$\kappa\cot(\kappa r) - r^{-1} = -\kappa^2 r/3 + O(\kappa^4 r^3), \quad (\sin(\kappa r)/\kappa)^{-2} - r^{-2} = \kappa^2/3 + O(\kappa^4 r^2)$$

show the new positive radial and tangential contributions. In the tangential calculation the leading coefficient is $1 + (n-2)/3 + 2(n-2)/3 = n-1$. The mixed term also changes: in normalized frames its difference is $(\kappa r/\sin(\kappa r) - 1)$ times the old mixed term, hence is bounded by $C\theta\kappa^2$. Combining all blocks, not just the diagonal ones, gives

$$\mathrm{Ric}_{G_\kappa} \geq (n - 1 - C\theta - C\kappa^2)\kappa^2 G_\kappa.$$

These estimates are proved in [M, Proposition 6.3].

**A value constraint that leaves a derivative margin.** Define

$$\lambda_\varepsilon = \sqrt{\sin^2\varepsilon - 8\varepsilon^4}, \qquad \chi(1)\sin\kappa_\varepsilon = \lambda_\varepsilon, \qquad c_\varepsilon = \frac{\sin\varepsilon\,\kappa_\varepsilon}{\lambda_\varepsilon}.$$

For sufficiently small $\varepsilon$, the square root is positive and $\varepsilon/2 \leq \kappa_\varepsilon \leq 4\varepsilon$. Let $a = a_{\kappa_\varepsilon}$. The rescaled metric $G_\varepsilon = c_\varepsilon^2 G_{\kappa_\varepsilon}$ has outer boundary metric $\sin^2\varepsilon\, g(1)$, since $c_\varepsilon a(1) = \sin\varepsilon$. The endpoint estimates $1 - \varepsilon^4 \leq \chi(1) \leq 1$, $|\chi'(1)| \leq \varepsilon^4$ imply

$$\chi(1)^2 - \lambda_\varepsilon^2 - \cos^2\varepsilon \geq 6\varepsilon^4.$$

Rationalization, with a positive denominator at most two, gives $\chi(1)\cos\kappa_\varepsilon - \cos\varepsilon \geq 3\varepsilon^4$. Since $\sin\kappa_\varepsilon/\kappa_\varepsilon \leq 1$,

$$a'(1) = \chi(1)\cos\kappa_\varepsilon + \frac{\sin\kappa_\varepsilon}{\kappa_\varepsilon}\chi'(1) \geq \cos\varepsilon + 2\varepsilon^4.$$

The slice is stationary at the boundary. With the outward normal, therefore,

$$\mathrm{II}_{\mathrm{ann,out}} = \frac{a'(1)}{\sin\varepsilon} g_\partial \geq \left(\cot\varepsilon + \frac{2\varepsilon^4}{\sin\varepsilon}\right) g_\partial. \tag{14.1}$$

The subtraction $8\varepsilon^4$ in the tuning equation has provided an explicit convexity reserve after matching the boundary value.

The same normalization keeps a useful Ricci coefficient because

$$\frac{\kappa_\varepsilon^2}{c_\varepsilon^2} = 1 - \frac{8\varepsilon^4}{\sin^2\varepsilon} \geq 1 - 32\varepsilon^2.$$

Thus $\mathrm{Ric}_{G_\varepsilon} \geq (n - 1 - C\varepsilon^2)G_\varepsilon$. Short radial paths and the bounded diameters of normalized links give $\mathrm{diam}_{\mathrm{int}} \leq (16 + \pi)\varepsilon$ and a volume bound $8^n\sigma_{n-1}\varepsilon^n$. Their constants are independent of the unrescaled path derivatives. [M, Proposition 6.3].

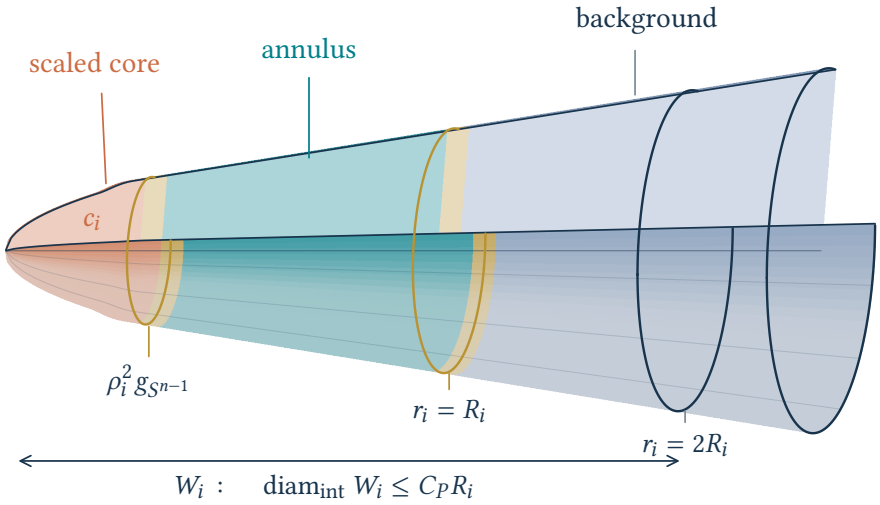



FIGURE 15. The core, annulus, and unchanged outer region form one retained domain $W_i$. Its boundary lies at the background coordinate sphere $r_i = 2R_i$. The coordinate $r_i$ labels the background only; no radial coordinate on the core is asserted. The two smoothing collars are separate.

**The surplus survives the actual background matching.** The normalized link is $h_i = (1 - \delta_i)\widetilde{h}_i$. Matching its scaled outer boundary to the removed sphere fixes

$$\sin \varepsilon_i = \frac{\sin R_i}{\sqrt{1 - \delta_i}}, \qquad R_i \le \varepsilon_i \le 2R_i.$$

More precisely, under the terminal Moser map $\phi_i$, the common tensor is

$$g_\partial = \sin^2 \varepsilon_i \, \phi_i^* h_i = \phi_i^* (\sin^2 R_i \, \widetilde{h}_i). \tag{14.2}$$

The exterior's outward second fundamental form is at least $-\cot R_i$ times its boundary metric. On the common boundary,

$$\cot R_i - \cot \varepsilon_i = \frac{\delta_i}{\sin R_i (\cos R_i + \sqrt{\cos^2 R_i - \delta_i})} \le \frac{R_i^3}{25}.$$

Here (11.6) and the previously fixed small $R_*$ are used. The surplus in (14.1) is at least $2R_i^3$, since $\varepsilon_i \ge R_i$ and $\sin \varepsilon_i \le \varepsilon_i$. Consequently

$$\mathrm{II}_{\mathrm{ann,out}} + \phi_i^* \mathrm{II}_{\mathrm{ext}} \ge \frac{49}{25} R_i^3 g_\partial > 0. \tag{14.3}$$

Both normals are outward from their respective pieces. The lower bound is small but strictly positive, and the smoothing width is chosen only after this margin is known. [M, Proposition 8.3].

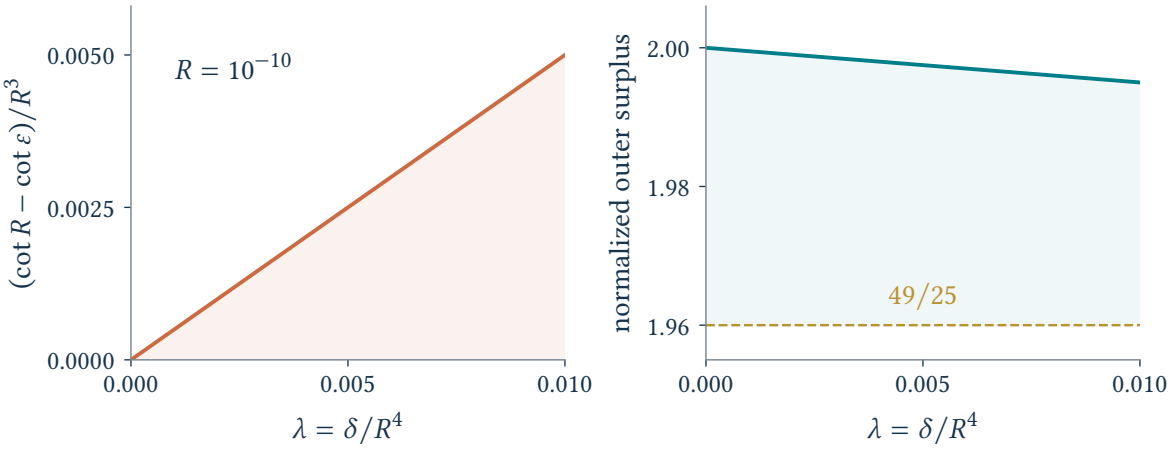
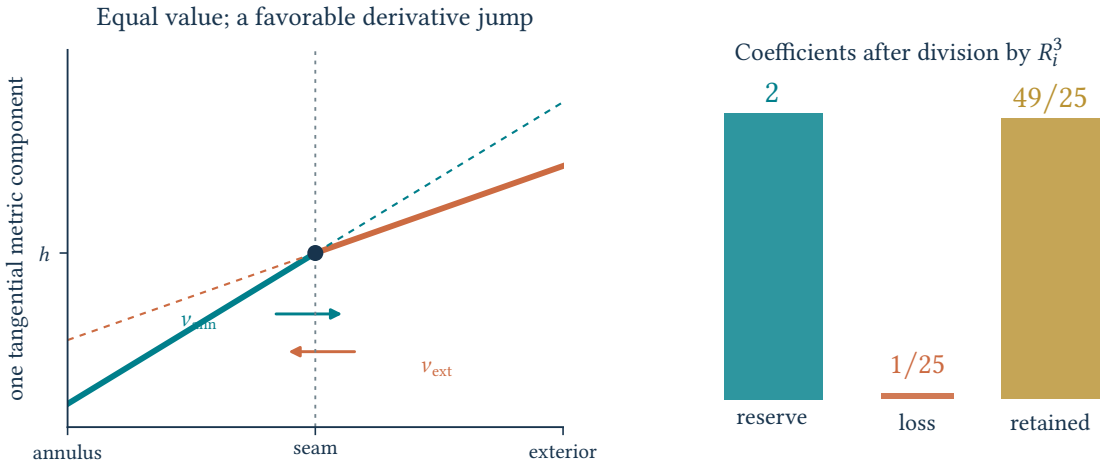



FIGURE 16. Exact matching algebra at the illustrative radius $R = 10^{-10}$. The cotangent difference is evaluated in rationalized form. The retained outer convexity, divided by $R^3$, exceeds $49/25$. The graph does not claim that the sample radius satisfies every global construction condition; $\delta = 0$ denotes the algebraic limiting endpoint.




$$h = \sin^2 \varepsilon_i \, \phi_i^* h_i = \phi_i^* (\sin^2 R_i \, \widetilde{h}_i)$$

$$\mathrm{II}_{\mathrm{ann,out}} + \phi_i^* \mathrm{II}_{\mathrm{ext}} \geq \frac{49}{25} R_i^3 h > 0$$

VISUAL Q. The common boundary tensor is fixed through the terminal Moser map, while the outward normal derivatives retain a positive sum. The left tangent-line schematic separates value matching from this derivative requirement; its slopes are not numerical annular data. The right columns show the proved bounds, in units of $R_i^3$: a reserve of at least 2, a loss of at most $1/25$, and a retained margin of at least $49/25$.

## 15. CLOSE THE MANIFOLDS, AND RETAIN EACH INSERTED DOMAIN

The boundary-link volume estimate gives $\tau_i^{n-1} \leq 10q < (\beta/2)^{n-1}$. The annulus is exactly round at its inner end, with limiting physical radial slope $\tau_i$. Choose its inner cut so that its radius $\rho_i$ satisfies

$0 < \rho_i < 2^{-i} R_i$ and its slope is less than $\beta$. The outward inner second fundamental form is at least $-\beta \rho_i^{-1} \rho_i^2 g_{S^{n-1}}$. Attach the core $(P_n, \rho_i^2 k)$. Constant scaling gives

$$\mathrm{II}_{\mathrm{core}} + \mathrm{II}_{\mathrm{ann,in}} \geq \frac{\alpha - \beta}{\rho_i} \rho_i^2 g_{S^{n-1}} > 0, \qquad \mathrm{Ric}_{\rho_i^2 k} \geq \frac{\mu}{\rho_i^2} (\rho_i^2 k).$$

The outer annular smoothing retains $\mathrm{Ric} \geq \frac{1}{2} g$; the inner smoothing can retain the common constant $\kappa = \frac{1}{8} \min\{1, \mu\} > 0$. The annulus connects to the background and the core has no remaining boundary, so every finite realization $M_j$ is smooth, connected, and closed. [M, Proposition 8.3; §9].

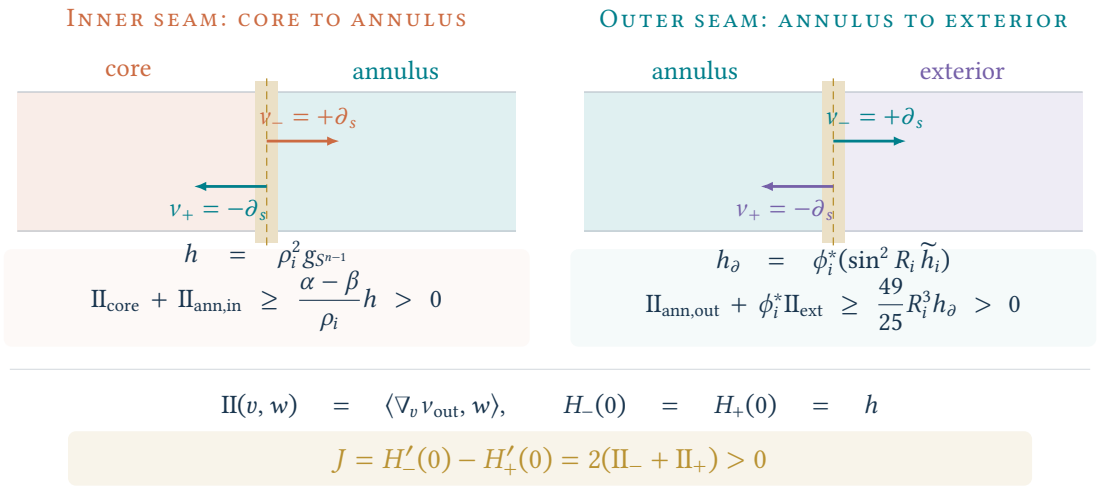



FIGURE 17. The common boundary metrics and outward normal signs are fixed at both seams. In signed coordinates increasing from left to right, $J = H_-'(0) - H_+'(0) = 2(\mathrm{II}_- + \mathrm{II}_+) > 0$. The outer tensor equality is interpreted through the terminal Moser attaching map.

**Why the seam can be smoothed.** In signed normal coordinates write the metrics as $ds^2 + H_-(s)$ and $ds^2 + H_+(s)$, with common boundary metric $h$. Their outward-normal convention gives $J = H_-'(0) - H_+'(0) = 2(\mathrm{II}_- + \mathrm{II}_+) > 0$. The cubic Hermite interpolation on $[-d, d]$ matches the metric and its first normal derivative at both ends. Its expansion is

$$H_d = h - \frac{d}{4} J + \frac{s}{2} (H_-'(0) + H_+'(0)) - \frac{s^2}{4d} J + R_d, \qquad \partial_s^r R_d = O_{C^2}(d^{2-r}), \quad 0 \leq r \leq 2.$$

Thus $H_d'' = -J/(2d) + O(1)$. In the Gaussian-coordinate Ricci formulas the leading normal and tangential terms are $\mathrm{tr}_h J/(4d)$ and $J/(4d)$, while the mixed terms stay bounded. For sufficiently small $d$, the whole quadratic form retains the prescribed slightly weaker Ricci lower bound. This argument uses the sign of the normal jump, not a claim that Ricci positivity is preserved by arbitrary convex combinations of metrics. [M, Lemma 3.1].

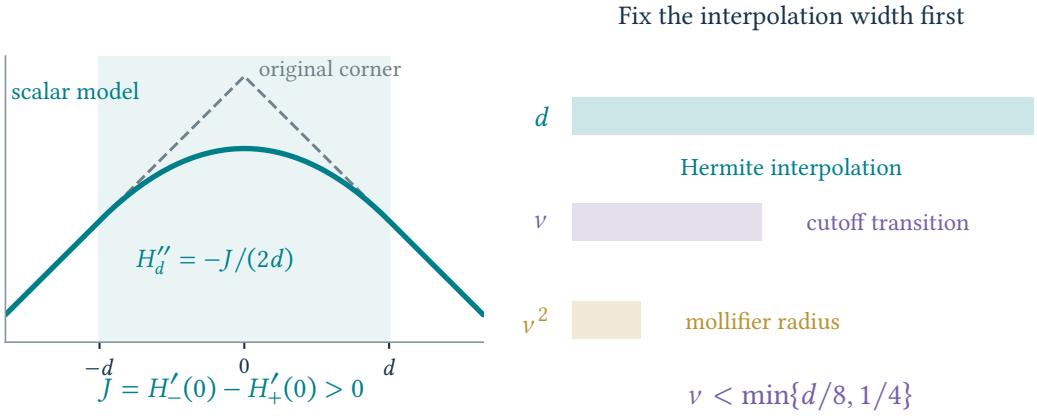
**There are two smoothing scales.** After $d$ has been fixed, smooth the two remaining joins using an even mollifier of radius $\nu^2$, supported through a cutoff of width $\nu$, with $\nu < \min\{d/8, 1/4\}$. The interpolated tensor is $C^1$, so its distributional second derivative has no point mass at these joins. Ricci is linear in that second normal derivative once the lower jets are fixed. The nonnegative kernel averages the strict Ricci inequality, and the commutator error tends to zero for fixed $d$. On the cutoff transition, evenness gives an initial $O_d(\nu^4)$ error; two cutoff derivatives leave an $O_d(\nu^2)$ error. Thus the smoothing preserves the required lower bound and is unchanged outside its chosen collar.

Tracing the same formula and integrating the interpolation gives

$$\int_{[-d,d]\times\Sigma} \mathrm{Scal}\ dv = \int_{\Sigma} \mathrm{tr}_h J\, dv_h + O(d).$$

Fix $d$ small enough for this error and the Ricci estimate, then choose $v$ for the remaining integral and tensor errors. Hence for any prescribed $\epsilon > 0$, the final collar cost is at most $\int_\Sigma \mathrm{tr}_h J\, dv_h + \epsilon$. A thinner collar alone is not an explanation of bounded scalar mass: the boundary flux is the controlled quantity. [M, Lemmas 3.1–3.2].

A favorable normal jump supplies a margin for two smoothing scales




leading Ricci terms: $\mathrm{tr}_h J/(4d)$ and $J/(4d)$; mixed terms remain bounded

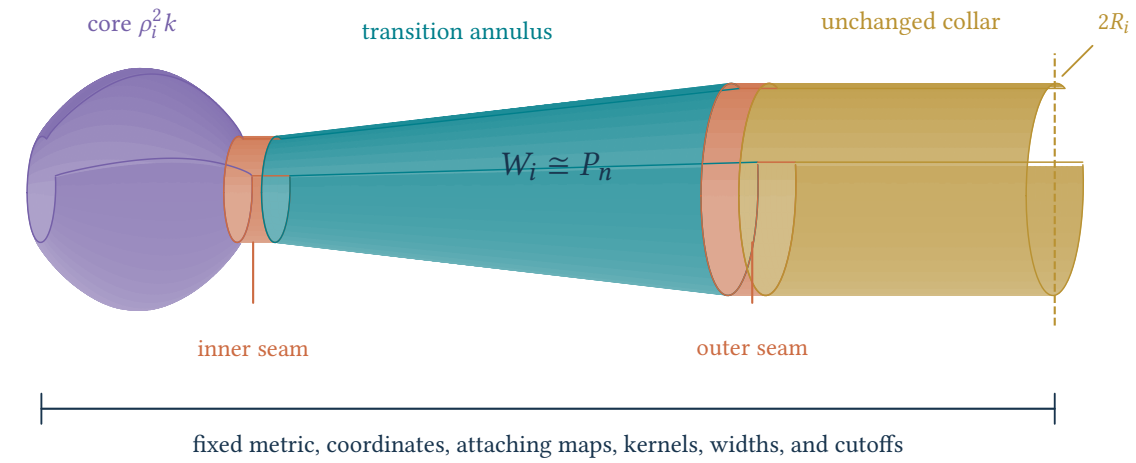
Scale bars are separately magnified; the complete tensor estimate uses the strict margin.

Visual R. The scalar graph displays the leading Hermite mechanism: a positive derivative jump becomes a negative second derivative on the interpolation strip. For the tensor metric, the resulting leading normal and tangential Ricci terms are positive, while mixed terms remain bounded. Fix the strip width first; then choose the cutoff width and smaller mollifier radius. Scalar concavity alone is not the Ricci proof.

Each seam smoothing is chosen once. The boundary parametrization, signed collar coordinates, kernels, widths, and cutoffs are reused whenever the same data recur. By (11.5), those data recur *identically* for all $j \geq i$. Include the core, its annulus, both smoothing regions, and the unchanged outer collar up to $r_i = 2R_i$ in $W_i$. Then

$$W_i \cong P_n, \qquad \mathrm{diam}_{\mathrm{int}}\, W_i \leq C_P R_i, \qquad (W_i, h_i^W) \text{ is fixed for every } j \geq i. \tag{15.1}$$

The diameter estimate is proved by paths staying inside $W_i$. The collars' supports are disjoint and do not touch later insertion regions. Thus the fixed target $W_i/\partial W_i$ used in Section 6 really is one compact metric space, not a sequence of diffeomorphic quotients with varying metrics. [M, §9; Proposition 2.3]. Figure 15 gives the assembly overview; Visual S identifies the complete metric domain retained here, including both smoothed seams and the outer collar.




VISUAL S. The retained domain contains the core, transition annulus, both smoothed seams, and the unchanged outer collar ending at $r_i = 2R_i$. All these metric data are fixed once $i$ is fixed, so the quotient $W_i/\partial W_i$ has one fixed metric in every later approximant. The cutaway suppresses dimension and depicts the decomposition, not an embedding of $P_n$.

**The accumulation set must not collapse through shortcuts.** A lower bound for distances between marked points is not automatic from the induced metric on the stratum. In its small chart, use the base projection $\Pi(t, x, \theta) = \cos t\, z(x)$, where $z$ is the Euclidean chart marking. On each unchanged annulus $2R_i \leq r_i \leq 4R_i$, replace it by

$$\Phi_j = \Pi(x_i) + \zeta(r_i)(\Pi - \Pi(x_i)), \qquad |\zeta'| \leq 2/R_i,$$

with $\zeta = 0$ near $2R_i$ and one near $4R_i$. Extend constantly over $W_i$, and keep $\Pi$ elsewhere. The difference $|\Pi - \Pi(x_i)| = O(R_i)$ cancels the cutoff derivative $O(R_i^{-1})$. The construction gives $\operatorname{Lip}\Phi_j \leq 32$ and $\Phi_j \iota_j(z) = z$ on $F$. Thus

$$d_{M_j}(\iota_j(z), \iota_j(z')) \geq \tfrac{1}{32}|z - z'|.$$

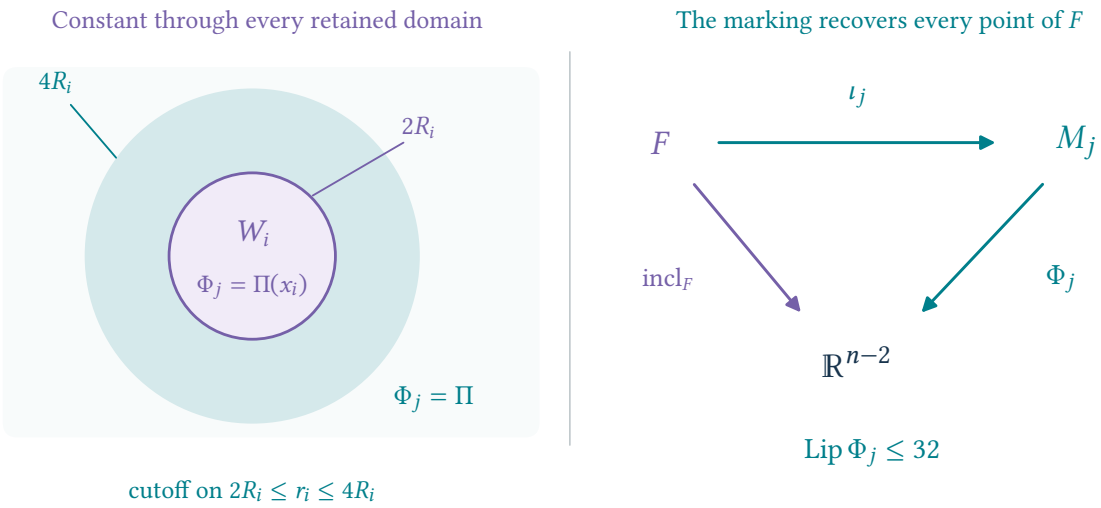
For the upper bound, replace portions of a short round path in the stratum that enter removed balls by arcs in their boundary spheres. Each detour costs at most $\pi/2$ times the replaced round length, and separation keeps it away from all other balls. The chart comparison gives $d_{M_j}(\iota_j(z), \iota_j(z')) \leq 2|z - z'|$. These uniform inequalities pass to the limit and give the bi-Lipschitz embedding $\iota$. [M, Lemmas 9.1–9.2]. The marking diagram in Figure 18 isolates the lower bound: the map recovers the chart coordinates on $F$, so its uniform Lipschitz constant limits ambient shortcuts between marked points.

The same detours and the intrinsic diameter estimate give

$$\sup_{w \in W_i} d_{M_j}(w, \iota_j(z_i)) \leq C_P R_i + 2 d_\gamma(x_i, z_i), \qquad j \geq i.$$

The right side tends to zero. Taking its decreasing tail supremum gives the $\eta_i \downarrow 0$ used in Section 6. The estimate controls every point of the domain, not only a chosen core center, and is uniform in the later approximating index. This completes the geometric input to the neighborhood homology argument. [M, Lemma 9.2; Proposition 2.4].

A uniform marking prevents the accumulation set from collapsing




$$|z - z'| = |\Phi_j \iota_j(z) - \Phi_j \iota_j(z')| \le 32\, d_{M_j}(\iota_j(z), \iota_j(z'))$$

FIGURE 18. For $i \le j$, the marking is constant on $W_i$ and returns to the base projection across the unchanged annulus. The triangle commutes in $\mathbb{R}^{n-2}$: $\Phi_j \circ \iota_j$ is the coordinate inclusion of $F$. Its uniform Lipschitz bound gives the lower ambient-distance estimate displayed in the diagram for $z, z' \in F$. Boundary detours give the separate upper bound $d_{M_j}(\iota_j(z), \iota_j(z')) \le 2|z - z'|$. The concentric circles depict cutoff regions; $r_i$ is a background coordinate and does not parametrize the core.

**Compactness and noncollapse.** An untouched common coordinate ball gives $\mathrm{Vol}\, B_1(p_j) \ge v > 0$. The common positive Ricci bound gives $\mathrm{diam}\, M_j \le D = \pi\sqrt{(n-1)/\kappa}$ by Bonnet–Myers. Since $B_1(p_j) \subset B_{D+1}(y)$, Bishop–Gromov implies

$$\mathrm{Vol}\, B_r(y) \ge \frac{v r^n}{(D+1)^n} \quad (0 < r < 1).$$

The upper comparison bounds the total volumes; disjoint small balls then give uniform bounds on the sizes of separated sets. Gromov–Hausdorff compactness supplies a convergent subsequence with compact noncollapsed limit. This is a limit of the already constructed closed smooth manifolds, not an unjustified assertion about an infinite smooth gluing. [M, Lemma 9.3].

## 16. Curvature concentration has a summable cost

Write $G = \bar{G} + b^2 d\theta^2$, where $\bar{G} = dt^2 + a^2\gamma$ and $m = n - 2$. The circle-warping formula gives $\mathrm{Scal}_G = \mathrm{Scal}_{\bar{G}} - 2b^{-1}\Delta_{\bar{G}} b$. Multiplication by $dv_G = b\, dv_{\bar{G}}\, d\theta$ cancels the inverse factor. Spatial divergence integrates to zero on the closed $S^m$. For $T$ in the local region, $a(0,x) = b_t(0,x) = 1$ gives

$$\int_{t<T} \mathrm{Scal}_G\, dv_G = 2\pi \int_0^T \int_{S^m} b\, \mathrm{Scal}_{\bar{G}}\, a^m\, dv_\gamma\, dt - 4\pi \int_{S^m} a(T,x)^m b_t(T,x)\, dv_\gamma + 4\pi\sigma_m. \tag{16.1}$$

The upper-boundary term is nonpositive. The remaining volume term is bounded by $C\int_0^T t(1+|\log t|^2)\,dt$, uniformly in the positive floor. For example,

$$\int_0^T t(\log t)^2\,dt = \frac{T^2}{2}\big((\log T)^2 - \log T + \tfrac{1}{2}\big) \longrightarrow 0.$$

The steep smooth opening contributes the finite endpoint flux $4\pi\sigma_m$, not an unbounded curvature maximum. The fixed outer strip is bounded separately. [M, Lemma 10.1].

**Why the slice scalar integrals are uniformly bounded.** Let $d = n-1$ and use the unscaled path $h(p)$ from Section 13. Its circle warping has derivative one at $u = 0$; the factor $\cos^{d-2} u$ kills the other endpoint flux at $u = \pi/2$. The same integration gives the exact identity

$$\int_{S^d} \mathrm{Scal}_{h(p)}\; dv_{h(p)} = (d-1)(d-2)\,\mathrm{Vol}(h(p)) + 4\pi\sigma_{d-2}. \tag{16.2}$$

Since its circle radius is at most $\sin u$, $\mathrm{Vol}(h(p)) \le \sigma_d$. If a metric is multiplied by a scalar $a \le 1$, its scalar integral is multiplied by $a^{(d-2)/2} \le 1$. The normalization in (13.2) therefore preserves a uniform upper bound, and Moser pullback does not change the integral. This is a bound independent of $i$ and of the derivatives of its metric path. [M, Lemma 10.2].

For an unscaled annulus $G = dr^2 + a^2 g(r)$, with $dv_{g(r)} = \nu$ fixed and $d = n-1$,

$$\mathrm{Scal}_G = a^{-2}\,\mathrm{Scal}_{g(r)} - d(d-1)(a'/a)^2 - 2da''/a - \tfrac{1}{4}|g'|_g^2.$$

Put $V = \int \nu$ and $S(r) = \int \mathrm{Scal}_{g(r)}\;\nu$. After integrating against $a^d\nu\,dr$, the two radial contributions are

$$d(d-1)V\int_{r_0}^1 a^{d-2}(a')^2\,dr - 2dV[a^{d-1}a']_{r_0}^1. \tag{16.3}$$

The slice term is $\int_{r_0}^1 a^{d-2}S(r)\,dr$. The uniform bound from (16.2), $a \le r$, and $|a'| \le 2$ control these terms independently of $r_0$. The remaining $-\frac{1}{4}|g'|_g^2$ has a favorable sign for this upper estimate. Nonnegativity of scalar curvature comes separately from the Ricci construction; it is not a consequence of discarding that term. Rescaling gives an annular cost $CR_i^{n-2}$. [M, Lemma 10.3].

A core scaled by $\rho_i^2$ has scalar integral multiplied by $\rho_i^{n-2}$. The outer seam has area $O(R_i^{n-1})$ and second fundamental form of size $O(R_i^{-1})$; its flux is $O(R_i^{n-2})$. At the inner seam the corresponding cost is $O(\rho_i^{n-2})$, hence $O(R_i^{n-2})$. Lemma 3.2 of [M] permits a separately prescribed integral error, chosen at each of the two seams to be at most $R_i^{n-2}/2$, while retaining the support and Ricci estimates.

Adding the background, annuli, cores, and collars yields

$$\int_{M_j} \mathrm{Scal}_{g_j}\; dv_{g_j} \le C_0 + C\sum_{i=1}^{j} R_i^{n-2} \le C_0 + \frac{CR_*^{n-2}}{2^{n-2}-1}. \tag{16.4}$$

Here $R_i < 2^{-i}R_*$. All Ricci eigenvalues are nonnegative, so their Hilbert–Schmidt norm is at most their sum: $|\mathrm{Ric}| \le \mathrm{Scal}$. This proves (2.4). [M, Proposition 10.4].

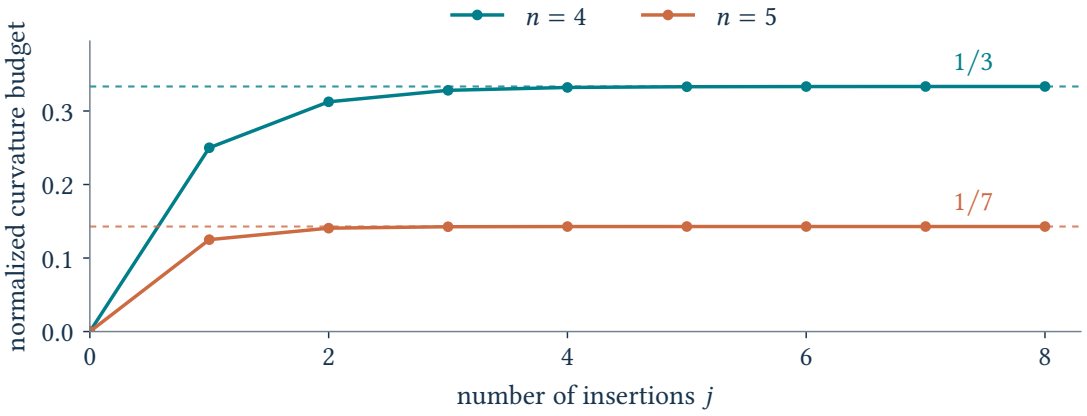
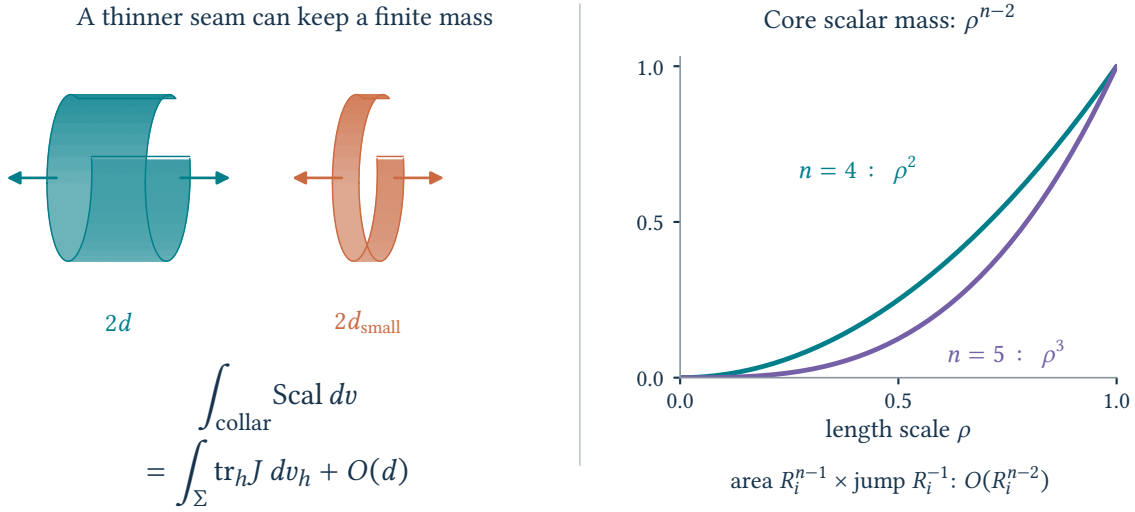



FIGURE 19. The normalized added scalar-curvature envelope has partial sums $(1 - 2^{-j(n-2)})/(2^{n-2} - 1)$. Its limits are $1/3$ for $n = 4$ and $1/7$ for $n = 5$. The graph is a summable upper bound, not the measured scalar mass of an approximant.

## Concentration is paid for by flux and scaling




VISUAL T. The Hermite strip contributes its boundary flux plus $O(d)$. After fixing $d$, the second smoothing error is chosen separately small; width alone does not force zero mass. The exact curves show scalar-integral scaling for a fixed core metric multiplied by $\rho^2$. The outer seam costs $O(R_i^{n-2})$, since area $O(R_i^{n-1})$ and normal jump $O(R_i^{-1})$ contribute complementary powers.

**Products preserve the required bounds.** For $N \geq 5$, take the five-dimensional sequence times a fixed flat $T^{N-5}$; for $N = 4$, keep the four-dimensional sequence. The product-distance formula carries converging correspondences to the product. Products of based radius-1/2 balls lie in the based unit ball, so noncollapse persists. Ricci remains nonnegative, but a positive-dimensional flat factor introduces zero Ricci directions. Scalar curvature and Ricci norm are those of the base, and their integrals are multiplied by the fixed torus volume. This proves the higher-dimensional metric and integral claims; the deletion threshold still uses the full product perturbation of Section 9. [M, proof of Corollary 1.2].

## 17. At one point, absolute angular error is enough

For the isolated example take $n = 4$ and $F = \{x_*\}$. Let $p = (x_*, 0)$. The unmodified positive-floor backgrounds converge smoothly on $S^4 \setminus \{p\}$ to a metric $g_0$. This requires the smooth parameter convergence of [M, Lemma 12.1], not just pointwise convergence of $b_\delta$. On the floor transition the derivatives are bounded by an exponential smallness factor times powers of its logarithm. For fixed $t > 0$, the expression

$$b_\delta(t) = q \sin t \exp\{-\tfrac{\delta}{2} \log(\sin^2 t + e^{-2c/\delta})\}$$

also has its smooth parameter extension at $\delta = 0$; the integrated correction kernel is treated in the same way.

The completion $Y_0$ is topologically $S^4$, and its distance $d_0$ satisfies the same two-sided round-distance comparison as (11.7). To get uniform distance convergence, first compare tensors outside a fixed small round ball about $p$. A path entering that ball can have its whole inner passage replaced at cost $O(\rho)$, because the background metrics are bounded above by the round metric. Let the approximating index tend to infinity, and then $\rho \downarrow 0$. This proves convergence to $Y_0$. It is the background $Y_0$, not yet the space $Y$ containing the persistent inserted cores. [M, Lemma 12.2].

**Rescale the coefficients that multiply lengths.** Use normal coordinates $y$ on the parameter sphere and write $y = rY$, $t = rT$. Then $\gamma(rY) \to |dY|^2$. Flatness of $\delta$ and $A$ at $x_*$ makes the other longitudinal coefficient tend to one uniformly on bounded rescaled sets. The useful angular error is

$$e_B(r) := \sup_{|Y| \le B,\ 0 \le T \le B} |r^{-1} b_{\delta(rY)}(rT) - qT| \longrightarrow 0. \tag{17.1}$$

Indeed, with $\delta_r = \sup_{|y| \le Br} \delta(y)$,

$$e_B(r) \le C_B \left[ r^2 + \delta_r (1 + |\log r|) e^{\delta_r |\log r|} \right].$$

For every $a > 0$, $\delta_r = O(r^a)$, so the right side vanishes. At $T = 0$, both coefficients in (17.1) are zero. An estimate obtained by dividing by $T$ would instead demand a false uniform relative opening-angle estimate down to every smooth microscopic pole.

A minimizing curve in $\mathbb{R}^2 \times \mathsf{C}(S^1_{2\pi q})$ can be chosen with total variation of the period-$2\pi$ angular coordinate at most $\pi$. Thus its extra angular length is at most $\pi e_B(r)$; the longitudinal error also tends to zero. Conversely, the lower coefficient comparison bounds lengths from below. The round-distance comparison confines any sufficiently short competitor to a larger fixed rescaled coordinate region, so an excursion outside that region cannot create an unexamined shortcut. These comparisons, with geodesic trimming to the actual closed balls, give

$$(Y_0, r^{-1} d_0, p) \longrightarrow (\mathbb{R}^2 \times \mathsf{C}(S^1_{2\pi q}), o) \quad \text{as } r \downarrow 0. \tag{17.2}$$

The estimates hold for the full limit, not for a selected sequence of scales. [M, Lemma 12.2].

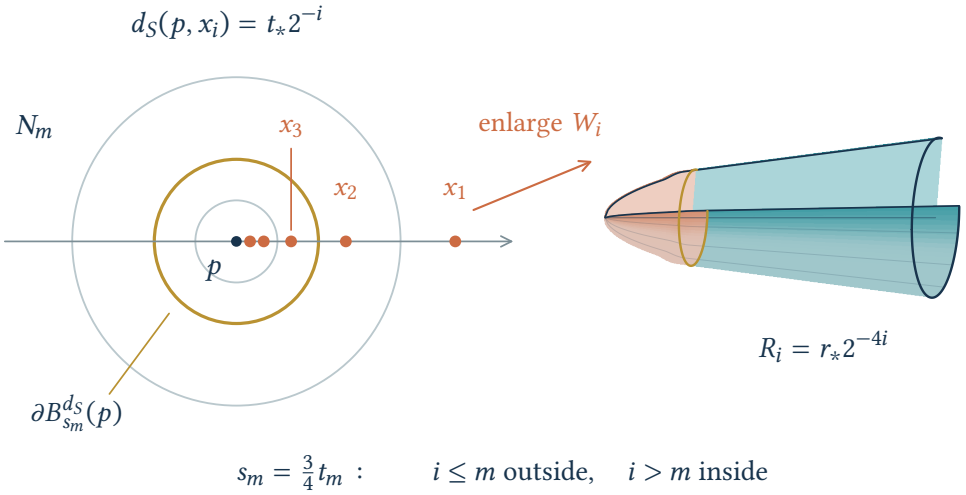
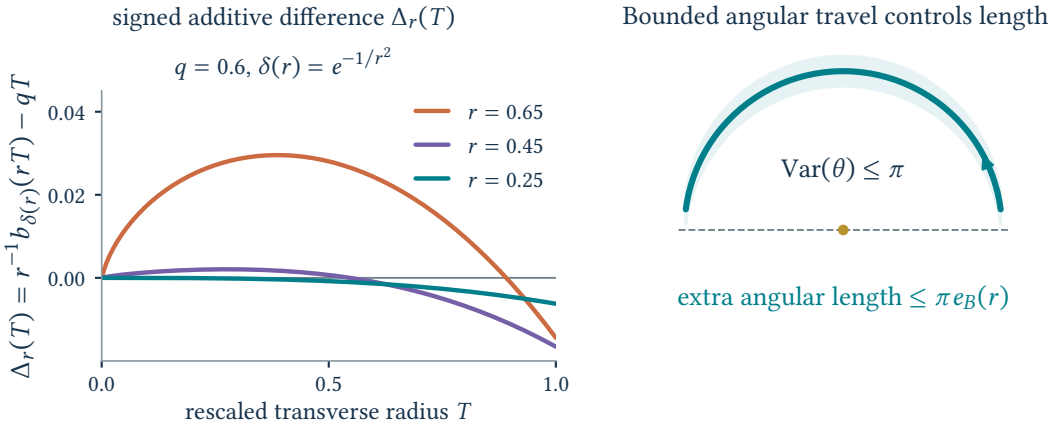



FIGURE 20. The isolated construction uses $t_i = t_* 2^{-i}$, $R_i = r_* 2^{-4i}$ and round coordinate radii $s_m = 3t_m/4$. With the prescribed small $r_*$, the first $m$ replacement regions lie outside that coordinate ball and all later ones inside. Dots show centers; the domain at right is separately enlarged.

The useful error is additive, even at the smooth pole




At $T = 0$ both length coefficients vanish; no division by $T$ is needed.

VISUAL U. The curves show the signed additive difference $\Delta_r(T) = r^{-1} b_{\delta(r)}(rT) - qT$ for the exact circle profile, with illustrative choices $q = 0.6$ and $\delta(r) = e^{-1/r^2}$. The proof instead bounds the absolute value uniformly in both $Y$ and $T$, with the constructed spatial parameter $\delta(rY)$; this bound is $e_B(r)$. A comparison curve in the product cone can have angular variation at most $\pi$, giving additional angular length at most $\pi e_B(r)$.

## 18. The replacement tail disappears under magnification

**Compare through the common exterior.** At a finite stage, let $M$ be the replaced manifold and $B$ its unmodified background. Their common exterior $E$ has boundary components $C_i$. Collapse each $C_i$ to a point, using the chain length induced by the *intrinsic* exterior metric. The resulting quotient is the same whether one fills the boundaries with the domains $W_i \subset M$ or the original balls $K_i \subset B$.

A quotient chain is a finite list of exterior segments whose successive endpoints may lie in the same collapsed component. If the chain visits a component twice, discard the intervening loop. This never increases length. Only finitely many component orders remain; compactness of their boundary sets gives minimizing representatives. The distances between distinct compact boundary components are positive. These facts justify the quotient metric and a minimizing chain with no repeated component, rather than assuming that a drawing of collapsed regions automatically defines the desired distance. [M, Lemma 13.1].

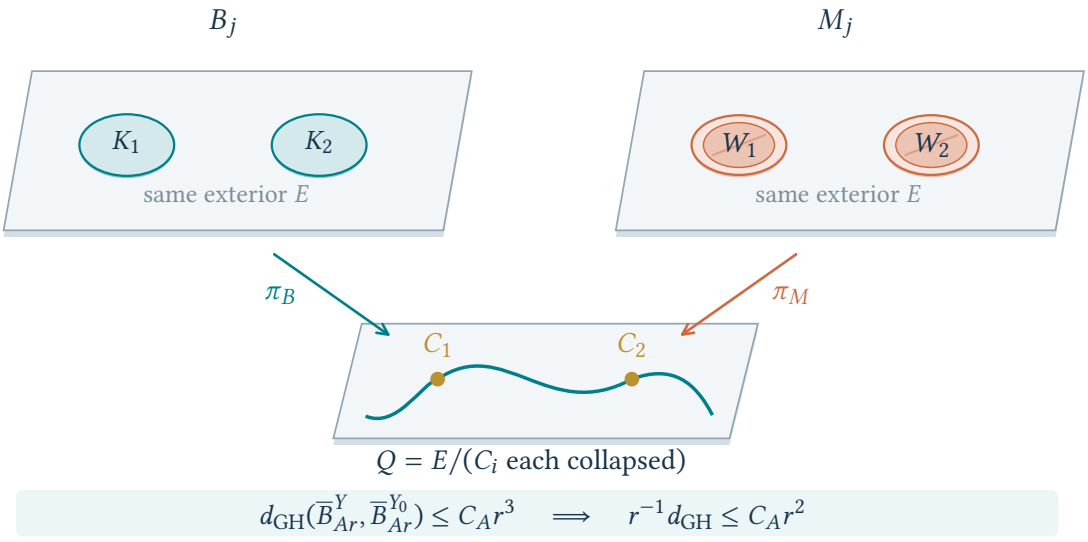


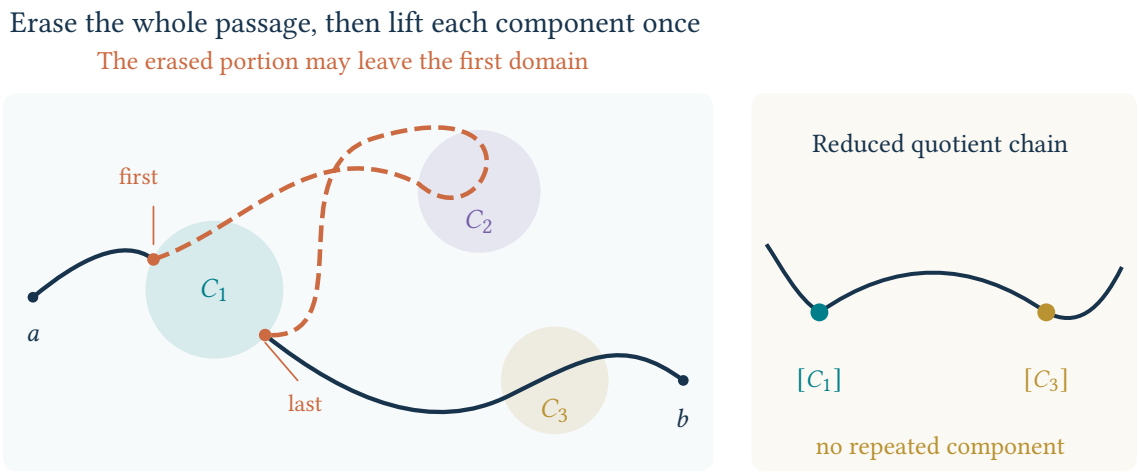

FIGURE 21. The common exterior with its boundary components collapsed gives a common metric quotient at each finite stage. Lifting a chain without repeated components costs a sum of intrinsic diameters. The pictured quotient need not be a Euclidean region. Closed balls in the comparison have their restricted ambient distances.

Both collapse projections are 1-Lipschitz. For a curve in either manifold, erase its entire passage between the first and last remaining visits to one domain; the endpoints have the same quotient image. The remaining exterior segments cost no more than the original curve. Conversely, lift a reduced quotient chain by joining its boundary representatives inside each $W_i$ or $K_i$. This costs at most their intrinsic diameters, including the possible endpoint domains.

If $(u, v), (u', v')$ have equal quotient images, [M, Lemma 13.1] gives

$$|d_M(u, u') - d_B(v, v')| \le 4 \sum_{t_i \le 4LR} \max\{\operatorname{diam}_{\text{int}} W_i, \operatorname{diam}_{\text{int}} K_i\} \tag{18.1}$$

when $u, u' \in \overline{B}^M_R(p)$, or symmetrically for $v, v'$ in the background ball. Here $L$ is the Lipschitz bound of the common localization function described next. Ambient diameters alone would not justify these lifted paths inside the domains.




$$\text{Lifting adds at most } \sum_i \max\{\operatorname{diam}_{\text{int}} W_i, \operatorname{diam}_{\text{int}} K_i\}.$$

Each lifted connector stays inside its filling; endpoint fillings are included.

> Visual V. Between the first and last remaining visits to one domain, erase the entire passage, including excursions through the exterior or other domains. Its endpoints have the same image after boundary collapse. A reduced quotient chain visits each remaining component at most once. Lifting it uses paths inside the fillings, so the cost is controlled by intrinsic diameters, with endpoint domains included.

**Localization excludes the distant domains.** Write $\rho(z) = d_{S^4}(p, z)$ and $t_i = \rho(x_i)$. On $2R_i \le r_i \le 4R_i$, define

$$\lambda_j = t_i + \zeta(r_i)(\rho - t_i),$$

with the same type of cutoff as in Section 15. Extend constantly with value $t_i$ over $W_i$, or over $K_i$ in the background, and use $\rho$ elsewhere. Since $|\rho - t_i| \le 4R_i$ on this annulus, the cutoff derivative has a uniform cost. [M, Lemma 13.2] gives

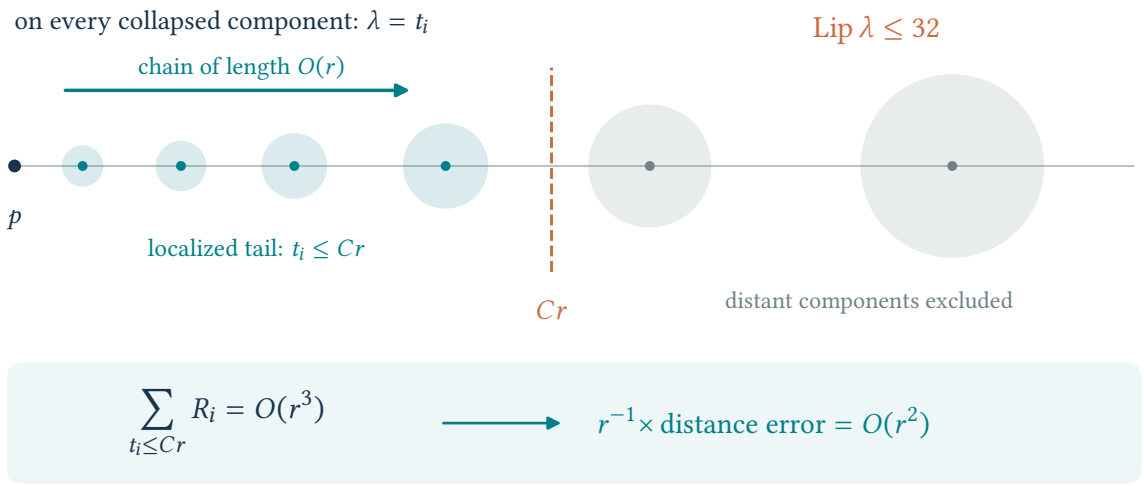
$$\operatorname{Lip} \lambda_j \le 32, \qquad \lambda_j \le 32 d(p, \cdot), \qquad d(p, \cdot) \le C_* \lambda_j. \tag{18.2}$$

It also gives $\rho/2 \le \lambda_j \le 2\rho$ on the exterior. The two versions agree there and are constant on every collapsed domain, so they descend to one localization function on the common quotient. A minimizing quotient chain starting $O(r)$ from the basepoint and of length $O(r)$ can only visit components with $t_i \le Cr$. This explains the restriction on the sum in (18.1). The function $\lambda_j$ is not claimed to equal metric distance. Visual W separates the full family from the components reachable by a short based chain. The required cubic rate concerns this localized tail; convergence of the global series alone does not supply it.

Assume

$$\sum_{t_i \le u} R_i \le C_0 u^3 \tag{18.3}$$

for small $u$. With $\operatorname{diam}_{\text{int}} W_i \le C_P R_i$ and $\operatorname{diam}_{\text{int}} K_i \le 4R_i$, the localized sum gives, for each fixed $A > 0$, a relation distortion $D_A(r) \le K_A r^3$ whenever either pair of corresponding endpoints is based within $2Ar$. Use the larger radius in the quotient estimate before fixing $K_A$, and take $r$ small enough that $D_A(r) < Ar/2$. These choices are independent of the approximating index. The common-quotient relation contains $(p, p)$. Consequently the distances to the basepoint of two related points differ by at most $D_A(r)$. This controls the radius error, but does not yet give a surjective relation between balls of exactly the same radius. [M, Proposition 13.3].




VISUAL W. The localization function is constant with value $t_i$ on each collapsed component and has a uniform Lipschitz bound. Hence a quotient chain beginning $O(r)$ from the basepoint and having length $O(r)$ cannot reach a component with $t_i > Cr$. Only the indicated tail contributes. Its cubic diameter budget becomes a quadratic error after distances are multiplied by $r^{-1}$.

**Correct the radius error in both directions.** For each $u \in \overline{B}^M_{Ar}(p)$, choose a related $v$ in the background. If $d_B(p, v) > Ar$, move $v$ along a minimizing geodesic from $p$ to its point at distance $Ar$; otherwise leave it unchanged. The displacement is at most $D_A(r)$. Repeat the construction in the opposite direction starting with each point of the background ball, and take the union of the two resulting relations.

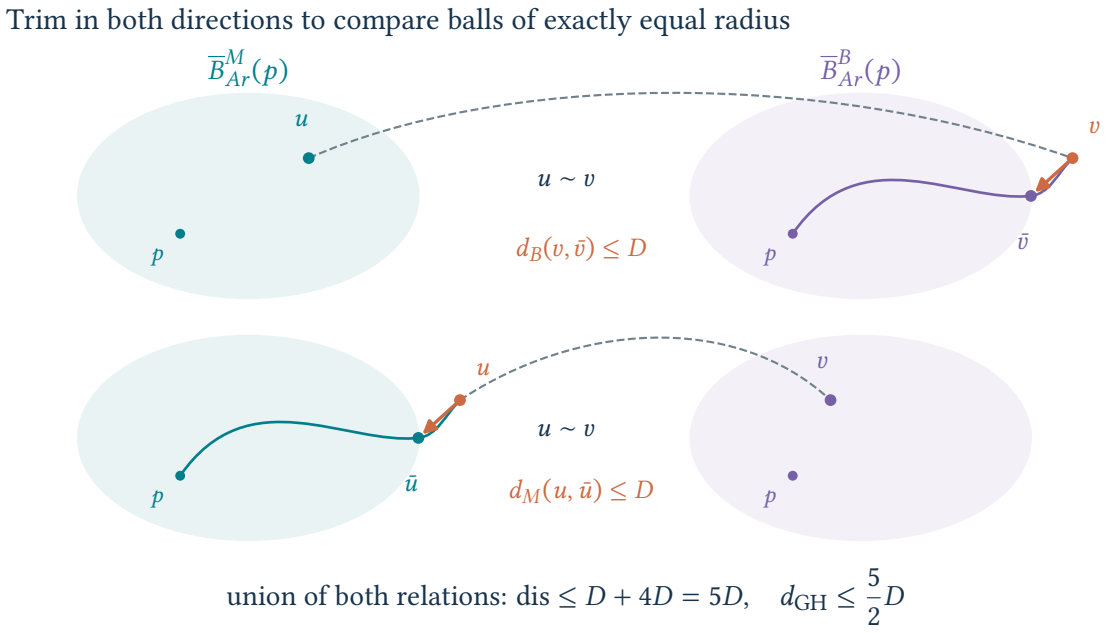
The union projects onto both closed balls and contains $(p, p)$. The original witnesses for each trimmed pair lie within $Ar + D_A(r) < 2Ar$, so the preceding localized estimate applies to them. When comparing two pairs, the original distance error is at most $D_A(r)$ and the endpoint motions contribute at most $4D_A(r)$. Hence its distortion is at most $5D_A(r)$. The standard correspondence bound gives

$$d_{\mathrm{GH}}(\overline{B}^M_{Ar}(p), \overline{B}^B_{Ar}(p)) \leq \tfrac{5}{2} D_A(r) \leq C_A r^3. \tag{18.4}$$

These are minimizing *geodesics*, not paths along an assumed radial coordinate inside the cores. No continuous or single-valued trimming map is needed. This is a correspondence argument with ambient ball distances.

First fix $r$ and let the approximating index tend to infinity. Closed balls of fixed radius in the compact geodesic spaces converge after correcting the small radius errors by the same geodesic trimming procedure. Replace each trimmed relation by its closure in the product of its two compact balls. Continuity of distances preserves its distortion bound; the two projections remain surjective and the basepoint pair is retained. Realize the two converging ball sequences in compact ambient spaces and take a Hausdorff limit of these closed relations. For any point of either limiting ball, choose approximating points and their partners; compactness in the other factor supplies a limiting partner. Thus both projections of the limiting relation are surjective. Its basepoint pair and distortion bound are retained as well. Thus (18.4) gives

$$d_{\mathrm{GH}}(\overline{B}^Y_{Ar}(p), \overline{B}^{Y_0}_{Ar}(p)) \leq C_A r^3$$




Visual X. A related point can lie just beyond the required ball. Move it inward along a minimizing geodesic by at most $D = D_A(r)$, and perform the same construction in the opposite direction. The union is surjective onto both closed balls and contains the basepoint pair. Comparing original witnesses and moved endpoints gives distortion at most $5D$; all ball distances are restricted ambient distances.

and a pointed correspondence with distortion at most $2C_A r^3$. Only now multiply distances by $r^{-1}$. Its distortion is at most $2C_A r^2 \to 0$. Composing with (17.2) yields

$$\mathrm{Tan}(Y, p) = \{\mathbb{R}^2 \times \mathsf{C}(S^1_{2\pi q})\}.$$

The argument applies for every fixed $A > 0$ and every sufficiently small $r$, so the cone does not depend on a blow-up subsequence. [M, Proposition 13.3; §14].

For the actual construction choose

$$t_i = t_* 2^{-i}, \qquad R_i = r_* 2^{-4i}, \qquad 0 < t_* < 1/100,$$
$$0 < r_* < \min\{R_*/4, t_*^3/4, t_*/10^8\}.$$

Then

$$\sum_{k \ge i} R_k = \frac{16}{15} r_* 2^{-4i} \le t_i^3.$$

For a small $u$, let $i$ be the first index with $t_i \le u$. The sum in (18.3) is this tail, hence at most $u^3$. The separation conditions also follow from the small ratio $r_*/t_*$. There is a stronger fourth-order envelope for this particular dyadic choice, but the proof uses the cubic estimate. Global summability by itself would not prove that the replacement error is $o(r)$ near the chosen point.

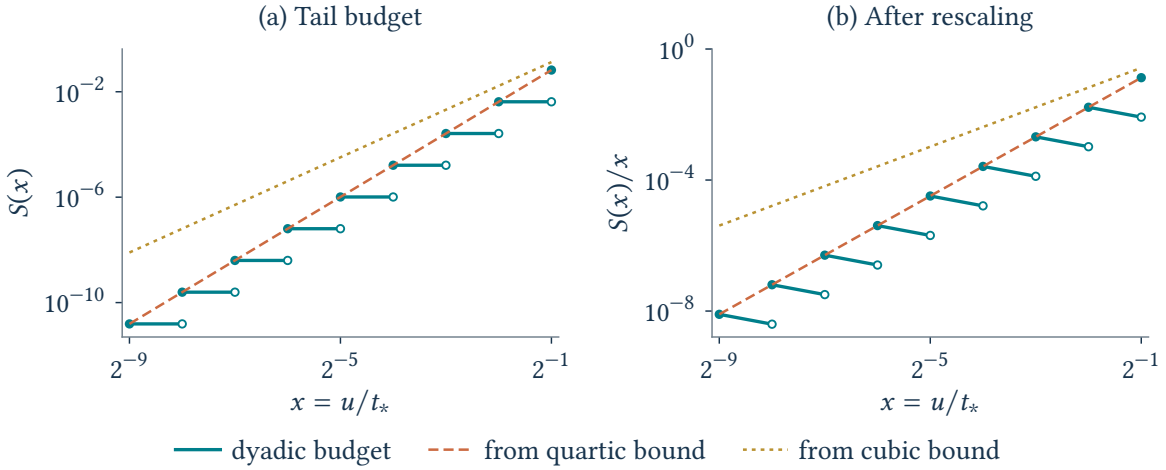



FIGURE 22. The normalized dyadic tail $S(x) = \sum_{2^{-i} \le x} 2^{-4i}$ lies below $(16/15)x^4$ and hence $(16/15)x^3$ for $0 < x \le 1/2$. Dividing by the observation scale changes a cubic bound into a quadratic one. The plotted quantities are diameter bounds, not measured Gromov–Hausdorff distances.

## 19. A SMOOTH PUNCTURE WITH NO CONICAL NEIGHBORHOOD

**Construct the candidate smooth puncture.** Put $s_m = 3t_m/4$. The first $m$ replacement regions and their enlarged collars lie outside the round coordinate ball of radius $s_m$; all later ones lie inside it. Define

$$N_m = M_m \setminus B^{d_{S^4}}_{s_m}(p).$$

The removed ball here is a *round coordinate ball in the unchanged background*, not an asserted distance ball for $g_m$. The common coordinates and fixed attaching maps identify $N_m \subset \operatorname{Int} N_{m+1}$ and embed each $N_m$ into every $M_j$, $j \ge m$. Thus $\mathcal{N} = \bigcup_m \operatorname{Int} N_m$ is a smooth four-manifold. Any two points lie in one member and can be separated there; the countable union of chart bases gives Hausdorffness and second countability.

The smoothly limiting background and the finitely many fixed replacement tensors agree on the collars. They define a smooth metric $g_\infty$ on $\mathcal{N}$, and the pulled-back metrics converge in $C^\infty(N_m)$ for every fixed $m$. Uniform limits of the compatible embeddings give a map $e : \mathcal{N} \to Y$. These facts alone do not yet say that it is an open embedding onto the entire puncture. [M, Lemma 14.1].

**First exit proves local metric agreement.** Fix a coordinate neighborhood with nested coordinate balls of radii $b < 2a < 3a$, where all relevant tensors satisfy $c\, g_{\mathrm{eucl}} \le g \le C\, g_{\mathrm{eucl}}$ on the middle ball. Choose $b < a\sqrt{c}/(4\sqrt{C})$. A path beginning in the small ball and leaving the middle one has length at least $\sqrt{c}(2a - b) > \sqrt{c}\, a$, while a coordinate segment between points of the small ball has length at most $2\sqrt{C}\, b < \sqrt{c}\, a/2$. Thus the distance infima are confined to the middle chart. Uniform tensor comparison there and passage to the limit give

$$d_Y(e(v), e(w)) = d_{g_\infty}(v, w)$$

for $v, w$ in the small chart. The same first-exit estimate, choosing the chart to exclude any second specified point, proves global injectivity of $e$. This controls ambient shortcuts, not just lengths of curves already assumed to remain in a chart.

**The image contains a whole neighborhood.** Take $y$ sufficiently close to $e(u)$ and points $y_j \in M_j$ converging to it. Their distances to the image of $u$ are eventually smaller than the uniform first-exit length of a smaller coordinate neighborhood. They must therefore lie in that coordinate neighborhood. A subsequence of their coordinates converges in its compact closure to $v$, and uniform convergence

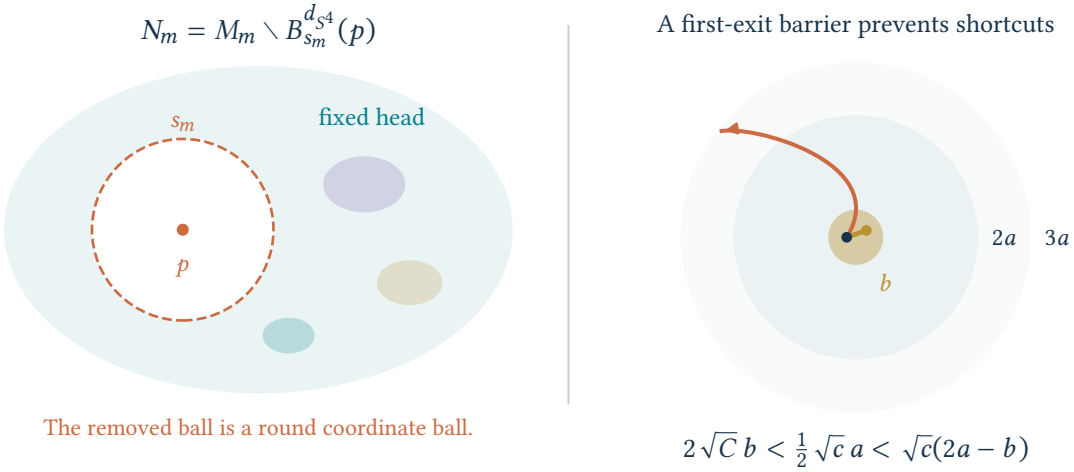
of the embeddings gives $e(v) = y$. This proves that an entire metric ball about $e(u)$ lies in the image. Together with the distance identity, it gives an isometry between small open balls. Thus $e(\mathcal{N})$ is open in $Y$. This extra argument is distinct from the local inverse-isometry statement used for avoidance in Section 7. [M, Lemma 14.1].

**No point of the puncture is missed.** The localization functions in (18.2) have a graph limit $\lambda : Y \to [0, \infty)$ with

$$\lambda \le 32 d_Y(p, \cdot), \qquad d_Y(p, \cdot) \le C_* \lambda, \qquad \lambda \circ e = \lambda_\infty > 0 \text{ on } \mathcal{N}.$$

Hence $\lambda^{-1}(0) = \{p\}$, so the smooth image avoids $p$. The construction also gives $\sup_{M_j \setminus N_m} \lambda_j \le 2 s_m$ for $j \ge m$, interpreted through the fixed embeddings. If $y \ne p$, take $a = \lambda(y)/2 > 0$ and choose $m$ with $2 s_m < a$. Any sequence $y_j \to y$ has $\lambda_j(y_j) > a$ eventually and is therefore trapped in this fixed compact $N_m$. Compactness and uniform convergence supply a preimage under $e$. Consequently $e(\mathcal{N}) = Y \setminus \{p\}$, proving the claimed smooth puncture. [M, Lemma 14.1].

A protected finite head becomes part of the smooth puncture



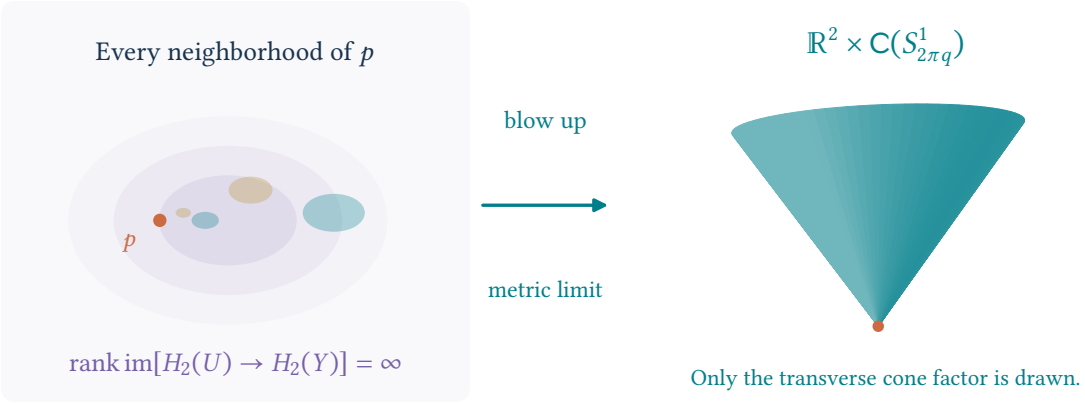

Local distance agreement, openness, and localization identify the entire puncture.

> Visual Y. The coordinate cut separates a fixed finite head from all later insertions, producing compatible smooth domains. In a small chart, a short internal segment is cheaper than any path that exits the middle chart; this rules out ambient shortcuts. Local distance agreement is followed by a separate openness argument and localization argument, which together identify the smooth union with the whole puncture.

At $p$, every open neighborhood contains all sufficiently late images $f_i(W_i)$. Their detector identities make the ambient classes independent, so the image of neighborhood $H_2$ in $H_2(Y; \mathbb{Z})$ has infinite rank. If $p \in V \subset U$ and $V \hookrightarrow U$ were null-homotopic, $H_2(V) \to H_2(U) \to H_2(Y)$ would vanish, a contradiction. Thus $p$ is not locally contractible.

An open cone $([0, a) \times K)/(\{0\} \times K)$ over a compact space contracts by $[r, z] \mapsto [(1 - s)r, z]$. Compactness supplies the uniform small-radius neighborhoods needed for continuity at the vertex. Its absolute second homology is zero. It therefore cannot be homeomorphic to an open neighborhood of $p$. No formula for $H_*(Y, Y \setminus \{p\})$ in terms of a fixed link is asserted or used. The metric tangent being a cone does not provide a cone chart on $Y$. [M, Theorem 1.3; §14].

A conical tangent does not provide a conical neighborhood




An open cone contracts to its vertex, so $H_2 = 0$.

It cannot be homeomorphic to an open neighborhood of $p$ in $Y$.

Visual Z. Every neighborhood of $p$ in $Y$ carries infinitely many classes detected in the ambient second homology, although the metric blow-ups at $p$ have one cone limit. The right-hand surface draws only the two-dimensional transverse cone factor, using $q = 0.6$ illustratively; the full tangent also contains $\mathbb{R}^2$. An open cone over a compact space is contractible and cannot model such a neighborhood.

For the splitting strata use

$$\mathcal{S}_k(Y) = \{y \,:\, \text{no tangent at } y \text{ splits } \mathbb{R}^{k+1}\}.$$

Every point other than $p$ has Euclidean tangent; at $p$ the unique tangent splits $\mathbb{R}^2$. Therefore $\mathcal{S}_1(Y) = \emptyset$, and $\mathcal{S}_0^{\mathrm{CC}}(Y) \subset \mathcal{S}_0(Y) \subset \mathcal{S}_1(Y)$ is empty as well. The interior of the complement of that prescribed zero-stratum is all of the nonmanifold space $Y$. Deleting the different set $\{p\}$, however, leaves a smooth manifold. The isolated example concerns the prescribed stratum, whereas the positive-measure examples obstruct *every* deletion below the stated Hausdorff dimension. [M, §1.3; proof of Theorem 1.3].

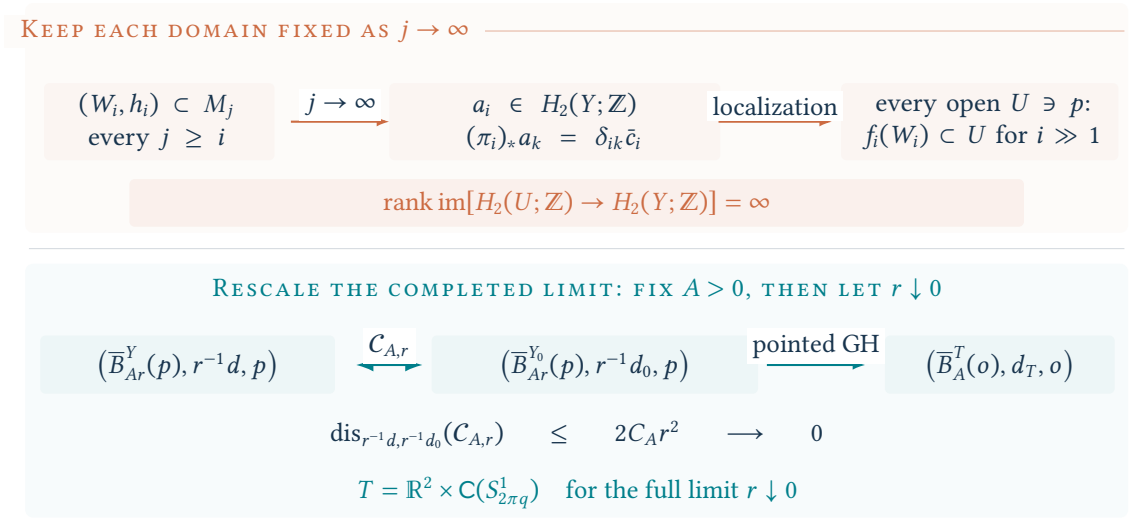



FIGURE 23. Fixing one domain and passing to the approximant limit preserves its homology detector. Computing a tangent instead starts with the completed spaces $Y, Y_0$, fixes a rescaled radius, and estimates the whole remaining tail as the scale tends to zero. The correspondence is a relation comparing distances, not an asserted continuous map between the spaces.

## 20. A GUIDE FROM THE PICTURES TO THE PROOF

**What the construction adds to the earlier mechanisms.** The comparison in [M, §1.4] attributes accumulation of second homology and an isolated nonmanifold point to earlier examples of Menguy, and ambient detection of homology also to Hupp–Naber–Wang in their different-dimensional limit setting. The distinctions emphasized here are the combinations stated in [M]: same-dimensional noncollapse, the positive-measure codimension-two obstruction with arbitrary-deletion persistence, and, separately, a compact four-dimensional isolated example with the specified metric tangent and empty prescribed zero-stratum. No claim that persistence or tangent uniqueness alone is new is needed for these explanations.

**Source map.** The labels below refer to [M]. The difficult geometric estimates remain explicitly located there; the companion spells out their hypotheses, the calculations connecting them, and the conclusions for which they are used.

| Question | Mechanism | Source in [M] |
|---|---|---|
| What detects one class? | Intrinsic boundary quotient and ambient Lipschitz collapse | Lemmas 2.1–2.2 |
| Why do classes survive? | Uniform inclusion limits, graph limits, finite coefficient detection | Propositions 2.3–2.4 |
| Why can a sphere avoid $E$? | Local inverse isometries and a submersion dimension estimate | Lemmas 2.5–2.8 |
| Why the threshold $N-2$? | Perturb the sphere in the full product domain | Proposition 2.9 |
| Why finite upper measure? | Ambient charts outside a finite-measure exceptional set | Lemma 11.2 |
| Why a uniform Ricci bound? | Parameter concavity, integrated compensation, full Schur test | Lemmas 4.1–4.2; Proposition 5.2 |
| Why a global background? | Separated fixed cutoffs and smooth endpoint conditions | Lemma 5.3 |

| Question | Mechanism | Source in [M] |
|---|---|---|
| Why exact earlier regions? | Flat scale, compensated positive floors, threshold choice | Lemmas 7.1–7.3 |
| Which cutting spheres? | Base projection and first-exit control of ambient distances | Lemma 8.1 |
| Why do boundaries match? | Fixed-$e_i$ path, Moser attaching map, spherical tuning | §6; Lemma 8.2; Proposition 8.3 |
| Why do collars persist? | Identical data, fixed smoothing choices, protected supports | Lemma 3.1; §9 |
| Why finite total curvature? | Boundary flux, bounded slice scalar integrals, summation | Lemma 3.2; §10 |
| Why is $F$ still visible? | Uniform Lipschitz marking and boundary detours | Lemmas 9.1–9.2 |
| Why the same tangent? | Absolute angular error, localized quotient chains, same-radius correspondences | §§12–13 |
| Why a smooth puncture? | Compatible exhaustion, first exit, openness, localization | Lemma 14.1 |

The two orders of limiting operations should remain separate. For topology, fix a domain and pass to the metric limit while retaining its maps; then localize infinitely many resulting classes. For a tangent, fix an observation radius after rescaling and bound the combined metric effect of all the small domains. Their homology need not vanish when that metric error tends to zero.

## References

[M] A. Nifa, *Homological obstructions to local contractibility in noncollapsed Ricci limits*, preprint, 25 August 2026. doi: 10.5281/zenodo.23049905.

Toulouse, France
*Email address*: nifa.anass@gmail.com