Introduction to Riemannian convergence theory

Simon Donaldson[1]

[1]Imperial College

March 25, 2022

We want to study limits of sequences of Riemannian manifolds (M_i, g_i), in various senses.

There is one obvious notion. On a fixed smooth manifold M, a Riemannian metric is a section of s^2T^*M so we can consider $g_i \to g_\infty$, as sections of this bundle, in standard functions spaces ($C^\infty_{\rm loc}, C^{k,\alpha}, L^\infty \dots$).

But the questions are more involved because:

- The diffeomorphisms of M act on the metrics.
- We need to consider cases where the limit is a structure on a different manifold, or singular space.

Section 1: Hyperbolic surfaces: the Deligne-Mumford moduli space

We consider a compact oriented 2-dimensional manifold Σ with a metric g of constant Gauss curvature -1.

The universal cover is isometric to the hyperbolic plane.

Taking the upper half-space model H, the (oriented) isometry group of H is $PSL(2, \mathbf{R})$, acting by Möbius maps.

The isometries of H act simply transitively on unit tangent vectors.

Any local isometry of H extends uniquely to a global one.

We have $\Sigma = H/\Gamma$ where $\Gamma \subset PSL(2, \mathbf{R})$.

For each point $p \in \Sigma$ we have the exponential map

$$\exp_p : T\Sigma_p \to \Sigma.$$

The *injectivity radius* $\mathrm{inj}(\Sigma)$ is the largest number r_0 such that, for all p, the map $\exp_p$ is injective on the open r_0 disc in the tangent space.

Theorem 1 Let (Σ_i, g_i) be a sequence of such Riemannian manifolds with fixed genus $\gamma \geq 2$. Suppose that $\mathrm{inj}(\Sigma_i) \geq \rho$ for some fixed ρ. Then there is a subsequence $\{i'\}$, a compact hyperbolic surface $(\Sigma_\infty, g_\infty)$ and diffeomorphisms $\phi_{i'} : \Sigma_\infty \to \Sigma_{i'}$ such that $\phi_{i'}^*(g_{i'}) \to g_\infty$ in C^∞.

First observation : by Gauss-Bonnet the area of Σ_i is $4\pi(\gamma - 1)$.

We take the point of view that a hyperbolic surface is given by gluing together a collection of discs in H by isometries in $PSL(2, \mathbf{R})$. (There are many other approaches to proving Theorem 1.)

Given Σ as above, choose a maximal set of points $p_1, \ldots p_N$ such that $d(p_\alpha, p_\beta) > \rho/10$. So any point $q \in \Sigma$ is within distance $\rho/10$ of some p_α. Thus the discs with centre p_α and radius $\rho/5$ cover Σ, while the discs with same centres and radii $\rho/20$ are disjoint. These latter discs have the same fixed area, determined by ρ.

So we get a bound on the number of discs in terms of ρ *and the genus* γ*.* : $N \leq N(\rho, \gamma)$
This obviously implies a bound on the *diameter* of Σ,

$$\mathrm{Diam}(\Sigma) \leq D.$$

Going back to the sequence in the Theorem, for each i take some cover of Σ_i by $N = N(\rho, \gamma)$ discs of fixed radius $\rho/5$ with centres $p_{\alpha,i} \in \Sigma_i$.
Let $D_{\alpha,i}$ be the open $\rho/4$ disc in Σ_i with centre $p_{\alpha,i}$.

Let $d_{\alpha,\beta,i}$ be the distance between $p_{\alpha,i}, p_{\beta,i}$ in the hyperbolic surface Σ_i.

Since $0 \leq d_{\alpha,\beta,i} \leq D$, we can suppose (passing to a subsequence) that these converge: $d_{\alpha,\beta,i} \to d_{\alpha,\beta}$.

So the discs $D_{\alpha,i}$ cover Σ_i and moreover each point of Σ_i lies well inside at least one $D_{\alpha,i}$.
Also, when a pair of these discs intersect the intersection is a connected set.

Fix a standard disc D of radius $\rho/4$ in H.
Second observation. The set of $F \in PSL(2, \mathbf{R})$ such that $F(D) \cap D$ is non-empty is *compact*.
Let

$$A = \{(\alpha, \beta) : d(\alpha, \beta) < \rho/2\}.$$

Without loss of generality we can the same set A if we replace $d(\alpha, \beta)$ by $d(\alpha, \beta, i)$ for any i.

For each Σ_i we have an atlas of isometric charts $\chi_{\alpha,i} : D \to D_{\alpha,i} \subset \Sigma_i$.

When $(\alpha, \beta) \in A$ the discs $D_{\alpha,i}, D_{\beta,i}$ intersect and we have an overlap map $f_{\alpha,\beta,i}$ defined on a subset of the standard disc D, such that

$$\chi_{\alpha,i} = \chi_{\beta,i} \circ f_{\alpha,\beta,i}.$$

The fact that the intersections are connected implies that $f_{\alpha,\beta,i}$ is the restriction of a global isometry $F_{\alpha,\beta,i} \in PSL(2, \mathbf{R})$.
So by the observation we can suppose that these have a limit $F_{\alpha,\beta}$ as $i \to \infty$.

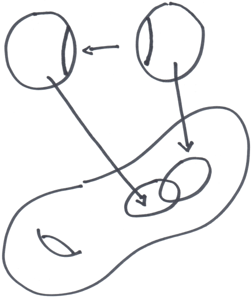

We construct a limiting surface as follows.

Start with the product $X = D \times \{1, \dots, N\}$ and define a relation by $(\alpha, z) \sim (\beta, w)$ if $(\alpha, \beta) \in A$ and $w = F_{\alpha,\beta}(z)$.

The compatability condition for the charts $\chi_{\alpha,i}$ in Σ_i implies, passing to the limit, that this is an equivalence relation. Let Σ_∞ be the quotient $X/\sim$.

The inclusion $D \times \{\alpha\} \subset X$ induces a map $\chi_\alpha : D \to \Sigma_\infty$.

One checks that these form a system of charts defining a compact hyperbolic structure on Σ_∞.

To get the statement of the Theorem we can use a general results about manifolds, which we state imprecisely

Let M be a compact n-dimensional manifold defined by an atlas of charts $\chi_\alpha : B^n \to M$ with overlap maps $f_{\alpha\beta} : U_{\alpha\beta} \to U_{\beta\alpha}$ where

$$U_{\alpha\beta} = \chi_\alpha^{-1} \left(\chi_\alpha(B^n) \cap \chi_\beta(B^n) \right) .$$

Let M' be another such manifold defined by data $\chi'_\alpha, f'_{\alpha\beta}$ with (for simplicity) the same domains $U_{\alpha\beta}$.

Proposition

If $f'_{\alpha\beta}$ is sufficiently close in C^∞ to $f_{\alpha\beta}$ then, after perhaps slightly shrinking the charts, we can find diffeomorphisms $g_\alpha : B^n \to B^n$ close to the identity such that $f'_{\alpha\beta} = g_\beta^{-1} \circ f_{\alpha\beta} \circ g_\alpha$.

This collection g_α is just the data needed to define a diffeomorphism $g : M \to M'$.

The proposition is a nonlinear version of the fact that $H^1(M, \mathbf{T}) = 0$, where $\mathbf{T}$ is the sheaf of C^∞ vector fields.

Theorem 1 is of interest in Riemann surface theory because the moduli space of genus γ hyperbolic surfaces modulo diffeomorphism can be identified with the moduli space $\mathcal{M}_\gamma$ of genus γ Riemann surfaces.

We are interested in compactifying this moduli space.

What happens in a sequence Σ_i with $\mathrm{inj}(\Sigma_i) \to 0$?

For $l > 0$ define the "standard collar" C_l to be the quotient of H identifying z with $e^l z$. The imaginary axis projects to a closed geodesic γ_l of length l in C_l.

Define the standard cusp E to be the quotient of H identifying z with $z+1$.

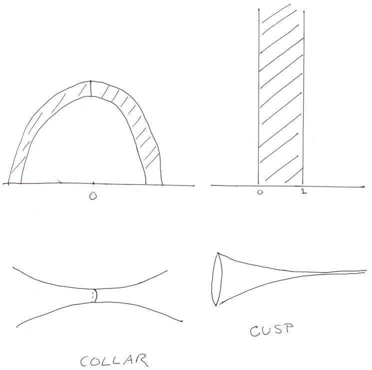
0
0
1
CUSP
COLLAR

The injectivity radius $\mathrm{inj}(\Sigma)$ is half the length of the shortest closed geodesic in Σ.

Suppose that we have a closed geodesic $\Gamma \subset \Sigma$, of length l. A neighbourhood of Γ in Σ is isometric to a neighbourhood of γ_l in C_l. More generally, there is a locally isometric covering map $C_l \to \Sigma$ taking γ_l to Γ.
Clearly, for fixed R the area of the R-neighbourhood of γ_l in C_l tends to 0 as $l \to 0$.

In our situation, if $\mathrm{inj}(\Sigma_i) \to 0$ we must have $\mathrm{Diam}(\Sigma_i) \to \infty$.

The converse is also true. as we saw.

When i is large the surface Σ_i contains a long "thin" region, isometric to part of the standard collar C_l, with l small.
The natural (subsequential) "limit"as $i \to \infty$ is a complete non-compact (possibly disconnected surface) Σ_∞.
The surface Σ_∞ has ends which are cusps, modelled on E.

This can be used to define a compactification $\overline{M}_\gamma$, which agrees with the algebraic geometers Deligne-Mumford compactification via stable nodal curves.

Section 2: bounded local geometry

Recall that at each point p in a Riemannian manifold there is a Riemann curvature tensor which is equivalent to the sectional curvature function on 2-planes in TM_p.

The sectional curvature gives strong control of geodesics, hence of the exponential map and in turn of the geometry, expressed in geodesic coordinates.

For any $\kappa \in \mathbf{R}$ let $N(\kappa, n)$ be the simply connected n-manifold of constant sectional curvature κ, with base point O.

The basic comparison theorem (somewhat roughly stated): Suppose that the sectional curvatures K of a complete manifold (M, g) satisfy $\kappa_1 \leq K \leq \kappa_2$. Fix $p \in M$ and an isometry $TM_p = TN_O$.

Compare M with the $N(\kappa, n)$, pulling back by the exponential maps from p and 0.

Then the metric on M is *smaller* than $N(\kappa_1, n)$ and *bigger* than $N(\kappa_2, n)$.

(More precisely, if $\kappa_i > 0$ we should restrict to points with distance $\leq \pi\sqrt{\kappa_i}$ from p.)

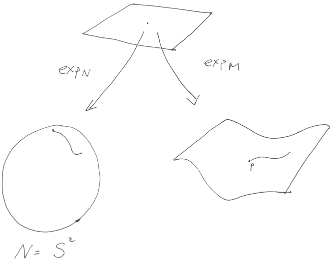
$\exp_N$
$\exp_M$
p
$N = S^2$

Fundamental fact Under the scaling $g \mapsto \lambda^2 g$ the curvature multiplies by λ^{-2}.

If f $r << \mathrm{inj}(M)$ and $r << \sup|\mathrm{Riem}|^{1/2}$ then the geometry of M on the scale of r is very close to Euclidean.

Theorem 2(Cheeger)
Fix $r_0, V, C > 0$. If (M_i, g_i) are compact n-dimensional Riemannian manifolds with $\mathrm{inj} \geq r_0, \mathrm{Vol} \leq V, |\mathrm{Riem}| \leq C$ then there is a subsequence which converges in $C^{1,\nu}$ to a $C^{1,\nu}$ limiting metric g_∞ on a manifold M.

Explanation:

- Here convergence is in the same sense as Theorem 1, via maps $\phi_{i'} : M \to M_{i'}$.
- $C^{1,\nu}$ refers to the usual Hölder space, with any exponent $\nu < 1$.
- Strictly, M is a $C^{2,\nu}$ manifold, but it is known that any such has a smooth structure.
- A $C^{1,\nu}$ metric has Christoffel symbols in $C^{,\nu}$ so we have a good theory of geodesics.
- We could also say that we have weak convergence in $L^p_{2,\mathrm{loc}}$ for any $p < \infty$. The limiting metric has a curvature tensor in L^p.

Equivalent hypotheses in Theorem 2 are:
$\mathrm{Diam} \leq D, \mathrm{Volume} \geq V', |\mathrm{Riem}| \leq C$.

A proof of Theorem 2 can follow the same strategy as we used for Theorem 1. We need to:

1. Choose charts so that the metric tensors converge in these charts.
2. Arrange that the overlap maps between the charts converge.

(There is an extra complication involving the domains of the overlap maps, but this can be dealt with.)

As a first attempt at a proof of Theorem 2 we could use use geodesic charts on small metric balls and follow the plan of proof of Theorem 1.
A bound on the curvature tensor gives a bound on the metric tensor in geodesic co-ordinates. This gives a C^1 bound on the overlap maps.
The compact inclusion $C^1 \to C^{,\nu}$ means that we can suppose that the overlap maps converge in $C^{,\nu}$.

We get a $C^{,\nu}$ limiting manifold M but this is two derivatives less than what we want, and it is not clear if there is a useful Riemannian limit.

To prove Theorem 2 we need a better choice of local coordinates. One approach uses harmonic coordinates, which we discuss in Section 6. Here we will outline a "gauge theory" approach.

In general the problem can be seen as seeking a "quantitative" version of a standard differential geometric result.
Standard Theorem: $\mathrm{Riem} = 0$ implies $\mathrm{Local\ \ Euclidean\ coordinates}$ $g_{ij} = \delta_{ij}$

Quantitative version: $\mathrm{Riem\ small}$ implies local co-ordinates with $g_{ij} - \delta_{ij}$ small.

Let g be a metric on the unit ball B^n with $|g_{ij} - \delta_{ij}| < \epsilon$ and $|\mathrm{Riem}| < \epsilon$. Regard the Levi-Civita connection of g as an $SO(n)$ connection ∇_g on a bundle over B. Parallel transport along rays gives an orthonormal frame of 1-forms η_i with $|\nabla_g \eta_i| \leq c\epsilon$. This implies that $|d\eta_i| \leq c\epsilon$ and so

$$\|d\eta_i\|_{L^p} \leq c\epsilon.$$

The usual proof of the Poincaré Lemma constructs a chain homotopy operator on differential forms over the ball $T_0 : \Omega^p \to \Omega^{p+1}$ such that $\eta = dT_0\eta + T_0 d\eta$. This construction is based on radial dilation about the origin. The operator T_0 does not behave very well on function spaces: it takes L^p to L^p. Averaging over different choices of origin (and being somewhat vague about the domain of definition of the forms) we get a better chain homotopy T which maps L^p to L^p_1, by elliptic theory.

,

In the situation above we get (on a slightly smaller ball) functions h_i such that $\eta_i = dh_i + \theta_i$ where $\|\theta_i\|_{L^p_1} \leq c\epsilon$. Write $G_{ij} = (dh_i, dh_j)$ then we have

$$\|G_{ij} - \delta_{ij}\|_{L^p_1} \leq c\epsilon.$$

Now use h_i as new coordinates. In these coordinates the metric tensor differs from Euclidean in L^p_1 norm from Euclidean by at most $c\epsilon$.

This has gained one derivative over the first attempt. To gain one more derivative one can apply a well-known theorem of Uhlenbeck to the connection ∇_g (once ϵ is small). This gives an orthonormal frame η_i with $\nabla_g \eta_i$ small in L^p_1, and we get the similar conclusion, but now with the metric tensor controlled in L^p_2.

For large enough p, this gives control in $C^{1,\nu}$, by Sobolev embedding theorems.

Uhlenbeck's theorem can be seen as a nonlinear version of the preceding discussion.

Section 3: Gromov-Hausdorff convergence

Let A, B be compact metric spaces. The Gromov-Hausdorff distance $D_{GH}(A,B)$ is the infimum of ϵ such that there is a metric d_ϵ on the disjoint union $A \sqcup B$ extending the given metrics on A, B and such that both A, B are ϵ-dense.

Given d_ϵ we define for $a \in A$ a set $N_\epsilon(a) \subset B$ of points within ϵ of a. Then for $b_1 \in N_\epsilon(a_1), b_2 \in N_\epsilon(a_2)$ we have

$$|d(b_1, b_2) - d(a_1, a_2)| \leq 2\epsilon.$$

Similarly for $b \in B$ we have $N_\epsilon(b) \subset A$.

On scales $>> \epsilon$ the metrics on A, B are almost equivalent.

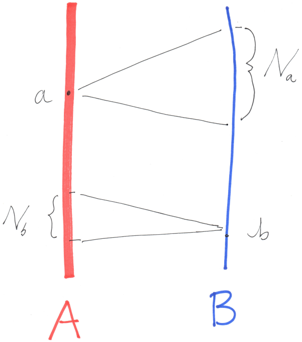
N_a
a
N_b
b
A
B

Theorem 3 (Gromov's compactness theorem).
Let X_i be a sequence of compact metric spaces with diameters $\leq D$. Suppose that for each $\eta > 0$ there is a number $N(\eta)$ such that each X_i can be covered by $N(\eta)$ η-balls. Then there is a subsequence which converges in the Gromov-Hausdorff metric.

This is analogous to the fact that a bounded sequence in a separable Hilbert space has a weakly convergent subsequence.

The proof of Gromov's Theorem is "elementary". The general idea is that if the X_i are *finite* sets with at most a fixed number of points then the result is clear. The hypothesis in the theorem means that for all ϵ each X_i can be approximated to within distance ϵ in D_{GH} by a finite set of fixed size.

Based convergence. Let X_i be locally compact metric spaces and $x_i \in X_i$. The pairs (X_i, x_i) have a *based GH limit* (X_∞, x_∞) if for each R the sequence of R-balls centred at x_i converge to that centred at x_∞.

Example Let Σ_i be a sequence of hyperbolic surfaces splitting into a union of two infinite-diameter pieces Σ_I, Σ_{II} by developing a long neck, as we discussed in Section 1. For different sequences of base points $p_i \in \Sigma_i$ we can arrange that the based limit is:

- (Σ_I, p_I) for some $p_I \in \Sigma_I$;
- (Σ_{II}, p_{II}) for some $p_{II} \in \Sigma_{II}$;
- $(\mathbf{R}, 0)$

In the last case the dimension drops in the limit.

We can also consider *scaled limits*. For a sequence of real numbers $\lambda_i > 0$ and based metric spaces $((X_i, \lambda_i d_{X_i}, x_i)$.
In the example above we can get additional scaled limits (for different choices of λ_i, p_i):

- $\mathbf{R}^2$,
- One point.
- $S^1 \times \mathbf{R}$.

Section 4: collapsing with bounded curvature and almost-flat manifolds

What can happen if we drop the injectivity radius assumption in Theorem 2? We clearly do not necessarily have convergence in the same sense. For example take $M_i = Y \times T^m_{\epsilon_i}$ where $T^m_{\epsilon_i}$ is a flat torus scaled to have diameter ϵ_i and $\epsilon_i \to 0$. In this section we look at some variations on this phenomenon.

Let P be the circle bundle over $\mathbf{R}^2$ with a connection having curvature $\omega = dx_1 dx_2$. Write V for the "vertical" vector field on P which generates the S^1 action and X_1, X_2 for the horizontal lifts of $\frac{\partial}{\partial x_1}, \frac{\partial}{\partial x_2}$. So X_1, X_2, V form a frame for the tangent bundle and the only non-vanishing Lie bracket between them is $[X_1, X_2] = V$.

For a fixed $\epsilon > 0$ define $X_0 = \epsilon^{-1} V$ so the only bracket is $[X_1, X_2] = \epsilon X_0$.
Define a Riemannian metric g_ϵ on the 3-manifold P by declaring that X_0, X_1, X_2 form an orthonormal frame. The formulae for the Levi-Civita connection and curvature show that the curvature of (P, g_ϵ) is $O(\epsilon^2)$.

The length of a circle fibre in the metric g_ϵ is $2\pi\epsilon$.

P carries a group structure, the *Heisenberg group*, such that the vector fields X_i are left-invariant. As a manifold we can identify $P = \mathbf{R}^2 \times S^1$ and the product is

$$(v, \lambda)(v', \lambda') = (v + v', \lambda\lambda' \exp(i\omega(v, v')).$$

Let L be a lattice in $\mathbf{R}^2$ with a fundamental domain of area 2π, for example $\mathbf{Z} \times 2\pi\mathbf{Z}$. Then L is a subgroup of P and the quotient $N = P/L$ is a circle bundle over a 2-torus $T^2 = \mathbf{R}^2/L$ with first Chern class 1.

The metric g_ϵ descends to N and clearly for ϵ small the diameter of N is essentially the same as that of T^2. We get a family of metrics on N with curvature tending to zero but diameter bounded below.

We have $H^1(N, \mathbf{R}) = \mathbf{R}^2$, so N is not homeomorphic to a 3-torus.

N is an example of an "almost-flat" manifold.

Take another parameter η and write $g_{\epsilon,\eta} = \eta^2 g_\epsilon$. Then $(N, g_{\epsilon,\eta})$ has diameter $O(\eta)$ and curvature $O(\eta^{-2}\epsilon^2)$,
Consider a Riemannian product $Y \times N$ with a fixed metric on Y and a sequence of metrics g_{ϵ_i,η_i} on N. If we choose the parameters so that $\eta_i \to 0$ but $\eta_i^{-2}\epsilon_i^2$ is bounded we get another example of a sequence of metrics with bounded curvature collapsing to a lower dimensional manifold Y (which is the Gromov-Hausdorff limit).

There are two different collapsing scales: η_i ("in the T^2 direction") and the much smaller $\eta_i \epsilon_i$ ("in the S^1 direction").

More generally we can construct fibrations

$$N \to M \to Y,$$

for example we could vary the lattice L as we move around Y.

There is a theory due to Cheeger, Gromov, Fukaya. . . of "F-structures" and "N-structures which, *very* roughly speaking, says that this is the general picture, for collapsing sequences with bounded curvature.

Example: graph manifolds

Let Γ_0 be the union of three line segments forming the letter Y. Let Z_0 be the 2-sphere with three discs removed and $\pi_0 : Z_0 \to \Gamma_0$ a map of the obvious kind taking the boundary circles to the end points.

Given a trivalent graph Γ we define a (family of) 3-manifolds M_Γ as follows. For each vertex take a copy of $Z_0 \times S^1$ and for each edge take a copy of $T^2 \times [0, 1]$.

Each boundary component of $Z_0 \times S^1$ is a 2-torus so we can glue these components together to get a compact 3-manifold M_Γ with a map $\pi : M_\Gamma \to \Gamma$.

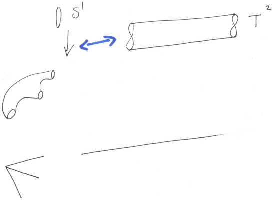
S^1
T^2

There are choices involved in gluing the T^2 boundaries and for generic choices there will be no free S^1 action on M_Γ, although such actions exist locally.

We can choose a sequence of metrics on M_Γ with bounded curvature such that based Gromov-Hausdorff limits are either $\mathbf{R}$ or Z_0 with its infinite-diameter hyperbolic metric.

These kind of metrics arise in the analysis of the Ricci flow on 3-manifolds and Perelman's proof of Thurston's Geometrisation conjecture.

The Ricci tensor (or Ricci curvature) of a Riemannian manifold (M, g) is a contraction of the full curvature tensor. It is a section of $S^2(T^*M)$, thus there is a notion of *positive* Ricci curvature etc.

The Ricci tensor is strongly related to the volume form vol_g of the metric. In geodesic coordinates about a point

$$\mathrm{vol}_g = \left(1 - \frac{1}{3}\sum R_{ij}x_i x_j + O(x^3)\right) \ dx_1 \dots dx_n.$$

So positive/negative Ricci curvature makes volumes smaller/larger in these coordinates (for sufficiently small x).

Digression

In complex differential geometry, a volume form on a complex manifold M is the equivalent to a Hermitian metric on the canonical line bundle $K_M = \Lambda^m T^*M$ (where m is the complex dimension).

This defines a connection on the line bundle with a curvature form $i\rho$ of type $(1,1)$.

In local complex co-ordinates $z_a = x_1 + iy_a$, if

$$|dz_1 \dots dz_m|^2 = h$$

then the volume form is

$$h^{-2} dx_1 dy_1 \dots dx_m dy_m$$

and the curvature form is

$$\rho = -i\overline{\partial}\partial \log h.$$

If M has a Kähler metric g and we use the volume form of g the tensor ρ is the same as the Ricci tensor, up to normalising factor.

Positive Ricci curvature $\Leftrightarrow K_M^{-1}$ ample $\Leftrightarrow$ Fano;

Negative Ricci curvature $\Leftrightarrow K_M$ ample $\Rightarrow$ General type.

The Ricci curvature is also strongly related to the Laplace operator.

The *Bochner identity* , for a 1-form α on a Riemannian manifold (M, g):

$$\Delta\alpha = \nabla^*\nabla\ \alpha + \mathrm{Ricci}.\alpha,$$

where $\Delta = d^*d + dd^*$ and

$$\nabla^* : \Gamma(T^*M) \otimes T^*M) \to \Gamma(T^*M)$$

is the formal adjoint of the covariant derivative

$$\nabla : \Gamma(T^*M) \to \Gamma(T^*M \otimes T^*M).$$

We will discuss two results about Ricci curvature:

- volume comparison (Bishop)
- splitting theorem (Cheeger-Gromoll)

A fundamental differential inequality

Let (M, g) be a Riemannian n-manifold with $\mathrm{Ricci} \geq 0$ and f be a function on M with $|df| = 1$. Write v for the gradient vector field of f. Then

$$\nabla_v(\Delta f) \geq |\nabla\nabla f|^2 \geq \frac{1}{n-1}(\Delta f)^2 . \qquad (**)$$

(In these lectures we use the "geometers" sign convention for the Laplacian $\Delta = d^* d$ on functions.)

This is a consequence of the Bochner identity.

Take $\alpha = df$ so $\Delta\alpha = d\Delta f$ and $\nabla_v \Delta f = (\Delta\alpha, \alpha)$. Since $\mathrm{Ricci} \geq 0$ we have $\nabla_v \Delta f \geq (\nabla^* \nabla \alpha, \alpha)$.
Now for any 1-form α

$$\frac{1}{2}\Delta|\alpha|^2 = (\nabla^*\nabla\alpha, \alpha) - |\nabla\alpha|^2.$$

In our case $|\alpha| = 1$ so $(\nabla^*\nabla\alpha, \alpha) = |\nabla\alpha|^2$ and we get

$$\nabla_v \Delta f \geq |\nabla\nabla f|^2.$$

For the second inequality in (**) notice that $|df| = 1$ implies that $\nabla_v v = 0$. (*i.e.* the integral curves of the vector field v are geodesics). This means v is in the kernel of the Hessian $\nabla\nabla f$. The last inequality follows from the fact that for a symmetric matrix H of rank $\leq n-1$ we have $\sum H_{ij}^2 \geq \frac{1}{n-1}\mathrm{Trace}(H)^2$.

Geometric significance The orthogonal complements of v are the tangent space of the level sets of f. The Hessian $\nabla\nabla f$ is the *second fundamental form* of the level set and $-\Delta f$ is the *mean curvature*.

Important supplement: if equality holds in (**) then the second fundamental form of the levels set is a multiple of the metric.

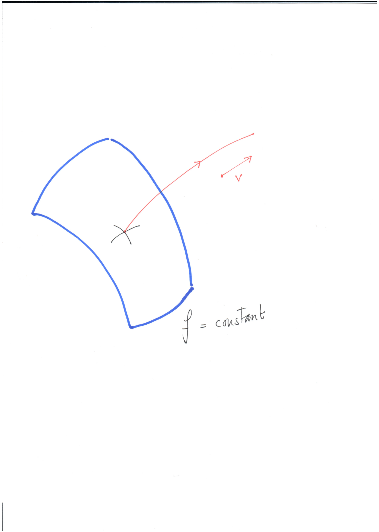
v
f = constant

Bishop comparison theorem

Let (M,g) be a complete Riemannian n-manifold with $\mathrm{Ricci} \geq 0$ and p be a point in M. Write $V(r)$ for the volume of the metric ball of radius r centred at p. Then $V(r)/r^n$ is a decreasing function of r.

(In particular, $V(r)$ is at most the volume of the Euclidean r-ball.)

There is a more general theorem under the hypothesis $\mathrm{Ricci} \geq \lambda$.

Outline proof.

For exposition, consider first the case when the metric has rotational symmetry about p and we take radius r less than the injectivity radius. Then $\frac{dV}{dr} = A(r)$ where $A(r)$ is the volume of the boundary of the r-ball. Let f be the distance to p, so $|df| = 1$. By considering the flux of the radial vector field, or otherwise, one sees that

$$\Delta f = -A^{-1}\frac{dA}{dr}.$$

Write $m(r) = -\Delta f$ and apply the inequality (**)

$$\frac{dm}{dr} \leq -\frac{1}{n-1}m^2.$$

This integrates to give $m \leq (n-1)/r$ which implies that A/r^{n-1} is decreasing and, integrating again, one sees that V/r^n is decreasing.

For the general case, take polar coordinates (r, σ) on $\mathbf{R}^n$, with $\sigma \in S^{n-1}$ and write the pull-back of the volume form by the exponential map as $A(r, \sigma) dr d\sigma$. Let $R(\sigma)$ be the largest distance such that the geodesic in the direction σ minimises up to distance $R(\sigma)$. The same arguments apply to $A(r, \sigma)$ for each fixed σ and for $r \leq R(\sigma)$.

The exponential map at p maps the union of these ray-segments *onto* M.

Important supplement If $s < r$ and $A(s)/s^{n-1} = A(r)/r^{n-1}$ then the metric is "conical" in the annulus $B(r) \setminus B(s)$ (see later).

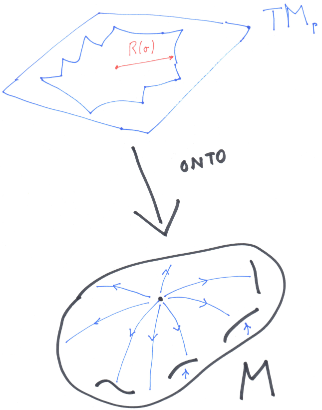

TM_p
$R(\sigma)$
ONTO
M

A geodesic $\gamma : \mathbf{R} \to M$ in a complete Riemannian manifold is called a "line" if $d(\gamma(a), \gamma(b)) = |a - b|$ for all $a, b \in \mathbf{R}$.

The splitting theorem
If M contains a line and $\mathrm{Ricci} \geq 0$ then M is a Riemannian product $\mathbf{R} \times N$.

Lemma If M is complete with $\mathrm{Ricci} \geq 0$ and if there is a smooth function $\phi : M \to \mathbf{R}$ with $|\nabla \phi| = 1$ then $M = N \times \mathbf{R}$.

Proof Let $c : \mathbf{R} \to M$ be an integral curve of the gradient vector field v and take $m(t) = \Delta\phi(c(t))$. So we get
$m'(t) \geq m^2/(n-1)$. It is an exercise to show that the only solution of this inequality defined for all t is $m = 0$.
The first inequality in (**) implies that $\nabla\nabla\phi = 0$ from which one easily sees that ϕ is the projection to $\mathbf{R}$ in a Riemannian product $M = \mathbf{R} \times N$.

The *Busemann function* associated to the line γ is

$$B(p) = \lim_{s \to \infty} d(\gamma(s), p) - s.$$

This function is not smooth *a priori*. The idea of (one) proof of the splitting theorem is to show that it is differentiable with $|\nabla B| = 1$ and to adapt the proof of the Lemma to this situation.

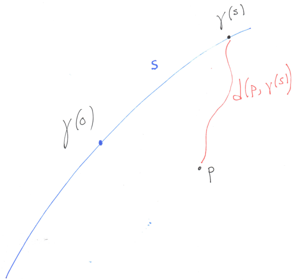
$\gamma(s)$
s
$d(p, \gamma(s))$
$\gamma(0)$
p

Section 6: Harmonic coordinates

Consider a local co-ordinate system x_i on a ball B in a manifold (M, g) with $\Delta_g x_i = 0$. Let $\eta_i = dx_i$. Then

$$\Delta(\eta_i, \eta_j) = (\nabla^* \nabla \eta_i, \eta_j + (\eta_i, \nabla^* \nabla \eta_j) + 2(\nabla \eta_i, \nabla \eta_j).$$

Since $\Delta \eta_i = 0$ the Bochner formula gives

$$\Delta(\eta_i, \eta_j) = 2(\mathrm{Ricci} \eta_i, \eta_j) + 2(\nabla \eta_i, \nabla \eta_j).$$

The covariant derivatives $\nabla \eta_i$ are given by the Christoffel symbols which depend on one derivative of the metric tensor so we get, schematically,

$$\Delta g^{ij} = R^{ij} + Q_{ij}(\partial g, g), \quad (****).$$

Fix p large. Suppose we have bounds

- $$C^{-1}\left(\delta_{ij}\right) \leq \left(g_{ij}\right) \leq C(\delta_{ij}),$$
- $$\|\partial g\|_{L^{2p}} \leq c$$

Then elliptic theory applied to (****) gives a bound on the L^p_2 norm of g_{ij} over an interior ball in terms of C, c and an L^p bound on the Ricci curvature.

The *harmonic radius* $r_h(M, g)$ of a compact Riemannian manifold is the supremum of numbers ρ such that there are harmonic coordinates satisfying the estimates above on every unit ball in the rescaled metric $(M, \rho^{-2}g)$.

In other words we get “good” coordinates on all balls of radius $\rho < r_h$ in the original metric.

Theorem 4 (Anderson)
Let (M, g) be a compact Riemannian n manifold with $|\mathrm{Ricci}| \leq \Lambda$ and $\mathrm{inj}(M) \geq \delta > 0$. Then $r_h(M, g) \geq \epsilon$ for some $\epsilon > 0$ depending only on n, Λ, δ.

A Corollary is that we can extend Theorem 2 to manifolds with a bound on $|\mathrm{Ricci}|$, rather than the full $|\mathrm{Riem}|$.

The proof of Theorem 4 goes by contradiction.

First show by elliptic theory that for a metric on the unit ball in $\mathbf{R}^n$ which is sufficiently close in L^p_1 to the Euclidean metric there are good harmonic coordinates.

Now if the statement is false there is a sequence of Riemannian manifolds (M_i, G_i) with $\mathrm{inj}(M_i) \to \infty$, $|\mathrm{Ricci}| \to 0$ and $r_h(M_i, G_i) = 1$.

Choose $p_i \in M_i$ such that **there are no good coordinates on the ball of radius** 2 **centred at** p_i**.**

We can suppose that there is a based limit $(M_\infty, G_\infty, p_\infty)$ with metrics converging strongly in L^p_2 on bounded subsets.

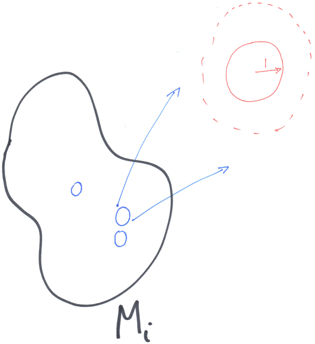
1
M_i

The limit has infinite injectivity radius and $\mathrm{Ricci} = 0$; it is a smooth metric by elliptic regularity.

If we know that $M_\infty = \mathbf{R}^n$, with Euclidean metric, we get a contradiction to the statement in bold above.

The fact that $M_\infty = \mathbf{R}^n$ can be established using the splitting theorem.

There is a related theorem of Anderson which, for simplicity, we state in the case $\mathrm{Ricci} = 0$.

Theorem 4' There is as a $\delta > 0$ such that if the volume of a unit ball of a metric with $\mathrm{Ricci} = 0$ is at least $(1 - \delta)$ times the volume of the Euclidean ball then there are coordinates on the half-sized ball in which the metric satisfies C^∞ estimates.

Section 7: Convergence under Ricci curvature bounds—theory

Theorem 5 (Gromov) Any sequence (M_i, g_i) of compact Riemannian n-manifolds with $\mathrm{Ricci} \geq \Lambda$, and diameter $\leq D$ for some Λ, D has a Gromov-Hausdorff convergent subsequence.

For simplicity suppose $\Lambda = 0$. Let M be one of the M_i and write $V = \mathrm{Vol}(M)$. Given $\epsilon > 0$, choose a maximal collection of disjoint $\epsilon/10$ balls B_α in M. The Bishop inequality gives

$$\mathrm{Vol}(B_\alpha) \geq \left(\frac{\epsilon}{10D}\right)^n V.$$

If the number of balls is N we get

$$N \left(\frac{\epsilon}{10D}\right)^n V \leq V,$$

which gives an upper bound on N. Then the statement follows from Theorem 3.

There is a similar compactness result for based limits (without diameter bound).

A sequence of based Riemannian manifolds $(M_i, g_i, p_i$ with $\mathrm{Ricci} \geq \Lambda$ is called *non-collapsing* if there is some $\delta > 0$ such that the volume of the unit ball in M_i centred at p_i is $\geq \delta$.

By the Bishop comparison theorem it is equivalent to take balls of any fixed radius.

There is a good structure theory, due to Cheeger, Colding, Naber... for non-collapsed limits of manifolds with a lower bound on Ricci.

Let (Z, d_Z) be the limit of (M_i, g_i). Initially, Z is a metric space.

Fix $q \in Z$ and let λ_i be a sequence $\lambda_i \to \infty$. By a variant of Theorem 5, the sequence of based metric spaces $(Z, \lambda_i d_Z)$ has a Gromov-Hausdorff convergent subsequence with limit W say.

W can also be obtained as a rescaled based limit of a subsequence of the M_i.

Such a space W is called a *tangent cone* of Z at q. (These are not always unique.)

For a point $w \in W$ and sequence $\mu_i \to \infty$ we can find a subsequential based limit of the rescaled metrics on W—an *iterated tangent cone* of Z. Then we can apply the same process again ....

Any iterated tangent cone can be obtained as a rescaled based limit of a subsequence of the original manifolds M_i.

Cones

Let (V, g_V) be a compact Riemannian manifold. The cone on V is the completion of the Riemannian metric $dr^2 + r^2 g_V$ on $(0, \infty) \times V$. Taking the completion adds a vertex point $r = 0$.

If (V, d_V) is a general metric space the cone can be defined, as a metric space, by the cosine formula:

$$d((v_1, r_1), (v_2, r_2))^2 = r_1^2 + r_2^2 - 2r_1 r_2 \cos\theta,$$

where $\theta = \min(d_V(v_1, v_2), \pi)$.

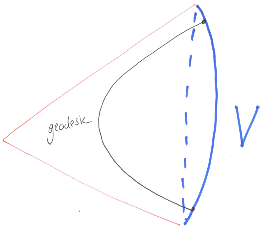
geodesic
V

Theorem 6 (Cheeger-Colding)
Any tangent cone, as above, is a metric cone.

(Also, the diameter of V is $\leq \pi$.)

Suppose that (M_i, g_i) have bounded Ricci curvature and non-collapsed Gromov-Hausdorff limit Z.
Let $R \subset Z$ be the set of points q such that some tangent cone is $\mathbf{R}^n$.
(R is the "regular set").

Theorem 6' R is an open subset of Z. It is an n-dimensional manifold and the metric on Z restricted to R is represented by a limiting $C^{1,\nu}$ (or L^p_2) Riemannian metric.

This follows from a result like Theorem 4' and some difficult facts about volumes and Gromov-Hausdorff limits.

The complement $\Sigma = Z \setminus R$ is called the *singular set* of the limit space.

We say that a tangent cone W splits off $\mathbf{R}^k$ if $W = \mathbf{R}^k \times W'$ for some W'.

For $k \leq n-1$ let $\Sigma_k \subset \Sigma$ be the set of points where some tangent cone splits off $\mathbf{R}^k$ but no tangent cone splits off $\mathbf{R}^{k+1}$.

So $\Sigma = \Sigma_{n-1} \sqcup \Sigma_{n-2} \sqcup \cdots \sqcup \Sigma_0$.

Theorem 6” (Cheeger-Colding)

$$\dim \Sigma_k \leq k$$

Here $\dim$ is the Hausdorff dimension.

In particular, R is *dense* in Z.

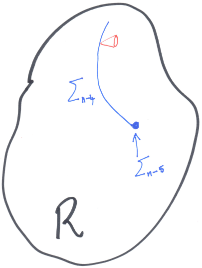
Σ_{n-4}
Σ_{n-5}
R

The proofs of Theorems 6, 6', 6" involve hard extensions of the volume comparison and splitting theorems to limits spaces, or equivalently "quantitative" versions of those theorems for Riemannian manifolds.

For example, if $p \in (M, g)$ with $\mathrm{Ricci} \geq 0$ and for some $s < r$ $r^{-n}\mathrm{Vol}(B_{p,r})$ is close to $s^{-n}\mathrm{Vol}(B_{p,s})$ then the metric is "close" to conical on the annulus $B_{p,r} \setminus B_{p,s}$.

Section 8: Convergence under Ricci curvature bounds—some examples

Hyperkähler 4-manifolds

A *hyperkähler triple* on an oriented 4-manifold X is a triple of closed 2-forms $\omega_1, \omega_2, \omega_3$ satisfying the orthogonality relations

$$\omega_i \wedge \omega_j = \delta_{ij} \mathrm{vol},$$

for some volume form vol.

Given such a triple, define a quadratic form g on tangent vectors by

$$g(v)\mathrm{vol} = \frac{1}{3} \sum i_v(\omega_a) \wedge i_v(\omega_b) \wedge \omega_c,$$

where the sum runs over cyclic permutations of (123).

Then it is a fact that g is a Riemannian metric on X with zero Ricci curvature.

Let $\beta_1, \beta_2, \beta_3$ be a standard basis of left-invariant 1-form on the the Lie group $SU(2)$ with

$$d\beta_1 = 2\beta_2 \wedge \beta_3,$$

etc.. Take a co-ordinate t on $\mathbf{R}$ (or an interval I in $\mathbf{R}$) and functions $f_1(t), f_2(t), f_3(t)$. Then define 2-forms on the product $I \times SU(2)$,

$$\omega_a = d\,(\,f(t)\beta_a\,)\,.$$

The orthogonality condition above for the volume form $2dt \wedge \beta_1 \wedge \beta_2 \wedge \beta_3$ is $2f_a f_a' = 1$, so $f_a(t) = \sqrt{t + \tau_a}$.

If all τ_a are 0 we get $|\frac{\partial}{\partial t}| = 1/(4t^{3/4})$. Take $r = t^{1/4}$ so $|\frac{\partial}{\partial r}| = 1$. We find that the metric is a cone

$$dr^2 + r^2(\beta_1^2 + \beta_2^2 + \beta_3^2).$$

Identify $SU(2)$ with S^3 (the unit quaternions). Then we see that what we have is the flat metric on $\mathbf{R}^4$.

Now take $f_1(t) = \sqrt{t+1}, f_2(t) = f_3(t) = \sqrt{t}$.

One finds that the metric is

$$\frac{1}{16t\sqrt{t+1}}dt^2 + \frac{t}{\sqrt{t+1}}\beta_1^2 + \sqrt{t+1}(\beta_2^2 + \beta_3^2).$$

Let v_1 be the left-invariant vector field dual to β_1. This generates an action of a circle by right multiplication on $SU(2)$. The quotient $SU(2)/S^1$ is $S^2 = \mathbf{CP}^1$.

As $t \to 0$ the lengths of the S^1 orbits tends to 0.

The choice of v_1 singles out a complex structure $\mathbf{R}^4 = \mathbf{C}^2$. Blow up the origin to get $\hat{\mathbf{C}}^2$, with a sphere E of self-intersection -1. We have a metric on $\hat{\mathbf{C}}^2 \setminus E$.

This does not quite extend to a smooth metric on $\hat{\mathbf{C}}^2$ because it has a cone angle 4π transverse to E.

Let X be the quotient of $\hat{\mathbf{C}}^2$ by the involution induced by $z \mapsto -z$ on $\mathbf{R}^4$.

We get a smooth metric g_X—the Eguchi-Hanson metric—on X.

For another point of view, we can see X as the cotangent bundle of the 2-sphere and consider the action of $SO(3) = SU(2)/\pm 1$. The zero section is a sphere S of self-intersection -2.

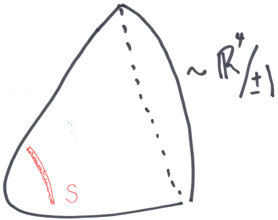
$\sim \mathbb{R}^4/\pm 1$
S

Fix a base point $x_0 \in X$ and consider the scaled metrics $\lambda^2 g_X$ with $\lambda \to 0$.

In the scaled metrics the area of S is $O(\lambda^2)$.

The based Gromov-Hausdorff limit of $(X, \lambda^2 g_X, x_0)$ is $\mathbf{R}^4 / \pm 1$, which is the cone over $S^3 / \pm 1$.

Going back to the general theory, it is known that (for non-collapsed limits with bounded Ricci curvature) $\Sigma_{n-1}, \Sigma_{n-2}, \Sigma_{n-3}$ are empty.

(This is a difficult theorem of Cheeger and Naber: in the Kähler case the proof is much easier.)

So the first situation to consider is that of “codimension 4” singularities, with tangent cone $\mathbf{R}^{n-4} \times W'$.

The 4-dimensional factor W' is the cone over a smooth Riemannian 3-manifold V with constant Ricci curvature. In dimension 3 this implies constant curvature, so $V = S^3/\Gamma$ for some finite subgroup $\Gamma \subset SO(4)$, acting freely on S^3.

Consider dimension $n = 4$ for simplicity.

There are abundant examples of sequences of 4-manifolds with bounded Ricci curvature converging to a limit with point singularities of this form.

The manifolds contain regions where the metric is well-approximated by $\lambda^2 g_X$ for small λ.

Example A smooth complex surface in $\mathbf{CP}^3$ of degree ≥ 3 has a Kähler-Einstein metric ($\mathrm{Ricci} = \lambda g$), unique up to scale.

Let S be a surface of degree ≥ 3 with one ordinary double point singularity, defined by an equation $P = 0$.

For generic p the surfaces S_t defined by equations $P + tp = 0$ are smooth, for small $t \in \mathbf{C}$, and the limit of their Kähler-Einstein metrics as $t \to 0$ has a singular point with tangent cone $\mathbf{C}^2 / \pm 1$.

Collapsing limits are more complicated.

Rough sketch of an example (Hein, Sun, Viaclovsky, Zhang 2021)

See also the article of Song Sun in Notices Amer. Math. Soc. March 2022.

Let V_1, V_2 be generic quadric surfaces in $\mathbf{CP}^3$ defined by polynomials Q_1, Q_2. Fix a generic polynomial p of degree 4; then the surface S_t defined by the equation $Q_1Q_2 + tp = 0$ are smooth for small t. They have Kähler-Einstein metrics with $\mathrm{Ricci} = 0$.

How do these metrics behave as $t \to 0$?

$V_1 \cap V_2$ is a curve of genus 1, a 2-torus T, which has a flat metric. The self-intersection number of T in V_i is 8.

The complements $V_i \setminus T$ have complete "Tian-Yau"metrics with $\mathrm{Ricci} = 0$.

The "end" of $V_i \setminus T$ is topologically $(0, \infty) \times N$, where N is an S^1 bundle over T of Chern class 8.

Taking account of orientations, these ends do *not* match up in the obvious way.

That would require gluing a bundle of Chern class 8 to one of Chern class -8

If the metrics on S_t are normalised to have diameter 1 the Gromov-Hausdorff limit is the interval $[0, 1]$.
A suitable choice of base points and scale factors gives a based, scaled limit $V_1 \setminus T$ with the Tian-Yau metric. Similarly for $V_2 \setminus T$.

The end of $V_1 \setminus T$ has the following asymptotic metric model. Let $g_{\rm Euc}$ be the flat Euclidean metric on $(0,\infty) \times T$. Take coordinates s on $(0,\infty)$ and x_1, x_2 on T.
Let $\pi : P \to (0,\infty) \times T$ be a circle bundle with Chern class 8 on T and let α be a connection 1-form on P with curvature a constant multiple of $dx_1 dx_2$.
The metric model (for $s >> 0$ of the end is

$$s\pi^* g_{\rm Euc} + s^{-1}\alpha^2$$

(for large s).

For each fixed s we have one the nilmanifold metrics we considered in Section 4: as s increases the T factor expands while the circle fibres shrink.

There are 16 special regions in the neck (topologically 4-balls), where the model is the "Taub-NUT" metric.

The Taub-NUT metric has an S^1-action with a fixed point. This provides a mechanism which allows matching of the two ends.

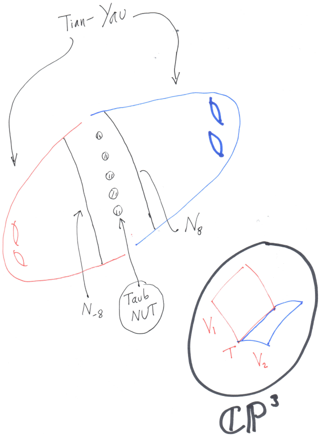
Tian-Yau
N_8
N_-8
Taub NUT
V_1
V_2
T
CP^3