Chapter 6

Limits of Metrics Spaces

A key ingredient in the proof of the Orbifold Theorem is the analysis of limits of metric spaces. In this chapter we give a short account of Gromov's theory of limits of metric spaces as re-interpreted using ϵ-approximations by Thurston. See Gromov's book [32] for a detailed treatment with many interesting applications.

Roughly, there is an "ϵ-approximation" between two metric spaces if the spaces look the same if we ignore details of size ϵ or smaller. From this we define the *Gromov-Hausdorff distance* between two compact metric spaces, and convergence of sequences of metric spaces.

This generalizes the classical notion of *Hausdorff distance* between two subsets A, B of a metric space X:

$$d_H(A, B) = \inf\{\epsilon > 0 : A \subset N(B, \epsilon; X) \text{ and } B \subset N(A, \epsilon; X)\},$$

where

$$N(A, r; X) = \{\ x \in X \ :\ \exists a \in A \ \ d(x, a) < r\ \}$$

denotes the (open) neighbourhood of radius r around A in X. (See [5] for a detailed discussion of the Hausdorff distance and many geometric applications).

The *Gromov-Hausdorff distance* also generalizes the notion of *Lipschitz distance* between homeomorphic metric spaces. A bijection $f : X \longrightarrow Y$ is *K-bilipschitz* if

$$\forall x \neq x' \in X, \quad 1/K \leq d_Y(fx, fx')/d_X(x, x') \leq K.$$

Two metric spaces are close in the Lipschitz sense if there is a $(1 + \epsilon)$-bilipschitz map between them with ϵ small.

Let X be a metric space. Then $A \subset X$ is an *ϵ-net* or *ϵ-dense* if for all $x \in X$ there exists $a \in A$ such that $d(x, a) < \epsilon$.

The basic idea is to approximate a *compact* metric space X by a *finite* ϵ-net A; we want information about a metric space accurate to within ϵ. If X, Y are compact metric spaces, we regard them as *"close"* if there are finite ϵ-nets $A \subset X$ and $B \subset Y$ and a $(1+\epsilon)$-bilipschitz map $f : A \longrightarrow B$.

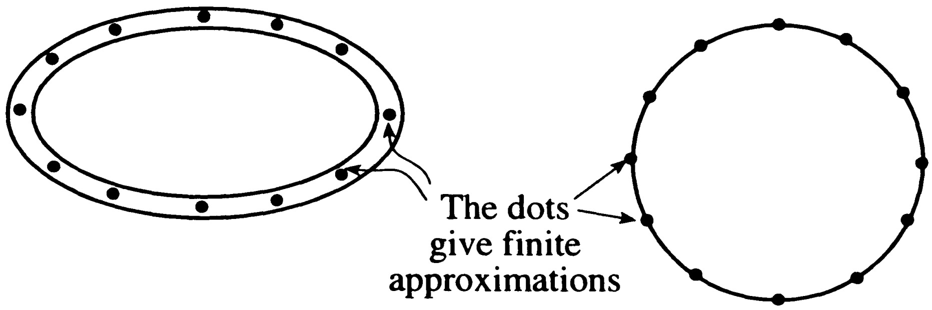

6.1 ϵ-approximations

Definition: An *ϵ-approximation* between metric spaces X and Y is a relation $R \subset X \times Y$ such that

(1) the projections $p_X : R \to X$ and $p_Y : R \to Y$ are both onto,

(2) if xRy and $x'Ry'$ then $|d_X(x, x') - d_Y(y, y')| \leq \epsilon$.

This defines a relation on metric spaces which is symmetric and almost transitive: if X is an ϵ-approximation to Y and Y is an ϵ'-approximation to Z, then X is an $(\epsilon + \epsilon')$-approximation to Z.

We begin with some examples.

Example 6.1. (a) Let A be an ϵ-net for X. Then we can define a 2ϵ-approximation $R \subset A \times X$ by: $aRx \Leftrightarrow d(a, x) < \epsilon$.
(b) A 0-approximation is an isometry.

Example 6.2. Suppose that X, Y are subsets of a metric space Z. Define a relation $R \subset X \times Y$ by xRy if $x \in X$, $y \in Y$ and $d_Z(x, y) \leq \epsilon$. This is a 2ϵ-approximation if $Y \subset N(X, \epsilon; Z)$ and $X \subset N(Y, \epsilon; Z)$. This leads to the Hausdorff metric on closed subsets of Z.

We also say that a sequence $\{X_n\}$ of metric spaces *converges* to Y and write $X_n \longrightarrow Y$ if there are ϵ_n-approximations between X_n and Y with $\epsilon_n \longrightarrow 0$ as $n \longrightarrow \infty$.

Example 6.3.
(a) $T_n = S^1 \times (1/n)S^1$ converges to the circle S^1.

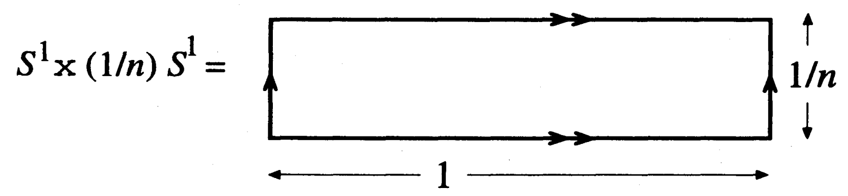

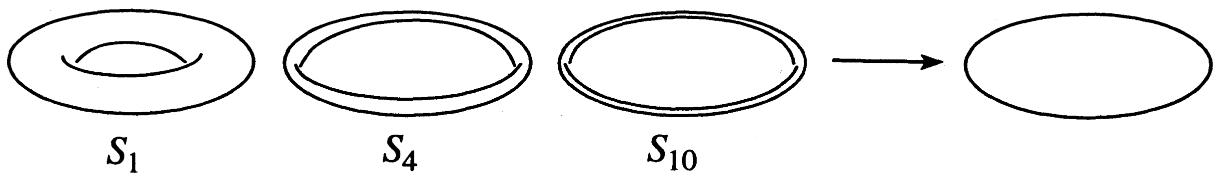

(b) X_n = Euclidean solid torus obtained from a cylinder of height $1/n$ by gluing ends with 180° twist converges to a 2-dimensional Euclidean cone with angle π.

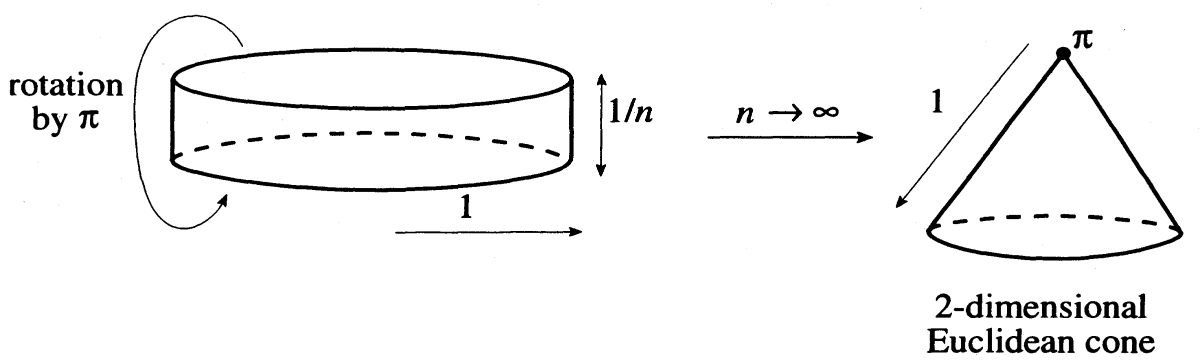

Exercise 6.4. What happens if we modify example (b) by varying the twist angles θ?

Proposition 6.5. *Let X, Y be compact metric spaces. Assume for all $\epsilon > 0$ there is an ϵ-approximation between X and Y. Then X is isometric to Y.*

Proof. Let $R_n \subset X \times Y$ be an $(1/n)$-approximation. There exists a countable dense set $A = \{x^k\} \subset X$. Choose $y_n^k \in Y$ with $x^k R_n y_n^k$. Choose a subsequence of y_n^1 converging to y^1, then a sub-subsequence of y_n^2 converging to y^2, etc. After repeating this process, we define $f : A \longrightarrow Y$ by $f(x^k) = y^k$.

We claim that f is an *isometry* onto $f(A)$ and therefore 1–1. This is because $x^i R_n y_n^i$ & $x^j R_n y_n^j$ implies

$$|d_X(x^i, x^j) - d_Y(y_n^i, y_n^j)| < 1/n.$$

Taking limits gives $d_X(x^i, x^j) = d_Y(y^i, y_n^j)$ which proves the claim. Since f is uniformly continuous and A is dense in X and Y is complete, it follows that there exists a unique continuous extension $f : X \longrightarrow Y$. Further, $f(A)$ is dense in Y because A is dense in X thus the image of A under R_n is $1/n$-dense in Y. Since X is compact it follows that $f(X)$ is closed. Hence f is onto. □

Corollary 6.6. *One can define a metric on the family $\mathcal{F}$ of isometry classes of compact metric spaces, by putting*

$$d(X, Y) = inf\{\epsilon \mid \text{there is an } \epsilon\text{−approximation between } X \text{ and } Y\}.$$

Further, metric spaces with finitely many points are dense.

(Note that diam $< \infty$ for compact spaces implies d is always finite.)

Remarks on related definitions:

For subsets of a fixed metric space, our definition of distance is closely related to the Hausdorff distance on closed subsets. Gromov uses this in [32] to define "Hausdorff convergence" of metric spaces. Gromov's definition of convergence of metric spaces is somewhat stronger than ours. Gromov's distance $d_G(X, Y)$ between metric spaces X and Y is the infimum of Hausdorff distance between $f(X)$ and $g(Y)$ over all isometric embeddings $f : X \to Z$, $g : Y \to Z$ in metric spaces Z. It's clear that $2d_G(X, Y) \geq d(X, Y)$ by example 6.2 above.

Exercise 6.7. Show that $2d_G(X, Y) = d(X, Y)$, as remarked by Bridson and Swarup in [13]. [Hint: Given an ϵ-approximation $R \subset X \times Y$, construct a suitable metric $\tilde{d}$ on the disjoint union $Z = X \cup Y$ which agrees with the given metrics d_X on X, d_Y on Y.]

6.2 Limits with basepoints

For non-compact space we introduce *basepoints*.

Definition: $(X_n, x_n) \to (Y, y)$ converges in the *Gromov-Hausdorff* topology if for all $r, \epsilon > 0$ and for all n sufficiently large, there is an ϵ-approximation R_n between $N(x_n, r; X_n)$ and $N(y, r; Y)$ such that $x_n R_n y$.

The idea is that "*big* neighbourhoods of the basepoint are almost isometric".

Example 6.8. C_n = hyperbolic torus with cone point π/n converges to a complete hyperbolic punctured torus.

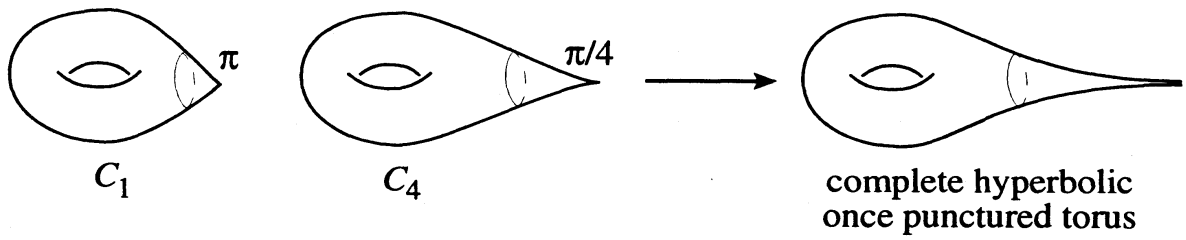

Limits will generally depend on choice of basepoint.

Example 6.9. A sequence of hyperbolic genus 3 surfaces with a long thin neck developing can converge to the three limits shown below for different choices of basepoints.

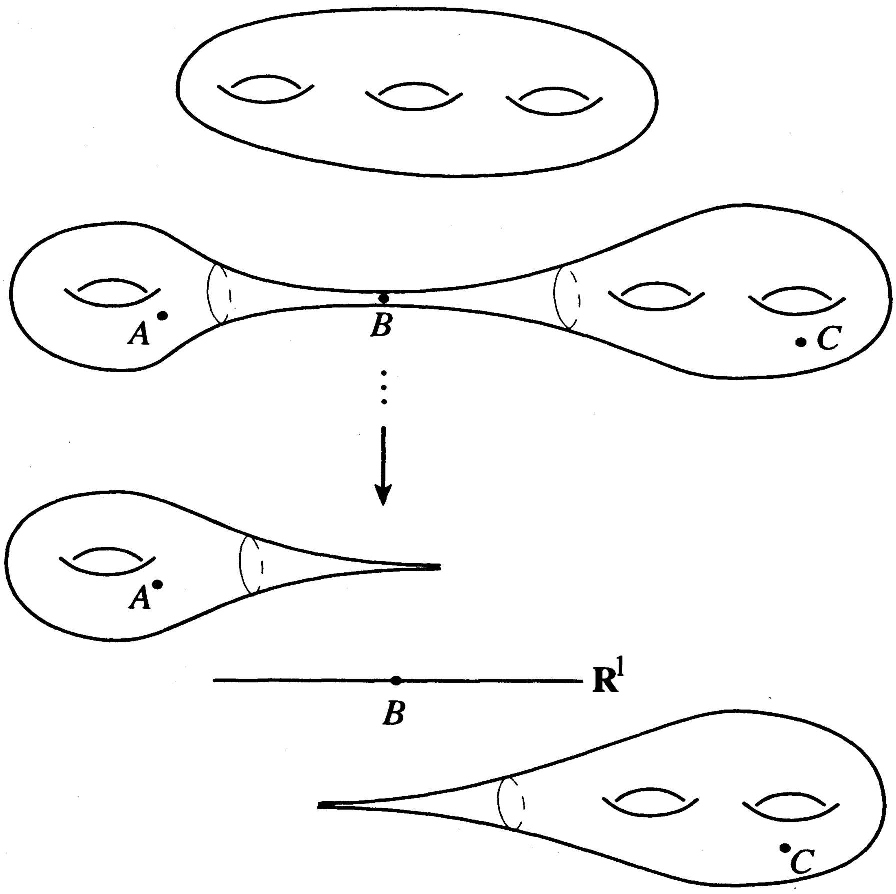

Given an ϵ-approximation $R \subset X \times Y$ if $x_0 R y_0$ we often wish to *restrict* R to an approximation between the r-neighbourhoods $N(x_0, r; X)$

and $N(y_0, r; Y)$. But there is a potential problem here: the projection onto the factors may not be *onto*.

The *smear of R* is the 3ϵ-approximation $R' \subset X \times Y$ given by: $xR'y \Leftrightarrow \exists x' \in X \quad y' \in Y$ such that

$$d(x, x') < \epsilon \text{ and } d(y, y') < \epsilon \text{ and } x'Ry'.$$

If x_0Ry_0 and $r > 0$ then R' restricts to a 3ϵ-approximation between $N(x_0, r; X)$ and $N(y_0, r; Y)$.

With this definition of Gromov-Hausdorff convergence we have: $n^{-1}\mathbb{Z} \to \mathbb{Q}$ and $n^{-1}\mathbb{Z} \to \mathbb{R}$. To get *unique* limits, we need to restrict the metric spaces involved. A metric space is called *proper* if *every closed ball is compact.* Note that every proper metric space is complete.

Corollary 6.10. *Let X, Y be proper metric spaces. If for all $r, \epsilon > 0$ there is an ϵ-approximation between $N(x, r; X)$ and $N(y, r; Y)$ then (X, x) is isometric to (Y, y).*

Corollary 6.11. *Let X_i, Y, Y' be proper metric spaces. If $(X_i, x_i) \to (Y, y)$ and $(X_i, x_i) \to (Y', y')$ then (Y, y) is isometric to (Y', y').*

Remark 6.12.
(a) One can always assume that limits are complete, since the distance between a metric space and its completion is zero.
(b) If closed balls in X_i are compact for all i, and $X_i \to Y$, with Y complete, then all closed balls in Y are also compact. This follows from the fact that a metric space is compact if and only it is complete and totally bounded (i.e. for all $\epsilon > 0$, there is a finite covering by ϵ-balls).
(c) The condition that closed balls are compact fails for infinite dimensional spaces. For example, consider $\mathbb{R}^\infty$ with the supremum metric. Then the ball of radius 1 is non-compact; e.g. the sequence $(1, 0, 0, \ldots), (0, 1, 0, \ldots), \ldots$ doesn't converge.

Example 6.13.
(a) $n^{-1}\mathbb{Z} \to \mathbb{E}^1$
(b) $nS^1 \to \mathbb{E}^1$
(c) $S^1 \times (n^{-1}S^1) \to S^1$.
(d) Let $\mathbb{H}^n(K)$ be "hyperbolic space" of constant curvature $K < 0$, obtained by *rescaling* the metric on $\mathbb{H}^n$. Then $\mathbb{H}^n(K) \to \mathbb{E}^n$ as $K \to 0$.
(e) Let C_n be the cone on n points with a path metric such that each edge

has length 1. Then the sequence $\{C_n\}$ does not converge. (There is no finite approximation to the ball of radius 1 in the "limit".)

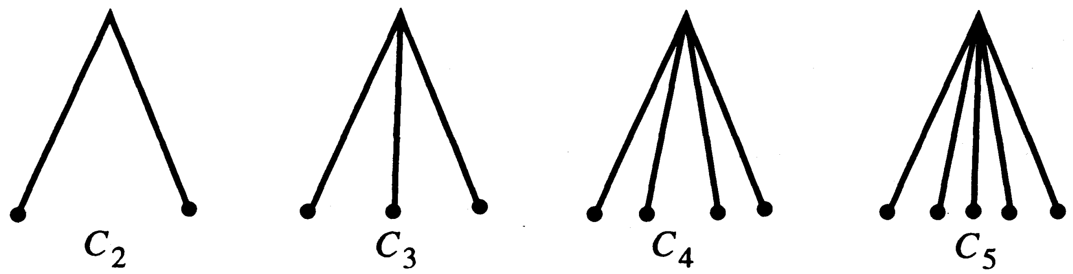

Exercise 6.14. Describe the possible geometric limits of sequences of 2-dimensional Euclidean tori. (Recall: these correspond to isometry classes of lattices in $\mathbb{R}^2$.)

6.3 Gromov's compactness theorem

We now examine the question: When does a sequence of metric spaces have a convergent subsequence?

Definitions: The *ϵ-count*, $\#(\epsilon, X)$, of a metric space X is the minimum number of balls of radius ϵ needed to cover X.

A collection of based metric spaces (X_n, x_n) is *uniformly totally bounded* if for all $\epsilon, r > 0$ there exists $K > 0$ such that $\#(\epsilon, N(x_n, r; X_n)) < K$.

Theorem 6.15 (Gromov's Compactness Theorem). ([32]) *If (X_i, x_i) is a sequence of proper metric spaces then the following are equivalent:*
(1) there is a subsequence $(X_{n_i}, x_{n_i}) \to (Y, y)$ with Y complete.
(2) there is a subsequence (X_{n_j}, x_{n_j}) which is uniformly totally bounded.

Proof. We show that (2) implies (1): Assume the sequence satisfies (2). For each ϵ and each i, we can choose K_ϵ ϵ-balls covering $B_{1/\epsilon}(x_i) \subset X_i$. Let $P_{\epsilon,i}$ be the set of centres of these balls together with the base point x_i. Then each $P_{\epsilon,i}$ is a finite set (containing $\leq K_\epsilon + 1$ points), and is an 3ϵ-approximation to $B_{1/\epsilon}(x_i)$. The metric on $P_{\epsilon,i}$ is described by a distance function

$$d_i : \{1, 2, \ldots, K_\epsilon + 1\}^2 \to [0, 2/\epsilon].$$

By compactness there is a convergent subsequence of d_i; hence there is a subsequence P_{ϵ,i_j} converging to a limiting (finite) metric space L_ϵ (containing $\leq K_\epsilon + 1$ points). (Possibly some points coalesce in the limit, but this won't matter.) Let l_ϵ be the limit of the base points x_{i_j}.

Now choose a collection of $K_{\epsilon/2}$ $\epsilon/2$-balls covering $B_{2/\epsilon}(x_i)$, and let

$$P_{\epsilon/2,i} = \{\text{centres of these } \epsilon/2 \text{ balls}\} \cup P_{\epsilon,i}.$$

We can choose a further subsequence such that $P_{\epsilon/2,i} \rightarrow L_{\epsilon/2}$; then L_ϵ embeds isometrically in $L_{\epsilon/2}$.

Continuing in this way, we obtain

$$L_\epsilon \subset L_{\epsilon/2} \subset L_{\epsilon/4} \subset \dots$$

Let L be the metric completion of $\bigcup_n^\infty L_{\epsilon/2^n}$ and $l \in L$ the limit of the base points. We claim that this is the limit of a (diagonal) subsequence of the X_i.

By our construction, we have a subsequence X_j such that for all $\epsilon > 0$, there is an ϵ-approximation between $B_{1/\epsilon}(x_j)$ and a subset $L_{\epsilon/3}$ of L. Further, each $L_{\epsilon/2^n}$ is an $\epsilon/2^n$-approximation to $L_{\epsilon/2^{n+1}}$, so is an $\epsilon/2^{n-1}$-approximation to L. Hence, $X_j \rightarrow L$. □

Exercise 6.16. Complete the above proof by showing that (1) implies (2).

6.4 Limits of hyperbolic cone-manifolds

We want to know :
(1) when a sequence of 3-dimensional hyperbolic cone-manifolds has a subsequence which converges to a complete metric space Y.
(2) when Y is a 3-dimensional hyperbolic cone-manifold.

Let M be a 3-dimensional hyperbolic cone-manifold; then in particular M is complete. Let D be a Dirichlet domain for M, and $x_0 \in D \subset \mathbb{H}^3$. If M has cone angles $< 2\pi$, then the natural quotient map $q : D \longrightarrow M$ is *onto*. But

$$\#(\epsilon, N(q(x_0), r; M)) \leq \#(\epsilon, N(x_0, r; D)) \leq \#(\epsilon/2, N(x_0, r; \mathbb{H}^3)).$$

Therefore, 3-dimensional hyperbolic cone-manifolds with cone angles in $(0, 2\pi]$ are uniformly totally bounded. Then Gromov compactness implies that (1) always holds.

Note: This is false if cone angles $> 2\pi$ are allowed. One may construct a hyperbolic surface with more and more cone points with cone angles larger than 2π in a small region. This can be done so that the area of this region increases without bound.

We have seen that a sequence of n-manifolds can converge to a space of dimension $< n$. We need a way to rule out this kind of behaviour.

If M is a Riemannian n-manifold and $x \in M$, the *injectivity radius at* x is the radius of the "largest embedded ball" in M with centre x.

Proposition 6.17. *Suppose* (M_k, x_k) *is a sequence of* n*-manifolds with constant curvature* K*. Suppose for all points* $x \in M_k$ *that* $inj(x) > inj_0$*. If* $(M_n, x_n) \to (Y, y_0)$ *then* Y *is an* n*-manifold of curvature* K*.*

Proof. Given $y \in Y$ choose $x_n \in M_n$ with $x_n R_n y$. After smearing we can assume that $N(x_n, inj_0; M_n) \to N(y, inj_0; Y)$. But $N(x_n, inj_0; M_n)$ is isometric to $N(p, r; \mathbb{H}^n(K))$, therefore $N(y, inj_0; Y)$ is isometric to $N(p, r; \mathbb{H}^n(K))$. □

We want to adapt this proof to *cone-manifolds.* This raises the question of what is an appropriate notion of *injectivity radius* in a cone-manifold? (With the usual definition $inj(x) \to 0$ as $x \to \Sigma$.) The role of injectivity radius in the above proof is to provide a *standard neighbourhood* of a certain size. The proof uses that the limit of such neighbourhoods is again such a neighbourhood. In a cone-manifold the local geometry is that of a *cone.* We will see that there is a compact family of such neighbourhoods, and this is what is used to extend the proof above.

A *cone* in a n-dimensional cone-manifold M is a subset isometric to a cone on a spherical $(n-1)$-dimensional cone-manifold, i.e. a "standard cone neighbourhood" as defined in section 3.2.

Definition: $inj(x)$ is the largest r for which $N(x, r; M)$ is contained in a cone:

$$inj(x) = \sup\{\ r > 0 \ :\ \exists\, x' \in M\ \exists r' > 0 \ s.t.\ N(x, r; M) \subset N(x', r'; M) \\ N(x', r'; M) \ = a\ cone\ \}.$$

Note that we *do not* assume that the standard neighbourhood is *centred* at the point x. This is to avoid difficulties near cone points: a point x near the cone locus has only a small standard ball centred at x; however there may be much larger standard *cones* centred at cone points which contain x.

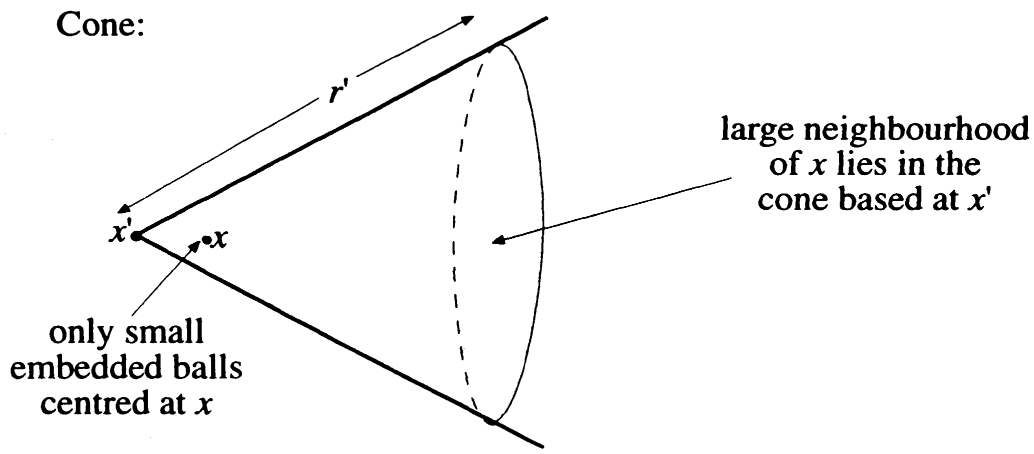

In general, the injectivity radius at a point x will be small if there is a short geodesic loop based at x, or if two *different* pieces of the singular locus are close to x.

Example 6.18. The injectivity radius at some points on an infinite Euclidean pillowcase are shown below. Maximal open balls about the points contained in standard cones are also indicated.

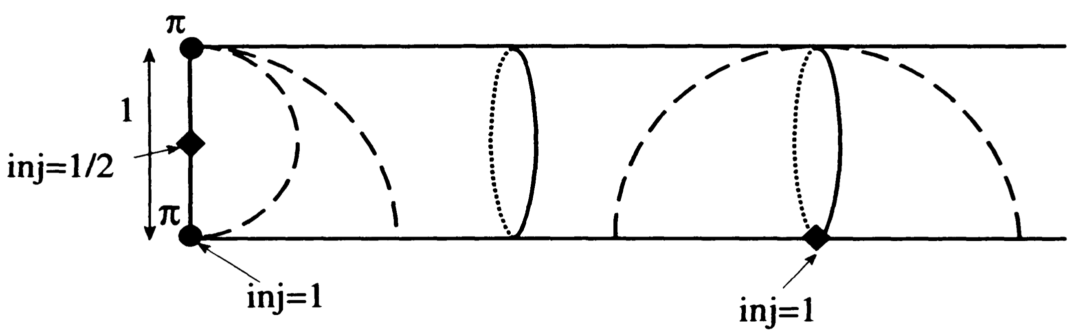

With this definition we obtain:

Proposition 6.19. *Let (M_n, x_n) be a sequence of cone-manifolds, with M_n of constant curvature $\kappa_n \in [-1, 0]$ and $\kappa_n \to \kappa_\infty$. Suppose that there are $\theta_0, inj_0 > 0$ such that all cone angles are in the range $[\theta_0, \pi]$ and $inj(x, M_n) > inj_0$ for all $x \in M_n$ and all n. Then there is a subsequence converging to a cone-manifold (M_∞, x_∞) of curvature κ_∞.*

The essential reason is that the cones of fixed radius r form a *compact* set of metric spaces, namely the cones on: $S^2(\alpha, \beta, \gamma)$, $S^2(\alpha, \alpha)$ or S^2 with $\alpha, \beta, \gamma \in [\theta_0, \pi]$.

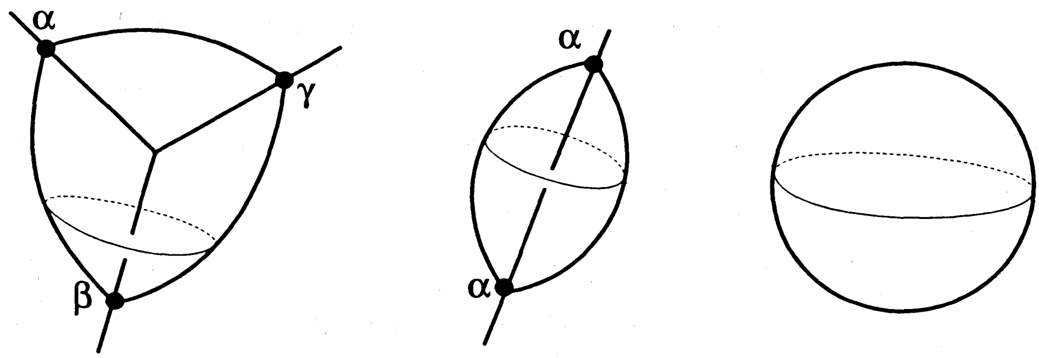

6.5 Bilipschitz convergence

Theorem 6.20. *Suppose compact hyperbolic* 3*-manifolds* M_n *converge in the Gromov-Hausdorff topology to a* compact *hyperbolic* 3*-manifold* M_∞. *Then for all* $\epsilon > 0$ *and for all sufficiently large n there is a* $(1+\epsilon)$-bilipschitz *map* $f : M_\infty \longrightarrow M_n$.

Proof. (Sketch) There exists $inj_0 > 0$ such that the injectivity radius of every point in every M_n and in M_∞ is bigger than inj_0. (There is a lower bound on injectivity radius in M_∞ because it is compact. Since $M_n \to M_\infty$, for large n the injectivity radius in M_n is nearly equal to that in the limit.) Let K be a *geodesic triangulation* of M_∞, such that every simplex $\sigma \in K$ is *small* compared to inj_0. Let $V = \{v_1, \cdots, v_k\}$ be the vertices of K. Let R_n be a $1/n$-approximation between M_n and M_∞. Choose $v_i^n \in M_n$ with $v_i^n R_n v_i$. If $v_a, v_b, v_c, v_d \in V$ span a 3-simplex $\sigma \in K$ then $v_a^n, v_b^n, v_c^n, v_d^n \in M_n$ span a *small geodesic simplex* in a standard metric ball. Thus we get a geodesic triangulation K_n of M_n combinatorially the same as K. Also, the corresponding edge lengths are nearly equal. Use ϵ_n-approximations with $\epsilon_n << \epsilon \cdot$ (edge lengths of K) to map the vertices into M_n. This extends to a simplicial map $f : K \longrightarrow K_n \subset M_n$ which is $(1+\epsilon)$-bilipschitz where $\epsilon \to 0$ as $n \to \infty$. □

Extensions:

(1) The previous result extends easily to cone-manifolds: this just requires some extra care constructing a *thin* geodesic triangulation near Σ. Note that bilipschitz convergence also implies convergence of cone angles, volume, etc.

(2) To extend the result to the case of a non-compact limit M_∞ requires estimates on the *decay of injectivity radius* described in theorem 7.8. In this case we get bilipschitz maps from any *compact subset* of M_∞ into the M_n.

(3) The result also extends easily to a sequence of cone manifolds M_n, where each has constant curvature K_n lying in the interval $[-1, 0]$.

Theorem 6.21. *Suppose that M_n is a sequence of complete hyperbolic cone 3-manifolds. Suppose that (M_n, x_n) converges in the Gromov-Hausdorff topology to a complete hyperbolic cone 3-manifold (M_∞, x_∞). Then given $\epsilon > 0$ and $R > 0$, for all sufficiently large n there is a $(1+\epsilon)$-bilipschitz map $f : N(x_\infty, R, M_\infty) \longrightarrow M_n$ with $d(f(x_\infty), x_n) < \epsilon$. Furthermore f maps singular set to singular set.*

6.6 Convergence of holonomy

If a sequence (M_n, x_n) of 3-dimensional hyperbolic cone-manifolds converges to a hyperbolic cone-manifold (M_∞, x_∞), we want convergence of the *holonomy* representations

$$h_n : \pi_1(M_n - \Sigma(M_n)) \longrightarrow PSL(2, \mathbb{C}).$$

Theorem 6.22. *Assume that for all large n there are $(1+\epsilon_n)$-bilipschitz homeomorphisms*

$$\phi_n : (M_\infty, \Sigma(M_\infty)) \longrightarrow (M_n, \Sigma(M_n))$$

with $\epsilon_n \to 0$ as $n \to \infty$. Let

$$\phi_n* : \pi_1(M_\infty - \Sigma(M_\infty)) \longrightarrow \pi_1(M_n - \Sigma(M_n))$$

be the induced homomorphism of fundamental groups, and assume that $\pi_1(M_\infty - \Sigma(M_\infty))$ is finitely generated. Then $h_n \circ \phi_{n*} \to h_\infty$ in the algebraic topology. *This means there are $A_n \in PSL(2, \mathbb{C})$ such that $A_n h_n(\phi_{n*}\alpha) A_n^{-1} \to h_\infty(\alpha)$ for all $\alpha \in \pi_1(M_\infty - \Sigma(M_\infty))$.*

Idea of Proof: Write $X_n = M_n - \Sigma_n$. Then the bilipschitz convergence implies that the developing maps $\mathrm{dev}_n : \tilde{X}_n \to \mathbb{H}^3$ can be adjusted by isometries g_n so that

$$g_n \circ \mathrm{dev}_n \circ \widetilde{\phi_n} \to \mathrm{dev}_\infty : \tilde{X}_\infty \to \mathbb{H}^3$$

uniformly on compact subsets. Applying this to a large compact subset of $\tilde{X}_\infty$ containing lifts of loops generating $\pi_1(X_\infty)$ gives the result. □