Gromov-Hausdorff convergence and homotopy type

Janko Latschev*

November 3, 1999

Abstract

We show that for every closed Riemannian manifold X there exists a positive number $\varepsilon_0 > 0$ such that for all $0 < \varepsilon \leq \varepsilon_0$ there exists some $\delta > 0$ such that for every metric space Y with Gromov-Hausdorff distance to X less than δ the geometric ε-complex $|Y_\varepsilon|$ is homotopy equivalent to X. In particular, this gives a positive answer to a question of Hausmann [Hau95].

1 Introduction

This note grew out of an attempt to understand which properties of a given Riemannian manifold X are reflected in metric spaces Y Gromov-Hausdorff close to it. Even thought the topology of Y will generally be quite complicated, we will prove that by a very simple procedure one can recover the homotopy type of X from it, provided it is close enough in a sense made explicit below.

Given a pseudo-metric space (X, d) and a positive number $\varepsilon > 0$, one can define an abstract simplicial complex X_ε whose vertices are the points of X and whose k-simplices consist of all (k+1)-tuples of points in X with pairwise distances less than ε. Apparently this complex was first introduced by Vietoris in the thirties in an attempt to define a homology theory for general metric spaces. It has more recently found applications in the theory of hyperbolic groups.

It was proved by Hausmann [Hau95] that for a closed Riemannian manifold X and ε sufficiently small the geometric realization $|X_\varepsilon|$ of this complex is homotopy

*Institut für Mathematik, Universität Zürich, Winterthurerstrasse 190, CH-8057 Zürich, Switzerland; *currently at:* Max-Planck-Institut für Mathematik, Bonn
email: `latschev@mpim-bonn.mpg.de`

equivalent to X. In fact, for $\varepsilon > 0$ sufficiently small he constructed an explicit homotopy equivalence $T : |X_\varepsilon| \longrightarrow X$ as follows. One chooses a total order on the points of X and defines the map T by induction on the dimension of the simplices, starting with the identity on the 0-skeleton. Assuming T is defined on all (k-1)-simplices, we define it on a k-simplex $\{x_0, ..., x_k\}$ with $x_0 < x_1 < ... < x_k$ by mapping the straight line between x_k and a point $p \in \{x_0, ..., x_{k-1}\}$ in the front face to the geodesic between the (already known) image points $T(x_k)$ and $T(p)$. This works for example when ε is smaller than the convexity radius of X. Because of the total order, these maps are compatible with each other and so T is continuous.

Hausmann then raised the question whether there would always be a finite subset $F \subset X$ and an $\varepsilon > 0$ such that the finite simplicial complex $|F_\varepsilon|$ is homotopy equivalent to X under the restriction of the map T. In this note, we give a positive answer to that question by proving the stronger

Theorem 1.1 *Let X be a closed Riemannian manifold. Then there exists $\varepsilon_0 > 0$ such that for every $0 < \varepsilon \leq \varepsilon_0$ there exists a $\delta > 0$ such that the geometric realization $|Y_\varepsilon|$ of the ε-complex of any metric space Y which has Gromov-Hausdorff distance less than δ to X is homotopy equivalent to X.*

When $Y \subset X$ is a subset, a homotopy equivalence is given by the map T as described above, and in general it is similar in spirit.

It was already observed by Hausmann that one can easily find $Y \subset X$ and $\varepsilon > 0$ such that the map $T : |Y_\varepsilon| \longrightarrow X$ is a domination, i.e. it admits a right homotopy inverse. The main difficulty in showing that this map can be chosen to be a genuine homotopy inverse is to show that $|Y_\varepsilon|$ is sufficiently locally contractible. The concept of ε-contraction as defined in the next section turns out to be central to our argument.

While the proof we present here is qualitative, it will be clear that it is possible to extract reasonable estimates for ε_0 and $\delta = \delta(\varepsilon_0)$ which depend only on the geometry of X. This observation should be useful in extending the result to apply to more general length spaces. We hope to return to this question elsewhere.

2 The flat case

In this section, we will prove the assertion in the very special case when X is a closed *flat* manifold and $F \subset X$ is a *finite subset*. In the next section we will repeat the argument along the same lines, but dealing with various technical issues arising in the more general context.

ε-contraction

ε-contraction is a variation of collapse as introduced by Hausmann [Hau95]. Let Y be a metric space and let $\varepsilon > 0$ be fixed. Suppose we are given a subset $Y' \subset Y$

and a point $y' \in Y \setminus Y'$ satisfying

$$\text{(C)} \qquad d(x,z) < \varepsilon \text{ implies } d(y',z) < \varepsilon \text{ for any } x \in Y' \text{ and } z \in Y.$$

Then the map $c : Y \longrightarrow Y \setminus Y'$ mapping Y' to the point $y' \in Y$ (and leaving all other points fixed) is called an *ε-contraction.* By (C), such a map induces a simplicial map on ε-complexes for Y and $Y \setminus Y'$, and using (C) with $x = z$ one sees that this map is contiguous to the identity of Y_ε. Recall that two maps $f, g : Y_\varepsilon \longrightarrow Y_\varepsilon$ are called contiguous if for every simplex σ the images $f(\sigma)$ and $g(\sigma)$ again span a simplex.

In particular $|(Y \setminus Y')_\varepsilon|$ is a strong deformation retract of $|Y_\varepsilon|$ (cf. [Spa66], p. 130). A metric space Y will be called *ε-contractible* if it can be mapped to a space of diameter less than ε through a sequence of ε-contractions. Clearly in this case $|Y_\varepsilon|$ is contractible in the usual sense.

In this section we will only deal with elementary ε-contractions, where Y' consists of only one point. The first key observation can now be stated as follows:

Lemma 2.1 *Let $B(0,r) \subset \mathbb{R}^n$ be the standard open ball of radius r, and let $0 < \varepsilon < r$ be given. Then there exists $\delta_1 > 0$ such that every finite δ_1-dense subset $F \subset B(0,r)$ is ε-contractible.*

Proof: Let any subset $F \subset B(0,r)$ be given. Let $p \in F$ have maximal distance $r' = d(p,0)$ to the origin, and consider the set

$$I_p = \bigcap_{x \in \overline{B(p,\varepsilon)} \cap \overline{B(0,r')}} \overline{B(x,\varepsilon)}.$$

If we want to move p by an elementary ε-contraction, than the points of I_p are precisely the allowed image points. We claim that I_p contains an open ball of definite radius, which we are trying to compute. It is easy to see that the only constraints on I_p come from points x in the intersection $S(p,\varepsilon) \cap S(0,r')$ of the ε-sphere around p with the r'-sphere around 0.

So let $x \in S(p,\varepsilon) \cap S(0,r')$ be given. Then the distance of x to the line $\overline{0p}$ is given by

$$h = \varepsilon \cdot \sqrt{1 - \frac{\varepsilon^2}{4r'^2}}.$$

The point $p' \in \overline{0p}$ closest to x has distance

$$\eta = \eta(\varepsilon, r) = \frac{\varepsilon^2}{2r'} \tag{1}$$

from p and distance

$$\delta_1 = \delta_1(\varepsilon, r') = \varepsilon \cdot \left(1 - \sqrt{1 - \frac{\varepsilon^2}{4r'^2}}\right) \tag{2}$$

from the boundary of $B(x,\varepsilon)$. Notice that for $\varepsilon \leq 2r'$ we have

$$\delta_1 \leq \varepsilon(1-(1-\frac{\varepsilon^2}{4r'^2})) \leq \eta. \tag{3}$$

Combining these results, we see that the ball $B(p',\delta_1)$ is completely contained in I_p and does not contain p.

Notice that both η and δ_1 increase with decreasing r'. Hence if F is δ_1-dense in $B(0,r)$ for $\delta_1 = \delta_1(\varepsilon, r)$, and $p' \in B(0,r)$ is constructed from some point $p \in F$ of maximal distance to 0 as above, then there exists a point $q \in F \cap B(p',\delta_1)$. Now by the definition of I_p the map $c_{pq} : F \longrightarrow F$ is an elementary ε-contraction. We can apply this same argument to $F' = F \setminus \{p\}$, and, continuing inductively, we will eventually move through a sequence of elementary ε-contractions to a subset $\widetilde{F} \subset F$ of diameter less than ε. Thus F is ε-contractible as claimed. $\square$

Remark 2.2 *Note that the same argument applies more generally to every finite set F δ_1-dense in some convex subset $D \subset B(0,r)$ containing 0.*

To be able to apply this remark, the following technical lemma is useful.

Lemma 2.3 *Let $B_1, ..., B_N$ be a finite cover of some Riemannian manifold X by open balls, and let δ_1 be given. Then there exists a number δ_2 such that every subset $F \subset X$ which is δ_2-dense in X has the property that $F \cap (B_{i_1} \cap ... \cap B_{i_k})$ is δ_1-dense in every non-empty intersection $B_{i_1} \cap ... \cap B_{i_k}$.*

Proof: We will argue by contradiction. So suppose such a δ_2 does not exist. Then for every $\delta_2 > 0$ there exists a set F_{δ_2} which is δ_2-dense in X and there exists some non-empty intersection $B_{i_1} \cap ... \cap B_{i_k}$such that $F \cap (B_{i_1} \cap ... \cap B_{i_k})$ is not δ_1-dense in $B_{i_1} \cap ... \cap B_{i_k}$.

In particular, we can find a sequence $\delta_{2,k} \stackrel{k\to\infty}{\longrightarrow} 0$ such that the "offending" intersection $B_{i_1} \cap ... \cap B_{i_k}$ is fixed. Thus we get a sequence of points $p_k \in B_{i_1} \cap ... \cap B_{i_k}$ such that no point of $F_k \cap (B_{i_1} \cap ... \cap B_{i_k})$ is closer than δ_1 to p_k. Since X is compact, the p_k have some subsequence converging to a point $p \in X$.

Now $B(p, \frac{\delta_1}{2}) \cap (B_{i_1} \cap ... \cap B_{i_k})$ is open and non-empty, since it contains infinitely many points of the sequence $\{p_k\}$. Therefore it contains a ball of definite size $B(q,\rho)$. For $\delta_{2,k} < \rho$ there has to be some point $x_k \in F_k$ in this ball. But then

$$d(x_k, p_k) \leq d(x_k, p) + d(p, p_k) < \frac{\delta_1}{2} + d(p,p_k)$$

thus contradicting the assumption on p_k as soon as $d(p,p_k) < \frac{\delta_1}{2}$. $\square$

We are now ready to answer Hausmann's question in the flat case.

Proposition 2.4 *Let X be a flat closed manifold. Then there exists $\varepsilon_0 > 0$ such that for every $0 < \varepsilon \leq \varepsilon_0$ there is some $\delta > 0$ such that the ε-complex of any finite δ-dense subset $F \subset X$ is homotopy equivalent to X under the map T.*

Proof: Consider a finite cover $B(a_1, r) \cup B(a_2, r) \cup ... \cup B(a_N, r)$ of X by balls of radius $r < conv(X)$, and define ε_0 to be the Lebesgue number of this cover. Here $conv(X)$ denotes the convexity radius of X. By making r (and hence ε) smaller if necessary, we may assume that the distance of any two disjoint balls is strictly positive. For any $\varepsilon \leq \varepsilon_0$, let $\delta_1 = \delta_1(\varepsilon, r)$ and δ_2 be constructed as in the previous two lemmas. By making them smaller if necessary, we may assume that $\delta_2 \leq \delta_1 < \frac{\varepsilon}{2}$. We claim that $\delta = \delta_2$ satisfies the conclusion of the proposition.

Namely, consider any finite δ_2-dense set $F \subset X$ and choose a total order on its points. As described in the introduction, the map $T : |F_\varepsilon| \longrightarrow X$ is given by mapping each simplex to the geodesic suspension of its front face. Notice also that since the balls $B(a_i, r)$ are strictly convex, T maps simplices of $|(F \cap B(a_i, r))_\varepsilon|$ into $B(a_i, r)$.

To define a map $S : X \longrightarrow |F_\varepsilon|$, we use an idea found in [Gro96], §1.4.D'. Consider a continuous partition of unity $\{\rho_x\}_{x \in F}$ subordinate to the cover of X by δ_1-balls centered at points of F, where we require that $p \in B(a_i, r)$ and $\rho_x(p) > 0$ implies $x \in B(a_i, r)$. For example, if $B(a_1, r), ..., B(a_k, r)$ are the balls in the cover containing $x \in F$, we can define

$$\rho'_x(p) := d(p, \bigcup_{a_j \notin \{a_1, ..., a_k\}} B(a_j, r)) \cdot d(p, X \setminus B(x, \delta_1)).$$

By Lemma 2.3, our choice of δ_2 guarantees that for every $p \in X$ there is some $x \in F$ with $\rho'_x(p) > 0$. As usual, ρ_x is now obtained by normalizing.

Next, we use the total order on F to define a bijective map from F to the vertices of the standard $(N-1)$-simplex $\Delta = \{t \in \mathbb{R}^N \mid \sum t_i = 1\} \in \mathbb{R}^N$, where N is the number of points in F. We can identify $|F_\varepsilon|$ with the appropriate subcomplex of this simplex. Now define $S : X \longrightarrow |F_\varepsilon| \subset \mathbb{R}^N$ by

$$S(p) = (\rho_{x_1}(p), \rho_{x_2}(p), ..., \rho_{x_N}(p)).$$

The condition $\delta_1 < \frac{\varepsilon}{2}$ guarantees that the image of this map is indeed contained in $|F_\varepsilon|$. Notice also that by construction $S(B(a_i, r)) \subset |(F \cap B(a_i, r))_\varepsilon|$.

Now again by construction $T \circ S : X \longrightarrow X$ maps each of the balls $B(a_i, r)$ into itself, so that this map is homotopic to the identity of X by a straight line homotopy. In the same way, $S \circ T : |F_\varepsilon| \longrightarrow |F_\varepsilon|$ maps each of the sets $|(F \cap B(a_i, r))_\varepsilon|$ into itself. Now by an induction on cells argument using Lemma 2.1 and Remark 2.2 we see that this map is homotopic to the identity of $|F_\varepsilon|$, and the proposition is established. $\square$

3 The general case

We will start the treatment of the general case by presenting a stronger version of Lemma 2.1.

Lemma 3.1 *Let $0 < \varepsilon \leq \frac{r}{2}$ be given. Then there exists a number $\widetilde{\delta_1} > 0$ such that every metric space (Y, d) whose Gromov-Hausdorff distance to some convex subset $D \subset B(0, r) \subset \mathbb{R}^n$ of the standard flat ball of radius r is less than $\widetilde{\delta_1}$ is ε-contractible.*

Here $\widetilde{\delta_1} = \widetilde{\delta_1}(\varepsilon, r)$ can be chosen as $\delta_1(\varepsilon, r)/5$ with $\delta_1(\varepsilon, r)$ given by equation (2).

Proof: Let $\widetilde{\delta_1}$ be defined as stated, and let (Y, d) be a given metric space $\widetilde{\delta_1}$-close to the convex subset D of the standard ball $B(0, r) \subset \mathbb{R}^n$. Without loss of generality D contains 0.

Let $\widetilde{d}$ be some pseudometric on $Y \sqcup D$ extending d on Y and d_{eucl} on D such that Y is contained in the $\widetilde{\delta_1}$-neighborhood of D and vice versa. The existence of this pseudometric is equivalent to our assumption that D and Y are $\widetilde{\delta_1}$-close ([Gro99], ch. 3).

Fix a finite subset $F \subset D$ which is $\widetilde{\delta_1}$-dense in D. Then the $\widetilde{d}$-balls $\widetilde{B}$ of radius $2\widetilde{\delta_1}$ around points of F will cover all of $Y \sqcup D$.

Now pick a point $y \in F$ of maximal distance to 0, and construct $y' \in D$ as in the proof of Lemma 2.1. Then there is a point $\widetilde{y'} \in Y$ with $\widetilde{d}(\widetilde{y'}, y') < \widetilde{\delta_1}$.

We want to show that mapping $Y' = \widetilde{B}(y, 2\widetilde{\delta_1}) \cap Y$ to $\widetilde{y'} \in Y$ is an ε-contraction. So fix any point $\widetilde{y} \in Y'$.

To verify condition (C), suppose that $\widetilde{x} \in Y$ is a point with $d(\widetilde{y}, \widetilde{x}) < \varepsilon$. Then there exists $x \in D$ with $\widetilde{d}(\widetilde{x}, x) < \widetilde{\delta_1}$ and therefore

$$d_{\text{eucl}}(y, x) \leq \widetilde{d}(y, \widetilde{y}) + d(\widetilde{y}, \widetilde{x}) + \widetilde{d}(\widetilde{x}, x) < \varepsilon + 3\widetilde{\delta_1}.$$

So there has to be some point $z \in \overline{yx} \subset D$ such that $d_{\text{eucl}}(y, z) < \varepsilon$ and $d_{\text{eucl}}(x, z) < 3\widetilde{\delta_1}$. From the proof of Lemma 2.1 we also know that

$$d_{\text{eucl}}(z, y') \leq \varepsilon - \delta_1(\varepsilon, r),$$

where $\delta_1(\varepsilon, r)$ is given by equation (2). Therefore

$$\begin{aligned} d(\widetilde{x}, \widetilde{y'}) &\leq \widetilde{d}(\widetilde{x}, x) + d_{\text{eucl}}(x, z) + d_{\text{eucl}}(z, y') + \widetilde{d}(y', \widetilde{y'}) \\ &< \varepsilon - \delta_1(\varepsilon, r) + 5\widetilde{\delta_1}. \end{aligned}$$

Hence if $\widetilde{\delta_1} \leq \delta_1/5$, then condition (C) holds. Note also that

$$\begin{aligned} d(\widetilde{y'}, 0) &\leq \widetilde{d}(\widetilde{y'}, y') + d_{\text{eucl}}(y', 0) \\ &< \widetilde{\delta_1} + (d_{\text{eucl}}(y, 0) - \eta(\varepsilon, r)) \\ &< d(\widetilde{y}, 0) - \eta(\varepsilon, r) + 3\widetilde{\delta_1} \\ &< d(\widetilde{y}, 0), \end{aligned}$$

which shows that $\widetilde{y'} \notin Y' = \widetilde{B}(y, 2\widetilde{\delta_1}) \cap Y$ as required. The procedure can now be iterated just as in the proof of Lemma 2.1, and since everything is controlled by the finite set F we are assured that it will terminate after a finite number of steps. $\Box$

Now let some closed Riemannian manifold X be given. Since the assertion of Theorem 1.1 is invariant under scaling, it is sufficient to prove it for some rescaled $\hat{X}$. By multiplying the metric on X by a large constant, we can achieve that any subset $D \subset B(p,1)$ of any ball $B(p,1)$ of radius 1 in the rescaled space $\hat{X}$ has distance less than $\widetilde{\delta_1}(\frac{1}{2},1)/2$ to the subset $D' \subset T_p\hat{X}$ in the flat ball of radius 1 corresponding to it under the exponential map at p.

We can cover $\hat{X}$ by finitely many balls $B(a_1,1),...,B(a_N,1)$ of radius 1 such that the centers are $\frac{1}{2}$-dense in $\hat{X}$. It follows that the Lebesgue number ε_0 of this cover is at least $\frac{1}{2}$.

For a metric space Y which is δ-close to $\hat{X}$ for some $\delta > 0$, again consider some metric $\widetilde{d}$ on the disjoint union $Y \sqcup \hat{X}$ as in the proof of Lemma 3.1 such that Y is contained in the δ-neighborhood of $\hat{X}$ and vice versa. It is easy to see that if δ is smaller than the Lebesgue number of the cover $\bigcup B(a_i,1)$ of $\hat{X}$, then the balls $\widetilde{B}(a_i,1)$ of radius 1 in the metric $\widetilde{d}$ on $Y \sqcup \hat{X}$ will cover all of $Y \sqcup \hat{X}$.

The analogue of Lemma 2.3 in the present setting is

Lemma 3.2 *Given $\hat{X} = \bigcup B(a_i,1)$ as above and $0 < \delta < \frac{1}{2}$, there exists $\widetilde{\delta_2} > 0$ such that every metric space Y which is δ_2-close to $\hat{X}$ has the property that $Y \cap (\widetilde{B}(a_{i_1},1) \cap ... \cap \widetilde{B}(a_{i_k},1))$ is δ-close to $B(a_{i_1},1) \cap ... \cap B(a_{i_k},1)$ whenever the later is non-empty.*

The proof is essentially the same as for Lemma 2.3 and will be ommitted.

To prove Theorem 1.1 for $\hat{X}$ with $\varepsilon_0 = \frac{1}{2}$, we proceed just as in the proof of Proposition 2.4. Given any $\varepsilon < \varepsilon_0$, we first find $\widetilde{\delta_1} = \widetilde{\delta_1}(\varepsilon, 1)$ as in Lemma 3.1 and then $\widetilde{\delta_2} = \widetilde{\delta_2}(\frac{\widetilde{\delta_1}}{2})$ as in Lemma 3.2.

Given a metric space Y which is $\frac{\widetilde{\delta_2}}{4}$-close to $\hat{X}$, we can find a finite subset $F \subset Y$ which is $\frac{\widetilde{\delta_2}}{4}$-close to $\hat{X}$ by covering $\hat{X}$ with finitely many balls of radius $\frac{\widetilde{\delta_2}}{4}$ and then picking one point from each of the corresponding balls in the metric $\widetilde{d}$ on $Y \sqcup \hat{X}$.

Notice also that the sets $Y \cap (\widetilde{B}(a_{i_1},1) \cap ... \cap \widetilde{B}(a_{i_k},1))$ are ε-contractible by Lemma 3.1 and the construction of $\hat{X}$ and $\widetilde{\delta_2}$. We now define the map $T : |Y_\varepsilon| \longrightarrow \hat{X}$ as follows. Each point $y \in Y$ is contained in a certain maximal intersection $\bigcap \widetilde{B}(a_i,1)$ of $Y \sqcup \hat{X}$. Now simply map y to a closest point $x \in \bigcap B(a_i,1) \subset X$, then proceed by induction on cells as before. To define the inverse map $S : \hat{X} \longrightarrow |Y_\varepsilon|$, use a partition of unity on $\hat{X}$ subordinate to the cover by balls $\widetilde{B}(y, \widetilde{\delta_1})$ around the points of $F \subset Y$ as in the proof of Proposition 2.4. Now the rest of the argument goes through just as in that case, using the Lemmas 3.1 and 3.2 in place of Lemmas 2.1 and 2.3.

Acknowlegdments

I would like to thank Th. Mautsch for pointing out [Gro96] to me and F. Schlenk for critical comments on an earlier version of this paper. I also wish to thank the Max-Planck-Institut für Mathematik in Bonn, where this work was completed, for its hospitality and support.

References

[Gro96] M. Gromov, *Carnot-Carathéodory spaces seen from within*, in: Sub-Riemannian Geometry (A. Belaiche and J.-J. Risler, eds.), Progress in Mathematics, vol. 144, Birkhäuser, 1996, pp. 85–324.

[Gro99] M. Gromov, *Metric structures for Riemannian and non-Riemannian spaces*, Progr. in Math. 152, Birkhäuser, 1999.

[Hau95] J.-C. Hausmann, *On the Vietoris-Rips Complexes and a Cohomology Theory for Metric Spaces*, in: Prospects in Topology: Proceedings of a conference in honour of William Browder, Princeton University Press, 1995, pp. 175–188.

[Spa66] E. Spanier, *Algebraic topology*, McGraw-Hill, 1966.